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 debt expands, an additional claim adds to rollover needs and raises the refinancing premium paid as the existing stock comes due. The resulting stock-wide marginal cost can overtake the finite-horizon wealth effect of further issuance, generating an aggregate-demand safe-debt Laffer curve. The equilibrium safe rate reaches its maximum at the same point. Empirically, I combine cash-Treasury and real-duration premia with estimates of how additional private supply reprices the outstanding Treasury stock. Over the past decade, the benchmark estimate of the marginal cost of producing safe Treasury claims more than doubled, rising from about 80 basis points in 2015Q1 to about 187 basis points in 2026Q1. This increase reduced the safe-debt margin—the rate-equivalent distance to the Laffer peak—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. With finite-horizon households, issuance raises financial wealth because current households do not fully capitalize the future taxes needed to service the debt. I call this gap the non-Ricardian wedge. It is positive at every debt level, so the usual model makes additional safe debt expansionary everywhere.[1] A chronic demand for stores of value, as emphasized in the safe-asset-shortage literature, only strengthens that 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 assessment required to mobilize support. Participation takes a threshold form. The marginal participant prices the claim by averaging these state-contingent assessments. Under the flat-prior benchmark, its zero-profit premium does not depend on the equilibrium participation cutoff. 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.

This marginal production cost enters aggregate demand through taxes. Perpetual-youth households count the government debt they hold as financial wealth but capitalize only the taxes they expect to pay while alive. Debt service and the cost of producing safe-asset services reduce their disposable income and human wealth. Finite horizons prevent this future tax burden from fully offsetting the value of the newly issued claim; that is the source of the non-Ricardian wedge. The safe-debt margin is the resulting demand benefit of another unit of debt minus its stock-wide marginal production cost. Debt raises demand while this margin is positive and lowers it after the margin turns negative. The crossing is an aggregate-demand peak reached while the sovereign remains solvent and its debt remains safe.

The equilibrium safe rate turns at the same point. Below the peak, issuance creates excess demand at the prevailing rate. The safe rate rises, lowering human wealth until the goods market clears again. Above the peak, the fiscal resources absorbed by safe-debt production exceed the wealth created by the marginal claim. Issuance then lowers demand at the prevailing rate, and the safe rate falls to restore equilibrium. The claims remain safe on this declining branch; the rate response reflects the rising cost of producing safety and the adjustment required to clear the goods market.

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 consequently pushes the Laffer peak to a larger debt stock and raises the equilibrium safe rate at any given stock in the stress region.

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.

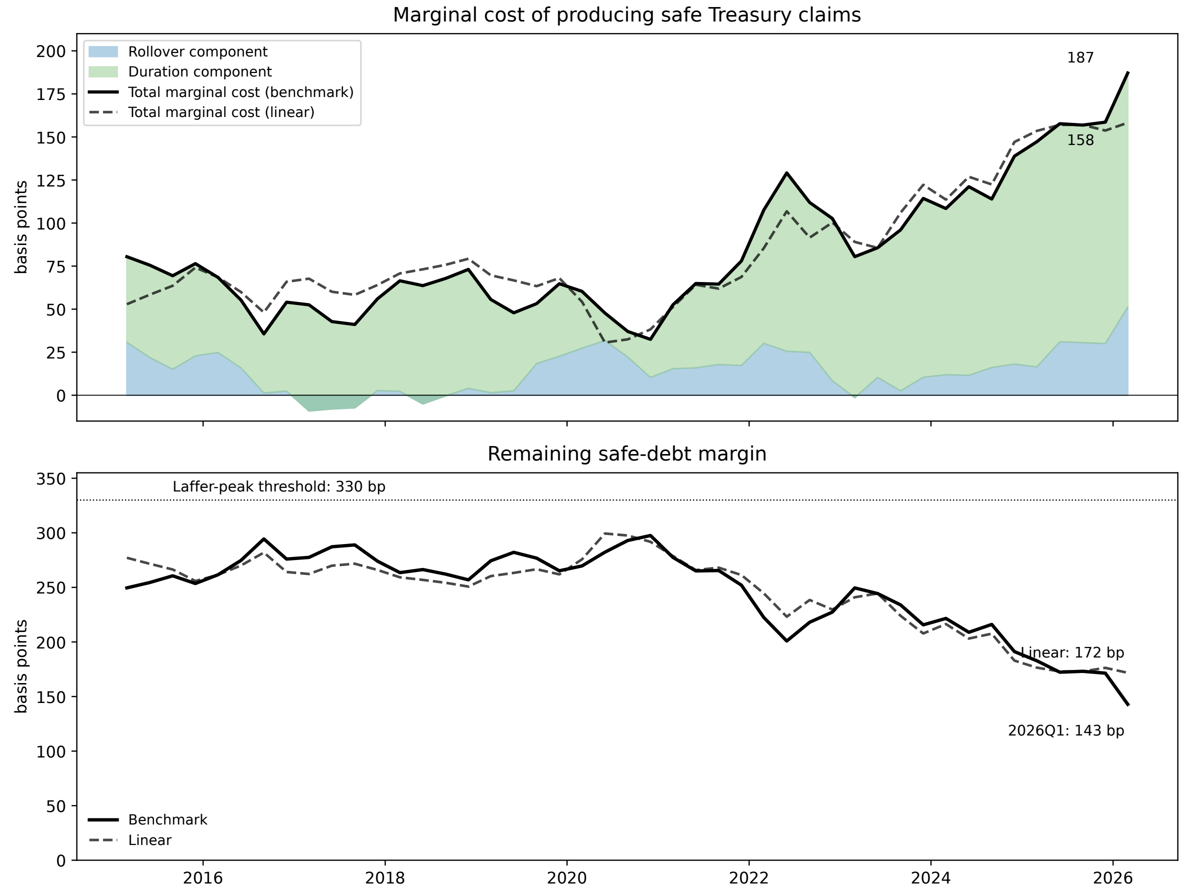

Figure 1: The marginal cost of producing safe Treasury claims and the remaining safe-debt margin. The solid series use the benchmark 2015+ OLS estimates, in which Treasury repricing varies with dealer-parent capacity under the 5 percent leverage mapping; dashed lines show the corresponding linear estimates. 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 330-basis-point household-horizon benchmark. In 2026Q1, the benchmark marginal cost is approximately 187 basis points and the remaining margin is approximately 143 basis points; the linear counterparts are 158 and 172 basis points.

For cash-Treasury repricing, I follow D'Amico and King (2013) and compare changes in the SOMA share 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 estimates allow Treasury repricing to depend on dealer-parent capacity. The interaction has the sign predicted by the model and is statistically significant under overlap-robust inference. The resulting marginal cost more than doubles, rising from about 80 basis points in 2015Q1 to about 187 basis points in 2026Q1. At the benchmark household horizon, the safe-debt margin falls from about 250 to 143 basis points. A linear decomposition helps identify the sources of the change: premium levels contribute 48.2 basis points, growth in the private coupon stock contributes 55.7 basis points through duration repricing, and cash–Treasury repricing contributes 1.5 basis points. The linear estimate rises by 105 basis points, from 53 to 158, and leaves a 2026Q1 margin of 172 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 make government debt private wealth. Mian, Straub, and Sufi (2021, 2025) develop related non-Ricardian demand mechanisms. 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. The object here is the demand contribution of the debt stock after the financial and fiscal resources needed to produce its safety have been paid.

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 derives aggregate demand while taking the cost of producing safe public debt as given. The environment is a stationary economy with output normalized to one and government purchases fixed at the share $\bar{g}$. The experiment varies the outstanding debt stock b, rather than current purchases, and lets the safe rate r adjust to clear the goods market. Households hold fully insured government claims and therefore receive the safe return. Section 3 endogenizes the premium that the government pays to support that insurance.

2.1 Perpetual-youth households

Households follow the perpetual-youth structure. They die with Poisson intensity $\lambda > 0$ and are replaced one-for-one by newborns who inherit none of their assets. Actuarially fair annuity markets transfer the assets of the deceased to survivors. Because a current household may die before future taxes are levied, it capitalizes only part of the tax burden that supports the debt it holds. The mortality hazard λ therefore governs the departure from Ricardian equivalence. With discount rate ρ and logarithmic flow utility, the marginal propensity to consume out of total wealth is $\rho+\lambda$.

Let χ denote the household's after-tax income flow. Its human wealth is the present value of that flow over the household's remaining lifetime,

$$h = \frac{\chi}{r+\lambda}. \tag{1}$$

Aggregate consumption is proportional to financial plus human wealth:

$$c = (\rho+\lambda)\big(b+h\big), \tag{2}$$

where b is the safe government-debt claim held by households.

Equations (1)–(2) are the standard log perpetual-youth demand system. Their role here is to separate the debt held by current households from the taxes borne partly by future households. Appendix A embeds this system in finite-risk-aversion Epstein–Zin preferences (Epstein and Zin, 1989). Households can insure the losses associated with rollover stress. In the high-risk-aversion limit used in the main text, they insure the full position: households receive r, while the government's excess yield finances the insurance premium derived below.

2.2 Fiscal closure and aggregate demand

Let $S \geq 0$ denote the aggregate flow premium paid to intermediaries for producing the safe-asset services of the outstanding debt stock. Section 3 derives $S(b,N)$, where $N>0$ is financial-system capacity. For now, treat S as a fiscal cost taken as given by households.

Households finance government purchases, safe-rate debt service, and the safe-asset premium through taxes. In the benchmark, intermediary owners allocate the premium income to global consumption or investment, so it does not return to domestic expenditure.[2] The government's flow budget requires

$$\tau(b,S;r) = \bar{g} + rb + S. \tag{3}$$

Disposable income is therefore

$$\chi(b,S;r) = 1-\tau(b,S;r) = 1-\bar{g}-rb-S. \tag{4}$$

Substitution into the household block gives human wealth and consumption as functions of the debt stock, its production cost, and the safe rate:

$$h(b,S;r) = \frac{1-\bar{g}-rb-S}{r+\lambda} \tag{5}$$

and

$$c(b,S;r) = (\rho+\lambda)\left(b+\frac{1-\bar{g}-rb-S}{r+\lambda}\right). \tag{6}$$

[2] Following the spirit of Gabaix and Maggiori (2015), intermediary balance-sheet capacity is the scarce input, while ownership and spending incidence are summarized by the share of premium income that returns to domestic expenditure. Appendix B allows this share to be positive.

Adding government purchases yields aggregate demand at a given safe rate:

$$c(b,S;r)+\bar{g}=(\rho+\lambda)\left(b+\frac{1-\bar{g}-rb-S}{r+\lambda}\right)+\bar{g}. \tag{7}$$

This expression isolates the two effects of additional debt. The new claim enters household financial wealth one-for-one, while debt service and safe-asset production reduce disposable income and human wealth. With infinitely lived households, the tax burden would offset the claim. With finite horizons, part of that burden falls beyond current households' expected lifetimes, leaving a positive wealth effect.

