The Safe-Debt Laffer Curve

Ricardo J. Caballero*

July 13, 2026

Link to the latest version

Abstract

This paper develops a model in which safe public debt can become a drag on aggregate activity even when the sovereign is solvent and its debt remains safe. Solvency rules out default, but it does not make safety costless: safe-asset services require rollover support, market-making, and balance-sheet capacity. As debt rises, safe-asset services become more costly to produce, and their marginal cost can overtake the wealth effect of additional issuance. The result is an aggregate-demand safe-debt Laffer curve: beyond the peak, additional debt remains safe but reduces demand and lowers the equilibrium safe rate. I measure the gap between the wealth effect of additional issuance and the stock-wide marginal cost of producing safe-asset services—the safe-debt margin. For the United States, this margin falls from about positive 280 basis points in 2015 to a 2026Q1 range of positive 39 to *negative* 58 basis points, placing the economy at the peak or on the contractionary side of the safe-debt Laffer curve.

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 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 stands at levels rarely seen outside major wars, and it continues to rise. A debt burden this large is widely expected to weigh on the economy eventually—to become, beyond some point, a substantial drag on aggregate activity. The most familiar version of this concern is about fiscal stability: whether fiscal positions are sustainable, and whether mounting debt ultimately ends in chaotic implicit or explicit default. This paper sets that version aside. Suppose the sovereign is solvent and its debt is, and remains, safe. Can high public debt still become a drag on demand? The canonical aggregate-demand model says no—and says so at any debt level.

The logic is familiar. Government debt is the economy's reference safe asset: the benchmark store of value and the collateral behind much private contracting. Additional issuance raises safe financial wealth, but it also raises the taxes needed to service the debt. With finite-horizon households, the wealth effect is only partly offset: current households do not fully capitalize the future tax burden. I call this unoffset component the non-Ricardian wedge. Since this wedge is present at every debt level, more safe debt is always expansionary, never a drag. The safe-asset-shortage literature goes further: with chronic excess demand for stores of value, abundant safe public debt is potentially beneficial rather than problematic.

This paper shows that once safe-asset services are produced rather than assumed, the canonical conclusion holds only up to a point: at high debt levels, additional safe public debt becomes a drag on aggregate demand. The mechanism is the rising cost of producing the services that give government debt its safe-asset role: market-making, warehousing, and rollover support when refinancing conditions deteriorate. The fiscal authority supplies the promise, while the financial system supplies the balance-sheet capacity that keeps outstanding claims liquid and safe. The excess yield the government pays for this service is a balance-sheet premium—the price of scarce intermediary capacity. As the debt stock rises, rollover pressure rises relative to financial capacity. Because the premium applies to the stock of outstanding claims, marginal issuance raises the support cost on the new unit and also raises the cost of supporting the existing stock. The safe-debt Laffer peak is reached when this stock-wide marginal cost exhausts the non-Ricardian wedge.

The model starts with a macro block. Output is fixed and aggregate demand determines the equilibrium safe rate. In this macro block, the safe rate is the economy-wide real risk-free rate that clears aggregate demand. The government's all-in cost of issuing and maintaining safe debt also includes the premium required to produce safe-asset services. Households live in a perpetual-youth economy, so public debt is not fully offset by the present value of future taxes. Additional safe debt therefore has a direct non-Ricardian demand benefit. The government also raises taxes to finance purchases, debt service, and the premium paid to produce the safe-asset services of public debt. The macro block keeps these two forces separate: the wealth effect of additional safe debt raises demand, while the fiscal cost of producing safe-asset services lowers demand.

The financial block determines the safety-production premium. In the benchmark, public debt is short-term, so the entire stock must be refinanced and kept safe through rollover support. Intermediaries provide that support using scarce balance-sheet capacity. When rollover exposure is small relative to financial capacity, the required premium is low. When rollover exposure presses against capacity, intermediaries require a higher premium to warehouse and support the debt stock in stress states. The premium therefore rises with debt pressure. Because the premium is paid on the outstanding stock, the cost of additional issuance is stock-wide.

Combining the macro and financial blocks yields the aggregate-demand safe-debt Laffer curve. At low debt levels, the non-Ricardian demand benefit dominates and additional safe debt is expansionary. At high debt levels, the stock-wide marginal premium cost of producing safe-asset services can dominate. The debt remains safe, but the marginal contribution of additional safe debt to aggregate demand turns negative. The same margin also determines the equilibrium safe-rate response: before the peak, more safe debt raises the safe rate; after the peak, more safe debt lowers it. Financial capacity shifts the peak outward by reducing the marginal cost of producing safe-asset services.

The positive-maturity extension then connects the model to the empirical measurement. The benchmark treats debt as short-term, so the only cost is rollover support. With longer maturities, additional issuance also changes the price of outstanding long bonds and the real term premium required to hold Treasury duration. The model therefore separates the total premium into a rollover or balance-sheet component and a duration component, the two premium margins measured in the empirical section.

The empirical part of the paper measures the two marginal premium effects that determine where the United States lies on the safe-debt Laffer curve. The key state variable is rollover exposure relative to dealer capacity: the stock of Treasury debt that must be refinanced in the near term, divided by a measure of dealer-parent balance-sheet capacity that I construct from regulatory capital and the balance-sheet room left under leverage limits. This ratio rises from about 0.5 in 2015 to 2.59 in 2026Q1. What makes additional debt contractionary is not a sharp rise in premium levels themselves—the rollover premium is only about 29 basis points in 2026Q1, and the baseline real-duration premium about 39—but their stock-wide repricing. An additional unit of debt raises the relevant premium on a much larger stock, so the marginal cost of keeping the debt safe is large even though premium levels remain moderate, and the rollover premium response steepens further as capacity tightens during 2025–26.

For the rollover or balance-sheet premium, I use the local-supply design of D'Amico and King (2013), which isolates the effect of Treasury supply from the many other forces that move with it: within a month, the Federal Reserve absorbs different fractions of the Treasury float in different maturity buckets, so month fixed effects remove aggregate news and curve-wide factors while maturity fixed effects remove persistent tenor differences. The estimated local-supply slope is then converted into the marginal cost of repricing the entire outstanding stock—the stock-wide rollover term used in the safe-debt margin.

For the real-duration premium—the second premium component, the compensation for holding Treasury duration—I use Federal Reserve asset-purchase and balance-sheet-runoff announcements together with the real term-premium measure of D'Amico, Kim, and Wei (2018), which jointly prices nominal Treasury securities and inflation-indexed Treasuries. The announcement response is converted into the corresponding stock-wide real-duration repricing term, so that the two premium responses are expressed on a common basis-point scale and can be added.

Combining the identified premium responses gives the safe-debt margin: the gap, in basis points, between the non-Ricardian demand benefit and the stock-wide marginal cost of producing safe-asset services, positive when additional debt is still expansionary and negative once it is contractionary. This margin falls from roughly positive 280 basis points in 2015 to a 2026Q1 range of positive 39 to *negative* 58 basis points, indicating that the economy is at or past the peak. The deterioration is not that dealer parents have become smaller. It reflects shrinking leverage room: rising leverage exposure presses against roughly flat Tier-1 capital, crowding out the balance-sheet space available for Treasury warehousing even as dealer parents remain well capitalized.

Literature Review. This paper is closely related to the safe-asset-shortage literature. Caballero, Farhi, and Gourinchas (2008, 2016, 2017, 2021) and Caballero and Farhi (2018) study the macroeconomic implications of scarce safe assets, including low safe rates, global imbalances, and recessionary pressure. Krishnamurthy and Vissing-Jorgensen (2012) document safety and liquidity premia on Treasury debt. Greenwood, Hanson, and Stein (2015) and D'Avernas and Vandeweyer (2024) study the supply and maturity structure of short safe claims. Maggiori (2017) and Farhi and Maggiori (2018) emphasize the international and intermediary-capacity dimensions of reserve-currency asset supply. Itskhoki and Mukhin (2026) provides a recent synthesis of the global-imbalances branch of this literature, and Brunnermeier, Merkel, and Sannikov (2024) a recent incomplete-markets treatment of safe-asset service flows. This paper turns that shortage logic around by asking when the marginal provision of safe public debt stops expanding aggregate demand.

A complementary literature places the safety and liquidity services of government bonds directly in preferences or liquidity constraints, generating a downward-sloping demand curve for Treasury debt; Krishnamurthy and Vissing-Jorgensen (2012) estimate it, with convenience yields falling as supply rises. In fiscal applications, the resulting Laffer curve is a budget object: issuance dilutes the premium on the outstanding stock, so revenue or sustainable debt peaks at an interior debt level (Brunnermeier, Merkel, and Sannikov, 2024; Mian, Straub, and Sufi, 2025; Li and Merkel, 2026). The curve developed in this paper shares the same stock-repricing logic but characterizes a different object: not the fiscal value or sustainable quantity of debt, but the government's marginal contribution to aggregate demand, which turns negative while the debt remains safe.

The paper also relates to the broader asset-shortage view. Caballero (2006, 2015) argue that low real rates, high asset valuations, bubbles, and macroeconomic fragility can arise when demand for stores of value exceeds the economy's capacity to produce financial assets. Auclert, Malmberg, Rognlie, and Straub (2025) provide a recent asset-supply and asset-demand framework for studying U.S. wealth, real interest rates, and fiscal sustainability. The present paper focuses on the safe-asset branch of this broader logic: public debt is a store of value only insofar as the fiscal and financial systems can produce its safety services.

The demand block builds on the non-Ricardian effects of government debt. Barro (1974) gives the Ricardian benchmark; Blanchard (1985) and Weil (1989) provide finite-horizon environments in which government debt is private wealth; and Mian, Straub, and Sufi (2021, 2025) study indebted demand and non-Ricardian fiscal deficits. Relatedly, Angeletos, Lian, and Wolf (2024) study whether deficit-financed expansions partly finance themselves through aggregate-demand effects, and Fornaro and Wolf (2025) study a feedback between public debt, fiscal distortions, and productivity growth. Relative to this literature, the paper adds the production side of safe debt:

households value insured public debt as safe wealth, but producing the insured claim consumes financial-system capacity and fiscal resources, generating the aggregate-demand safe-debt Laffer curve. The declining branch also relates to classic work on fiscal adjustment and aggregate demand (Giavazzi and Pagano, 1990; Alesina and Perotti, 1997), though the mechanism here is a steady-state one, in which fiscal adjustment finances the ongoing production of safe claims.

A separate empirical literature estimates the effect of debt and fiscal positions on long-term government yields: Laubach (2009), Gruber and Kamin (2012), and Gamber and Seliski (2019) find positive effects on advanced-economy long rates, and World Bank (2026) applies this to emerging-market borrowing costs. Recent high-frequency work identifies fiscal-news and supply effects on Treasury pricing through term premia, inflation, and convenience yields (Gómez-Cram, Kung, and Lustig, 2024, 2025; Jiang, Richmond, and Zhang, 2025). The empirical object in this paper is complementary: the marginal cost of producing safe Treasury claims and the resulting safe-rate Laffer condition.

The macro-finance channel relates to risk-centric models in which asset prices, financial conditions, and risk-bearing capacity shape aggregate demand. Caballero and Simsek (2020, 2021, 2023) develop models in which risk premia and asset prices affect demand and policy transmission, and Caballero, Caravello, and Simsek (2024) studies financial-conditions targeting. This paper studies the safe-asset services of public debt and treats the financial system's risk-bearing capacity as a macroeconomic state variable.

The Treasury-market intermediation literature gives institutional content to the financial-capacity block. Kashyap, Stein, Wallen, and Younger (2025) study a Treasury market with asset managers, hedge funds, and broker-dealers in which limited dealer balance-sheet capacity can generate market dysfunction as Treasury supply rises. Related work on the March 2020 episode and Treasury-market resilience includes 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). This literature provides empirical counterparts for the model's capacity variable and premium schedules; the contribution here is to embed that capacity margin in an aggregate-demand safe-debt Laffer curve.

The production block relates to sovereign rollover risk and safe-asset determination. Calvo (1988) and Cole and Kehoe (2000) develop self-fulfilling rollover-crisis mechanisms. Eaton and Gersovitz (1981), Arellano (2008), and Aguiar and Amador (2014) study sovereign borrowing with endogenous spreads and repayment incentives. He, Krishnamurthy, and Milbradt (2019) is especially close because the safe-asset status of debt is endogenous and debt-size dependent; their analysis centers on asset-market determination of safety, whereas here the safety-production cost enters the government budget constraint, aggregate demand, and the equilibrium safe rate.

The maturity extension result relates to theories of government debt maturity and financial intermediation. Angeletos (2002), Buera and Nicolini (2004), and Faraglia, Marcet, Oikonomou, and Scott (2019) study the maturity structure of government debt as a tool for managing fiscal risk and hedging shocks. Greenwood, Hanson, and Stein (2015) emphasize comparative advantage in the supply of short safe debt. Section 4 shows how positive maturity turns rollover risk into a price component through a local balance-sheet price-feedback multiplier; Appendix H uses the same premium schedule to study the maturity tradeoff between rollover exposure and term-premium exposure.

The remainder of the paper is organized as follows. Section 2 builds the model blocks: perpetual-youth demand, fiscal closure, the aggregate-demand schedule, and the rollover-insurance cost schedule. Section 3 combines these blocks, derives the aggregate-demand safe-debt Laffer curve, and studies how financial capacity shifts the peak. Section 4 derives the positive-maturity price and term-premium components used in the empirical margin. Section 5 identifies the U.S. safe-debt margin. Section 6 concludes. The appendices collect the empirical construction, the household and insurance derivation, the threshold-participation foundation for the rollover-insurance premium, the premium comparative statics, the closed-form illustration, the balance-sheet amplification block, and the maturity-choice extension.

2 The Model Blocks

This section builds the two model blocks used in the main result. The macro block maps debt and a safety-production premium into aggregate demand at a given safe rate. The financial block derives that premium from rollover-insurance pricing. Section 3 then combines the two blocks and imposes goods-market clearing.

2.1 Perpetual-Youth Demand

Time is continuous and output is normalized to one. The stock of outstanding public debt is denoted by b. Households hold insured government-debt claims and receive the safe return r. Government purchases are fixed throughout the benchmark; with output normalized to one, $\bar{g}$ denotes the constant government-purchases share.

Households are of the perpetual-youth type. Each household dies with Poisson intensity $\lambda > 0$, and newborn households replace dying households one-for-one. Actuarially fair annuity markets redistribute the assets of dying households to survivors as mortality credits; newborns enter without inheriting the assets of the deceased. The parameter λ governs the strength of non-Ricardian behavior: because households may die before taxes are levied, they internalize only part of future fiscal adjustments. Households discount utility at rate ρ. With logarithmic flow utility in this perpetual-youth economy, households consume a constant share of total wealth,

$$A \equiv \rho + \lambda. \tag{1}$$

Here ρ is the rate of impatience, while λ is the mortality component that makes wealth more valuable for current consumption.

Let χ denote the after-tax flow received by households. Human wealth is the present value of this flow at the mortality-adjusted discount rate,

$$h = \frac{\chi}{r + \lambda}. \tag{2}$$

Aggregate consumption is

$$c = A\big(b + h\big), \tag{3}$$

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

Appendix C derives (2)–(3) from the household problem. It starts from Epstein–Zin preferences (Epstein and Zin, 1989) with unit intertemporal elasticity and risk aversion σ, lets households choose government-bond holdings and rollover-insurance purchases, and records the finite-risk-aversion insurance conditions. The main text uses the full-insurance limit, $\sigma \to \infty$: government bonds pay the government yield, rollover insurance absorbs the risky component, and the net return on insured debt is the safe rate r. The insurance premium is covered by the excess yield paid by the government.

