Inelastic Demand Meets Optimal Supply of Risky Sovereign Bonds*

Matias Moretti
University of Rochester

Lorenzo Pandolfi
University of Naples Federico II and CSEF

Sergio L. Schmukler
World Bank

German Villegas-Bauer
International Monetary Fund

Tomas Williams
George Washington University

August 10, 2026

Abstract

We study how investor demand shapes bond prices and sovereign borrowing. We develop a sovereign debt model with long-term debt, endogenous default, and imperfectly elastic demand. Selling more bonds lowers the price the government receives, even holding expected repayment fixed. This price impact discourages current issuance, but also restrains *future* issuance, reducing dilution and raising long-term bond prices. We first characterize when the dilution-curbing force outweighs the price-impact cost and raises borrowing, showing that the condition is easier to satisfy when default risk alone provides limited discipline against dilution. We then estimate the elasticity of investor demand from passive portfolio flows induced by bond-index rebalancing and map it into the model. Quantitatively, the dilution-curbing force dominates: less elastic demand lowers spreads and default risk while supporting *higher* debt. Since the market itself imposes discipline, smaller default costs are needed to match observed spreads, and fiscal rules have milder effects.

Keywords: sovereign debt, imperfectly elastic demand, debt dilution, market discipline, fiscal rules

JEL Codes: F34, F41, G15

*We thank Mark Aguiar, Manuel Amador, Cristina Arellano, Yan Bai, Anmol Bhandari, Javier Bianchi, Gaston Chaumont, Masao Fukui, Robin Greenwood, Chang He, Rohan Kekre, Rishabh Kirpalani, Marco Pagano, Diego Perez, Walker Ray, Alessandro Rebucci, and Elia Sartori for useful discussions and feedback. We also thank participants at the 2024 ASSA Annual Meetings, 2024 Barcelona SED, George Washington University, ifo Dresden Workshop on Macroeconomics and International Finance, IMF Sovereign Debt Workshop, LACEA-LAMES 2024 Annual Meeting, NBER IFM 2024 Spring Conference, NBER 2024 Emerging Markets Conference, Minneapolis Fed, Richmond Fed, and University of Virginia. We are grateful to Joseph Lansley, Jefferson Martinez Saavedra, Xinyu Wei, and Patricio Yunis for research assistance. The World Bank Chile Research and Development Center and Knowledge for Change Program and the George Washington University Facilitating Fund provided financial support for this paper. Lorenzo Pandolfi gratefully acknowledges financial support from the European Union–NextGenerationEU–and the Italian Ministry of Universities and Research (PRIN Project No. 20227N3BMK), and from the UniCredit Foundation. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the IMF, the World Bank, those of their Executive Directors, or the governments they represent. Moretti: matias.moretti@rochester.edu. Pandolfi: lorenzo.pandolfi@unina.it. Schmukler: sschmukler@worldbank.org. Villegas-Bauer: gvillegas-bauer@imf.org. Williams: tomaswilliams@email.gwu.edu.

1 Introduction

Governments in emerging markets rely heavily on bonds issued in international capital markets for financing. Policymakers frequently emphasize the importance of who buys these bonds and whether investors can easily absorb additional issuances. Yet most quantitative sovereign debt models abstract from the role of investor demand, despite a growing asset pricing literature documenting downward-sloping demand schedules across asset markets. In this paper, we ask how investor demand shapes bond prices, default risk, and sovereign borrowing.

To answer this question, we formulate a sovereign debt model with defaultable long-term bonds and imperfectly elastic investor demand. Bond prices depend not only on default risk, but also on the quantity issued: selling more bonds lowers the price the government receives even when expected repayment is unchanged. We call this the *price impact* of issuance, and it affects borrowing through two opposing forces. First, issuing more today lowers the price received, discouraging current borrowing. Second, and specific to risky long-term debt, the same price impact makes *future* issuance less attractive, curbing dilution. When the government cannot commit to future issuance, investors anticipate that future borrowing will dilute the bonds sold today, so they require higher spreads up front. By restraining future borrowing, price impact lowers expected dilution and raises today's bond prices.

Analytically, we characterize these forces and derive conditions under which the dilution-curbing force outweighs the price-impact cost, so that less elastic demand leads the government to borrow *more* today. We then estimate the demand elasticity from variation in passive benchmark-tracking flows and map it into the model. Quantitatively, we find that the dilution-curbing force dominates. Three results follow. First, relative to the perfectly elastic case, spreads are lower, default risk falls, and the government sustains a *larger* stock of debt. Second, because price impact provides a source of restraint beyond default risk, our model matches comparable spreads and debt with smaller default costs. Third, fiscal rules such as debt ceilings have more limited effects, as the market already restrains future borrowing.