A constant premium per unit of debt would reduce the positive wealth effect of additional debt but could not make it negative. A reversal requires the premium itself to rise with issuance, so that an additional unit changes the cost of supporting the stock as it rolls over. The next section derives exactly this schedule. All of the nonlinearity behind the Laffer curve will come from the production of safety, rather than from the household block.

3 Producing Safe Public Debt

Public debt is safe in the household block because the financial system absorbs the losses associated with rollover stress. This section describes that production process. Some debt comes due; the mass of balance sheets willing to support the rollover varies across states; a shortfall is divided among those that participate; and the marginal provider prices the compensation required to enter. The resulting premium rises when the refinancing need grows relative to financial capacity.

3.1 The rollover episode and the insurance premium

Let $r^G(b,N)$ denote the yield paid by the government on insured public debt. Households receive the safe return r. The difference

$$s(b,N)\equiv r^G(b,N)-r. \tag{8}$$

I refer to $s(b,N)$ as the Treasury absorption premium: the compensation paid to the intermediary system for supplying rollover support. Thus the government pays $r+s(b,N)$, households receive r, and the aggregate premium payment is

$$S(b,N)=s(b,N)b \tag{9}$$

in steady state. This expression annualizes the sequence of payments made as the debt stock matures and is refinanced. A higher premium therefore affects the terms on which existing claims will be rolled over; it is not charged retroactively to claims that remain outstanding.

The benchmark treats the full debt stock as coming due during the rollover episode. Financial-system capacity $N>0$ is the amount of this exposure that the full intermediary sector can absorb without stress. Section 5 restores finite maturity after deriving the participation mechanism in this short-maturity case.

3.2 Capacity, participation, and exposure

Normalize the mass of potential support providers to one. A share $\underline{\alpha} \in (0,1)$ remains available in every rollover state; it represents private and official capacity that is inelastic over the episode, including reserves, repurchase-agreement facilities, and market-functioning purchases. The remaining providers may or may not enter. Let $\alpha \in [\underline{\alpha}, 1]$ denote the total mass active in a particular state. Active capacity is then αN.

If $b > \alpha N$, the active providers collectively absorb the shortfall $b - \alpha N$, or $(b - \alpha N)/\alpha$ per provider. Dividing by α matters because a given refinancing gap is more demanding when fewer balance sheets share it. The support schedule Γ maps exposure per active provider, $z \geq 0$, into the flow assessment per unit of insured debt required to mobilize that support. It satisfies

$$\Gamma(0) = 0, \qquad \Gamma'(z) > 0 \text{ for } z > 0, \qquad \Gamma'_{+}(0) > 0, \qquad \Gamma''(z) \geq 0. \tag{10}$$

The assessment rises with individual exposure. Convexity captures the increasing cost of using balance sheets more intensively as leverage constraints and price impact become more severe. The state-contingent assessment is therefore

$$a(b,N,\alpha) = \Gamma\left(\frac{\max\{b - \alpha N, 0\}}{\alpha}\right). \tag{11}$$

When capacity is ample, the assessment vanishes. When capacity is scarce, lower participation raises exposure per active provider and hence the compensation required from the government.

The ratio of refinancing needs to full-system capacity is

$$q(b,N) = \frac{b}{N}. \tag{12}$$

A shortfall occurs exactly when $\alpha < q(b,N)$. This divides the state space into three regions. If $q \leq \underline{\alpha}$, the inelastic base alone can absorb the rollover and the assessment is zero in every state. If $\underline{\alpha} < q < 1$, stress arises only when participation is sufficiently low. If $q > 1$, even full participation leaves a shortfall. The first region ends at the debt threshold $\widehat{b}(N) = \underline{\alpha} N$.

3.3 Marginal-participant pricing

The elastic providers decide whether to participate after observing private signals about an aggregate fundamental that shifts their willingness or ability to support the rollover. An entrant receives the common premium s per unit of insured debt and bears the assessment $a(b,N,\alpha)$. Since the assessment is larger when fewer providers participate, a provider's payoff depends on how many others enter. The private-signal game in Appendix C has the familiar monotone structure: providers below a signal cutoff enter, those above it stay out, and the provider at the cutoff is indifferent. Its zero-profit condition is

$$s = \mathbb{E}\left[a(b,N,\alpha) \mid \text{marginal signal}\right]. \tag{13}$$

The premium is set before the marginal provider knows the realized mass of entrants. It therefore equals that provider's expected assessment across the participation states it considers possible.

Lemma 1 (Marginal-participant pricing)**.** *Suppose elastic support providers use threshold participation strategies and have a flat prior over the aggregate fundamental. Conditional on the marginal signal, the active mass α is uniformly distributed on $[\underline{\alpha}, 1]$. The marginal-participant premium is independent of the equilibrium cutoff and equals*

$$s(b,N) = \frac{1}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{1} a(b,N,\alpha)\, d\alpha. \tag{14}$$

The uniform distribution in the lemma is a property of the marginal provider's posterior, not an assumption about realized participation. Under the flat prior, transforming the aggregate fundamental into the implied rank of the marginal signal produces a uniform variable. Since the active mass is an affine transformation of that rank, the marginal provider assigns equal density to every $\alpha \in [\underline{\alpha}, 1]$. The cutoff shifts which fundamental produces a given active mass but does not change the distribution of active mass conditional on being marginal. This is why different threshold equilibria can have different realized participation and the same zero-profit premium. Appendix C gives the signal structure and the formal change-of-variable argument.

3.4 The premium schedule

Substituting the state-contingent assessment into the marginal provider's expectation gives the premium schedule. The assessment is positive only when $\alpha < q(b,N)$, so (14) becomes

$$s(b,N) = \frac{1}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{\min\{q(b,N),1\}} \Gamma\left(\frac{b-\alpha N}{\alpha}\right) d\alpha, \tag{15}$$

with $s(b,N) = 0$ when $q(b,N) \leq \underline{\alpha}$. The integral has a direct economic interpretation. For every participation state in which capacity is insufficient, it records the assessment associated with exposure per active provider; the marginal provider averages those assessments over the possible active masses. Debt and capacity affect both the exposure within a stressed state and the set of states in which stress occurs.

In the interior stress region $\underline{\alpha} < q(b,N) < 1$, differentiating gives

$$s_b(b,N) = \frac{1}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{q(b,N)} \frac{1}{\alpha} \Gamma'\left(\frac{b-\alpha N}{\alpha}\right) d\alpha > 0, \tag{16}$$

while

$$s_N(b,N) = -\frac{1}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{q(b,N)} \Gamma'\left(\frac{b-\alpha N}{\alpha}\right) d\alpha < 0. \tag{17}$$

The changing upper limit contributes nothing because the assessment is exactly zero at the boundary $\alpha = q(b,N)$. The remaining terms capture the economically relevant changes. More debt raises exposure for every active provider in every stressed state; it also moves the stress boundary outward. More capacity lowers exposure within each stressed state and moves the boundary inward. Appendix D supplies the derivative calculations and the finite-δ formulas.

3.5 Average cost and stock-wide marginal cost

The premium $s(b,N)$ is an average cost per unit of insured debt. Multiplying it by the debt stock gives the steady-state fiscal cost $S(b,N) = bs(b,N)$. This distinction matters because issuance changes both the quantity that must eventually be refinanced and the premium at which refinancing takes place. Differentiating gives

$$S_b(b,N) = s(b,N) + bs_b(b,N). \tag{18}$$

The term s is the direct cost of supporting the marginal unit at the prevailing premium. The term bs_b is the stock effect: the additional refinancing pressure raises the premium paid as the existing stock subsequently comes due. Nothing in this calculation requires the government to reprice a bond before maturity. It is the present steady-state value of more expensive future rollovers. A moderate premium response can therefore generate a large marginal fiscal cost when it applies to the refinancing flow of a large debt stock.

Marginal cost steepens with debt for two reasons. More debt increases exposure within every stressed participation state, and it brings additional states into the stress region. Convexity of Γ strengthens the first effect. Appendix D shows that $s_{bb} \geq 0$, and hence

$$S_{bb}(b,N) = 2s_b(b,N) + bs_{bb}(b,N) > 0 \tag{19}$$

throughout the active stress region. The premium and the refinancing base over which it is paid rise together.

Financial capacity works through the same two margins in reverse. It lowers the assessment within each stressed state and reduces the set of states that require support. It therefore lowers both the premium and its response to debt. The resulting total- and marginal-cost effects are

$$S_N(b,N) = bs_N(b,N) < 0, \qquad S_{bN}(b,N) = s_N(b,N) + bs_{bN}(b,N) < 0. \tag{20}$$

Appendix D establishes the negative cross-derivative. The next section places these production costs in the household demand system. The central comparison will be between the finite-horizon wealth created by another unit of debt and the stock-wide marginal cost S_b of keeping it safe.

4 Equilibrium and the Safe-Debt Laffer Curve

The household block tells us how debt and the cost of supporting it affect demand. The financial block tells us how that cost varies with debt and capacity. Combining them yields the paper's main result. The relevant marginal object is $S_b = s + bs_b$: the demand benefit of a new safe claim must be compared with the cost of supporting the marginal unit and the refinancing stock.

4.1 Goods-market clearing and the marginal effect of debt

For a given debt stock and financial capacity, the safe rate adjusts until aggregate demand equals normalized output:

$$(\rho + \lambda)\left(b + \frac{1 - \bar{g} - rb - S(b,N)}{r + \lambda}\right) + \bar{g} = 1. \tag{21}$$

This condition has the closed-form solution

$$r(b,N) = \rho + \frac{\rho+\lambda}{1-\bar{g}}\left[\lambda b - S(b,N)\right]. \tag{22}$$

The bracketed term is the net demand contribution of the safe-debt stock in annual-rate units. The component λb captures the finite-horizon wealth effect: current households own the debt but do not expect to pay all of the taxes that support it. Safe-debt production subtracts $S(b,N)$. A larger value of this net contribution requires a higher safe rate to reduce human wealth and keep total demand equal to output.

To locate the turning point, first hold the safe rate fixed and differentiate aggregate demand with respect to debt:

$$\left.\frac{\partial(c+\bar{g})}{\partial b}\right|_{r,N} = (\rho+\lambda)\left[1-\frac{r+S_b(b,N)}{r+\lambda}\right] = \frac{\rho+\lambda}{r+\lambda}\left[\lambda - S_b(b,N)\right]. \tag{23}$$

The decomposition inside the last bracket is the economic core of the result. One additional unit raises household financial wealth by one. It also requires taxes for safe debt service r and for the stock-wide production cost S_b. Current households capitalize those taxes at $r+\lambda$. Ordinary debt service therefore leaves the familiar finite-horizon benefit $\lambda/(r+\lambda)$, while safe-debt production subtracts $S_b/(r+\lambda)$. The marginal propensity to consume, $\rho+\lambda$, converts the change in wealth into consumption demand.

The safe-debt margin is the rate-equivalent difference $\lambda - S_b(b,N)$. It is positive when the wealth created by the marginal claim exceeds the resources needed to produce its safety and negative when the production cost is larger.