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 attached to the outstanding public-debt stock. In this subsection, S is an input into the fiscal block; the financial block below derives the schedule $S(b, N) = s(b, N)b$ from rollover-insurance pricing and intermediary balance-sheet capacity.

Under the benchmark incidence convention, taxes finance this premium inside the domestic demand block, while the corresponding intermediary income is not recycled into that block. The tax share adjusts across steady states to finance government purchases, safe-rate service on outstanding debt, and this safety-production premium. The fiscal closure is

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

The after-tax flow is therefore

$$\chi(b, S; r) = 1 - \tau(b, S; r) = 1 - \bar{g} - rb - S. \tag{5}$$

Substituting this fiscal closure into the household block gives

$$h(b, S; r) = \frac{1 - \bar{g} - rb - S}{r + \lambda} \tag{6}$$

and

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

Thus aggregate demand at safe rate r, taking the safety-production premium S as an input, is

$$\mathcal{D}(b, S; r) \equiv A\left(b + \frac{1 - \bar{g} - rb - S}{r + \lambda}\right) + \bar{g}. \tag{8}$$

At a fixed safe rate, an additional unit of safe public debt raises safe financial wealth directly. It also raises taxes, both through safe-rate debt service and through the cost of producing safe-asset services. If safe-asset services were costless, finite horizons would make the direct wealth effect dominate the capitalized tax burden. Even a constant per-unit premium would lower the demand contribution of debt, by s per unit. But a constant premium cannot generate a peak: the turn requires the stock-repricing term bs_b, which is why the premium must rise with debt.

The financial block below derives the safety-production premium S from rollover-insurance pricing. Section 3 then combines this safety-premium schedule with (8) and imposes goods-market clearing.

2.3 Rollover Exposure and Financial Capacity

The financial block derives the safety-production premium $S(b, N)$ from rollover-insurance pricing. Financial-system capacity $N > 0$ measures the maximum rollover exposure the potential rollover-insurance pool can absorb without stress. It is not a hard safety limit. When rollover needs exceed active capacity, the debt can remain safe, but producing its safe-asset services requires compensation for warehousing, market-making, and balance-sheet support.

Debt is refinanced gradually. Each unit of debt matures according to a Poisson process with hazard $\delta > 0$. The parameter δ is the rollover intensity, so average maturity is $1/\delta$. Over one model period—the rollover-insurance episode over which the premium is priced—the rollover share of the outstanding stock is

$$\mu(\delta) \equiv 1 - e^{-\delta}. \tag{9}$$

Thus the stock of debt exposed to rollover during the period is $\mu(\delta)b$. The benchmark below takes the short-maturity limit, $\delta \to \infty$, so $\mu(\delta) \to 1$ and the entire debt stock is exposed. Section 4 studies finite rollover intensity, $\delta < \infty$, and the positive-maturity price component of rollover risk.

2.4 The Rollover-Insurance Premium

Let $r^G(b, N)$ denote the gross yield paid by the government on insured public debt. Households hold the insured safe claim and receive the safe return r. The financial system absorbs the rollover risk embedded in the government bond and is compensated by the yield spread

$$s(b, N) \equiv r^G(b, N) - r. \tag{10}$$

The government pays $r + s(b, N)$, households receive the safe return r, and the payment

$$S(b, N) = s(b, N)b \tag{11}$$

compensates specialized financial balance sheets for providing rollover insurance.

The benchmark incidence assumption is that fiscal resources used to produce safe-asset services are borne inside the domestic non-Ricardian demand block, while the premium income paid to the intermediary sector is not recycled into that block.[1] Under this incidence, the support premium is a fiscal-financial cost of safe-debt production: it reduces the domestic demand contribution of public debt even though it is income to intermediary owners.

The key economic point is stock-wide pricing. The premium is a single market price paid on the entire insured stock, so an additional unit of debt raises the cost of the marginal unit and raises the cost of supporting the outstanding stock. The section derives the premium $s(b, N)$ and the marginal fiscal cost of producing safe-asset services, $S_b = s + bs_b$, in the benchmark short-maturity limit.

Let $\alpha \in [\underline{\alpha}, 1]$ denote the active fraction of potential absorption capacity. The state $\alpha = 1$ is full participation by the rollover-insurance pool; lower values describe stress states in which

[1] The tractable intermediary treatment follows the spirit of Gabaix and Maggiori (2015): intermediary balance-sheet capacity is the scarce input, while ownership and spending incidence are summarized by a demand-block convention. Appendix A makes the incidence convention explicit.

fewer balance sheets are willing or able to absorb rollover exposure. The lower bound $\underline{\alpha} \in (0, 1)$ is the inelastic investor base that is always present in the market. Active absorption capacity in participation state α is

$$B(\alpha, N) = \alpha N. \tag{12}$$

In the benchmark limit $\mu(\delta) = 1$, the debt stock exposed to rollover is b. The rollover shortfall in participation state α is

$$X(b, N, \alpha) = \max\{b - \alpha N, 0\}. \tag{13}$$

The aggregate shortfall is shared across active absorption capacity. The resulting per-active-unit stress is

$$\zeta(b, N, \alpha) = \frac{X(b, N, \alpha)}{\alpha}. \tag{14}$$

The support schedule Γ maps per-active-unit stress into the state-contingent assessment required to produce safe-asset services, where

$$\Gamma(0) = 0, \qquad \Gamma'(z) > 0 \text{ for } z > 0, \qquad \Gamma'_{+}(0) > 0, \qquad \Gamma''(z) \geq 0. \tag{15}$$

The assessment represents the compensation required to draw scarce balance-sheet capacity into warehousing, market-making, and supporting the debt stock in stress states. The schedule is increasing because larger individual stress requires more compensation. It is convex because additional stress is borne by progressively more constrained balance sheets: leverage limits bind more tightly and price impact rises. This individual-stress convexity is the theoretical counterpart of the capacity-interacted premium responses measured below.

The per-active-unit assessment is

$$a(b, N, \alpha) = \Gamma\big(\zeta(b, N, \alpha)\big). \tag{16}$$

Low participation states are costly because a given aggregate shortfall is borne by fewer active units, raising per-active-unit stress.

It is useful to summarize rollover pressure by the participation threshold required to absorb the exposed debt stock,

$$q(b, N) = \frac{b}{N}. \tag{17}$$

The threshold q indexes the stress states: $X(b, N, \alpha) > 0$ exactly when $\alpha < q(b, N)$. If $q(b, N) \leq \underline{\alpha}$, the inelastic investor base is sufficient, and no rollover-insurance premium is required. If $q(b, N) > \underline{\alpha}$, some participation states have too little active capacity to absorb the rollover exposure. When $q(b, N) > 1$, even full participation leaves a rollover shortfall, so every participation state contributes to the premium. In the benchmark short-maturity case, the stress threshold at which active rollover support first becomes required is $\widehat{b}(N) = \underline{\alpha} N$. Let $\bar{q}(b, N) = \min\{q(b, N), 1\}$.

Appendix D gives a threshold-participation foundation for the pricing measure. The main text uses its pricing implication. The marginal insurer prices the participation states in which its entry is pivotal; with a uniform participation state on $[\underline{\alpha}, 1]$, the competitive premium is the unweighted average of the per-active-unit assessment, which is zero outside the stress region. The per-unit rollover-insurance premium is therefore

$$s(b,N) = \frac{1}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{\bar{q}(b,N)} \Gamma\left(\frac{b-\alpha N}{\alpha}\right) d\alpha, \tag{18}$$

with $s(b,N) = 0$ when $q(b,N) \leq \underline{\alpha}$. In the interior stress region $\underline{\alpha} < q(b,N) < 1$, differentiating gives

$$s_b(b,N) = \frac{1}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{q(b,N)} \frac{1}{\alpha}\Gamma'\left(\frac{b-\alpha N}{\alpha}\right) d\alpha > 0, \tag{19}$$

while

$$s_N(b,N) = -\frac{1}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{q(b,N)} \Gamma'\left(\frac{b-\alpha N}{\alpha}\right) d\alpha < 0. \tag{20}$$

The upper-boundary term in the differentiation is zero because per-active-unit stress is zero at $\alpha = q(b,N)$. Higher debt raises the premium by increasing individual balance-sheet stress in each already-stressed participation state; since $\Gamma' > 0$, the per-active-unit assessment rises state by state. Higher financial capacity lowers the premium because it reduces stress in every participation state. Appendix E derives these formulas in the general finite-δ case; equations (18)–(20) are the benchmark limit $\mu(\delta) = 1$.

2.5 The Stock-Wide Safety Cost

In the benchmark stress region, the same forces make the stock-wide marginal cost convex. As debt rises, it raises individual balance-sheet stress in already-stressed states and brings additional participation states into stress; because Γ is convex, each additional unit of individual stress is more expensive than the previous one. The marginal cost of producing safe-asset services therefore steepens with the debt stock, and Appendix E gives the curvature condition $S_{bb} > 0$ used below.

The total fiscal cost of producing safe-asset services is $S(b,N) = s(b,N)b$, as in (11). Its marginal cost is

$$S_b(b,N) = s(b,N) + bs_b(b,N). \tag{21}$$

This decomposition is the central financial mechanism. The first term, $s(b,N)$, is the safety-production cost of the marginal unit of debt. The second term, $bs_b(b,N)$, is the repricing of the outstanding stock. The premium is a common market-clearing insurance price for the entire stock, and the government pays it through an excess yield on all insured debt. The marginal unit therefore makes the outstanding stock more expensive to support. This is the financial input used below: the safe-debt Laffer peak occurs when the marginal stock-wide cost of producing safe-asset services reaches the non-Ricardian wedge.

3 Equilibrium and the Safe-Debt Laffer Curve

This section combines the model blocks, derives the aggregate-demand safe-debt Laffer curve, and studies how financial capacity shifts the Laffer peak.

3.1 The Laffer Peak and the Equilibrium Safe Rate

Given $S(b, N)$, the steady-state safe rate $r(b, N)$ is determined by goods-market clearing:

$$\mathcal{D}(b, S(b, N); r) = A \left(b + \frac{1 - \bar{g} - rb - S(b, N)}{r + \lambda} \right) + \bar{g} = 1. \tag{22}$$

The key marginal object is the effect of an additional unit of debt on aggregate demand, evaluated at the prevailing safe rate. Government purchases are fixed, so the relevant object is

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

The first term inside brackets is the direct safe-wealth effect of an additional unit of debt. The second term is the tax burden required to service that unit and produce safe-asset services, capitalized by households at the mortality-adjusted rate. Using (21), the safety-cost component is the marginal stock-wide cost $S_b = s + bs_b$.

The aggregate-demand safe-debt Laffer curve is the relationship between the debt stock and the aggregate-demand contribution of insured public debt. 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-hand side is the non-Ricardian wedge created by finite horizons. The right-hand side is the marginal stock-wide cost of producing safe-asset services.

Proposition 1 (Safe-debt Laffer curve)**.** *Let $r(b, N)$ be the equilibrium safe rate that clears the goods market in (22), and define*

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

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, that $S_{bb}(b, N) > 0$ at any interior zero of H, 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 $b > b^L(N)$.

The proposition captures the main economic force. Below the peak, the direct safe-wealth effect of insured debt exceeds the capitalized tax burden created by ordinary safe-rate service and the stock-wide cost of producing safe-asset services. Above the peak, the marginal safety-production cost is large enough that the net marginal contribution of additional safe debt turns negative.

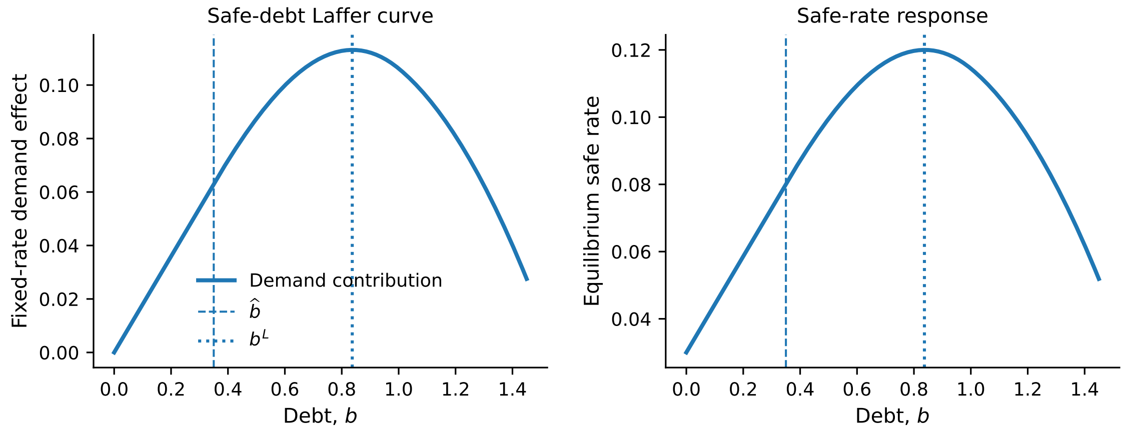

Figure 1: The aggregate-demand safe-debt Laffer curve and the equilibrium safe rate. The figure uses the closed-form linear-support-cost case of Appendix F. The vertical dashed line marks the stress threshold $\widehat{b}$; the dotted line marks the Laffer peak b^L. Above $\widehat{b}$, the marginal fiscal cost of producing safe-asset services rises with debt. At b^L, the marginal stock-wide cost of producing safe-asset services equals the non-Ricardian wedge.

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 the sign of the debt derivative of the equilibrium safe rate identifies the branch of the aggregate-demand safe-debt Laffer curve.*

Implicit differentiation of (22) gives

$$\operatorname{sign} r_b(b,N) = \operatorname{sign}\big[\lambda - S_b(b,N)\big]. \tag{27}$$

Thus the rate response has the same economic interpretation as the Laffer curve. When $\lambda - S_b > 0$, debt creates excess demand at the old safe rate, and the safe rate rises to reduce human wealth and restore goods-market clearing. When $\lambda - S_b < 0$, debt creates deficient demand at the old safe rate, and the safe rate falls to raise human wealth and restore goods-market clearing. Debt can therefore become contractionary even while households continue to hold safe government-debt claims: producing those claims has become costly enough to absorb more demand than the marginal unit creates. The equilibrium safe rate and the government's all-in cost of safe debt therefore move through different components. After the peak, the equilibrium safe rate falls because aggregate demand is deficient, while the safety-production premium embedded in the fiscal cost of debt is rising.

Figure 1 illustrates the mechanism in the closed-form case of Appendix F. The stress threshold $\widehat{b}$ is the debt level at which rollover pressure first exceeds the inelastic investor base. The Laffer peak b^L occurs later, when the marginal stock-wide cost of producing safe-asset services has risen enough to reach the non-Ricardian wedge. The lower panel plots the equilibrium safe rate and shows the same turning point.

3.2 Financial Capacity and Fiscal Capacity

The Laffer peak depends on the marginal fiscal cost of producing safe-asset services. This subsection shows how financial capacity shifts that peak in the benchmark short-maturity limit.

Financial capacity lowers rollover pressure. Differentiating (17) gives

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

A higher debt stock increases rollover exposure relative to capacity. A higher N expands absorption capacity and lowers the participation share required to absorb the same rollover exposure.

Appendix E shows that $S_N < 0$ and $S_{bN} < 0$. Higher financial capacity lowers both the total fiscal cost of the rollover-insurance premium and the marginal fiscal cost of producing safe government-debt claims. Since the Laffer peak is pinned down by the marginal stock-wide safety cost, the effect of financial capacity on fiscal capacity follows directly from differentiating the peak condition.