Our baseline framework adds a stylized demand side to an otherwise standard infinite-horizon sovereign debt model. Investor demand is downward sloping, and its slope governs the price impact of issuance, which the government internalizes. Several microfoundations can generate a demand schedule of this form. Our preferred specification follows the mandate-based framework of Gabaix and Koijen (2021). In this interpretation, mandates specify reference positions, such as buy-and-hold exposures and passive benchmark tracking, and govern how strongly portfolios tilt toward bonds with higher expected excess returns. The strength of these tilts determines the demand elasticity.[1] This specification links directly to our empirical design and provides a parsimonious way to match the large price responses we estimate.

We formalize the dilution-curbing mechanism in a tractable three-period version of the model. Here, long-term bonds issued at $t = 0$ can be diluted by additional borrowing at $t = 1$. We first characterize the perfectly elastic case, in which default risk is the only force curbing dilution. We then perturb the demand slope around that solution and derive a necessary and sufficient condition under which a steeper demand schedule curbs dilution enough for the government to raise period-0 borrowing. The condition is most easily satisfied when default risk provides little discipline against dilution on its own, and when the demand schedule steepens with borrowing, so that additional period-1 issuance pushes prices down more sharply. Because the benefit comes from shielding long-term claims against dilution, maturity is central. As period-0 bonds become shorter lived, less of that debt remains exposed to period-1 dilution, and the dilution-curbing force weakens.

To quantify these forces, we first need an estimate of the demand elasticity. The benchmark-tracking component of mandates maps naturally to the identifying variation we exploit. When index weights change, passive investors adjust their holdings mechanically. These adjustments shift the quantity held outside passive benchmark-tracking portfolios while leaving government debt unchanged, so the associated price response is informative about the demand elasticity.

We measure these shifts using the J.P. Morgan Emerging Markets Bond Index Global Diversified (EMBIGD), the largest benchmark index for emerging market sovereign bonds denominated in U.S. dollars. Specifically, we construct *flows induced by index rebalancing* (FIR) by combining assets passively tracking the EMBIGD with monthly changes in its composition. Because actual index weights depend on bond prices, we instrument FIR using the EMBIGD's diversification rule, which caps country weights by face value. We use

[1] Analogous pricing functions can arise from preferred-habitat demand (Vayanos and Vila, 2021), non-pecuniary benefits (Greenwood et al., 2024; Choi et al., 2026), or other limits to arbitrage. We discuss these alternatives in Appendix A.1.

this rule to construct synthetic weights based only on face amounts and the resulting cap adjustments. We restrict attention to countries whose own eligible amount is unchanged between two consecutive rebalancing events. The identifying variation therefore comes from *other* countries' issuances and retirements, which alter the index's composition and shift each country's weight differently depending on how tightly its cap binds.

We combine this instrument with the known timing of rebalancing dates and estimate reduced-form price responses in narrow windows around them. Our estimates imply that a 1 percentage point (p.p.) increase in FIR, scaled by non-passive holdings, raises bond prices by 0.30%, corresponding to an inverse demand elasticity of -0.30. In yield terms, this implies a decline of 5.5 basis points in the yield to maturity. Responses are stronger for riskier countries, reaching nearly 0.40%, and close to zero for safer ones. The magnitude of these responses is large relative to what frictionless asset-pricing models would imply. Our model reproduces responses of this magnitude by anchoring demand around reference positions and limiting the sensitivity of portfolio tilts to expected excess returns.

Mapping the reduced-form price responses into the structural demand slope, however, is not straightforward. Because the sovereign's problem is dynamic and its debt is long term, passive-demand shifts unrelated to *current* fundamentals can still affect *future* borrowing and expected dilution. These policy responses feed back into expected repayment and, in turn, into current prices. To address this, we use indirect inference: we extend the model with secondary markets to replicate the empirical exercise and pin down the parameter governing the demand slope to match the reduced-form estimates. For the quantitative exercise, we use a mean-variance portfolio rule to govern how portfolios tilt around reference positions, so that the demand slope increases with repayment variance, consistent with the stronger estimated price responses for riskier bonds. Calibrating the model to Argentina, a common case study in the sovereign debt literature, we obtain an average model-implied inverse elasticity of -0.18. The gap between this value and the reduced-form estimate implies that over one-third of the response reflects endogenous changes in repayment prospects.