4.2 The Laffer peak and the equilibrium safe rate

Debt is expansionary when

$$S_b(b,N) < \lambda, \tag{24}$$

and contractionary when

$$S_b(b,N) > \lambda. \tag{25}$$

The peak satisfies

$$\lambda = s(b^L,N) + b^L s_b(b^L,N), \tag{26}$$

where b^L is the debt stock at the peak. The left side expresses the finite-horizon debt benefit in annual-rate units; the right side is the stock-wide marginal cost of safe-asset production. The non-Ricardian wedge itself remains the part of the future tax burden that current households do not capitalize.

The resulting schedule is a Laffer curve for debt's contribution to aggregate demand. Before rollover stress begins, the production premium is zero and debt is expansionary. Inside the stress region, S_b rises with debt and gradually compresses the safe-debt margin. At b^L, the two sides are equal. Beyond that point, the claims remain safe, but the fiscal resources required to support another unit and refinance the stock exceed the wealth created by the marginal claim.

Proposition 1 (Safe-debt Laffer curve)**.** *Let $r(b,N)$ be the equilibrium safe rate that clears the goods market in* (21)*, and define the safe-debt margin*

$$H(b,N) \equiv \lambda - S_b(b,N).$$

Suppose the economy enters the stress region at $\widehat{b}(N)$, with $S_b(\widehat{b},N) = 0$. Assume that $S_b(b,N)$ is continuous and strictly increasing on the stress region and that $S_b(b,N) \to \infty$ as $b \to \infty$. Then there is a unique Laffer peak $b^L(N) > \widehat{b}(N)$ satisfying (26)*. At the prevailing safe rate, debt raises the aggregate-demand contribution of public debt for $b < b^L(N)$ and lowers it for $b > b^L(N)$.*

Strictly increasing S_b makes the crossing unique. The proposition concerns debt's contribution to demand at the prevailing rate. The equilibrium safe rate inherits the same turning point because it is the price that clears the goods market.

Differentiating (22) gives the exact response

$$r_b(b,N) = \frac{\rho + \lambda}{1 - \bar{g}} \big[\lambda - S_b(b,N)\big]. \tag{27}$$

If $\lambda - S_b > 0$, additional debt creates excess demand at the old safe rate. The rate rises and reduces human wealth until demand again equals output. Once $\lambda - S_b < 0$, issuance lowers demand at the old rate; the safe rate must fall to raise human wealth and restore clearing. The aggregate-demand effect and the equilibrium-rate response are therefore governed by the same margin. Past the peak, the premium paid by the government can continue to rise even as the safe return received by households falls.

Corollary 1 (Coincident aggregate-demand and safe-rate peaks)**.** *The aggregate-demand Laffer peak and the peak of the equilibrium safe-rate schedule coincide at $b^L(N)$. Hence whether the safe rate rises or falls with debt identifies the branch of the aggregate-demand safe-debt Laffer curve.*

Figure 2 illustrates the closed-form case of Appendix E. Rollover stress begins at $\widehat{b}$; the Laffer and safe-rate peaks coincide at the higher debt level b^L.

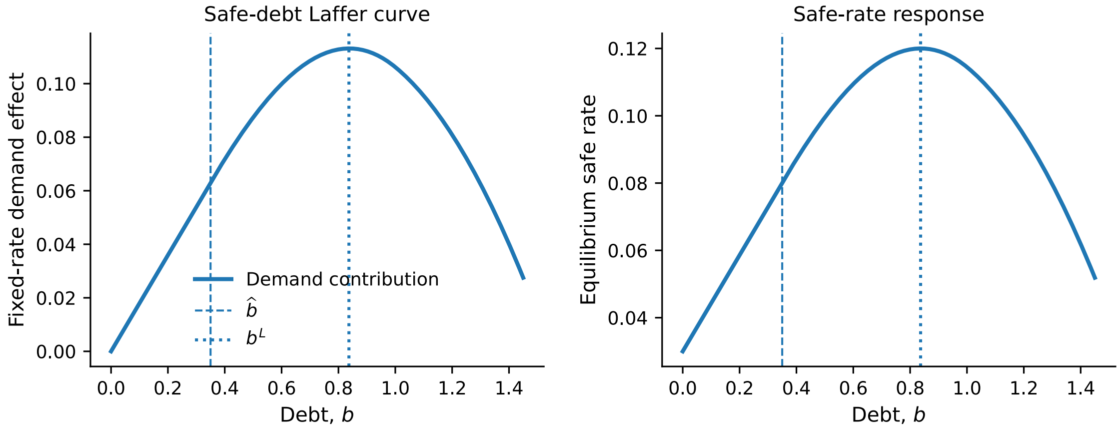

Figure 2: The aggregate-demand safe-debt Laffer curve and the equilibrium safe rate. The figure uses the closed-form linear-support-cost case of Appendix E. The vertical dashed line marks the stress threshold $\widehat{b}$; the dotted line marks the Laffer peak b^L. Above $\widehat{b}$, the marginal fiscal cost of producing safe-asset services rises with debt. At b^L, this cost equals the finite-horizon debt benefit expressed in annual-rate units by λ.

4.3 Financial capacity and the Laffer peak

Financial capacity affects the model at two distinct points. It determines when rollover stress begins, and, once stress is active, it determines how rapidly marginal production cost rises. The first effect follows directly from the debt-capacity ratio. Differentiating (12) gives

$$q_b(b,N) = \frac{1}{N} > 0, \qquad q_N(b,N) = -\frac{1}{N}q(b,N) < 0. \tag{28}$$

Debt raises rollover exposure relative to capacity; higher N lowers the active share needed to absorb a given refinancing flow. Since the inelastic base can absorb debt up to $\widehat{b}(N) = \underline{\alpha}N$, greater capacity moves the onset of stress one-for-one with N.

Within the stress region, capacity lowers both total and marginal safety-production costs: $S_N < 0$ and $S_{bN} < 0$ in (20). At a given debt stock, more balance-sheet capacity reduces the assessment in each stressed state and makes the premium less sensitive to debt. It therefore raises $\lambda b - S(b,N)$ and slows the compression of the safe-debt margin. The Laffer peak moves outward because the marginal-cost schedule reaches λ at a larger debt stock.

Proposition 2 (Financial capacity and the Laffer peak)**.** *Let $b^L(N)$ denote an interior Laffer peak satisfying*

$$\lambda - S_b(b^L,N) = 0.$$

Suppose $S_{bb}(b^L,N) > 0$. Then

$$\frac{db^L}{dN} = -\frac{S_{bN}(b^L,N)}{S_{bb}(b^L,N)} > 0. \tag{29}$$

The derivative follows from the implicit-function theorem applied to the peak condition. The denominator is positive because marginal cost rises with debt; the numerator is negative because capacity lowers marginal cost. Figure 3 shows both effects: higher N moves the onset of stress and the Laffer peak to the right.

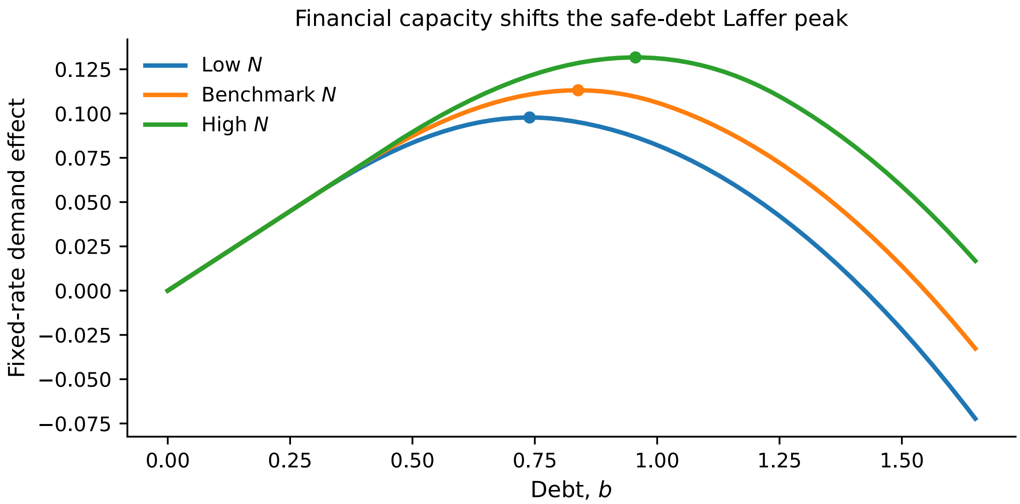

Figure 3: Financial capacity shifts the Laffer peak. The figure plots the aggregate-demand contribution of insured public debt for three levels of financial capacity in the closed-form case. Higher N lowers rollover pressure and shifts the peak outward.

Capacity also shifts the equilibrium safe-rate schedule at a fixed debt stock. Differentiating (22) with respect to N gives

$$r_N(b,N) = -\frac{\rho+\lambda}{1-\bar{g}} S_N(b,N) > 0. \tag{30}$$

Greater capacity also raises the equilibrium safe rate at a given debt stock in the stress region. By lowering the fiscal cost of producing safety, it increases demand at the old rate. The safe rate rises to reduce human wealth and restore equilibrium. Capacity therefore lifts the rate schedule locally while extending its increasing branch to a larger debt stock.

The short-maturity benchmark exposes the full debt stock to rollover. Finite maturity changes both the amount refinanced and the price response of the claims that remain outstanding.

5 Maturity: Quantity and Price Dimensions of Rollover Risk

The short-maturity benchmark exposes the full debt stock to each rollover episode. Positive maturity changes the mechanism in two ways. First, only part of the stock comes due during a given interval, reducing the quantity that requires immediate support. Second, the claims that remain outstanding have longer duration, so their prices respond more strongly to a change in the premium associated with that same rollover exposure. The second effect amplifies the first in prices; it does not introduce an independent source of risk.

These two effects cannot be added as separate schedules. The total premium enters the yield used to value a long claim and therefore affects the claim's duration. Duration, in turn, determines how strongly the rollover assessment affects its price. The premium and its price multiplier must consequently be solved together. The analysis below first restores the quantity of debt that comes due, then adds the duration-price fixed point. Throughout the section, the safe rate is held fixed when differentiating the premium; its equilibrium adjustment continues to follow the goods-market condition in Section 4.

5.1 The quantity exposed to rollover

Let δ be the Poisson intensity at which each unit of debt matures. Average maturity is $1/\delta$, and the share of the stock that comes due during the unit-length episode is

$$\mu(\delta) = 1 - e^{-\delta}.$$

Higher δ means shorter maturity and a larger refinancing flow $\mu(\delta)b$. The participation mechanism of Section 3 is unchanged; finite maturity simply replaces the full stock b with the amount that actually comes due.

Rollover pressure and the upper boundary of the stress region become

$$q(b,N,\delta) = \frac{\mu(\delta)b}{N}. \tag{31}$$

For $q(b,N,\delta) > \underline{\alpha}$, the direct rollover-insurance premium is

$$s^R(b,N,\delta) = \frac{1}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{\min\{q(b,N,\delta),1\}} \Gamma\left(\frac{\mu(\delta)b - \alpha N}{\alpha}\right) d\alpha, \tag{32}$$

and $s^R(b,N,\delta) = 0$ when $q(b,N,\delta) \leq \underline{\alpha}$. The superscript R identifies the direct rollover component and will carry unchanged into the empirical section.