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

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

Then

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

The proposition states the financial-capacity comparative static. Higher N lowers rollover pressure and reduces the marginal fiscal cost of producing safe-asset services, captured by $S_{bN} < 0$. Because the peak occurs when this marginal cost reaches the non-Ricardian wedge, greater financial capacity shifts the Laffer peak outward.

Figure 2 illustrates this comparative static in the same closed-form case as Figure 1. A higher N shifts the stress threshold and the Laffer peak to the right. A lower N brings the stress region forward and makes the aggregate-demand reversal occur at a lower debt stock.

The benchmark comparative static holds maturity at the short-maturity limit and therefore studies how debt strains financial capacity when the full stock is exposed to rollover. The next section studies finite rollover intensity and the positive-maturity price component of rollover risk.

4 Positive Maturity and Term-Premium Compensation

The benchmark takes the short-maturity limit $\delta \to \infty$, in which the full stock is exposed to rollover and maturity collapses. This section restores positive maturity by considering a finite rollover intensity $\delta < \infty$. Positive maturity matters for the empirical margin because debt pressure need not be priced only through a short-rollover wedge. Longer claims also carry price risk, and the compensation for that risk can absorb part of the market response to high public debt. The section therefore constructs the total positive-maturity premium used in Section 5: a rollover component, a term-premium component, and a local balance-sheet multiplier linking the two.

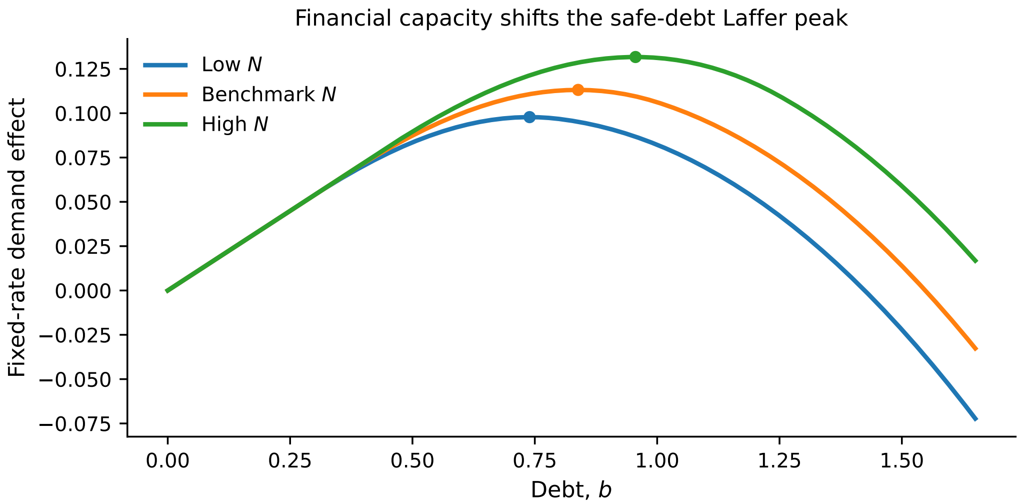

Figure 2: Financial capacity shifts the Laffer peak. The figure plots the aggregate-demand contribution of insured public debt for three levels of financial capacity in the closed-form case. Higher N lowers rollover pressure and shifts the peak outward.

4.1 Finite rollover intensity and direct rollover pricing

For a fixed rollover intensity δ, the share of the outstanding stock exposed to rollover over the insurance episode is

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

A higher δ means shorter average maturity and larger rollover exposure $\mu(\delta)b$. A lower δ means longer average maturity and less frequent rollover.

The rollover-pressure state and active-support upper limit are

$$q(b, N, \delta) = \frac{\mu(\delta)b}{N}, \qquad \bar{q}(b, N, \delta) = \min\{q(b, N, \delta), 1\}. \tag{30}$$

For $q(b, N, \delta) > \underline{\alpha}$, the direct rollover-insurance premium, before the price feedback derived in Appendix G, is

$$s^{base}(b, N, \delta) = \frac{1}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{\bar{q}(b,N,\delta)} \Gamma\left(\frac{\mu(\delta)b - \alpha N}{\alpha}\right) d\alpha, \tag{31}$$

and $s^{base}(b, N, \delta) = 0$ when $q(b, N, \delta) \leq \underline{\alpha}$. Thus finite rollover intensity scales the benchmark rollover-pressure variable by $\mu(\delta)$. For fixed δ, higher debt raises rollover pressure and higher financial capacity lowers it. Hence the direct rollover-insurance premium is increasing in public debt and decreasing in financial capacity on the stress region where the support schedule is active.

4.2 Term-premium compensation and balance-sheet feedback

Positive maturity turns rollover risk into a price object. A positive-maturity government bond is exposed to the current rollover-insurance premium and to the valuation consequences of future rollover-insurance premia. Let $s^T(b, N, \delta)$ denote the total excess-yield premium for a claim with rollover intensity δ, and define the associated effective discount-plus-hazard rate by

$$\rho^T(b, N, \delta) = r + s^T(b, N, \delta) + \delta. \tag{32}$$

The premium s^T is determined as a fixed point because the price component depends on duration, and duration depends on the total premium.

The term-premium component summarizes compensation for the price volatility of longer bonds. Let ς_P^2 denote the variance of the persistent premium or rate shock that moves the price of the long-bond portfolio. Then

$$s^X(b, N, \delta) = \vartheta(N)\frac{\varsigma_P^2}{[\rho^T(b, N, \delta)]^2}, \qquad \vartheta_N(N) < 0. \tag{33}$$

A lower δ lengthens duration and raises price sensitivity. The function $\vartheta(N)$ captures the price of bearing term-premium risk; higher financial capacity lowers that price. Appendix H derives this form from the duration $1/\rho^T$ of a Poisson-maturity bond.

The multiplier below is a local amplification device. A higher required premium lowers the price of outstanding long bonds. The price loss reduces effective balance-sheet capacity. Lower effective capacity then raises the rollover-insurance premium again. I measure N in mark-to-market balance-sheet-capacity units, so a valuation loss can be read locally as a capacity loss. With this normalization, the price sensitivity of long debt, $1/\rho^T$, and the capacity sensitivity of the premium enter the same local feedback equation.

Define the direct sensitivity of the term-premium component to financial capacity by

$$\psi^X_{\text{dir}}(b, N, \delta) \equiv -\left.\frac{\partial s^X(b, N, \delta)}{\partial N}\right|_{\rho^T} > 0, \tag{34}$$

where the derivative holds the effective discount-plus-hazard rate fixed and isolates the direct price-of-risk effect of financial capacity.

The multiplier combines two local sensitivities: the direct effect of capacity on the rollover-insurance premium and the direct effect of capacity on the term-premium component. Solving the local feedback loop with these two sensitivities gives the balance-sheet multiplier

$$M(b, N, \delta) \equiv \frac{\rho^T(b, N, \delta) + \psi^X_{\text{dir}}(b, N, \delta)}{\rho^T(b, N, \delta) + s_N^{base}(b, N, \delta)}. \tag{35}$$

Appendix G derives (35) as the local fixed point of the balance-sheet price-feedback loop. The rollover component of the positive-maturity premium is

$$s^R(b, N, \delta) = M(b, N, \delta)s^{base}(b, N, \delta), \tag{36}$$

and the total positive-maturity premium is

$$s^T(b, N, \delta) = s^R(b, N, \delta) + s^X(b, N, \delta). \tag{37}$$

Proposition 3 (Positive maturity and the price multiplier)**.** *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:*

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

Let $M(b, N, \delta)$ be the amplification multiplier in (35). Under the admissibility condition

$$\rho^T(b, N, \delta) + s_N^{base}(b, N, \delta) > 0,$$

the multiplier satisfies $M(b, N, \delta) > 1$. Moreover, $M(b, N, \delta) \to 1$ as $\delta \to \infty$.

The proposition separates the two positive-maturity channels needed for the empirical exercise. Shorter maturity raises direct rollover exposure through $\mu(\delta)$. Longer maturity creates price exposure, so premium changes move bond prices and feed back into effective financial capacity through M. In the short-maturity limit, duration vanishes and the multiplier converges to one, returning the benchmark rollover-insurance case.

4.3 The positive-maturity margin

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

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

The steady-state marginal cost is stock-wide:

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

The positive-maturity safe-debt margin is $\lambda - S_b^T(b, N, \delta)$. Thus positive maturity changes the premium schedule, not the logic of the aggregate-demand safe-debt Laffer curve. At a fixed δ, the Laffer peak solves

$$\lambda = s^T(b, N, \delta) + b\frac{\partial s^T(b, N, \delta)}{\partial b}. \tag{41}$$

The right-hand side now includes both rollover compensation and term-premium compensation. This is the bridge to Section 5: the empirical margin uses the total premium because the market consequences of high public debt can be absorbed through term premia as well as through short-rollover balance-sheet wedges.

4.4 Maturity choice as an extension

The same total premium schedule can be used to study optimal maturity. Appendix H considers the choice of rollover intensity δ, equivalently the average maturity $1/\delta$. At an interior optimum, the government balances the rollover cost of shorter maturity against the term-premium saving from shorter duration. The appendix shows that, under the maturity-capacity cross-partial condition stated there, higher financial capacity permits a higher optimal rollover intensity, while scarce rollover-insurance capacity tilts issuance toward longer maturity. This maturity-choice result is an application of the positive-maturity premium schedule; the empirical margin below uses the total premium s^T itself.

5 Identifying the U.S. Safe-Debt Laffer Margin

This section identifies the U.S. empirical counterpart of the model's safe-debt margin. Proposition 1 defines the theoretical margin as the gap between the non-Ricardian demand benefit of safe debt and the stock-wide marginal cost of producing safe-asset services. The empirical task is to measure the two premium margins that enter that stock-wide cost. The first is a rollover or balance-sheet margin, measured from Treasury-swap spreads. The second is a real-duration margin, measured from real term premia. The section then combines these margins with U.S. rollover exposure and dealer-parent Treasury capacity to construct a time series for the empirical margin from 2015 to 2026Q1. Appendix B records the data construction and regression details.

I denote the empirical time-series counterpart of the model's margin by H_t. The notation matches the empirical figures, while $H(b, N)$ remains the theoretical object. The economy is on the increasing branch of the safe-debt Laffer curve when $H_t > 0$ and on the declining branch when $H_t < 0$. By Corollary 1, the same margin determines the sign of the aggregate-demand contribution of debt and the sign of the equilibrium safe-rate response.

5.1 Rollover exposure and financial capacity

Rollover exposure is measured relative to Treasury-intermediation capacity:

$$Q_t = \frac{\mu_t b_t}{N_t}, \tag{42}$$

where b_t is privately held marketable Treasury debt, μ_t is the share of that debt maturing within one year, and N_t is measured Treasury-intermediation capacity. The numerator $\mu_t b_t$ is therefore the privately held marketable Treasury debt stock coming due within one year. This is the finite-maturity analogue of the rollover pressure in the model.

The baseline capacity measure is dealer-parent holdable Treasury capacity. "Dealer-parent" refers to the parent banking organizations that own the dealers active in Treasury intermediation; "holdable" means the amount of Treasury inventory those organizations could hold, given regulatory leverage capacity, plus the Treasury inventory they already hold. The measure is

$$N_t^{\text{hold}} = \sum_i \left(\frac{T1_{i,t}}{\kappa_{\text{SLR}}} - LE_{i,t} \right) + \sum_i UST_{i,t}, \qquad \kappa_{\text{SLR}} = 6\%. \tag{43}$$

Here $T1$ is Tier-1 capital, LE is total leverage exposure, and UST is Treasury securities already held by dealer-parent institutions. The first term measures unused balance-sheet space under a Supplementary Leverage Ratio constraint. The second term adds Treasuries already held, so the measure captures both unused room to absorb Treasuries and Treasury inventory already warehoused by dealer-parent institutions.

The capacity series starts in 2015Q1 because the leverage-exposure input is available in FR Y-9C filings from that point. Rollover exposure relative to dealer-parent capacity rises from about 0.5 in 2015 to 2.59 in 2026Q1. The transition is not simply that banks become smaller. During the 2025–2026Q1 squeeze, dealer-parent equity rises by about 28 percent, while holdable Treasury capacity falls by about 11 percent over 2025 and by another 19 percent in 2026Q1. The relevant constraint is leverage room: total leverage exposure rises by roughly \$800 billion against roughly flat Tier-1 capital, reducing the balance-sheet space available for Treasury warehousing even as dealer parents remain well capitalized.

The 2026Q1 observation is unusually far into scarce-capacity territory. For that reason, estimates below that allow the rollover response to steepen when capacity is scarce are treated as endpoint-sensitive evidence rather than as the baseline estimate.

This distinction is important for the mechanism. The safe-debt margin depends on the capacity available to absorb Treasury supply in stress states, not on bank equity in isolation. If dealer parents use leverage capacity for higher-return exposures, Treasury-warehousing capacity can contract while the institutions themselves look equity-rich. The 2025–26 episode is therefore informative about the model's prediction that the rollover premium steepens when Treasury-intermediation capacity becomes scarce.

For the monthly capacity state used in the rollover-margin interaction, I anchor the level of N_t to the quarterly regulatory measure and use dealer-parent market equity only to allocate movements within the quarter.

No high-frequency proxy is ideal. Dealer holdings and maturity-sector size are especially problematic because they are equilibrium quantities in the same local Treasury markets used for identification: they move with local supply, demand, dealer positioning, and maturity-sector scale. Dealer-parent market equity is also endogenous, but it is broader and less directly tied to maturity-local Treasury positions. The empirical construction therefore uses equity only as a high-frequency scaler for the quarterly regulatory capacity measure; the quarterly mean remains equal to regulatory holdable Treasury capacity by construction.

5.2 Premium components

The total premium used in the empirical margin is

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

This decomposition mirrors the maturity section's separation between the rollover component and the term-premium component in (37). I use s_t^D in the empirical section because the measured object is a real-duration premium. It is the empirical counterpart of the theoretical term-premium component s^X in the positive-maturity model. Both premium levels are aggregated across maturity buckets using outstanding Treasury debt weights.

The first component, s_t^R, is the rollover or balance-sheet premium. It is measured using a Treasury-swap spread: the Treasury par yield minus the maturity-matched par swap rate. A swap provides interest-rate exposure without requiring an investor to finance and warehouse a cash Treasury security. Thus, when cash Treasuries trade at a higher yield than matched swaps, the spread captures the balance-sheet cost of holding and intermediating the cash Treasury. I choose the sign convention so that a higher s_t^R means a higher rollover or balance-sheet premium. The debt-weighted rollover premium is about 29 basis points in 2026Q1.

The second component, s_t^D, is the real-duration premium. The baseline measure is the real term premium of D'Amico, Kim, and Wei (2018), which jointly prices nominal Treasury securities and inflation-indexed Treasuries. I use this real term-premium measure because the duration margin in the model is real duration compensation.

This distinction matters because nominal term premia also contain inflation-risk-premium (IRP) news, and asset-purchase announcements can move that component. Removing the inflation-risk-premium component is conservative for the safety-production object studied here. If high public debt raises inflation-risk premia because investors assign greater probability to inflationary debt dilution, that is a conventional fiscal-risk channel rather than the safe-debt-production cost isolated in the model.

The DKW series is native at the five- and ten-year maturities; the debt-weighted level interpolates through those anchors and extrapolates to the longer-maturity stock. In 2026Q1, the DKW real-duration premium is about 39 basis points. The older construction that subtracts an inflation-risk-premium estimate from the nominal term premium of Adrian, Crump, and Moench (2013) gives about 8 basis points and is reported as the ACM-minus-IRP level robustness case.