With the calibrated model, we quantify the aggregate implications of downward-sloping demand. Relative to a perfectly elastic case, our baseline features spreads roughly 350 basis points lower, reduced default risk, and an average debt-to-output ratio about

2 percentage points higher. Together, these differences imply a certainty-equivalent consumption gain of about one-third of a percentage point. In our calibration, we normalize the average *inelasticity premium*—the spread differential relative to a perfectly elastic case, holding policies fixed—to zero. As we increase the average premium, debt falls relative to our baseline. Yet the dilution-curbing force remains strong enough that, even when the premium accounts for 30% of spreads, the government still sustains more debt than in the perfectly elastic case. Finally, the differences in debt and spreads relative to the perfectly elastic case narrow as debt maturity shortens, consistent with the three-period model.

We then compare our model to a *recalibrated* perfectly elastic model that matches the same target moments for spreads and debt. With perfectly elastic demand, borrowing is disciplined by default risk, which depends on default costs such as output losses and exclusion from international markets. In our model, price impact provides an additional source of discipline. As a result, we find that sustaining similar spreads and debt requires smaller average output losses in default. This market discipline reshapes the timing of default, leading to earlier deleveraging and defaults that occur only after a prolonged downturn. It also alters the effects of fiscal rules: under a debt-to-output ceiling, the response in our baseline is modest, since price impact already curbs future issuance. In contrast, in the perfectly elastic model, the ceiling causes spreads and default risk to fall sharply, prompting the government to borrow more.

Related literature. We build on quantitative models of sovereign default (Aguiar and Gopinath, 2006; Arellano, 2008; Chatterjee and Eyigungor, 2012), and especially on work with long-term debt and dilution (Hatchondo et al., 2016; Aguiar et al., 2019; Aguiar and Amador, 2020; DeMarzo et al., 2023). We link these models to demand-based asset pricing (Koijen and Yogo, 2019), which studies how asset prices respond to quantities. Within fixed income, the demand-based asset pricing literature has focused mainly on safe government bonds.[2] Most closely related to our analysis, Choi et al. (2026) quantify how downward-sloping demand contributes to the underprovision of U.S. Treasuries.[3] For

[2]See, among others, Krishnamurthy and Vissing-Jorgensen (2012), Greenwood et al. (2015), Jiang et al. (2024), and Ray et al. (2024). Relatedly, Mian et al. (2025) study the fiscal implications of safe-debt demand and convenience yields.

[3]Kaldorf and Röttger (2025) study how collateral services shape prices and issuance of European sovereign bonds. Costain et al. (2025) add default risk to a preferred-habitat model of the euro-area yield curve but take supply as exogenous, remaining silent on dilution.

risky long-term debt, we show instead that this effect can reverse. By curbing dilution, less elastic demand can lead the government to issue *more* debt, and we characterize when this occurs.

Our emphasis on market discipline connects our framework to sovereign debt models featuring risk-averse lenders. In Aguiar et al. (2016), a higher price of default risk reduces the government's future incentive to issue debt in states where default is possible, lowering expected dilution.[4] We focus instead on a broader mechanism at the core of demand-based asset pricing: the elasticity of investor demand and the resulting price impact. Frictionless models, including those with risk aversion, typically imply demand that is far too elastic to match observed price responses (Petajisto, 2009; Gabaix and Koijen, 2021). Our demand specification is designed to reproduce the large responses we estimate while remaining agnostic about their source, which can include risk aversion, balance-sheet or regulatory constraints, liquidity frictions, or other limits to arbitrage.[5]

Because imperfectly elastic demand disciplines future issuance, our analysis relates to work on monetary and fiscal rules as disciplining devices in sovereign debt models (Aguiar et al., 2013; Alfaro and Kanczuk, 2017; Dovis and Kirpalani, 2020; Hatchondo et al., 2022a; Bianchi et al., 2023), and to a broader literature on rules under limited commitment or enforcement (Amador et al., 2006; Halac and Yared, 2021, 2022). In these studies, rules restrain future borrowing or other policy choices. In our model, part of this restraint comes from the market through the price impact of future issuance, making rules such as debt ceilings less effective.