Longer maturity lowers s^R through both channels already present in the benchmark. It reduces exposure per active provider in every stressed state and may remove high-participation states from the stress region. Thus s^R rises with debt and rollover intensity and falls with financial capacity. Capacity also weakens the additional rollover pressure created by shorter maturity: Appendix F establishes $s^R_{\delta N} < 0$.

5.2 The price effect of duration

Let $s^D(b,N,\delta)$ denote the duration-price component and let

$$s(b,N,\delta) = s^R(b,N,\delta) + s^D(b,N,\delta) \tag{33}$$

be the total premium. A Poisson-maturity claim that pays a unit flow while outstanding has duration $1/[r + s(b,N,\delta) + \delta]$. Lower δ lengthens the claim; a higher total premium raises its discount rate and shortens it.

Let $\vartheta(N) > 0$ scale the transmission of rollover-support risk into bond prices, with $\vartheta_N(N) < 0$. The duration-price component is

$$s^D(b,N,\delta) = \frac{\vartheta(N)}{[r+s(b,N,\delta)+\delta]^2} s^R(b,N,\delta). \tag{34}$$

Longer duration magnifies the price effect of the direct rollover premium. Capacity operates twice within this mechanism: it lowers s^R and dampens its transmission into bond prices.

Combining the two components gives the fixed point

$$s(b,N,\delta) = s^R(b,N,\delta)\left\{1+\frac{\vartheta(N)}{[r+s(b,N,\delta)+\delta]^2}\right\}. \tag{35}$$

Greater rollover stress raises the direct premium and the amount magnified by duration. The resulting increase in s raises the discount rate, shortens duration, and partially offsets the initial amplification. The right-hand side is therefore continuous and decreasing in s, while the left-hand side is increasing and unbounded. The fixed point is unique for finite s^R and $r+\delta > 0$.

Implicit differentiation with respect to capacity gives

$$s_N = \frac{\left[1+\frac{\vartheta(N)}{(r+s+\delta)^2}\right]s_N^R + \frac{s^R\vartheta_N(N)}{(r+s+\delta)^2}}{1+\frac{2s^R\vartheta(N)}{(r+s+\delta)^3}} < 0, \tag{36}$$

where all functions are evaluated at (b,N,δ). Capacity lowers both the direct rollover premium and its duration-price amplification. The denominator captures the offset from the endogenous shortening of duration.

Differentiation with respect to rollover intensity gives

$$s_\delta = \frac{\left[1+\frac{\vartheta(N)}{(r+s+\delta)^2}\right]s_\delta^R - \frac{2s^R\vartheta(N)}{(r+s+\delta)^3}}{1+\frac{2s^R\vartheta(N)}{(r+s+\delta)^3}}. \tag{37}$$

Shorter maturity raises rollover exposure, so $s_\delta^R > 0$. It also shortens duration and reduces price amplification. Either force can dominate the response of the total premium.

Proposition 3 (Positive maturity and rollover-price amplification)**.** *Fix $\delta < \infty$ and consider the active stress region $\underline{\alpha} < q(b,N,\delta) < 1$. The direct rollover-insurance component is increasing in rollover intensity:*

$$s_\delta^R(b,N,\delta) = \frac{\mu'(\delta)b}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{1}{\alpha}\Gamma'\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)d\alpha > 0. \tag{38}$$

The total premium is uniquely determined by (35)*. Its capacity derivative is negative as in* (36)*, and its maturity derivative is given by* (37)*. As $\delta \to \infty$, the duration-price component vanishes and the short-maturity premium is recovered.*

Quantity and price are therefore two dimensions of the same rollover exposure, joined through the fixed point.

5.3 The positive-maturity margin

For a fixed maturity structure, the total fiscal cost of producing safe-asset services is

$$S(b,N,\delta) = b\,s(b,N,\delta). \tag{39}$$

The steady-state marginal cost is stock-wide:

$$S_b(b,N,\delta) = s(b,N,\delta) + b\frac{\partial s(b,N,\delta)}{\partial b}. \tag{40}$$

Using $s = s^R + s^D$, the same object can be written as

$$S_b = \left(s^R + bs_b^R\right) + \left(s^D + bs_b^D\right). \tag{41}$$

Each bracket contains a premium level and a stock-repricing term. The first reflects direct rollover pricing; the second reflects its duration-price amplification. The empirical implementation estimates the corresponding cash-Treasury and real-duration responses directly.

The safe-debt margin is $\lambda - S_b(b,N,\delta)$, so at fixed δ the Laffer peak solves

$$\lambda = s(b,N,\delta) + b\frac{\partial s(b,N,\delta)}{\partial b}. \tag{42}$$

In the empirical implementation, s_t^R is the observed cash-Treasury counterpart of the rollover component and s_t^D is the observed real-duration component. Their premium levels and stock-repricing terms correspond to the two brackets in (41).

5.4 Maturity choice as an extension

Appendix F allows the government to choose rollover intensity δ, or average maturity $1/\delta$. At an interior optimum,

$$\left[1 + \frac{\vartheta(N)}{(r+s+\delta)^2}\right] s_\delta^R = \frac{2s^R\vartheta(N)}{(r+s+\delta)^3}. \tag{43}$$

The amplified marginal rollover cost of shortening maturity equals the reduction in duration-price amplification. If capacity sufficiently reduces the marginal rollover cost, the optimum shifts toward faster rollover; scarcity tilts issuance toward longer maturity. The empirical exercise instead evaluates the premium at the observed maturity structure.

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

The positive-maturity model implies a stock-wide marginal cost $S_b = s + b\,\partial s/\partial b$. I construct its U.S. counterpart for privately held marketable Treasury notes and bonds and compare 2015Q1 with 2026Q1. The empirical decomposition uses the same rollover and duration notation as the model:

$$s_t = s_t^R + s_t^D. \tag{44}$$

Let $\mathscr{M}_t^R$ and $\mathscr{M}_t^D$ denote the repricing of the outstanding private Treasury stock through these two components. The empirical marginal cost is

$$S_{b,t} \equiv s_t^R + s_t^D + \mathscr{M}_t^R + \mathscr{M}_t^D. \tag{45}$$

The premium levels s_t^R and s_t^D are the measured components of the Treasury absorption premium. The terms $\mathscr{M}_t^R$ and $\mathscr{M}_t^D$ capture the additional cost created when issuance reprices the outstanding private Treasury stock.

6.1 Premium levels

The cash-Treasury component s_t^R is the Treasury par yield minus the maturity-matched par swap rate. A swap supplies interest-rate exposure without financing and warehousing a cash Treasury, so the spread measures cash-specific Treasury pricing. The measure is signed: its negative 2015Q1 value indicates that cash-specific benefits then outweighed balance-sheet costs. I construct a debt-weighted measure from the two-, three-, five-, seven-, ten-, and thirty-year maturities. The swap-spread panel uses continuous benchmark-versus-IRS chains in each bucket, with a LIBOR-consistent seven-year series constructed from Treasury and swap par yields. The resulting measure rises from −18.2 basis points in 2015Q1 to −0.5 basis points in 2026Q1.

The real-duration component s_t^D is the D'Amico–Kim–Wei (DKW) real term premium (D'Amico, Kim, and Wei, 2018). DKW jointly prices nominal and inflation-indexed Treasuries, separating real-duration from inflation-risk compensation. I linearly interpolate or extrapolate its five- and ten-year estimates to the benchmark maturities and weight them by private debt. The observed real term premium is the signed equilibrium price associated with placing longer claims in private hands. In the model it maps to the duration-price component generated when maturity magnifies rollover-support exposure. It rises from 8.4 basis points in 2015Q1 to 38.9 basis points in 2026Q1.

Together, the two premium levels rise from −9.9 basis points in 2015Q1 to 38.3 basis points in 2026Q1. Their 48.2-basis-point increase is the contribution of measured premium levels to the change in marginal cost.

6.2 Cash-Treasury repricing

Following D'Amico and King (2013), I use a monthly panel of the six benchmark maturity buckets to relate changes in Treasury–swap spreads to changes in the SOMA share. At each horizon h, I estimate

$$\Delta_h s_{m,t}^R = \phi_m + \psi_t + \beta_h \Delta_h share_{m,t} + \varepsilon_{m,t}, \tag{46}$$

where m denotes a maturity bucket and $\Delta_h x_{m,t} \equiv x_{m,t+h} - x_{m,t}$. I define $share_{m,t} = 100 \times SOMA_{m,t}/out_{m,t}$, using System Open Market Account holdings and outstanding Treasury debt in the bucket. Maturity effects absorb persistent differences across buckets; start-month effects absorb aggregate news and curve-wide repricing. The sample excludes the first half of 2020 so that the estimates capture persistent repricing outside the emergency-market episode.

The coefficient β_h is the basis-point change in the spread following a one-percentage-point increase in the SOMA share. A negative value means that a larger SOMA share lowers the

Treasury–swap spread in that maturity segment. Let $b_{m,t}$ denote private float, $b_t = \sum_m b_{m,t}$, $w_{m,t} = b_{m,t}/b_t$, and $\varpi_{m,t} = SOMA_{m,t}/out_{m,t}$. A one-percentage-point increase in the SOMA share removes one percent of the bucket's outstanding stock from private portfolios. Reversing the sign gives the response to greater private supply; multiplication by $100(1-\varpi_{m,t})$ scales that response to the bucket's private float, and $w_{m,t}$ aggregates across maturities. The stock-wide repricing term is

$$\mathscr{M}^R_{h,t} = -100\,\beta_h \sum_m w_{m,t}(1-\varpi_{m,t}). \tag{47}$$

This mapping distributes marginal issuance in proportion to current private maturity weights. After the start-month effects remove common curve movements, it retains the maturity-specific cash-Treasury response. In the Monthly Statement of the Public Debt (MSPD) notes-and-bonds universe, the weighting sum $\sum_m w_{m,t}(1-\varpi_{m,t})$ rises from 0.781 in 2015Q1 to 0.837 in 2026Q1.

I average $\mathscr{M}^R_{h,t}$ over $h = 12$–24 to capture persistent repricing. Because adjacent starting months share most of their outcome windows, inference uses Driscoll–Kraay/Newey–West standard errors with 24 monthly lags. The average linear slope is approximately -0.27 in the sample beginning in 2015 and -0.48 in the sample beginning in 2010. They imply 2026Q1 cash-Treasury repricing terms of 23 and 40 basis points. Table 1 also reports an IV sensitivity estimate.

Table 1: 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. These linear estimates provide comparisons for the nonlinear benchmark. Inference allows for serial dependence across overlapping starting months. The IV row uses lagged SOMA share and trailing issuance, an on-the-run proxy for recently issued securities. Appendix G gives the specification and inference details.