5.3 Identifying the rollover margin

The rollover-margin design uses maturity-local supply variation created by Federal Reserve Treasury purchases. Within a month, the Federal Reserve absorbs different fractions of the Treasury float in different maturity buckets. Month fixed effects absorb aggregate news, Federal Reserve balance-sheet-policy shifts, regulatory conditions, and curve-wide term-premium movements. Maturity fixed effects absorb persistent differences across tenors. The remaining maturity-local variation identifies how removing Treasury supply from a local segment affects the Treasury-swap spread, producing the rollover-premium margin used in the empirical Laffer calculation.

The design follows the local-supply logic of D'Amico and King (2013) and estimates long differences on a native maturity grid:

$$\Delta_h s^R_{m,t} = \alpha_m + \tau_t + \beta_h \Delta_h share_{m,t} + \varepsilon_{m,t}, \tag{45}$$

where m indexes maturity buckets, $\Delta_h x_{m,t} \equiv x_{m,t+h} - x_{m,t}$, α_m are maturity fixed effects, τ_t are month fixed effects, and $share_{m,t}$ is the Federal Reserve's share of the outstanding Treasury float in bucket m, measured in percentage points. The 2020H1 crisis window is excluded because the design is meant to identify slow-moving stress accumulation rather than emergency-market dysfunction and crisis-intervention pricing.

The slope β_h is measured in basis points per percentage point of local Federal Reserve absorption. With the sign convention above, a negative β_h means that Federal Reserve absorption lowers the rollover or balance-sheet premium; equivalently, private Treasury supply raises it. The rollover marginal term is constructed on the native maturity grid of the local-supply design, using the estimated slope and the private share of each bucket's float. The conversion from the local-supply coefficient to the stock-wide marginal term is

$$Q s^R_Q = -\beta_h \sum_m \left(100 - share_{m,t}\right) \simeq -496\, \beta_h. \tag{46}$$

At 2026Q1, the sum of private-float shares across the six native maturity buckets is about 496, which gives the reported conversion.

Averaged over the persistent twelve- to twenty-four-month window, the no-state slope is $\beta \simeq -0.27$ in the 2015+ sample and $\beta \simeq -0.48$ in the full sample. These estimates imply

$Qs_Q^R \simeq 133$ to 236 basis points. Figure 3 plots the horizon profile of the full-sample no-state rollover-margin estimate. The response builds gradually and stabilizes over the persistent twelve- to twenty-four-month window.

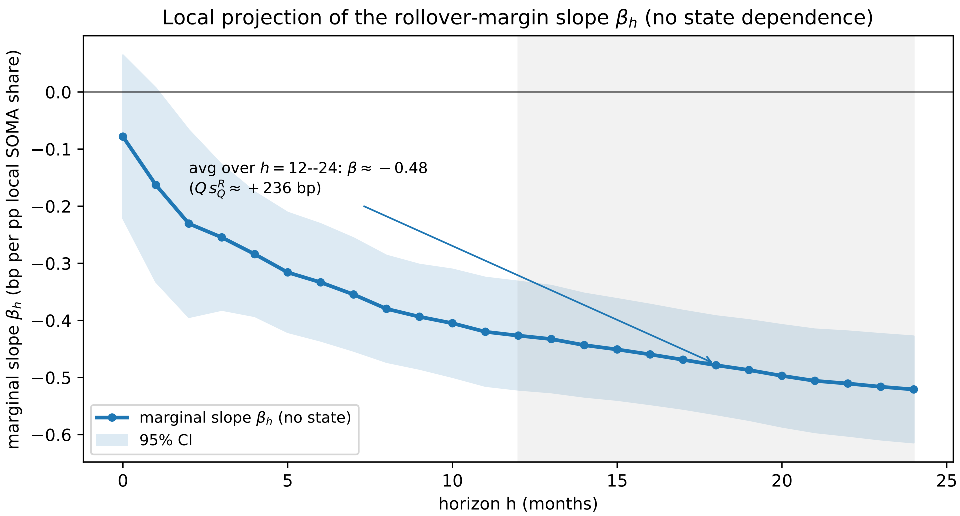

Figure 3: Local projection of the non-state-dependent rollover-margin slope. The figure reports the full-sample no-state response by horizon, excluding 2020H1. The slope deepens and settles over the twelve- to twenty-four-month window, with an average coefficient around −0.48, corresponding to $Qs_Q^R \simeq 236$ basis points.

The model predicts that the rollover-margin slope should steepen when capacity is scarce. Greater financial capacity lowers the marginal cost of producing safe-asset services and shifts the safe-debt Laffer peak outward. The empirical counterpart is that the local-supply effect should be smaller when dealer-parent capacity is ample and larger when dealer-parent capacity is scarce. I therefore estimate the interacted specification

$$\Delta_h s_{m,t}^R = \alpha_m + \tau_t + \beta_h \Delta_h share_{m,t} + \gamma_h \left(\Delta_h share_{m,t} \times \widetilde{N}_t\right) + \varepsilon_{m,t}, \tag{47}$$

where $\widetilde{N}_t$ is standardized dealer-parent capacity, with higher values denoting greater capacity. Since $\beta_h < 0$, a positive γ_h means the supply effect flattens when capacity is ample and steepens when capacity is scarce. The capacity interaction builds over the horizon. It is modest at twelve months and becomes stronger by twenty-four months, where the estimated interaction is about 0.17 with a t-statistic of about 2.5. Figure 4 plots this state-dependent local projection.

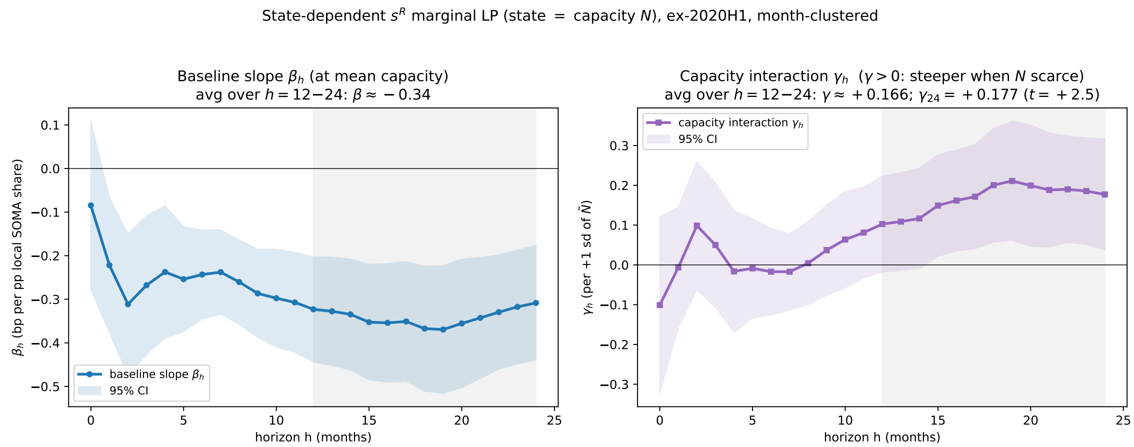

Figure 4: State-dependent local projection of the rollover-margin slope with dealer-parent capacity as the state variable. The left panel shows the baseline slope at mean capacity; the right panel shows the capacity interaction. A positive interaction means the local-supply slope steepens when capacity is scarce. The interaction builds over the horizon and is strongest over the twelve- to twenty-four-month window.

Figure 5 plots the linear rollover marginal term and the version that allows the rollover response to steepen when capacity is scarce. The paths move together when capacity is ample and fan apart when capacity becomes scarce.

Evaluated at today's record-scarce capacity, $\widetilde{N}_{\text{today}} = -3.06$, the interaction raises Qs_Q^R from the linear 133–236 basis-point range to 230 basis points under the eight- to twelve-month average, 308 basis points at twelve months, 408 basis points at twenty-four months, and 409 basis points for the twelve- to twenty-four-month average. The twelve- to twenty-four-month average slightly exceeds both the twelve- and twenty-four-month endpoints because the interaction path is hump-shaped and peaks inside the averaging window.

The eight- to twelve-month reading gives the smaller capacity-interaction effect and is used as the baseline interacted reading. The twelve- to twenty-four-month window stabilizes the coefficient more cleanly, but applying it to the 2026Q1 capacity state uses a record-scarce observation outside the 2015–25 capacity range. I therefore report those estimates as longer-horizon coefficient-window sensitivity rows in the margin table.

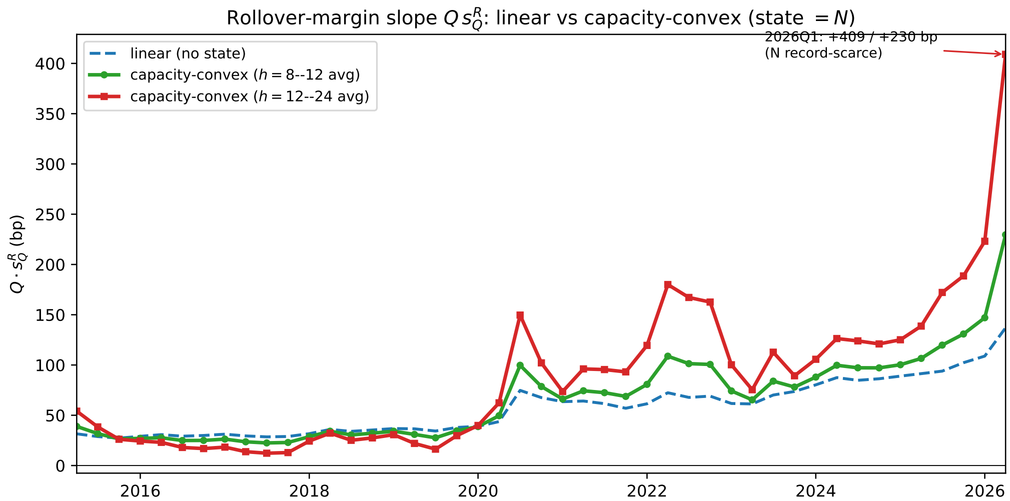

Figure 5: Rollover-margin slope over time: linear versus capacity-interacted. The figure plots the no-state linear marginal term, the eight- to twelve-month capacity-interacted average, and the twelve- to twenty-four-month capacity-interacted average. The paths coincide when capacity is ample and fan out in scarce-capacity episodes, especially in 2026Q1 when dealer-parent holdable Treasury capacity reaches a record-scarce state.

Table 1 reports the local-supply estimates underlying the rollover margin. Column 1 gives the full-sample no-state estimate, column 2 gives the same specification over the 2015+ sample used for comparison with the capacity interaction, and column 3 adds the interaction with standardized dealer-parent capacity. The full-sample estimate is the most precise linear estimate and does not use the capacity interaction. The 2015+ no-state estimate is the same-sample benchmark for comparison with the interacted specification.

Table 1: Long-difference local-supply estimates for the rollover margin

	(1) no state Full sample	(2) no state 2015+	(3) + capacity N 2015+
$\Delta_h share_{m,t}$ (β, $h = 24$)	-0.521^{***}	-0.232^{***}	-0.293^{***}
	(0.048)	(0.054)	(0.062)
$\Delta_h share_{m,t} \times \widetilde{N}_t$ (γ)			$+0.174^{**}$
			(0.070)
Maturity fixed effects	yes	yes	yes
Month fixed effects	yes	yes	yes
Cluster by month	yes	yes	yes
Sample months	2010-02–2024-03	2015-01–2024-03	2015-01–2024-03
Observations	996	642	642
Month clusters	166	107	107
Within R^2	0.180	0.027	0.036
Qs_Q^R, avg. $h = 8$–12	+201	+134	+230
Qs_Q^R, $h = 12$	+211	+137	+308
Qs_Q^R, $h = 24$	+258	+115	+408
Qs_Q^R, avg. $h = 12$–24	+236	+133	+409

Notes: Standard errors clustered by month are in parentheses. Stars denote 10, 5, and 1 percent significance. The coefficient β is measured in basis points per percentage point of local Federal Reserve Treasury absorption. The sample excludes 2020H1. The displayed coefficients are the $h = 24$ estimates. The bottom panel converts each column's horizon-specific slopes to Qs_Q^R at the 2026Q1 stock using the conversion in (46); the interacted column evaluates the capacity interaction at $\widetilde{N}_{\text{today}} = -3.06$. Capacity is observed only from 2015Q1, so the interacted rows extrapolate the interaction to today's record-scarce state.

5.4 Identifying the real-duration margin

The duration-margin design uses announcements because real term premia reprice when investors learn about changes in the expected path of Federal Reserve Treasury purchases or runoff, before the holdings path is implemented. I therefore use twelve Treasury-relevant asset-purchase and balance-sheet-runoff announcements and measure the two-trading-day response of the DKW real term premium. Appendix B lists the events and records how announcement amounts are constructed.

The event-level specification relates the change in the real term premium to the announced change in expected cumulative Treasury purchases:

$$\Delta s_e^D = a + \beta_D \, Amount_e + u_e, \tag{48}$$

where $Amount_e$ is measured in billions of dollars, positive for purchase increases and negative for tapering or runoff.

The conversion to the stock-wide repricing term is the duration analogue of the rollover conversion. The event slope is applied to the privately held Treasury float and reported as the stock-wide duration repricing term $Q_t s_{Q,t}^D$.

For the ten-year DKW real term premium, the event slope is about -0.0066 basis points per billion dollars. With privately held marketable debt of about \$16.67 trillion, this implies $Q_t s_{Q,t}^D \simeq 110$ basis points. The baseline outstanding-weighted real response gives $Q_t s_{Q,t}^D = 90$ basis points, while a share specification gives about 53 basis points and provides a unit-clean lower bound. I therefore treat 90 basis points as the baseline and 50–110 basis points as the relevant range.

This real-duration response is roughly three times smaller than the nominal term-premium responses commonly reported in the quantitative-easing literature because it strips out inflation-risk-premium news and uses a whole-sample event response rather than a QE1-heavy average.

The duration estimates show no reliable state dependence. For consistency with the rollover-premium specification, I estimate the state interaction using the same capacity variable N. The interaction is imprecise and statistically weak, so the empirical exercise carries a constant duration margin and assigns the capacity-interaction channel to the rollover component. Figure 6 shows the local-projection pattern: the pooled real-duration response is small and persistent, and the capacity interaction is not statistically distinguishable from zero.

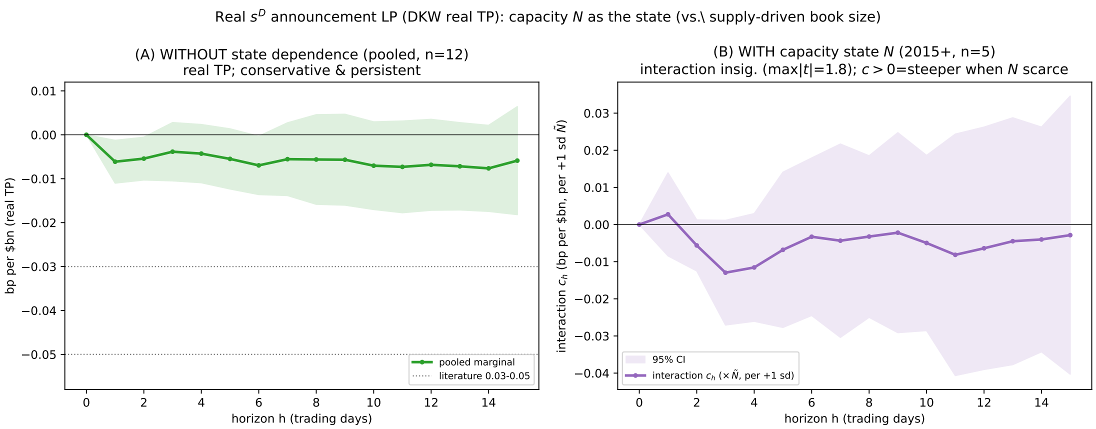

Figure 6: Announcement-identified real-duration margin. Panel A shows the pooled DKW real-term-premium response around Federal Reserve asset-purchase and balance-sheet-runoff announcements. Panel B interacts the announcement response with the same capacity state N used in the rollover-premium regression. The pooled response is small and persistent, while the capacity interaction is imprecisely estimated and statistically weak.