Lastly, our empirical analysis builds on work showing that investor composition shapes bond pricing (Faia et al., 2025; Zhou, 2024; Fang et al., 2025) and on studies using index rebalancing to estimate price responses and demand elasticities (Harris and Gurel, 1986; Shleifer, 1986; Greenwood, 2005; Hau et al., 2010; Chang et al., 2014; Raddatz et al., 2017; Pandolfi and Williams, 2019; Pavlova and Sikorskaya, 2022; Beltran and He, 2025). Unlike these studies, we pair bond-level price responses with a structural model that endogenizes issuance and repayment. This allows us to disentangle the demand slope from endogenous changes in expected repayment and recover the elasticity relevant for

[4] Borri and Verdelhan (2010), Lizarazo (2013), and Arellano et al. (2017) also introduce risk-averse investors, but with one-period bonds, and so abstract from dilution.

[5] Passadore and Xu (2022), Chaumont (2025), and Moretti (2020) consider sovereign debt models in which liquidity premia rise with default risk through search frictions in over-the-counter secondary markets; Bigio et al. (2023) study related liquidity costs of issuance.

risky long-term sovereign debt.

The paper is structured as follows. Section 2 presents a sovereign debt model with endogenous default and imperfectly elastic demand. Section 3 outlines the empirical strategy and main estimates. Section 4 reports the quantitative results. Section 5 concludes.

2 A Sovereign Debt Model with Imperfectly Elastic Demand

We formulate a sovereign debt model with two main blocks. On the issuer side, the setup features long-term debt, lack of commitment, and endogenous default. On the demand side, we introduce a downward-sloping demand schedule whose slope measures how strongly prices respond to quantities, holding expected repayment fixed. We first present the infinite-horizon framework and then analyze a simplified three-period version that characterizes the key mechanisms analytically.

2.1 Environment

Consider a small open economy with incomplete markets. Output, y, is exogenous and follows a continuous Markov process with transition function $P_y(y_{t+1} \mid y_t)$. The economy is inhabited by a representative risk-averse household whose expected discounted utility is given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u\left(c_t\right), \tag{1}$$

where $\beta \in (0,1)$ is the discount factor, c denotes consumption, and $u(\cdot)$ is strictly increasing and concave.

The government enters each period with an outstanding stock of bonds B_{t-1}. Each unit matures with probability λ and pays a coupon ν otherwise. In each period, the government chooses next-period debt B_t to maximize Equation (1), subject to the economy's resource constraint. The government cannot commit to future issuance decisions and may also choose to default. Let $d_t \in \{0,1\}$ denote the default decision, where $d_t = 1$ indicates default. Default leads to temporary exclusion from international markets and an exogenous output loss, $\Phi(y_t)$.

If the country is currently not in default, the resource constraint is

$$c_t = y_t + q_t\left(B_t - (1-\lambda)B_{t-1}\right) - \left(\lambda + (1-\lambda)\nu\right)B_{t-1}. \tag{2}$$

Here, q_t denotes the bond price, which is endogenous and depends on the government's policies and investor demand. The term $B_t - (1-\lambda)B_{t-1}$ captures new bond issuance, while $(\lambda + (1-\lambda)\nu)\, B_{t-1}$ represents current debt service. If the government is in default, it is excluded from debt markets, and the resource constraint simplifies to $c_t = y_t - \Phi(y_t)$.

2.2 Investor Demand

Bonds are priced by global investors in international markets. We introduce a minimal departure from standard sovereign debt models by adding a downward-sloping demand schedule.

We allow part of investor demand to be price-insensitive (perfectly inelastic) and tied to holdings that do not respond to contemporaneous expected excess returns. We denote this component by $\mathcal{B}_t$ and call it *reference demand*. It can capture buy-and-hold exposures, benchmark tracking, and preferred-habitat motives, among others. The term $B_t - \mathcal{B}_t$ is bond supply net of these reference positions. For brevity, we refer to it as *residual supply*. The demand slope, $\mathcal{S}_t$, is the central object of our analysis because it governs price impact. It measures how much the bond price must adjust per additional unit of residual supply, holding expected repayment fixed.

With this notation, the bond price is

$$q_t = \frac{\mathbb{E}_t(\mathcal{R}_{t+1})}{r^f} - \mathcal{S}_t\,(B_t - \mathcal{B}_t), \tag{3}$$

where $\mathcal{R}_{t+1}$ denotes next-period repayment per unit and r^f is the gross risk-free rate. The first term is the discounted expected repayment, the price that would prevail under perfectly elastic demand. The second term is the quantity-based price adjustment implied by the demand schedule. Our formulation is minimal by design. It adds a single term beyond discounted expected repayment, keeping the model tractable while capturing the price impact of issuance.