6.3 Nonlinear repricing and financial capacity

The model predicts that scarce capacity steepens cash-Treasury repricing. I measure capacity at the holding companies that own Treasury dealers, which I call dealer parents. The supplementary leverage ratio requires Tier 1 capital to be at least a fraction κ of total leverage exposure:

$$\frac{T1_{i,t}}{LE_{i,t}} \geq \kappa. \tag{48}$$

For given Tier 1 capital, $T1_{i,t}/\kappa$ is the maximum leverage exposure consistent with the ratio. The difference $T1_{i,t}/\kappa - LE_{i,t}$ is unused leverage-exposure room. Adding Treasury inventory already held converts this unused room into a measure of total Treasury absorption—current holdings plus additional room—and aggregating across dealer parents gives

$$N_t(\kappa) = \sum_i \left(\frac{T1_{i,t}}{\kappa} - LE_{i,t}\right) + \sum_i UST_{i,t}. \tag{49}$$

The index i runs over dealer parents, and $UST_{i,t}$ is the Treasury inventory they already hold. The baseline sets $\kappa = 0.05$: Tier 1 capital is at least 5 percent of total leverage exposure, the historical holding-company threshold used for affected U.S. global systemically important bank (GSIB) parents. The tight sensitivity mapping sets $\kappa = 0.06$, based on the historical well-capitalized threshold for their insured depository subsidiaries. Appendix G describes the data, monthly interpolation, and 2026 regulatory change.

To estimate how capacity changes the cash-Treasury response, I interact the change in the SOMA share with standardized dealer-parent capacity:

$$\Delta_h s^R_{m,t} = \phi_m + \psi_t + \beta_h \Delta_h share_{m,t} + \gamma_h \left(\Delta_h share_{m,t} \times \widetilde{N}_t \right) + \varepsilon_{m,t}, \tag{50}$$

where $\widetilde{N}_t$ measures capacity in standard deviations from its 2015–2024 mean, with higher values denoting greater capacity, and γ_h is the coefficient on the capacity interaction. The estimated $\gamma_h > 0$ has the model's predicted sign: abundant capacity flattens the negative within-bucket supply slope, while scarce capacity steepens it. Under the baseline 5 percent mapping, the average direct and interaction coefficients are -0.24 and 0.19. The Driscoll–Kraay/Newey–West t-statistics on the interaction range from 2.4 to 2.9 over the $h = 12$–24 window, with a mean of 2.6. The nonlinear response is therefore statistically significant under inference that accounts for the overlapping horizons.

Table 2: Nonlinear cash-Treasury repricing estimates

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 evaluate $-100(\beta_h + \gamma_h \widetilde{N}_{2026Q1}) \sum_m w_{m,2026Q1}(1 - \varpi_{m,2026Q1})$.

The 2026Q1 capacity observation is about -2.0 standard deviations under the baseline mapping. Evaluating the interaction at that state gives a 52-basis-point cash-Treasury repricing term, compared with 23 basis points in the linear specification. Under the tight 6 percent mapping, the capacity state is -3.1 standard deviations and the repricing term is 72 basis points. The 2026Q1 observations lie below the 2015–2024 estimation range. Appendix G reports the full coefficient paths, the comparison across mappings, and IV sensitivities.

6.4 Real-duration repricing

The duration design uses twelve Treasury-relevant asset-purchase and balance-sheet-runoff announcements. I regress the two-trading-day change in the DKW real term premium on the announced change in expected cumulative Treasury purchases:

$$\Delta s^D_e = c_D + \beta_D Amount_e + u_e, \tag{51}$$

where e runs over announcements, c_D is an intercept, and $Amount_e$ is measured in billions of dollars. Positive amounts increase expected Federal Reserve purchases and reduce private

Treasury supply; tapering and runoff enter with negative amounts. Appendix G lists the events and sizing assumptions.

For the debt-weighted, average-maturity DKW real term premium, the with-intercept estimate is $\beta_D \simeq -0.0055$ basis points per billion dollars. Let b_t^{priv} denote privately held notes and bonds in the MSPD universe. Since purchases reduce private supply, the duration repricing term is

$$\mathscr{M}_t^D = -\beta_D b_t^{priv}. \tag{52}$$

The private coupon float rises from \$7.525 trillion in 2015Q1 to \$17.642 trillion in 2026Q1, giving duration repricing terms of 41.4 and 97.1 basis points.

The baseline uses the debt-weighted, average-maturity DKW measure and scales announcements in dollars. A specification using the ten-year DKW premium as the dependent variable gives 116 basis points in 2026Q1. A specification that scales announcement amounts by contemporaneous private float gives a stock-wide repricing term of 53 basis points. These specifications provide sensitivity to the duration measure and the scaling of purchase amounts.

The twelve-event sample and the sizing of several open-ended programs leave substantial uncertainty around the slope. Using HC3 heteroskedasticity-robust inference and a small-sample t critical value, the 95 percent interval for β_D is $[-0.0109, -0.0001]$; this maps into a 1–193 basis point interval for $\mathscr{M}_{2026Q1}^D$. Leave-one-event-out estimates imply 74–116 basis points, while excluding the two March 2020 announcements gives 151 basis points. The point estimate measures the average response of the duration-price component to private Treasury supply—the price response that magnifies rollover-support exposure in the model.

6.5 The benchmark marginal cost

Combining the four components in (45), the benchmark nonlinear OLS specification gives a 2026Q1 marginal cost of approximately 187 basis points and a safe-debt margin of approximately 143 basis points. Figure 1 plots the full quarterly path. The marginal cost rises from about 80 basis points in 2015Q1, so the remaining margin falls from about 250 to 143 basis points over the period.

The linear specification provides a transparent decomposition of this movement. Its marginal cost rises from 52.9 basis points in 2015Q1 to 158.3 basis points in 2026Q1, an increase of 105.4 basis points. The nonlinear benchmark places the marginal cost at a higher level in both endpoints because capacity is scarce in both 2015Q1 and 2026Q1. Table 3 reports the preferred 2026Q1 calculation and the linear comparisons.

Table 3: 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
2026Q1, nonlinear OLS benchmark	−0.5	38.9	52	97.1	≈ 187	≈ 143
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 basis points. The first row evaluates the OLS capacity interaction under the benchmark 5 percent mapping. The remaining rows report linear specifications. The repricing terms use the date-specific mappings in (47) and (52) on the MSPD notes-and-bonds universe. Sums use unrounded components; $H_t = 330 - S_{b,t}$.

In the linear decomposition, measured premium levels contribute 48.2 basis points to the increase. Growth in the private coupon float raises the duration repricing term by 55.7 basis points, and cash-Treasury repricing contributes another 1.5 basis points. The three components sum to 105.4 basis points. The 2010+ and IV linear specifications produce increases of 106.6 and 106.0 basis points.

Table 4 shows how the nonlinear benchmark varies with the duration measure and purchase-amount scaling.

Table 4: Duration sensitivity of the 2026Q1 marginal cost

Benchmark nonlinear specification	$S_{b,2026Q1}$	Margin H_{2026Q1} ($\lambda = 330$)
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 52-basis-point cash-Treasury repricing term from the benchmark 5 percent capacity mapping. Differences of one basis point can arise from unrounded regression coefficients and premium levels. The share specification scales announcement amounts by contemporaneous private float; the average-maturity and ten-year specifications scale them in dollars. All use the same twelve-event design.

6.6 The safe-debt margin

The benchmark sets $\lambda = 330$ basis points, or a 3.3 percent annual mortality hazard, giving an expected household horizon of about 30 years. The safe-debt margin

$$H_t = \lambda - S_{b,t} \tag{53}$$

is in annual-rate units: λ expresses the finite-horizon debt benefit in those units, and $S_{b,t}$ is the stock-wide marginal cost of producing safe Treasury claims. The Laffer peak occurs when this gap reaches zero. In the nonlinear benchmark, H_t falls from about 250 basis points in 2015Q1 to approximately 143 basis points in 2026Q1. The linear specification gives 277 and 172 basis points at the same dates. The tight 6 percent capacity mapping places the 2026Q1 margin at approximately 123 basis points.

At any date t, the horizon parameter that places the economy at the peak is $\lambda_t^* = S_{b,t}$. Holding λ fixed,

$$\Delta H_t = -\Delta S_{b,t}, \tag{54}$$

so the increase in marginal cost compresses the margin one-for-one. The level of H_t is a benchmark translation of the household calibration and the duration construction; the calculation does not propagate the event-study uncertainty into that level.

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 this cost. At low debt, the wealth created by an additional claim exceeds the fiscal cost of supporting it. As rollover pressure builds, the premium rises and so does the cost of refinancing the stock. The turning point is reached when this stock-wide marginal cost exhausts the finite-horizon wealth effect. Beyond it, further issuance lowers both debt's contribution to aggregate demand and the equilibrium safe rate, even though the claims remain safe.

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. 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 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 Households, Insurance, and Perpetual-Youth Demand

The full-insurance limit used in the main text follows from finite-risk-aversion insurance choices. The derivation also recovers the perpetual-youth demand system in Section 2.

A.1 Preferences, states, and insurance choices

Time is continuous. Households die with Poisson hazard λ, receive the actuarial mortality return on surviving wealth, and discount at rate ρ. Newborns replace them but do not inherit their assets, which the annuity market redistributes to survivors. Epstein–Zin preferences (Epstein and Zin, 1989) have unit intertemporal elasticity and risk aversion σ across aggregate payoff states. If $\omega \in \Omega_t$ indexes rollover conditions and $c_i(\omega)$ is state-contingent consumption, the within-date certainty equivalent is

$$C_t^{\sigma} = \left(\mathbb{E}_t\left[c_i(t,\omega)^{1-\sigma}\right]\right)^{1/(1-\sigma)}. \tag{55}$$

The logarithmic time aggregator delivers unit intertemporal elasticity; σ prices within-date rollover risk.

With financial wealth w_i, the household chooses government-bond holdings b_i, insured units $\iota_i \in [0, b_i]$, and uninsured exposure

$$u_i \equiv b_i - \iota_i \geq 0. \tag{56}$$

Any financial wealth not held in government debt is invested in a safe annuitized asset earning r. Government debt pays r^G, insurance costs π per unit, and uninsured debt suffers loss $\ell(\omega) \geq 0$. Disposable private income χ follows the fixed-purchases tax rule in (4). The state-contingent wealth law is

$$\dot{w}_i(\omega) = (r+\lambda)w_i + \chi - c_i(\omega) + \left(r^G - r - \pi\right)\iota_i + \left(r^G - r - \ell(\omega)\right)u_i. \tag{57}$$

The final two terms separate insured from uninsured exposure.