5.5 The U.S. safe-debt margin

The empirical marginal safety-production premium at date t combines the observed premium level with the stock-wide marginal repricing terms. The terms $Q_t s^R_{Q,t}$ and $Q_t s^D_{Q,t}$ denote the rollover and duration repricing margins, each expressed in basis points and in the same stock-wide marginal-cost units. The empirical counterpart of the theoretical stock-wide marginal object $S_b = s + bs_b$ is

$$S_{b,t} \equiv s^T_t + Q_t s^T_{Q,t} = s^R_t + s^D_t + Q_t s^R_{Q,t} + Q_t s^D_{Q,t}. \tag{49}$$

The corresponding safe-debt margin is

$$H_t \equiv \lambda - S_{b,t}. \tag{50}$$

The level term $s^T_t = s^R_t + s^D_t$ is the average premium paid on the outstanding stock. The two stock-wide repricing terms are measured differently: $Q_t s^R_{Q,t}$ comes from maturity-local Treasury supply variation, while $Q_t s^D_{Q,t}$ comes from the response of real term premia to Treasury-supply announcements. Both are reported in basis points and enter the same marginal-cost object $S_{b,t}$.

The macro anchor is the non-Ricardian wedge λ, which makes safe public debt expansionary in the standard benchmark. In the benchmark calibration, I set the household discount rate equal to the safe rate, $\rho = r$, so $\lambda = A - r$. The baseline value is $\lambda = 330$ basis points. The empirical exercise is therefore a location exercise: it compares the measured marginal safety-production cost with the macro demand benefit of issuing one more safe claim.

Figure 7 plots H_t over time. The margin falls from roughly positive 280 basis points in 2015 toward and through zero as rollover exposure grows relative to dealer-parent capacity. The linear path and the eight- to twelve-month capacity-interacted path provide the baseline readings. The longer-horizon capacity-interacted path shows how the margin moves when the more stabilized twelve- to twenty-four-month coefficient is applied to the record-scarce 2026Q1 capacity observation.

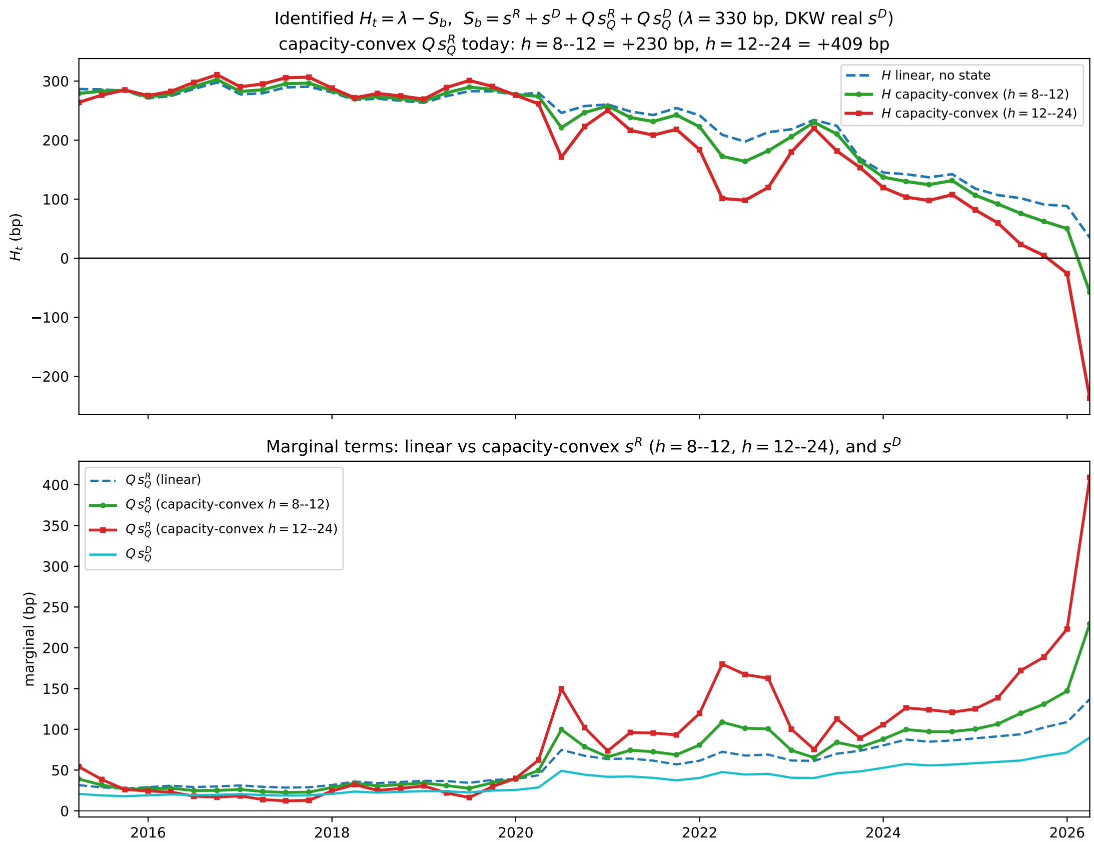

Figure 7: Identified safe-debt margin H_t. The top panel plots three DKW real-duration-level paths: the linear no-state rollover margin, the eight- to twelve-month capacity-interacted rollover margin, and the twelve- to twenty-four-month capacity-interacted rollover margin. The bottom panel shows the stock-wide repricing terms $Q_t s^R_{Q,t}$ and $Q_t s^D_{Q,t}$. The margin falls from roughly positive 280 basis points in 2015 toward and through zero by 2026Q1; the spread across paths reflects how far the capacity-interaction estimate is read into the record-scarce 2026Q1 state.

Table 2 reports the main 2026Q1 readings. The same-sample linear reading is positive 39 basis points under DKW real levels. The eight- to twelve-month capacity-interacted reading is *negative* 58 basis points. The full-sample linear estimate is the most precise estimate that does not use the capacity interaction at the 2026Q1 endpoint; it places the margin below zero, at negative 64 basis points.

Table 3 reports sensitivity to the real-duration level construction and to the horizon window used for the capacity interaction. The ACM-minus-IRP row changes only the premium level. The remaining rows keep the same 2026Q1 capacity state but use different horizon windows for the estimated capacity interaction. These rows show how far the margin moves when the more persistent twelve- to twenty-four-month rollover response is applied to today's capacity observation.

The twelve- to twenty-four-month capacity-interacted average is negative 237 basis points, and the single twelve-month capacity-interacted reading is negative 136 basis points.

Table 2: U.S. safe-debt margin: main readings, 2026Q1

Specification	s^T level	Rollover term	Duration term	H_t
Same-sample linear, DKW level	68	133	90	39
Capacity-convex $h = 8$–12, DKW level	68	230	90	-58
Full-sample linear, DKW level	68	236	90	-64

Notes: Entries are basis points. The baseline non-Ricardian wedge is $\lambda = 330$ basis points. The premium level is $s^R + s^D$. The rollover and duration terms are $Q_t s^R_{Q,t}$ and $Q_t s^D_{Q,t}$, respectively. The first two rows give the baseline 2026Q1 readings used in the text: the same-sample linear row, which does not use the capacity interaction, and the eight- to twelve-month interacted row. The third row reports the most precise linear estimate. Displayed entries are rounded; recomputing margins from displayed components can differ by at most one basis point from calculations based on the underlying empirical files.

Table 3: U.S. safe-debt margin: sensitivity readings, 2026Q1

Specification	s^T level	Rollover term	Duration term	H_t
Same-sample linear, ACM-minus-IRP level	37	133	90	70
Capacity-convex $h = 12$–24 average, ACM-minus-IRP level	37	409	90	-206
Single twelve-month capacity-interacted, DKW level	68	308	90	-136
Capacity-convex $h = 24$, DKW level	68	408	90	-236
Capacity-convex $h = 12$–24 average, DKW level	68	409	90	-237

Notes: Entries are basis points. The rollover and duration terms are $Q_t s^R_{Q,t}$ and $Q_t s^D_{Q,t}$, respectively. The rows use the same identified ingredients as Table 2 to show sensitivity to real-duration level construction and horizon choice, including the longer-horizon coefficient-window readings at the same 2026Q1 state. The identification sample ends in 2024-03, while the diagnostic evaluates the observed 2026Q1 state. The interacted rows extrapolate the interaction to today's record-scarce capacity, $\widetilde{N}_{\text{today}} = -3.06$, outside the 2015–25 sample range; the linear rows do not use the capacity interaction. The twelve- to twenty-four-month average can exceed both endpoint values because the interaction path is hump-shaped and peaks inside the averaging window.

Several remaining measurement choices make the calculation conservative. The capacity denominator is broad and treats measured leverage slack as usable for Treasury intermediation, which lowers measured rollover pressure. The duration component is real rather than nominal, which removes inflation-risk-premium news. The Treasury-swap spread nets out some of the liquidity and safety value of cash Treasuries, so it can understate the balance-sheet cost of warehousing those securities. Finally, the calculation uses observed premium levels rather than fitted high-stress levels. These choices bias the measured marginal safety-production cost downward. Even so, the U.S. moves from a comfortably positive margin in 2015 to the neighborhood of zero under linear rollover pricing and just past the peak under the eight- to twelve-month estimate that allows the rollover response to steepen when capacity is scarce.

6 Conclusion

The paper's central result is an aggregate-demand safe-debt Laffer curve: beyond a point, additional safe public debt lowers rather than raises its contribution to aggregate demand and the equilibrium safe rate. Corollary 1 links these two turning points: the aggregate-demand peak and the safe-rate peak coincide. The drag raised in the introduction is therefore not a default or solvency drag; it is the aggregate-demand drag that appears when the cost of producing safe-asset services absorbs the wealth effect of additional debt.

The mechanism is fiscal-financial. The safety quality of public debt is endogenous: safe-asset production depends jointly on sovereign solvency, debt-service arithmetic, and the financial system's capacity to roll over government debt and produce its safety services. When that capacity is abundant, additional debt expands the stock of safe claims and raises the aggregate-demand contribution of public debt and the equilibrium safe rate. When it is scarce, the marginal safety-production cost can be large enough that the added claims remain safe while their marginal contribution to aggregate demand and the safe rate turns negative.

The empirical section identifies the premium margins that locate the United States on this curve. Rollover exposure relative to dealer-parent Treasury capacity rises sharply from 2015 to 2026Q1, while the rollover premium response steepens when capacity becomes scarce. The resulting empirical margin falls from roughly positive 280 basis points in 2015 to a range of positive 39 to *negative* 58 basis points in 2026Q1. The additional estimates that allow the rollover response to steepen with balance-sheet scarcity show how the 2025–26 amplification of rollover costs can push the margin further into negative territory. The empirical conclusion is that the United States has moved from the comfortably increasing branch of the safe-debt Laffer curve to the neighborhood of the peak and, even under the conservative interacted rollover estimate, just past it.

The location of the peak is itself an economic object. Financial capacity shifts it outward: greater balance-sheet capacity for Treasury intermediation raises the debt level at which the safe-debt margin turns negative. Past the peak, additional debt lowers the equilibrium safe rate; when that rate is free to fall, the drag is absorbed through the rate, but at a binding lower bound it falls directly on demand (Caballero and Farhi, 2018). The peak of the safe-debt Laffer curve is therefore not a fixed feature of the fiscal landscape, but a moving boundary set by the financial system's capacity to produce safe-asset services.

Appendix

A Global-Owner Incidence and Partial Recycling

This appendix gives one ownership structure that delivers the incidence convention used in the main text. The convention is that fiscal resources used to produce safe-asset services are borne by households inside the domestic non-Ricardian demand block, while the premium income received by intermediary owners is not recycled into that block.

Recall from Section 3 and Proposition 1 that the marginal aggregate-demand contribution of public debt is governed by the margin

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

where λ is the non-Ricardian wedge and $S_b(b, N)$ is the marginal fiscal-financial cost of producing safe-asset services. The peak of the safe-debt Laffer curve occurs when this margin is zero.

To make the incidence explicit, suppose that the intermediary sector is owned by investors who are marginally global. This investor class may include foreign investors as well as domestic wealthy investors. What matters is marginal behavior: additional intermediary income is either consumed outside the domestic goods block or allocated to a global residual asset outside the domestic goods and local safe-bond demand block.

Let $\xi \in [0, 1]$ denote the fraction of intermediary premium income that is recycled into domestic aggregate demand. The benchmark case in the main text is $\xi = 0$. In that case, the fiscal cost $S(b, N)$ is paid by households inside the domestic demand block, while the corresponding intermediary income generates no offsetting domestic demand term. With partial recycling, the net marginal drag from safe-asset-service production is only the unrecycled component, $(1 - \xi)S_b(b, N)$. The margin becomes

$$H_\xi(b, N) = \lambda - (1 - \xi)S_b(b, N).$$

The Laffer peak is therefore characterized by

$$(1 - \xi)S_b(b, N) = \lambda.$$

Thus the benchmark incidence convention is a transparent limiting case rather than a knife-edge requirement. Partial recycling weakens the safety-production drag, but the same mechanism applies whenever $\xi < 1$ and the effective marginal cost $(1 - \xi)S_b(b, N)$ eventually exceeds the non-Ricardian wedge.

This formulation also clarifies the empirical interpretation of financial capacity N. The dealer-parent capital and leverage-room variables used in Section 5 measure balance-sheet capacity available for U.S. Treasury intermediation. They need not be interpreted as wealth owned by the marginal domestic consumer in the demand block. Indeed, the relevant intermediary balance sheets may be owned by globally diversified investors, including domestic investors whose marginal portfolio decisions are global. This is consistent with the incidence convention above: N measures the capacity used to produce safe Treasury services, while ξ governs how much of the resulting premium income returns to domestic aggregate demand.

B Empirical Construction

This appendix records the measurement details behind Section 5. The body uses identified rollover and duration premium margins as the empirical inputs.

B.1 Data construction

The rollover-stress statistic is $Q_t = \mu_t b_t / N_t$, as in (42). The baseline construction uses a one-year rollover horizon. Thus $\mu_t b_t$ is the dollar amount of privately held marketable Treasury debt maturing within one year, net of Federal Reserve System Open Market Account (SOMA) holdings. The current value used in Section 5, $Q = 2.59$, is the 2026Q1 value of this one-year privately held rollover exposure divided by dealer-parent holdable Treasury capacity.

Dealer-parent holdable Treasury capacity is constructed from FR Y-9C filings as in (43). The measure is deliberately broad: it treats all measured leverage slack as potentially usable for Treasury intermediation and then adds Treasury inventory already held. This broad denominator makes the baseline rollover-stress statistic conservative. Because the margin calculation is invariant to constant rescalings of N_t, alternative denominators matter through the time-series shape of Q_t, especially during the post-2020 rise in privately held short-maturity exposure.

The monthly capacity state used in the interaction specification is a high-frequency interpolation of this quarterly regulatory object. Dealer-parent market equity is not treated as an exogenous measure of N_t. The narrower claim is that it is a broad high-frequency scaler for dealer-parent balance-sheet franchise value and is less mechanically tied to maturity-local Treasury positions than dealer holdings or maturity-sector size. The interpolation preserves the quarterly regulatory mean of holdable capacity by construction, so the level of N_t remains pinned down by regulatory leverage-room data rather than by equity valuation.