Equation (3) can be obtained from several demand structures, as analyzed in Appendix A.1. Our preferred interpretation follows the mandate-based framework of Gabaix and Koijen (2021). In this interpretation, portfolio mandates determine the reference positions summarized by $\mathcal{B}_t$. They also govern how strongly portfolios tilt around those reference positions in response to expected excess returns, which determines $\mathcal{S}_t$. Preferred-habitat motives, risk constraints, and non-pecuniary holding motives can generate an analogous pricing equation.

For the quantitative analysis in Section 4, we impose a mean-variance specification for the portfolio tilts, so that they respond to expected excess returns relative to repayment variance. This yields the following demand slope:

$$\mathcal{S}_t = \frac{\kappa}{r^f}\,\mathbb{V}_t(\mathcal{R}_{t+1}), \tag{4}$$

where $\kappa \geq 0$ captures how tightly mandates limit these portfolio tilts. When $\kappa = 0$, the tilts are unconstrained, so demand is perfectly elastic and bond prices equal discounted expected repayment. When $\kappa > 0$, mandates constrain the tilts, demand is imperfectly elastic, and the price impact of a given change in residual supply is stronger when repayment variance is high, consistent with the evidence in Section 3.

Reference demand $\mathcal{B}_t$, in turn, plays two roles in the quantitative analysis. First, it maps directly to our identification strategy: index rebalancing shifts benchmark-tracking holdings, providing the variation we use in Section 3 to estimate the reduced-form demand elasticity. Second, it helps separate the price impact of issuance from the average compensation investors require for holding the debt, allowing the model to match the large price responses we document. Market discipline and the dilution-curbing force we emphasize, however, do not hinge on a positive reference demand. In the three-period environment, we explicitly characterize the dilution-curbing force for any $\mathcal{B}$, including $\mathcal{B} = 0$.

2.3 Government Problem

We now embed the pricing equation (3) into the government's infinite-horizon problem. To keep the state space parsimonious, we specify reference demand as $\mathcal{B}(\tau, B')$, where B' is the end-of-period debt stock. The dependence on B' allows reference demand to scale with outstanding debt. The state variable τ is an exogenous shifter of these reference positions and follows a continuous Markov process with transition function $P_\tau(\tau' \mid \tau)$.

The government problem is as follows. If not in default, its value function is

$$V\left(y, \tau, B\right) = \max\left\{V^r\left(y, \tau, B\right), V^d\left(y, \tau\right)\right\}, \tag{5}$$

where $V^r(\cdot)$ denotes the value under repayment and $V^d(\cdot)$ the value under default. Under repayment, it chooses the next-period debt stock B' to solve:

$$V^r\left(y, \tau, B\right) = \max_{B'} \;\; u(c) + \beta\,\mathbb{E}_{s'|s} V\left(y', \tau', B'\right), \tag{6}$$

$$\text{subject to} \quad c = y + q(y, \tau, B')\left(B' - (1-\lambda)B\right) - \left(\lambda + (1-\lambda)\nu\right)B,$$

where $q(y, \tau, B')$ is the bond price schedule as a function of current states and the next-period debt choice. The expectation operator $\mathbb{E}_{s'|s}$ is taken over next-period exogenous states $s' \equiv (y', \tau')$, conditional on the current exogenous state $s \equiv (y, \tau)$.

While in default, the government is excluded from debt markets and cannot issue new debt. It exits default with probability χ, with no recovery value. The default value is

$$V^d(y, \tau) = u(y - \Phi(y)) + \beta\, \mathbb{E}_{s'|s}\left[\chi V(y', \tau', 0) + (1-\chi)\, V^d(y', \tau')\right]. \tag{7}$$

Given the next-period default choice, repayment per unit of debt B' is

$$\mathcal{R}(y', \tau', B') \equiv [1 - d(y', \tau', B')]\,[\lambda + (1-\lambda)(\nu + q(y', \tau', B''))], \tag{8}$$

where $d(y', \tau', B')$ is next period's default decision and $B'' \equiv B'(y', \tau', B')$ is the next-period debt choice under repayment.

Substituting the repayment function into Equation (3) gives the following recursive bond price schedule:

$$q(y, \tau, B') = \frac{1}{r^f}\, \mathbb{E}_{s'|s}\mathcal{R}(y', \tau', B') \;-\; \mathcal{S}(y, \tau, B')\,(B' - \mathcal{B}(\tau, B')), \tag{9}$$

where $\mathcal{S}(y, \tau, B')$ is the demand slope evaluated at the current state and debt choice. The first term is discounted expected repayment, which falls with default risk. The second term is the quantity-based price adjustment.