A.2 Finite-σ insurance conditions

The certainty equivalent in (55) induces the normalized within-date state-price weight

$$\Lambda_i^{\sigma}(\omega) \equiv \frac{c_i(\omega)^{-\sigma}}{\mathbb{E}[c_i(\omega)^{-\sigma}]}. \tag{58}$$

A marginal uninsured position pays $r^G - r - \ell(\omega)$ relative to the safe asset:

$$\mathbb{E}\left[\Lambda_i^{\sigma}(\omega)(r^G - r - \ell(\omega))\right] = 0, \qquad \text{if } u_i > 0, \tag{59}$$

$$\mathbb{E}\left[\Lambda_i^{\sigma}(\omega)(r^G - r - \ell(\omega))\right] \leq 0, \qquad \text{if } u_i = 0. \tag{60}$$

A marginal insured position has state-independent excess payoff $r^G - r - \pi$:

$$r^G - r - \pi = 0, \qquad \text{if } \iota_i > 0, \tag{61}$$
$$r^G - r - \pi \leq 0, \qquad \text{if } \iota_i = 0. \tag{62}$$

Holding total bond holdings fixed, a marginal unit of insurance pays $\ell(\omega) - \pi$ and satisfies

$$\mathbb{E}\left[\Lambda_i^\sigma(\omega)(\ell(\omega) - \pi)\right] = 0, \qquad \text{if } 0 < \iota_i < b_i, \tag{63}$$
$$\mathbb{E}\left[\Lambda_i^\sigma(\omega)(\ell(\omega) - \pi)\right] \leq 0, \qquad \text{if } \iota_i = 0, \tag{64}$$
$$\mathbb{E}\left[\Lambda_i^\sigma(\omega)(\ell(\omega) - \pi)\right] \geq 0, \qquad \text{if } \iota_i = b_i. \tag{65}$$

Households therefore compare the insurance premium with the state-price-weighted rollover loss.

A.3 Full-insurance benchmark

Let

$$\ell^{\max} \equiv \operatorname*{ess\,sup}_{\omega \in \Omega} \ell(\omega) \tag{66}$$

be the worst-state rollover loss. In the equilibrium studied in the main text, the insurance premium π is the competitive average support cost across rollover states. With a nondegenerate distribution of rollover losses,

$$\pi < \ell^{\max}. \tag{67}$$

Assumption 1 (Full insurance)**.** *Households fully insure their government-bond holdings against rollover-stress losses: $\iota_i = b_i$, $u_i = 0$.*

The main text studies the full-insurance limit as $\sigma \to \infty$. When $\pi < \ell^{\max}$, worst-state consumption governs portfolio choice in this limit, and households insure the entire government-bond position as in Assumption 1.

Under full insurance, equation (61) implies

$$r^G = r + \pi, \tag{68}$$

so a fully insured unit of government debt earns the safe net return,

$$r^G - \pi = r. \tag{69}$$

Substituting $u_i = 0$ and $r^G - r - \pi = 0$ into (57) gives the state-independent wealth law

$$\dot{w}_{i,t} = (r_t + \lambda) w_{i,t} + \chi_t - c_{i,t}. \tag{70}$$

Full insurance makes $C_t^\sigma = c_{i,t}$, reducing the problem to the standard log perpetual-youth case.

A.4 Human wealth and consumption under full insurance

Human wealth is the present value of disposable income while alive:

$$h_t = \int_t^{\infty} \exp\left[-\int_t^{\tau} (r_v + \lambda)\, dv\right] \chi_\tau \, d\tau. \tag{71}$$

It evolves according to

$$\dot{h}_t = (r_t + \lambda) h_t - \chi_t, \tag{72}$$

and in a steady state with constant safe rate r and constant χ,

$$h(r) = \frac{\chi}{r + \lambda}. \tag{73}$$

For total wealth $\mathscr{W}_{i,t} = w_{i,t} + h_t$, equations (70) and (72) imply

$$\dot{\mathscr{W}}_{i,t} = (r_t + \lambda)\mathscr{W}_{i,t} - c_{i,t}. \tag{74}$$

The log perpetual-youth Euler equation is $\dot{c}_{i,t}/c_{i,t} = r_t - \rho$. In steady state, $\mathscr{W}_i = w_i + h$, and imposing the no-Ponzi condition on (74) yields

$$c_i = (\rho + \lambda)\mathscr{W}_i. \tag{75}$$

Aggregating over the unit-mass population, with fully insured public debt as households' safe financial asset, gives

$$c = (\rho + \lambda)(b + h). \tag{76}$$

Finite horizons make public debt net wealth because current households only partly capitalize taxes paid by future households. Under fixed purchases, additional debt raises financial wealth but also the tax burden embedded in human wealth. Combining (2) and (4) yields (23).

A.5 Goods-market clearing and the safe rate

In equilibrium, $\pi = s(b, N)$, the marginal-participant premium derived in Appendix C. The full-insurance household block and fixed-purchases fiscal closure give (21). Solving that condition gives (22); differentiation then gives (27).

B Global-Owner Incidence and Partial Recycling

Let ξ denote the share of intermediary premium income spent on domestic goods. The benchmark sets this share to zero; this section allows some of the income financed by households' payments for safe-asset services to return to domestic expenditure.

The marginal aggregate-demand contribution of public debt is governed by

$$H(b, N) = \lambda - S_b(b, N),$$

where λ is the mortality hazard, which expresses the non-Ricardian debt benefit in annual-rate units in this normalization, and $S_b(b,N)$ is the stock-wide marginal cost of safe-debt production. The Laffer peak occurs at $H=0$.

Intermediary owners are marginally global: they may be foreign or wealthy domestic investors, and they allocate premium income between domestic goods and global consumption or investment.

The benchmark sets $\xi=0$, so intermediary premium income has zero weight in domestic expenditure. With a positive domestic expenditure share, aggregate demand is

$$\mathscr{D}_\xi(b,S;r)=(\rho+\lambda)\left(b+\frac{1-\bar{g}-rb-S}{r+\lambda}\right)+\bar{g}+\xi S. \tag{77}$$

At a fixed safe rate,

$$\left.\frac{\partial \mathscr{D}_\xi}{\partial b}\right|_{r,N}=\frac{\rho+\lambda}{r+\lambda}\left\{\lambda-\left[1-\frac{\xi(r+\lambda)}{\rho+\lambda}\right]S_b(b,N)\right\}. \tag{78}$$

The corresponding safe-debt margin is therefore

$$H_\xi(b,N)=\lambda-\left[1-\frac{\xi(r+\lambda)}{\rho+\lambda}\right]S_b(b,N). \tag{79}$$

Provided $\rho+\lambda>\xi(r+\lambda)$, the Laffer peak satisfies

$$S_b(b,N)=\frac{\lambda}{1-\xi(r+\lambda)/(\rho+\lambda)}=\frac{(\rho+\lambda)\lambda}{\rho+\lambda-\xi(r+\lambda)}. \tag{80}$$

The benchmark $\xi=0$ recovers $S_b=\lambda$. Partial recycling weakens the marginal drag through the coefficient $1-\xi(r+\lambda)/(\rho+\lambda)$. If $\rho+\lambda\le\xi(r+\lambda)$, rising premium income does not generate a peak through this incidence channel.

This ownership structure separates production capacity from spending incidence. Dealer-parent capital and leverage room determine the Treasury-intermediation capacity N, while ξ determines the share of premium income that returns to domestic demand.

C 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

$$s^R(b,N,\delta)=\frac{1}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{\min\{q(b,N,\delta),1\}}\Gamma\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)d\alpha, \tag{81}$$

where

$$q(b,N,\delta)=\frac{\mu(\delta)b}{N}. \tag{82}$$

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 (81) gives (15). The derivation first defines active insurers' state-contingent assessment, then prices it 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).

C.1 Assessment in a rollover-stress state

The government has debt stock b. Over the unit-length insurance episode 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 insurers to one. A fraction $\underline{\alpha} \in (0,1)$ is inelastic and always participates. Among the remaining insurers, the participating share is $\tilde{\alpha} \in [0,1]$, so the active mass is

$$\alpha = \underline{\alpha} + (1 - \underline{\alpha})\tilde{\alpha} \in [\underline{\alpha}, 1]. \tag{83}$$

With aggregate financial-system capacity N, the active insurers can absorb αN. A rollover-stress shortfall arises when $\mu(\delta)b > \alpha N$, or equivalently when $\alpha < q(b,N,\delta)$. Each active insurer is assigned $\max\{\mu(\delta)b - \alpha N, 0\}/\alpha$. The support schedule maps this exposure into the flow assessment per unit of insured debt required to preserve government debt's safety:

$$\Gamma(0) = 0, \quad \Gamma'(z) > 0 \text{ for } z > 0, \quad \Gamma'_+(0) > 0, \quad \Gamma''(z) \geq 0. \tag{84}$$

The schedule Γ gives the compensation generated by each active insurer's assigned exposure.

The state-contingent assessment is

$$a(b,N,\alpha) = \Gamma\left(\frac{\max\{\mu(\delta)b - \alpha N, 0\}}{\alpha}\right). \tag{85}$$

For $\alpha < q(b,N,\delta)$, the assessment is $\Gamma((\mu(\delta)b - \alpha N)/\alpha)$; otherwise it is zero. Low participation raises exposure per active insurer. All premium and assessment objects are flow spreads per unit of insured debt; multiplying a premium by the debt stock gives aggregate premium payments.

C.2 The participation game

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

$$x_i = \theta + \eta \varepsilon_i, \tag{86}$$

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 insurers.

Each elastic insurer chooses

$$e_i \in \{0,1\}, \tag{87}$$

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

$$\Pi_i(e_i, \alpha; b, N) = e_i \left[s^R - a(b,N,\alpha)\right]. \tag{88}$$

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

$$e_i(x_i) = 1 \quad \Longleftrightarrow \quad x_i \leq x^*. \tag{89}$$

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

$$\tilde{\alpha}(\theta; x^*) = F\left(\frac{x^* - \theta}{\eta}\right), \tag{90}$$

and total participation is

$$\alpha(\theta; x^*) = \underline{\alpha} + (1 - \underline{\alpha}) F\left(\frac{x^* - \theta}{\eta}\right). \tag{91}$$

C.3 Threshold pricing and price invariance

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

$$V(x_i; x^*, s^R, b, N) = s^R - \mathbb{E}\left[a\big(b, N, \alpha(\theta; x^*)\big) \mid x_i\right]. \tag{92}$$

The marginal insurer is indifferent when

$$V(x^*; x^*, s^R, b, N) = 0, \tag{93}$$

or equivalently

$$s^R = \mathbb{E}\left[a\big(b, N, \alpha(\theta; x^*)\big) \mid x_i = x^*\right]. \tag{94}$$

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

$$f\left(\frac{x^* - \theta}{\eta}\right). \tag{95}$$

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

$$z = F\left(\frac{x^* - \theta}{\eta}\right). \tag{96}$$

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

$$\frac{dz}{d\theta} = -\frac{1}{\eta} f\left(\frac{x^* - \theta}{\eta}\right). \tag{97}$$

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

$$\alpha = \underline{\alpha} + (1-\underline{\alpha})z, \tag{98}$$

α is uniform on $[\underline{\alpha},1]$, independently of x^*. The zero-profit premium is therefore the uniform average of the state-contingent assessment, as in (81); setting $\mu(\delta)=1$ gives (14). 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 (85) into the uniform-average pricing equation gives

$$s^R(b,N,\delta) = \frac{1}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{1}\Gamma\left(\frac{\max\{\mu(\delta)b-\alpha N,0\}}{\alpha}\right)d\alpha \tag{99}$$

$$= \frac{1}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{\min\{q(b,N,\delta),1\}}\Gamma\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)d\alpha, \tag{100}$$

This is the finite-rollover-intensity counterpart of (15).