The rollover component s^R is measured using the Treasury-swap spread: the Treasury par yield minus the maturity-matched par swap rate. The native maturity grid for the local-supply design uses the non-interpolated benchmark tenors two, three, five, seven, ten, and thirty years. The sign convention is chosen so that a higher s^R means a higher rollover or balance-sheet premium.

The duration component s^D is measured using the DKW real term premium of D'Amico, Kim, and Wei (2018). The series jointly prices nominal Treasury and inflation-indexed Treasury securities and therefore nets out the inflation-risk-premium component inside a single term-structure model. This removal is conceptually important: an inflation-risk premium that rises with the stock of debt would partly reflect inflationary debt dilution or fiscal-risk compensation, while the paper's empirical object is real duration compensation associated with safe-debt production. The real term premium is native at the five- and ten-year maturities; the debt-weighted level interpolates through these anchors and extrapolates to longer maturities. The ACM-minus-inflation-risk-premium construction used in earlier drafts is retained as the ACM-minus-IRP lower level robustness case.

Table 4 lists the data sources used to construct the debt, capacity, premium, Federal Reserve holdings, and dealer-parent equity series.

Table 4: 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 CUSIP	Federal Reserve SOMA holdings	Weekly/monthly
Tier-1 capital, leverage exposure, and Treasuries held by dealer parents	FR Y-9C regulatory filings	Quarterly
Treasury par yields and maturity-matched par swap rates	Treasury and swap-curve data	Daily
Real term premia	D'Amico, Kim, and Wei (2018) term-premium estimates	Daily
Nominal term premia and robustness checks	Adrian, Crump, and Moench (2013) estimates and inflation-risk-premium estimates	Daily/monthly
Dealer-parent equity values for high-frequency capacity proxy	CRSP dealer-parent equity-market data	Daily

B.2 Rollover-margin identification details

The rollover-margin design is a monthly maturity panel. For horizon h, the no-state specification is (45). The state-dependent specification is (47). Both use maturity fixed effects and month fixed effects, with standard errors clustered by month. The sample excludes 2020H1 to focus on slow-moving stress accumulation rather than emergency-market dysfunction and crisis-intervention pricing.

The unit conversion in (46) uses the fact that the Federal Reserve share is measured in percentage points. A one-percentage-point increase in the Federal Reserve share of a maturity bucket removes one percent of that bucket's outstanding float from private balance sheets. Summing $100 - share_{m,t}$ across the native maturity grid gives the scale factor used in the body. This converts a local-supply coefficient in basis points per percentage point of Federal Reserve share into the rollover marginal term $Q_t s^R_{Q,t}$.

The capacity interaction is estimated on the 2015+ sample because regulatory capacity is observed from 2015Q1. The 2026Q1 capacity observation is outside the historical range. The eight- to twelve-month interacted row is the baseline interacted reading; the twelve- to twenty-four-month average is reported as a longer-horizon coefficient-window sensitivity row. The linear rows require no extrapolation of the capacity interaction.

B.3 Duration-margin identification details

The duration-margin event study uses Treasury-relevant asset-purchase and balance-sheet-runoff announcements. The outcome is the two-trading-day change in the DKW real term premium. The treatment is the announced change in expected cumulative Treasury purchases or runoff. The Maturity Extension Program is excluded because it is approximately face-neutral for Treasury supply. The resulting duration marginal is an announcement-identified real-duration response. Its capacity interaction, estimated with the same capacity state N as in the rollover-margin regression, is statistically weak; the baseline empirical margin therefore treats s^D as constant.

The baseline outstanding-weighted stock-wide duration repricing term is $Q_t s^D_{Q,t} = 90$ basis points. A ten-year real-term-premium specification gives about 110 basis points, while a share specification gives about 53 basis points. Section 5 carries 90 basis points as the baseline and interprets 50–110 basis points as the relevant range. Table 5 reports the Treasury-relevant asset-purchase and balance-sheet-runoff announcements used in the duration-margin event study. Open-ended purchase and pace-change announcements require judgment in translating policy news into an expected cumulative purchase amount; the table records the event-size assumptions used in the baseline.

Table 5: 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. The Maturity Extension Program is excluded because it is approximately face-neutral for Treasury supply. Open-ended and pace-change announcements require judgment in translating statements into cumulative amounts; this is the main soft input in the duration-margin construction.

C Households, Insurance, and Perpetual-Youth Demand

This appendix records the household block behind the main text. The first part gives the finite-risk-aversion insurance conditions. The second states the full-insurance benchmark and shows that, under that benchmark, the household problem reduces to the perpetual-youth demand system used in Section 2.

C.1 Preferences, states, and insurance choices

Time is continuous. A household dies with Poisson hazard λ, receives the actuarial mortality return on surviving wealth, and discounts at rate ρ. Newborns replace dying households one-for-one, but they do not inherit the assets of the deceased; the annuity market redistributes those assets to surviving households through the mortality return. As in the main text, $A \equiv \rho + \lambda$. Preferences are Epstein–Zin (Epstein and Zin, 1989) with 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

$$C^{\sigma}_t = \left(\mathbb{E}_t \left[c_i(t,\omega)^{1-\sigma}\right]\right)^{1/(1-\sigma)}. \tag{51}$$

The logarithmic aggregator over time gives unit intertemporal elasticity; σ prices aggregate rollover-state risk within each date.

The household state variable is financial wealth w_i. The household chooses government-bond holdings b_i, insured units $\iota_i \in [0, b_i]$, and uninsured exposure

$$u_i \equiv b_i - \iota_i \geq 0. \tag{52}$$

The residual safe annuitized asset earns r. Government debt pays r^G; rollover insurance costs π per insured unit; an uninsured unit suffers rollover-stress loss $\ell(\omega) \geq 0$. Disposable private income is denoted by χ. In the benchmark steady state of the main text, χ is determined by the fixed-purchases tax rule in (5). The state-contingent wealth law can then be written as

$$\dot{w}_i(\omega) = (r + \lambda)w_i + \chi - c_i(\omega) + \left(r^G - r - \pi\right)\iota_i + \left(r^G - r - \ell(\omega)\right)u_i. \tag{53}$$

This separates the insured claim from the uninsured rollover-risk claim.

C.2 Finite-σ insurance conditions

The certainty equivalent in (51) induces the normalized within-date state-price weight

$$\Lambda_i^\sigma(\omega) \equiv \frac{c_i(\omega)^{-\sigma}}{\mathbb{E}[c_i(\omega)^{-\sigma}]}. \tag{54}$$

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

$$\mathbb{E}\left[\Lambda_i^\sigma(\omega)(r^G - r - \ell(\omega))\right] = 0, \qquad \text{if } u_i > 0, \tag{55}$$

$$\mathbb{E}\left[\Lambda_i^\sigma(\omega)(r^G - r - \ell(\omega))\right] \leq 0, \qquad \text{if } u_i = 0. \tag{56}$$

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

$$r^G - r - \pi = 0, \qquad \text{if } \iota_i > 0, \tag{57}$$

$$r^G - r - \pi \leq 0, \qquad \text{if } \iota_i = 0. \tag{58}$$

Equivalently, holding total bond holdings fixed, a marginal unit of insurance has state-contingent payoff $\ell(\omega) - \pi$ and satisfies

$$\mathbb{E}\left[\Lambda_i^\sigma(\omega)(\ell(\omega) - \pi)\right] = 0, \qquad \text{if } 0 < \iota_i < b_i, \tag{59}$$

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

$$\mathbb{E}\left[\Lambda_i^\sigma(\omega)(\ell(\omega) - \pi)\right] \geq 0, \qquad \text{if } \iota_i = b_i. \tag{61}$$

These conditions are the finite-risk-aversion discipline behind the benchmark used in the main text: households compare the insurance premium with the rollover loss under their state-price weights.

C.3 Full-insurance benchmark

Let

$$\ell^{\max} \equiv \operatorname*{ess\,sup}_{\omega\in\Omega} \ell(\omega) \tag{62}$$

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,

$$\pi < \ell^{\max}. \tag{63}$$

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

Assumption 1 is the benchmark allocation. With high aversion to aggregate payoff risk and $\pi < \ell^{\max}$, households prefer the insured claim to worst-state rollover exposure. The finite-risk-aversion conditions above record the insurance margins that motivate this benchmark.

Under full insurance, equation (57) implies

$$r^G = r + \pi, \tag{64}$$

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

$$r^G - \pi = r. \tag{65}$$

Substituting $u_i = 0$ and $r^G - r - \pi = 0$ into (53) gives the state-independent wealth law

$$\dot{w}_{i,t} = (r_t + \lambda)w_{i,t} + \chi_t - c_{i,t}. \tag{66}$$

Thus $C_t^{\sigma} = c_{i,t}$ in the insured benchmark, and the household problem collapses to the standard log perpetual-youth problem.

C.4 Human wealth and consumption under full insurance

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

$$h_t = \int_t^{\infty} \exp\left[-\int_t^{\tau} (r_v + \lambda)\, dv\right] \chi_\tau \, d\tau. \tag{67}$$

It evolves according to

$$\dot{h}_t = (r_t + \lambda)h_t - \chi_t, \tag{68}$$

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

$$h(r) = \frac{\chi}{r + \lambda}. \tag{69}$$

Let total wealth be $\mathcal{W}_{i,t} = w_{i,t} + h_t$. Combining (66) and (68) gives

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

The log perpetual-youth Euler equation is $\dot{c}_{i,t}/c_{i,t} = r_t - \rho$. In steady state, $\mathcal{W}_i = w_i + h$, and imposing the no-Ponzi condition on (70) yields

$$c_i = A\mathcal{W}_i. \tag{71}$$

Aggregating over the unit-mass population and using the fact that fully insured public debt is the aggregate safe financial asset held by households gives

$$c = A(b + h), \tag{72}$$

which is equation (3). Public debt is net wealth because finite horizons imply partial capitalization of taxes paid by future households. In the fixed-purchases benchmark, an additional unit of debt raises safe financial wealth directly and also raises the tax burden capitalized in human wealth. Combining (3) with (5) gives the fixed-rate marginal effect recorded in (23).

C.5 Goods-market clearing and the safe rate

In equilibrium, the household insurance premium equals the marginal-participant rollover-insurance premium from Appendix D: $\pi = s(b, N)$. Combining the full-insurance household block with the fixed-purchases fiscal closure gives the goods-market condition (22); differentiating it gives the safe-rate response (27). This is the rate result used in the main text.

D Threshold Participation and the Rollover-Insurance Premium

This appendix derives the pricing measure used in Section 2.3. Throughout this appendix, δ is fixed unless stated otherwise. To keep notation light, I write $s(b, N)$ for the finite-rollover premium $s^{base}(b, N, \delta)$. For orientation, the premium schedule derived below is the following. On the active stress region $q(b, N, \delta) > \underline{\alpha}$,

$$s(b, N) = \frac{1}{1 - \underline{\alpha}} \int_{\underline{\alpha}}^{\bar{q}(b,N,\delta)} \Gamma\left(\frac{\mu(\delta)b - \alpha N}{\alpha}\right) d\alpha, \qquad \bar{q}(b, N, \delta) = \min\{q(b, N, \delta), 1\}, \tag{73}$$

where

$$q(b, N, \delta) = \frac{\mu(\delta)b}{N}. \tag{74}$$

When $q(b, N, \delta) \le \underline{\alpha}$, the inelastic investor base absorbs rollover exposure and $s(b, N) = 0$. The derivation has two pieces. The first defines the state-contingent assessment paid by active rollover insurers. The second shows that the marginal participant prices that assessment under the flat-prior Laplacian distribution, which is uniform over active participation states. The private-signal threshold environment follows Morris and Shin (1998).

D.1 Assessment in a rollover-stress state

The government has debt stock b. The rollover share over the insurance episode is $\mu(\delta) = 1 - e^{-\delta}$, so rollover need is

$$R(b, \delta) = \mu(\delta)b. \tag{75}$$

A fraction $\underline{\alpha} \in (0, 1)$ of potential rollover insurers is inelastic and always participates. The remaining elastic mass has participation share $\tilde{\alpha} \in [0, 1]$, so total participation is

$$\alpha = \underline{\alpha} + (1 - \underline{\alpha})\tilde{\alpha} \in [\underline{\alpha}, 1]. \tag{76}$$

Active absorption capacity is

$$B(\alpha, N) = \alpha N, \tag{77}$$

where N is aggregate financial-system capacity. A rollover-stress shortfall is

$$X(b, N, \alpha) = \max\{\mu(\delta)b - \alpha N, 0\}. \tag{78}$$

Equivalently, stress occurs when $\alpha < q(b, N, \delta)$. The per-active-unit stress is

$$\zeta(b, N, \alpha, \delta) = \frac{X(b, N, \alpha)}{\alpha}. \tag{79}$$

The support schedule maps this individual balance-sheet stress into the state-contingent assessment required to preserve the safety of government debt:

$$\Gamma(\zeta), \qquad \Gamma(0) = 0, \quad \Gamma'(\zeta) > 0 \text{ for } \zeta > 0, \quad \Gamma'_+(0) > 0, \quad \Gamma''(\zeta) \geq 0. \tag{80}$$

Throughout, Γ is a per-active-unit stress-cost schedule. It maps the stress borne by an active balance sheet into the compensation required in that state.

Per active unit, quoted per unit of insured debt, the assessment is

$$a(b, N, \alpha) = \Gamma\big(\zeta(b, N, \alpha, \delta)\big). \tag{81}$$

Equivalently, for $\alpha < q(b, N, \delta)$, the assessment is $\Gamma((\mu(\delta)b - \alpha N)/\alpha)$, and otherwise it is zero because $\Gamma(0) = 0$. Low participation raises the assessment by increasing the rollover stress borne by each active unit. All quantities are in flow-spread units per unit of insured government debt, so aggregate premium payments in the macro block are $S = sb$.

D.2 The participation game

An aggregate fundamental $\theta \in \mathbb{R}$ shifts elastic insurers' willingness or ability to participate; larger θ is a worse participation environment. Elastic insurer i observes

$$x_i = \theta + \eta\varepsilon_i, \tag{82}$$

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

$$e_i \in \{0, 1\}, \tag{83}$$

where $e_i = 1$ means entry into the rollover-insurance pool. The payoff per unit of insured debt is

$$\Pi_i(e_i, \alpha; b, N) = e_i \left[s - a(b, N, \alpha)\right]. \tag{84}$$

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

$$e_i(x_i) = 1 \quad \Longleftrightarrow \quad x_i \le x^*. \tag{85}$$

The inequality follows from the convention that higher signals indicate worse participation conditions. Under this strategy, elastic participation is

$$\tilde{\alpha}(\theta; x^*) = F\left(\frac{x^* - \theta}{\eta}\right), \tag{86}$$

and total participation is

$$\alpha(\theta; x^*) = \underline{\alpha} + (1 - \underline{\alpha})F\left(\frac{x^* - \theta}{\eta}\right). \tag{87}$$

D.3 Threshold pricing and price invariance

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

$$V(x_i; x^*, s, b, N) = s - \mathbb{E}\left[a\big(b, N, \alpha(\theta; x^*)\big) \mid x_i\right]. \tag{88}$$

The marginal insurer is indifferent when

$$V(x^*; x^*, s, b, N) = 0, \tag{89}$$

or equivalently

$$s = \mathbb{E}\left[a\big(b, N, \alpha(\theta; x^*)\big) \mid x_i = x^*\right]. \tag{90}$$

Monotonicity of $a(b, N, \alpha)$ in α, together with the fact that $\alpha(\theta; x^*)$ is decreasing in θ, gives the single-crossing property that supports threshold equilibria.

Lemma 1 (Laplacian price invariance and threshold equilibria)**.** *Fix any threshold x^*. Under the flat prior over θ, conditional on the marginal signal $x_i = x^*$, the total participation state $\alpha(\theta; x^*)$ is uniformly distributed on $[\underline{\alpha}, 1]$. Hence the marginal-participant premium implied by (90) is independent of x^* and equals*

$$\bar{a}(b, N) = \frac{1}{1 - \underline{\alpha}} \int_{\underline{\alpha}}^{1} a(b, N, \alpha)\, d\alpha. \tag{91}$$

For $s = \bar{a}(b, N)$, each cutoff x^* satisfies the marginal indifference condition, and single-crossing makes the cutoff a best response to itself. Monotone threshold equilibria therefore exist for every cutoff x^*, and all carry the same marginal-participant price $\bar{a}(b, N)$.

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

$$f\left(\frac{x^* - \theta}{\eta}\right). \tag{92}$$

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

$$z = F\left(\frac{x^* - \theta}{\eta}\right). \tag{93}$$

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

$$\frac{dz}{d\theta} = -\frac{1}{\eta} f\left(\frac{x^* - \theta}{\eta}\right). \tag{94}$$

The likelihood term is exactly offset by the Jacobian in the change of variables from θ to z: the transformed density is proportional to

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

which is constant. Therefore z is uniform on $[0, 1]$ conditional on $x_i = x^*$. Since

$$\alpha = \underline{\alpha} + (1 - \underline{\alpha})z, \tag{95}$$

participation α is uniform on $[\underline{\alpha}, 1]$. The distribution is independent of x^*, so the right-hand side of (90) is the same for every threshold. Substitution gives (91). With $s = \bar{a}(b, N)$, the marginal indifference condition holds at every cutoff x^*. The single-crossing property described above then implies that insurers with $x_i < x^*$ enter and insurers with $x_i > x^*$ stay out, so each cutoff is a monotone threshold equilibrium and all such equilibria carry the common price $\bar{a}(b, N)$. □

Substituting (81) into (91) gives

$$s(b, N) = \frac{1}{1 - \underline{\alpha}} \int_{\underline{\alpha}}^{1} \Gamma\big(\zeta(b, N, \alpha, \delta)\big)\, d\alpha \tag{96}$$

$$= \frac{1}{1 - \underline{\alpha}} \int_{\underline{\alpha}}^{\bar{q}(b,N,\delta)} \Gamma\left(\frac{\mu(\delta)b - \alpha N}{\alpha}\right) d\alpha, \tag{97}$$

where $\bar{q}(b, N, \delta) = \min\{q(b, N, \delta), 1\}$. This is the finite-rollover-intensity counterpart of (18).

D.4 Interpretation

The threshold-participation foundation makes $s(b, N)$ the zero-profit compensation required by the marginal rollover insurer for the state-contingent assessment $a(b, N, \alpha)$. Low-participation states create individual balance-sheet stress because a given rollover shortfall is borne by fewer active units, and the flat-prior Laplacian belief averages the active-unit assessment uniformly over $[\underline{\alpha}, 1]$. In the macro block, the resulting stationary spread is paid on insured government debt. When b rises, the market-clearing assessment changes for the rolled-over stock, which is why the marginal fiscal cost of producing safe-asset services contains both the direct premium term and the stock-wide rollover term bs_b.

E Comparative Statics of the Rollover-Insurance Premium

This appendix differentiates the premium formula from Appendix D. The finite-maturity formulas use $\mu(\delta) = 1 - e^{-\delta}$; the benchmark in Section 2.3 sets $\mu(\delta) = 1$. In this appendix, the finite-rollover premium $s(b, N)$ is the base premium denoted $s^{base}(b, N, \delta)$ in Section 4, with δ held fixed. Derivatives with respect to b and N hold δ fixed unless indicated otherwise. The active stress region is $\underline{\alpha} < q(b, N, \delta) < 1$, with $q(b, N, \delta) = \mu(\delta)b/N$. Write

$$\zeta \equiv \zeta(b, N, \alpha, \delta) \equiv \frac{\mu(\delta)b - \alpha N}{\alpha} \tag{98}$$

for per-active-unit stress in participation state α. Since per-active-unit stress is zero at the boundary, $\zeta = 0$ when $\alpha = q(b, N, \delta)$, and because $\Gamma(0) = 0$, the first derivative has no moving-boundary term:

$$s_b(b, N) = \frac{\mu(\delta)}{1 - \underline{\alpha}} \int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{\Gamma'(\zeta)}{\alpha} \, d\alpha > 0, \tag{99}$$

$$s_N(b, N) = -\frac{1}{1 - \underline{\alpha}} \int_{\underline{\alpha}}^{q(b,N,\delta)} \Gamma'(\zeta) \, d\alpha < 0. \tag{100}$$

More debt raises per-active-unit stress in every stress state; more capacity lowers it.

The total fiscal cost of producing safe-asset services is

$$S(b, N) = bs(b, N), \tag{101}$$

so

$$S_b(b, N) = s(b, N) + bs_b(b, N). \tag{102}$$

This is the stock-wide marginal-cost formula: the marginal unit pays the premium, and the new debt stock changes the premium paid on outstanding insured debt.

The curvature result follows from differentiating (99). The interior derivative is

$$\frac{\partial}{\partial b} \left[\frac{\mu(\delta)\Gamma'(\zeta)}{\alpha} \right] = \frac{\mu(\delta)^2 \Gamma''(\zeta)}{\alpha^2}, \tag{103}$$

and the moving-boundary term is

$$\frac{\mu(\delta)}{1 - \underline{\alpha}} q_b(b, N, \delta) \frac{\Gamma'_+(0)}{q(b, N, \delta)}. \tag{104}$$

Hence

$$s_{bb}(b,N) = \frac{1}{1-\underline{\alpha}}\left[\mu(\delta)^2\int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{\Gamma''(\zeta)}{\alpha^2}\,d\alpha + \mu(\delta)q_b(b,N,\delta)\frac{\Gamma'_+(0)}{q(b,N,\delta)}\right] \geq 0. \tag{105}$$

The first term is the curvature of the individual balance-sheet stress cost inside the stress region. The second is the boundary effect from expanding the set of stress states. Since

$$S_{bb}(b,N) = 2s_b(b,N) + bs_{bb}(b,N), \tag{106}$$

it follows that

$$S_{bb}(b,N) > 0. \tag{107}$$

The financial-capacity cross-partial uses

$$S_{bN}(b,N) = s_N(b,N) + bs_{bN}(b,N). \tag{108}$$

Differentiating (99) with respect to N, the interior derivative is

$$\frac{\partial}{\partial N}\left[\frac{\mu(\delta)\Gamma'(\zeta)}{\alpha}\right] = -\frac{\mu(\delta)\Gamma''(\zeta)}{\alpha}, \tag{109}$$

and the moving-boundary term is

$$\frac{\mu(\delta)}{1-\underline{\alpha}}\,q_N(b,N,\delta)\frac{\Gamma'_+(0)}{q(b,N,\delta)}. \tag{110}$$

Because

$$q_N(b,N,\delta) = -\frac{1}{N}q(b,N,\delta), \tag{111}$$

this gives

$$\begin{aligned} s_{bN}(b,N) = &-\frac{\mu(\delta)}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{\Gamma''(\zeta)}{\alpha}\,d\alpha \\ &- \frac{\mu(\delta)}{(1-\underline{\alpha})N}\Gamma'_+(0) < 0. \end{aligned} \tag{112}$$

Together with $s_N < 0$, this implies

$$S_{bN}(b,N) < 0. \tag{113}$$

Thus higher financial capacity lowers the marginal fiscal cost of producing safe government-debt claims.

Finally, let $\mu_\delta(\delta) \equiv e^{-\delta}$. In the active stress region,

$$s_\delta(b,N,\delta) = \frac{\mu_\delta(\delta)b}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{\Gamma'(\zeta)}{\alpha}\,d\alpha > 0, \tag{114}$$

and

$$s_{b\delta}(b,N,\delta) = \frac{\mu_\delta(\delta)}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{\Gamma'(\zeta)}{\alpha}\, d\alpha + \frac{\mu(\delta)\mu_\delta(\delta) b}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{\Gamma''(\zeta)}{\alpha^2}\, d\alpha + \frac{\mu_\delta(\delta)\Gamma'_+(0)}{1-\underline{\alpha}}. \tag{115}$$

The first two terms are interior effects; the last is the moving-boundary term. Since $S_b = s + bs_b$,

$$S_{b\delta}(b,N,\delta) = s_\delta(b,N,\delta) + bs_{b\delta}(b,N,\delta) > 0. \tag{116}$$

This derivative is the finite-rollover-intensity comparative static used in Section 4.

F Closed-Form Illustration

This appendix solves the premium block in closed form under a linear per-active-unit stress-cost schedule. The closed form draws the figures and illustrates the general comparative statics derived in Appendix E.

Assume

$$B(\alpha,N) = \alpha N, \qquad \Gamma(z) = \varphi z, \qquad \varphi > 0. \tag{117}$$

The stress threshold is the debt level at which the inelastic base just absorbs rollover exposure. The benchmark figures impose $\mu(\delta) = 1$, so $\widehat{b}(N) = \underline{\alpha}N$. The formulas below are written for the more general finite-δ case and collapse to the benchmark by setting $\mu(\delta) = 1$:

$$\widehat{b}(N) = \frac{\underline{\alpha}N}{\mu(\delta)}. \tag{118}$$

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

$$\widehat{b}(N) < b < \frac{N}{\mu(\delta)}, \tag{119}$$

the interior region satisfies $\underline{\alpha} < q(b,N,\delta) < 1$. Substituting the linear stress-cost schedule into (97) gives

$$\begin{aligned} s(b,N) &= \frac{\varphi}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{\mu(\delta)b - \alpha N}{\alpha}\, d\alpha \\ &= \frac{\varphi}{1-\underline{\alpha}} \left[\mu(\delta) b \int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{1}{\alpha}\, d\alpha - N \int_{\underline{\alpha}}^{q(b,N,\delta)} 1\, d\alpha \right] \\ &= \frac{\varphi}{1-\underline{\alpha}} \left[\mu(\delta) b \log\left(\frac{\mu(\delta) b}{\underline{\alpha}N} \right) - \mu(\delta) b + \underline{\alpha}N \right]. \end{aligned} \tag{120}$$

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

$$\begin{aligned} s_b(b,N) &= \frac{\varphi}{1-\underline{\alpha}}\left[\mu(\delta)\log\left(\frac{\mu(\delta)b}{\underline{\alpha}N}\right)+\mu(\delta)-\mu(\delta)\right] \\ &= \frac{\varphi\mu(\delta)}{1-\underline{\alpha}}\log\left(\frac{\mu(\delta)b}{\underline{\alpha}N}\right) > 0. \end{aligned} \tag{121}$$

The middle $+\mu(\delta)$ comes from differentiating $\mu(\delta)b\log(\mu(\delta)b/(\underline{\alpha}N))$, and it cancels the derivative of $-\mu(\delta)b$. The second derivative is

$$s_{bb}(b,N) = \frac{\varphi\mu(\delta)}{(1-\underline{\alpha})b} > 0. \tag{122}$$

The derivative with respect to financial capacity is

$$s_N(b,N) = -\frac{\varphi}{1-\underline{\alpha}}\left[q(b,N,\delta)-\underline{\alpha}\right] < 0. \tag{123}$$

The marginal safety-production cost is

$$\begin{aligned} S_b(b,N) &= s(b,N) + bs_b(b,N) \\ &= \frac{\varphi}{1-\underline{\alpha}}\left[\mu(\delta)b\log\left(\frac{\mu(\delta)b}{\underline{\alpha}N}\right)-\mu(\delta)b+\underline{\alpha}N\right] + \frac{\varphi}{1-\underline{\alpha}}\left[\mu(\delta)b\log\left(\frac{\mu(\delta)b}{\underline{\alpha}N}\right)\right] \\ &= \frac{\varphi}{1-\underline{\alpha}}\left[2\mu(\delta)b\log\left(\frac{\mu(\delta)b}{\underline{\alpha}N}\right)-\mu(\delta)b+\underline{\alpha}N\right]. \end{aligned} \tag{124}$$

This expression isolates the repricing force: the log term appears twice, once from the average premium on the stock and once from the effect of debt on the premium itself. When rollover exposure exceeds the full inelastic-support capacity, $b \geq N/\mu(\delta)$, the upper integration limit in (97) is one. The corresponding marginal cost is

$$S_b(b,N) = \frac{\varphi}{1-\underline{\alpha}}\left[2\mu(\delta)b\log\left(\frac{1}{\underline{\alpha}}\right)-N(1-\underline{\alpha})\right], \tag{125}$$

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

In the interior stress region, the cross-partial is

$$S_{bN}(b,N) = \frac{\varphi}{1-\underline{\alpha}}\left[-\frac{2\mu(\delta)b}{N}+\underline{\alpha}\right] < 0. \tag{126}$$

The sign follows because $\mu(\delta)b > \underline{\alpha}N$. In the saturated region $b \geq N/\mu(\delta)$, (125) gives $S_{bN}(b,N) = -\varphi < 0$. Hence higher financial capacity lowers the marginal fiscal cost of producing safe-asset services throughout the stress region.

If the Laffer peak lies in the interior stress region, it solves

$$\lambda = \frac{\varphi}{1-\underline{\alpha}}\left[2\mu(\delta)b^L\log\left(\frac{\mu(\delta)b^L}{\underline{\alpha}N}\right)-\mu(\delta)b^L+\underline{\alpha}N\right]. \tag{127}$$

If the peak lies in the saturated region, the right-hand side is instead given by (125) evaluated at $b = b^L$. The left-hand side is the non-Ricardian wedge. The right-hand side is the marginal fiscal cost of producing safe-asset services. A larger N or larger inelastic base $\underline{\alpha}$ lowers the right-hand side for a given b and pushes the peak outward. A larger unit support cost φ or rollover intensity δ raises the right-hand side through the rollover share $\mu(\delta)$ and pulls the peak inward.

G Derivation of the Balance-Sheet Multiplier

This appendix derives the local amplification formula discussed in Section 4. The state variables are b, N, and the rollover intensity δ. The object being priced is the premium required to compensate the financial system for rollover insurance. The local price-feedback block links stress assessments, effective capacity, future premia, and the value of outstanding long debt.

The baseline rollover assessment premium is $s^{base}(b, N, \delta)$ from (31). Its sensitivity to financial capacity is $s_N^{base}(b, N, \delta) < 0$, so a capacity loss $dN < 0$ raises the baseline premium by approximately $-s_N^{base}(-dN)$. Let $s_N^T \equiv \partial s^T / \partial N < 0$ denote the local sensitivity of the total premium to financial capacity.

A stress assessment reduces future financial capacity. In the price-loss loop below, N is measured in units of mark-to-market long-bond capacity, so a unit valuation loss is a unit capacity loss. This is a units convention for the local representation; equivalently, any conversion coefficient for that valuation channel is absorbed into the local sensitivity $|s_N^T|$.

G.1 Price sensitivity of long debt

Let P denote the price of a Poisson-maturity long bond. The bond pays a unit flow while alive and matures with rollover intensity δ. If the safe discount rate is r and the required total premium is s^T, the effective discount-plus-hazard rate is $\rho^T(b, N, \delta) = r + s^T(b, N, \delta) + \delta$, as in (32). The price is the present value of the unit flow:

$$P = \int_0^\infty e^{-\rho^T u}\, du = \frac{1}{\rho^T}. \tag{128}$$

A persistent increase ds^T in the premium raises ρ^T one-for-one. Therefore

$$d \log P = -\frac{1}{\rho^T}\, ds^T. \tag{129}$$

The local log-price sensitivity to a persistent premium increase is $1/\rho^T$. The persistence qualifier matters: the valuation formula applies to the component of the stress-induced premium increase that investors expect to persist over the effective life of the long bond.