C.4 Interpretation

The resulting $s^R(b,N,\delta)$ is the marginal insurer's zero-profit compensation for $a(b,N,\alpha)$. Low participation concentrates a given rollover shortfall on fewer balance sheets, and the marginal insurer averages their assessments 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.

D Comparative Statics of the Rollover-Insurance Premium

To obtain the comparative statics, differentiate the direct rollover premium $s^R(b,N,\delta)$ derived in Appendix C. 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 insurer 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:

$$s_b^R(b,N,\delta) = \frac{\mu(\delta)}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)}\frac{1}{\alpha}\Gamma'\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)d\alpha > 0, \tag{101}$$

$$s_N^R(b,N,\delta) = -\frac{1}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)}\Gamma'\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)d\alpha < 0. \tag{102}$$

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

The fiscal cost of the direct rollover component is

$$S^R(b,N,\delta) = bs^R(b,N,\delta), \tag{103}$$

so

$$S^R_b(b,N,\delta) = s^R(b,N,\delta) + bs^R_b(b,N,\delta). \tag{104}$$

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

The curvature result follows from differentiating (101). The interior derivative is

$$\frac{\partial}{\partial b}\left[\frac{\mu(\delta)}{\alpha}\Gamma'\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)\right] = \frac{\mu(\delta)^2}{\alpha^2}\Gamma''\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right), \tag{105}$$

and the term from the changing upper integration limit is

$$\frac{\mu(\delta)}{1-\underline{\alpha}}q_b(b,N,\delta)\frac{\Gamma'_+(0)}{q(b,N,\delta)}. \tag{106}$$

Hence

$$\begin{aligned} s^R_{bb}(b,N,\delta) = \frac{1}{1-\underline{\alpha}}\Bigg[\mu(\delta)^2\int_{\underline{\alpha}}^{q(b,N,\delta)}\frac{1}{\alpha^2}\Gamma''\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)d\alpha \\ +\mu(\delta)q_b(b,N,\delta)\frac{\Gamma'_+(0)}{q(b,N,\delta)}\Bigg] \geq 0. \end{aligned} \tag{107}$$

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

$$S^R_{bb}(b,N,\delta) = 2s^R_b(b,N,\delta) + bs^R_{bb}(b,N,\delta), \tag{108}$$

it follows that

$$S^R_{bb}(b,N,\delta) > 0. \tag{109}$$

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

$$S^R_{bN}(b,N,\delta) = s^R_N(b,N,\delta) + bs^R_{bN}(b,N,\delta). \tag{110}$$

Differentiating (101) with respect to N, the interior derivative is

$$\frac{\partial}{\partial N}\left[\frac{\mu(\delta)}{\alpha}\Gamma'\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)\right] = -\frac{\mu(\delta)}{\alpha}\Gamma''\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right), \tag{111}$$

and the term from the changing upper integration limit is

$$\frac{\mu(\delta)}{1-\underline{\alpha}}q_N(b,N,\delta)\frac{\Gamma'_+(0)}{q(b,N,\delta)}. \tag{112}$$

Because

$$q_N(b,N,\delta) = -\frac{1}{N}q(b,N,\delta), \tag{113}$$

this gives

$$s^R_{bN}(b,N,\delta) = -\frac{\mu(\delta)}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{1}{\alpha}\Gamma''\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right) d\alpha$$
$$-\frac{\mu(\delta)}{(1-\underline{\alpha})N}\Gamma'_+(0) < 0. \tag{114}$$

Together with $s^R_N < 0$, this gives

$$S^R_{bN}(b,N,\delta) < 0. \tag{115}$$

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

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

$$s^R_\delta(b,N,\delta) = \frac{e^{-\delta}b}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{1}{\alpha}\Gamma'\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right) d\alpha > 0, \tag{116}$$

and

$$s^R_{b\delta}(b,N,\delta) = \frac{e^{-\delta}}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{1}{\alpha}\Gamma'\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right) d\alpha$$
$$+\frac{\mu(\delta)e^{-\delta}b}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{1}{\alpha^2}\Gamma''\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right) d\alpha + \frac{e^{-\delta}\Gamma'_+(0)}{1-\underline{\alpha}}. \tag{117}$$

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

$$S^R_{b\delta}(b,N,\delta) = s^R_\delta(b,N,\delta) + bs^R_{b\delta}(b,N,\delta) > 0. \tag{118}$$

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.

E 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 D.

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

$$\Gamma(z) = \varphi z, \qquad \varphi > 0. \tag{119}$$

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:

$$\widehat{b}(N) = \frac{\underline{\alpha}N}{\mu(\delta)}. \tag{120}$$

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

$$\widehat{b}(N) < b < \frac{N}{\mu(\delta)}, \tag{121}$$

the interior region satisfies $\underline{\alpha} < q(b,N,\delta) < 1$. Substituting the linear stress-cost schedule into (100) gives

$$\begin{aligned}
s^R(b,N,\delta) &= \frac{\varphi}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{\mu(\delta)b-\alpha N}{\alpha}\, d\alpha \\
&= \frac{\varphi}{1-\underline{\alpha}} \left[\mu(\delta)b \int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{1}{\alpha}\, d\alpha - N \int_{\underline{\alpha}}^{q(b,N,\delta)} 1\, d\alpha \right] \\
&= \frac{\varphi}{1-\underline{\alpha}} \left[\mu(\delta)b \log\left(\frac{\mu(\delta)b}{\underline{\alpha}N}\right) - \mu(\delta)b + \underline{\alpha}N \right].
\end{aligned} \tag{122}$$

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

$$\begin{aligned}
s_b^R(b,N,\delta) &= \frac{\varphi}{1-\underline{\alpha}} \left[\mu(\delta) \log\left(\frac{\mu(\delta)b}{\underline{\alpha}N}\right) + \mu(\delta) - \mu(\delta) \right] \\
&= \frac{\varphi\mu(\delta)}{1-\underline{\alpha}} \log\left(\frac{\mu(\delta)b}{\underline{\alpha}N}\right) > 0.
\end{aligned} \tag{123}$$

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

$$s_{bb}^R(b,N,\delta) = \frac{\varphi\mu(\delta)}{(1-\underline{\alpha})b} > 0. \tag{124}$$

The derivative with respect to financial capacity is

$$s_N^R(b,N,\delta) = -\frac{\varphi}{1-\underline{\alpha}} \left[q(b,N,\delta) - \underline{\alpha}\right] < 0. \tag{125}$$

The marginal safety-production cost is

$$\begin{aligned}
S_b^R(b,N,\delta) &= s^R(b,N,\delta) + b s_b^R(b,N,\delta) \\
&= \frac{\varphi}{1-\underline{\alpha}} \left[\mu(\delta)b \log\left(\frac{\mu(\delta)b}{\underline{\alpha}N}\right) - \mu(\delta)b + \underline{\alpha}N \right] + \frac{\varphi}{1-\underline{\alpha}} \left[\mu(\delta)b \log\left(\frac{\mu(\delta)b}{\underline{\alpha}N}\right) \right] \\
&= \frac{\varphi}{1-\underline{\alpha}} \left[2\mu(\delta)b \log\left(\frac{\mu(\delta)b}{\underline{\alpha}N}\right) - \mu(\delta)b + \underline{\alpha}N \right].
\end{aligned} \tag{126}$$

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

$$S_b^R(b,N,\delta) = \frac{\varphi}{1-\underline{\alpha}} \left[2\mu(\delta)b \log\left(\frac{1}{\underline{\alpha}}\right) - N(1-\underline{\alpha}) \right], \tag{127}$$

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

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

$$S^R_{bN}(b,N,\delta) = \frac{\varphi}{1-\underline{\alpha}}\left[-\frac{2\mu(\delta)b}{N} + \underline{\alpha}\right] < 0. \tag{128}$$

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

$$\lambda = \frac{\varphi}{1-\underline{\alpha}}\left[2b^L \log\left(\frac{b^L}{\underline{\alpha}N}\right) - b^L + \underline{\alpha}N\right]. \tag{129}$$

If the peak is saturated, its right-hand side is (127) at $b = b^L$ and $\mu(\delta) = 1$. Greater N or $\underline{\alpha}$ lowers marginal cost at a given b and pushes the peak outward; greater unit support cost φ pulls it inward. At finite maturity, rollover intensity raises the direct marginal cost in (126); the total Laffer peak must also include the duration-price component from Section 5.

F Duration-Price Amplification and Maturity Choice

The capacity and maturity comparative statics below complete the analysis of the total-premium fixed point in (35). The government then chooses rollover intensity δ; higher δ means shorter maturity, faster rollover, and lower duration.

F.1 Rollover component

The direct rollover component is $s^R(b,N,\delta)$ from (32). Because exposure is $\mu(\delta)b$, $s^R_\delta > 0$ in the stress region. Differentiating with respect to N gives

$$\begin{aligned} s^R_{\delta N}(b,N,\delta) = &-\frac{e^{-\delta}b}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{1}{\alpha}\Gamma''\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)d\alpha \\ &-\frac{e^{-\delta}b}{(1-\underline{\alpha})N}\Gamma'_+(0) < 0. \end{aligned} \tag{130}$$

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

F.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

$$s^D(b,N,\delta) = \frac{\vartheta(N)}{[r+s(b,N,\delta)+\delta]^2}s^R(b,N,\delta). \tag{131}$$

For the capacity dependence, parameterize

$$\vartheta(N) = \bar{\vartheta} N^{-\eta_\vartheta}, \qquad \bar{\vartheta} > 0, \qquad \eta_\vartheta > 0. \tag{132}$$

where $\bar{\vartheta}$ scales duration-price transmission and η_ϑ is its capacity elasticity. Higher capacity dampens the price effect of a given rollover assessment. Combining (131) with the direct rollover premium gives

$$s = s^R(b,N,\delta)\left[1 + \frac{\vartheta(N)}{(r+s+\delta)^2}\right]. \tag{133}$$

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 (36) and (37). Because $s_N^R < 0$ 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 (37) with respect to N yields the full cross-derivative of the premium with respect to rollover intensity and capacity:

$$\begin{aligned} s_{\delta N} = \Bigg\{ &\left[1 + \frac{\vartheta}{(r+s+\delta)^2}\right] s^R_{\delta N} + \left[\frac{\vartheta_N}{(r+s+\delta)^2} - \frac{2\vartheta s_N}{(r+s+\delta)^3}\right] s^R_\delta \\ &- (1+s_\delta)\frac{2s^R\vartheta}{(r+s+\delta)^3}\left[\frac{\vartheta_N}{\vartheta} + \frac{s^R_N}{s^R} - \frac{3s_N}{r+s+\delta}\right]\Bigg\}\left[1 + \frac{2s^R\vartheta}{(r+s+\delta)^3}\right]^{-1}. \end{aligned} \tag{134}$$