G.2 The feedback loop

Let s^R denote the rollover component after balance-sheet feedback. Suppose a unit loss of capacity occurs. The direct effect is to raise the baseline rollover-insurance premium by $-s_N^{base}$. If long bonds are outstanding, this premium increase lowers their price. From (129), the valuation loss per unit of persistent premium increase is proportional to $1/\rho^T$. This valuation loss further erodes capacity and triggers another premium increase.

Let $|s_N^T|$ denote the absolute value of the local capacity derivative of the total premium. If a unit capacity loss raises the total premium by $|s_N^T|$, then a persistent premium increase of this size lowers the mark-to-market value of the long-bond position by $|s_N^T|/\rho^T$, using (129).

Since N is measured in mark-to-market capacity units, this valuation loss is a further capacity loss. The exposure scale has already been absorbed into the local derivative $|s_N^T|$; the remaining amplification is the price sensitivity $1/\rho^T$. The price-loss channel therefore has local size

$$\mathfrak{m} = \frac{|s_N^T|}{\rho^T}. \tag{130}$$

The multiplier $\mathfrak{m}$ measures the extra capacity loss created by the bond-price response to a premium increase. The self-reference enters through the $(1 + \mathfrak{m})$ term in the capacity derivative of the premium schedule. Thus the converged rollover component is multiplied by $1 + \mathfrak{m}$,

$$s^R = (1 + \mathfrak{m})s^{base}. \tag{131}$$

Holding $\mathfrak{m}$ fixed in the local derivative gives

$$-s_N^R = (1 + \mathfrak{m})(-s_N^{base}). \tag{132}$$

The first term $-s_N^{base}$ is the direct effect of capacity on the rollover-insurance premium. The term $\mathfrak{m}(-s_N^{base})$ is the additional effect generated by the price-loss feedback on outstanding long debt. The fixed point for $\mathfrak{m}$ below uses the total sensitivity $|s_N^T|$, so the feedback is already incorporated in the converged multiplier $1 + \mathfrak{m}$.

Let s^X denote term-premium compensation and define its direct sensitivity to financial capacity, holding ρ^T fixed, by

$$\psi_{\text{dir}}^X(b, N, \delta) \equiv -\left.\frac{\partial s^X(b, N, \delta)}{\partial N}\right|_{\rho^T} > 0. \tag{133}$$

Appendix H derives this object from the term-premium component. The total sensitivity of the required premium to financial capacity is therefore

$$|s_N^T| = (1 + \mathfrak{m})(-s_N^{base}) + \psi_{\text{dir}}^X. \tag{134}$$

Combining (134) and (130) gives the fixed point

$$\mathfrak{m} = \frac{(1 + \mathfrak{m})(-s_N^{base}) + \psi_{\text{dir}}^X}{\rho^T}. \tag{135}$$

Solving for $\mathfrak{m}$ yields

$$\mathfrak{m} = \frac{-s_N^{base} + \psi_{\text{dir}}^X}{\rho^T + s_N^{base}}, \qquad \rho^T + s_N^{base} > 0. \tag{136}$$

The regularity condition $\rho^T + s_N^{base} > 0$ keeps the feedback finite. This is the same local admissibility condition used in Proposition 3: the feedback from capacity to prices amplifies the base premium without making the local fixed point explosive. Substituting (136) into (131) gives

$$s^R(b, N, \delta) = \frac{\rho^T(b, N, \delta) + \psi_{\text{dir}}^X(b, N, \delta)}{\rho^T(b, N, \delta) + s_N^{base}(b, N, \delta)} s^{base}(b, N, \delta). \tag{137}$$

The numerator captures the direct term-premium sensitivity to financial capacity. The denominator captures the balance-sheet feedback: a higher premium lowers bond prices, reduces effective financial capacity, and raises the premium further through the direct premium sensitivity s_N^{base}. The formula is a local linearization around a steady state that shows how balance-sheet fragility magnifies the safety-production cost.

Differentiating the multiplier with respect to rollover intensity along the fixed-point premium schedule gives

$$M_\delta = \frac{\left(s_N^{base} - \psi_{\mathrm{dir}}^X\right)\rho_\delta^T + \left(\rho^T + s_N^{base}\right)\psi_{\mathrm{dir},\delta}^X - \left(\rho^T + \psi_{\mathrm{dir}}^X\right)s_{N\delta}^{base}}{(\rho^T + s_N^{base})^2}. \tag{138}$$

Here $\rho_\delta^T = 1 + s_\delta^T$, since the derivative is evaluated along the fixed-point total premium. The identity separates duration compression, the direct term-premium sensitivity, and the rollover-capacity sensitivity of the direct premium. As $\delta \to \infty$, $\rho^T \to \infty$, $s^X \to 0$, and $\psi_{\mathrm{dir}}^X \to 0$, so the multiplier in (137) converges to one.

H Term-Premium Compensation and Maturity Choice

This appendix gives the maturity-choice extension referenced in Section 4. The government chooses the rollover intensity δ, equivalently the frequency with which debt must be rolled over. A higher δ means shorter average maturity, faster rollover, and lower duration. A lower δ means longer average maturity, slower rollover, and higher duration.

H.1 Rollover component

The rollover component uses the same pricing formula as $s^{base}(b, N, \delta)$ in (31), now treating δ as the government's choice. Since rollover exposure is $\mu(\delta)b$, the derivative s_δ^{base} in (38) is positive in the stress region. Let $\mu_\delta(\delta) \equiv e^{-\delta}$. Differentiating (38) with respect to N, and using ζ from (98), gives an interior term and a moving-boundary term:

$$\begin{aligned} s_{\delta N}^{base}(b, N, \delta) = &-\frac{\mu_\delta(\delta)b}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{\Gamma''(\zeta)}{\alpha}\, d\alpha \\ &- \frac{\mu_\delta(\delta)b}{(1-\underline{\alpha})N}\Gamma'_+(0) < 0. \end{aligned} \tag{139}$$

The first term says that higher capacity lowers the marginal per-active-unit stress generated by shorter maturity in every stress state. The second term is the boundary effect: higher capacity shrinks the set of states in which a marginal increase in δ creates rollover stress.

H.2 Term-premium component

For term risk, consider a Poisson-maturity bond with effective discount-plus-hazard rate $\rho^T(b, N, \delta) = r + s^T(b, N, \delta) + \delta$, as in (32). As shown in Appendix G, the price of this bond is locally proportional to $1/\rho^T$, so its duration is

$$\mathcal{D}\sqcap\nabla = \frac{1}{\rho^T}. \tag{140}$$

If the variance of the persistent premium or rate shock is ς_P^2, local price variance is proportional to $\mathcal{D}\sqcap\nabla^2\varsigma_P^2$. If the market price of bearing term risk is $\vartheta(N)$, the term-premium component is

$$s^X(b,N,\delta) = \vartheta(N)\frac{\varsigma_P^2}{[\rho^T(b,N,\delta)]^2}, \tag{141}$$

with

$$\vartheta(N) = \bar{\vartheta}N^{-\eta_\vartheta}, \qquad \bar{\vartheta} > 0, \qquad \eta_\vartheta > 0. \tag{142}$$

Higher financial capacity lowers the market price of bearing term risk. Holding ρ^T fixed, the direct sensitivity is

$$\psi^X_{\text{dir}}(b,N,\delta) = -\left.\frac{\partial s^X}{\partial N}\right|_{\rho^T} = \frac{\eta_\vartheta}{N}s^X(b,N,\delta). \tag{143}$$

Holding s^T fixed inside ρ^T, the direct cross-partial of the term-premium term is

$$s^{X,dir}_{\delta N}(b,N,\delta) = \frac{2\eta_\vartheta}{N}\frac{\vartheta(N)\varsigma_P^2}{(\rho^T)^3} > 0. \tag{144}$$

This term has the opposite sign from (139): higher capacity lowers the price of term risk, which reduces the term-premium saving from shortening maturity.

H.3 Fixed point and maturity choice

Throughout this subsection, $M \equiv 1 + \mathfrak{m} = (\rho^T + \psi^X_{\text{dir}})/(\rho^T + s^{base}_N)$ is the balance-sheet amplification multiplier derived in Appendix G. The total premium solves the fixed point

$$s^T(b,N,\delta) = \frac{\rho^T + \psi^X_{\text{dir}}}{\rho^T + s^{base}_N}s^{base}(b,N,\delta) + \vartheta(N)\frac{\varsigma_P^2}{(\rho^T)^2}, \qquad \rho^T = r + s^T + \delta. \tag{145}$$

The first term is rollover insurance with balance-sheet amplification. The second term is compensation for term risk. The local contraction condition is maintained:

$$\sup_{s^T}\left|\frac{\partial}{\partial s^T}\left[\frac{\rho^T + \psi^X_{\text{dir}}}{\rho^T + s^{base}_N}s^{base} + \vartheta(N)\frac{\varsigma_P^2}{(\rho^T)^2}\right]\right| < 1, \tag{146}$$

which ensures that (145) defines a unique local premium schedule for each (b,N,δ). The government chooses δ to minimize that well-defined schedule:

$$\delta^*(b,N) \in \arg\min_{\delta>0} s^T(b,N,\delta). \tag{147}$$

At an interior optimum,

$$s^T_\delta(b,N,\delta^*) = 0, \qquad s^T_{\delta\delta}(b,N,\delta^*) > 0. \tag{148}$$

The first-order condition equates the marginal rollover cost of shorter maturity with the marginal term-premium saving from shorter maturity.

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* (148) *holds. If*

$$s^T_{\delta N}(b, N, \delta^*) < 0, \tag{149}$$

then

$$\frac{d\delta^*(b, N)}{dN} > 0. \tag{150}$$

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:

$$s^T_{\delta\delta}(b, N, \delta^*)\frac{d\delta^*}{dN} + s^T_{\delta N}(b, N, \delta^*) = 0. \tag{151}$$

Thus

$$\frac{d\delta^*}{dN} = -\frac{s^T_{\delta N}}{s^T_{\delta\delta}}. \tag{152}$$

The sign of $s^T_{\delta N}$ decomposes into the rollover-capacity term and the term-premium-price term:

$$s^T_{\delta N} = s^R_{\delta N} + s^X_{\delta N}. \tag{153}$$

Because $s^R = M s^{base}$, the rollover part of the cross-partial is

$$s^R_{\delta N} = M s^{base}_{\delta N} + M_N s^{base}_{\delta} + M_\delta s^{base}_{N} + M_{\delta N} s^{base}. \tag{154}$$

The first term is the direct rollover-capacity force. It is negative because $M > 0$ and $s^{base}_{\delta N} < 0$, as shown in (139). The remaining terms capture how the balance-sheet multiplier itself changes with maturity and capacity. They are part of the full cross-partial.

The direct term-premium cross-partial in (144) is positive. The total maturity-capacity cross-partial is therefore

$$s^T_{\delta N} = M s^{base}_{\delta N} + M_N s^{base}_{\delta} + M_\delta s^{base}_{N} + M_{\delta N} s^{base} + s^X_{\delta N}. \tag{155}$$

The maturity comparative static holds when this full expression is negative:

$$M s^{base}_{\delta N} + M_N s^{base}_{\delta} + M_\delta s^{base}_{N} + M_{\delta N} s^{base} + s^X_{\delta N} < 0. \tag{156}$$

Equivalently,

$$-\left(M s^{base}_{\delta N} + M_N s^{base}_{\delta} + M_\delta s^{base}_{N} + M_{\delta N} s^{base}\right) > s^X_{\delta N}. \tag{157}$$

Under this condition,

$$s^T_{\delta N} < 0. \tag{158}$$

A checkable benchmark fixes the amplification multiplier in the (δ, N) comparison. In that case $M_N = M_\delta = M_{\delta N} = 0$, and (154) becomes

$$s^R_{\delta N} = M s^{base}_{\delta N}. \tag{159}$$

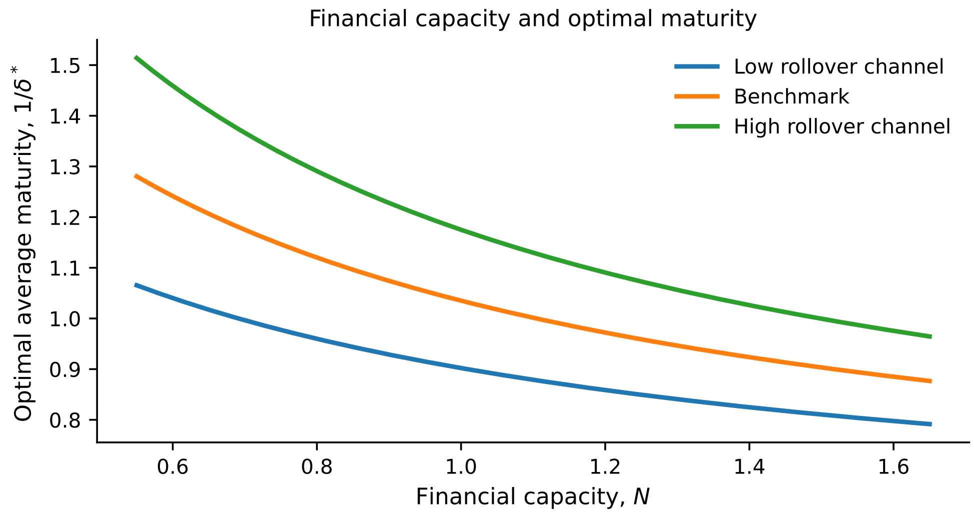

Figure 8: Financial capacity and optimal maturity. The figure plots the optimal average maturity $1/\delta^*$ against financial capacity N. The downward slope means that stronger financial capacity permits a higher optimal rollover intensity, while scarce financial capacity tilts issuance toward longer maturity. The three curves vary the strength of the rollover channel and the term-premium compensation to show the robustness of the comparative static.

Since $M > 0$ and $s_{\delta N}^{base} < 0$, this rollover-capacity term is negative. The full condition (156) then reduces to

$$-M s_{\delta N}^{base} > s_{\delta N}^{X}. \tag{160}$$

This benchmark records the primitive comparison obtained when the direct rollover-capacity channel is isolated. In the full model, the cross-partial in (155) includes the multiplier-derivative terms.

Together with the second-order condition, implicit differentiation gives

$$\frac{d\delta^*}{dN} = -\frac{s_{\delta N}^{T}}{s_{\delta\delta}^{T}} > 0. \tag{161}$$

Equivalently, lower rollover-insurance capacity lowers δ^*, which means longer maturity. Scarce rollover-insurance capacity therefore induces the government to term out the debt, trading off reduced rollover pressure against higher term-premium compensation.

Figure 8 illustrates the financial-capacity comparative static. The vertical axis reports optimal average maturity, $1/\delta^*$. As financial capacity rises, the government can sustain a higher rollover intensity and therefore chooses shorter average maturity. When financial capacity is scarce, rollover exposure is more costly and the optimum shifts toward longer maturity.

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.

Farhi, Emmanuel, and Iván Werning (2016). "A Theory of Macroprudential Policies in the Presence of Nominal Rigidities." *Econometrica*, 84(5), 1645–1704.

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.

Korinek, Anton, and Alp Simsek (2016). "Liquidity Trap and Excessive Leverage." *American Economic Review*, 106(3), 699–738.

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.

Spence, A. Michael (1975). "Monopoly, Quality, and Regulation." *The Bell Journal of Economics*, 6(2), 417–429.

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.