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

$$\begin{aligned} s_{\delta N} = \Bigg\{ &\left[1 + \frac{\vartheta}{(r+s+\delta)^2}\right] s^R_{\delta N} + \left[\frac{\vartheta_N}{(r+s+\delta)^2} - \frac{2\vartheta s_N}{(r+s+\delta)^3}\right] s^R_\delta \\ &- \frac{2s^R\vartheta}{(r+s+\delta)^3}\left[\frac{\vartheta_N}{\vartheta} + \frac{s^R_N}{s^R} - \frac{3s_N}{r+s+\delta}\right]\Bigg\}\left[1 + \frac{2s^R\vartheta}{(r+s+\delta)^3}\right]^{-1}. \end{aligned} \tag{135}$$

F.3 Maturity choice

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

$$\delta^*(b,N) \in \arg\min_{\delta > 0} s(b,N,\delta). \tag{136}$$

At an interior optimum,

$$s_\delta(b,N,\delta^*) = 0, \qquad s_{\delta\delta}(b,N,\delta^*) > 0. \tag{137}$$

Using (37), the first-order condition is

$$\left[1 + \frac{\vartheta(N)}{(r+s+\delta)^2}\right] s^R_\delta = \frac{2s^R\vartheta(N)}{(r+s+\delta)^3}. \tag{138}$$

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

Proposition 4 (Maturity choice and financial capacity)**.** *Let $\delta^*(b,N)$ be an interior solution of the maturity problem and suppose the second-order condition in (137) holds. If*

$$s_{\delta N}(b,N,\delta^*) < 0, \tag{139}$$

then

$$\frac{d\delta^*(b,N)}{dN} > 0. \tag{140}$$

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:

$$s_{\delta\delta}(b,N,\delta^*)\frac{d\delta^*}{dN} + s_{\delta N}(b,N,\delta^*) = 0. \tag{141}$$

Consequently,

$$\frac{d\delta^*}{dN} = -\frac{s_{\delta N}}{s_{\delta\delta}}. \tag{142}$$

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.

G Empirical Construction

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

G.1 Data construction

Dealer-parent Treasury capacity comes from FR Y-9C filings and the leverage-ratio mapping in (49). 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 5 lists the data sources used to construct the debt, capacity, premium, Federal Reserve holdings, and dealer-parent equity series.

Table 5: 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

G.2 Cash-Treasury repricing details

The cash-Treasury design is a monthly maturity panel. Equations (46) and (50) 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 (47) 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.

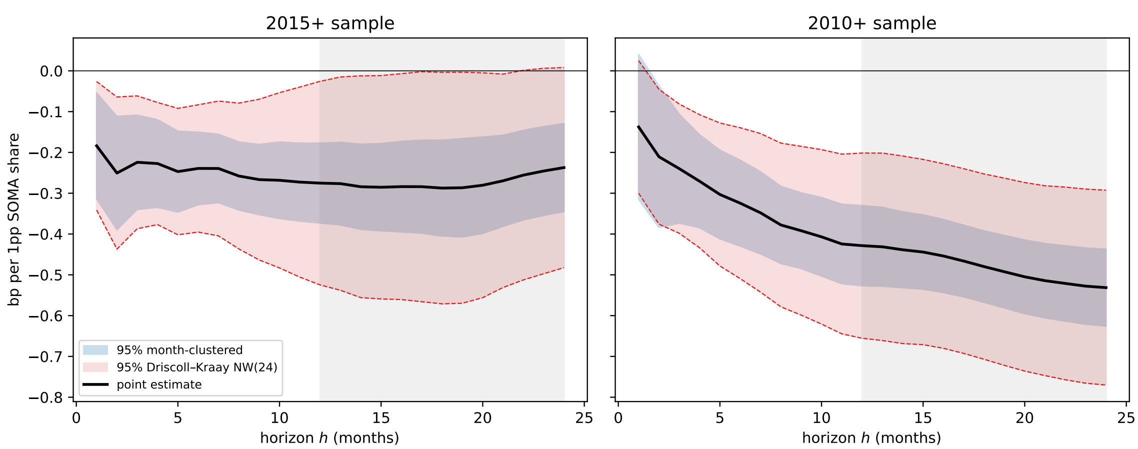

Figure 4: Linear cash-Treasury repricing estimates. The panels report β_h for the 2015+ and 2010+ samples. Black lines are point estimates; blue bands are month-clustered 95 percent confidence intervals; red dashed bands use Driscoll–Kraay/Newey–West inference with 24 monthly lags. The shaded region 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.

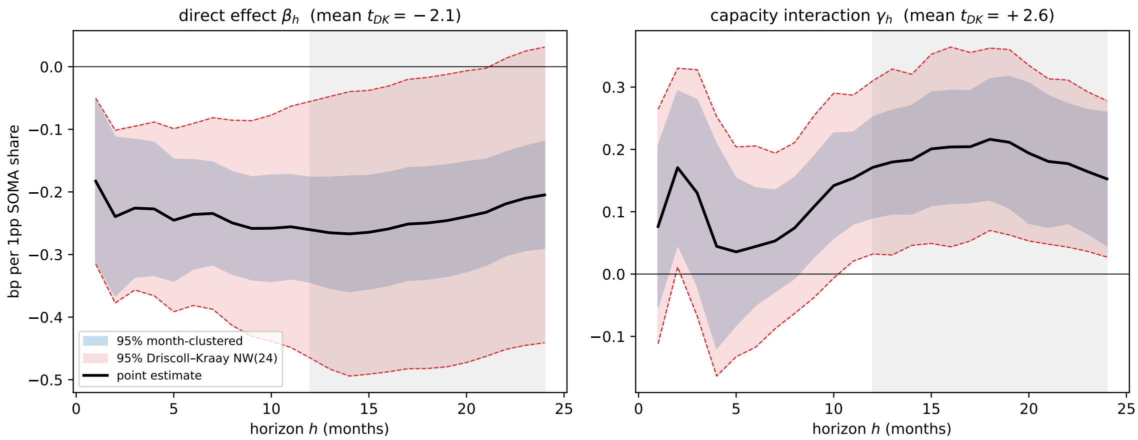

Figure 5: Benchmark nonlinear cash-Treasury repricing estimates. The left panel reports the direct slope β_h, and the right panel reports the capacity interaction γ_h, using the 5 percent leverage mapping. Black lines are point estimates; blue bands are month-clustered 95 percent confidence intervals; 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.

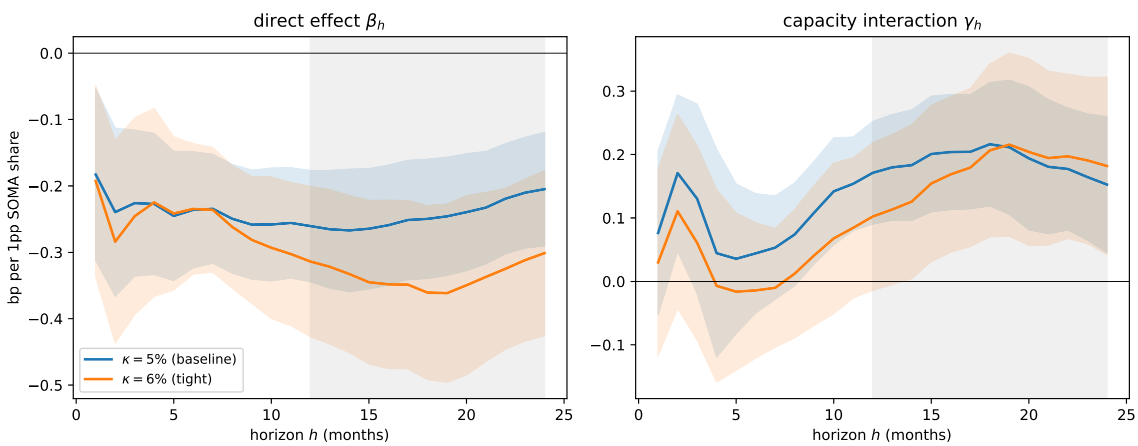

Figure 6: Nonlinear estimates under 5 and 6 percent leverage-ratio mappings. Point estimates and bands are shown for the direct slope β_h and capacity interaction γ_h; the shaded region is the $h = 12$–24 averaging window. The interaction paths are similar, while the standardized 2026Q1 capacity states 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.

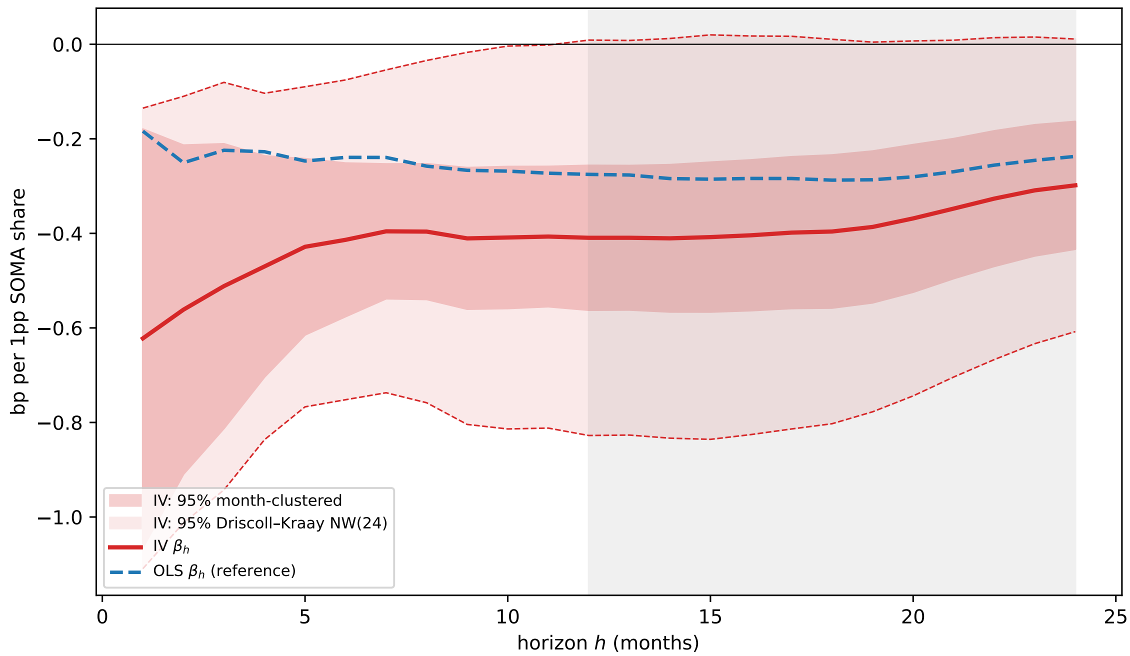

Figure 7: Linear IV sensitivity with overlap-robust inference. The instruments, measured at $t-6$, are lagged maturity-bucket SOMA share and a trailing-issuance on-the-run proxy. The solid red line is the IV estimate and the dashed blue line is the OLS benchmark; the bands show month-clustered and Driscoll–Kraay/Newey–West 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.

G.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 (51).

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 7 reports statistical uncertainty. Table 6 lists the announcements and sizing assumptions.

Table 6: 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 7 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 7: 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.

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.

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.

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.