Together, the government bond policy, default rule, repayment function, and pricing schedule characterize the recursive Markov equilibrium (Appendix A.2 provides the equilibrium definition). This infinite-horizon model is the environment we take to the data.

2.4 A Three-Period Illustration of Debt Dilution and Market Discipline

We specialize the framework to three periods to characterize analytically how imperfectly elastic demand shapes debt issuance and dilution. For tractability, we assume that the government has linear preferences $\sum_{t=0}^{2} \beta^t c_t$ and deterministic endowments in $t \in \{0, 1\}$ (normalized to 0). The only uncertainty is the terminal endowment in period $t = 2$.[6] We also fix reference demand at a constant level, $\overline{\mathcal{B}}$, to better isolate the role of the demand slope.[7] Appendix A.3 provides all derivations.

In periods $t \in \{0, 1\}$, the government issues bonds that promise one unit of consump-

[6]Linear preferences remove the consumption-smoothing motive, so borrowing reflects the desire to front-load consumption (since $\beta < 1/r^f$). This simplifies the analytical characterization while preserving the key forces we study below: the price-impact cost and the dilution-curbing force.

[7]We relax this assumption in our quantitative analysis, where reference demand scales with outstanding debt, given a calibrated process for τ.

tion in $t = 2$. Issuance at $t = 0$ consists of two-period (long-term) bonds, denoted B_0, while issuance at $t = 1$ consists of one-period bonds in amount $B_1 - B_0$, making B_1 the total stock of debt due at $t = 2$. Since $(y_0, y_1) = (0, 0)$, the consumption levels are $c_0 = q_0 B_0 \geq 0$ and $c_1 = q_1(B_1 - B_0) \geq 0$, where q_0 and q_1 are the bond prices. Although the model structure is heavily simplified, the key dilution channel is preserved: debt issued at $t = 0$ can be diluted by additional borrowing at $t = 1$ before repayment is realized.

In the final period, the endowment y_2 is stochastic with distribution function $F(y)$. After observing y_2, the government chooses between repayment and default. We impose a proportional output cost of default, $\Phi(y_2) \equiv y_2/\phi$ with $\phi \geq 1$, which implies the default policy $d(B_1, y_2) = \mathbf{1}\{y_2 < \phi B_1\}$ and the default probability $\Pr(\text{default} \mid B_1) = F(\phi B_1)$. The terminal repayment per unit is $\mathcal{R}(B_1, y_2) = 1 - d(B_1, y_2)$.

We normalize the gross risk-free rate to $r^f = 1$, so that the period-1 bond price schedule is

$$q_1(B_1; \kappa) = 1 - F(\phi B_1) \; - \; \mathcal{S}(B_1; \kappa)\,(B_1 - \overline{\mathcal{B}}), \tag{10}$$

where $\mathcal{S}(B_1; \kappa)$ denotes the demand slope, with $\mathcal{S}(B_1; 0) = 0$ corresponding to the perfectly elastic case. For the purposes of the three-period model, we treat $\mathcal{S}(B_1; \kappa)$ as a general, smooth function of repayment risk. This lets us illustrate the implications of a downward-sloping demand schedule in a flexible way, without taking a stand on a specific microfoundation.[8] Since $\mathcal{S}(B_1; 0) = 0$, to first order in κ around $\kappa = 0$, we can write $\mathcal{S}(B_1; \kappa) \approx \kappa \mathcal{S}_\kappa(B_1)$, where $\mathcal{S}_\kappa(B_1) \equiv \frac{\partial \mathcal{S}(B_1;\kappa)}{\partial \kappa}\Big|_{\kappa=0}$, which we assume to be strictly positive. The period-0 price is $q_0(B_0; \kappa) = q_1(B_1(B_0; \kappa); \kappa)$, since no uncertainty is resolved between $t = 0$ and $t = 1$.[9]

As a first step, we characterize the perfectly elastic case, $\kappa = 0$, and describe the role of default risk in restraining borrowing and curbing dilution.[10] After that, we consider

[8]Given that the default cutoff is a linear function of B_1, we write $\mathcal{S}(\cdot)$ as a function of B_1 to simplify notation. We also keep the dependence on κ explicit to facilitate local comparative statics around the perfectly elastic case.

[9]In the infinite-horizon model, price impact applies separately to current and future issuance through each period's demand schedule. The three-period model simplifies this structure, collapsing the two pricing dates into the single schedule $q_1\,(B_1(B_0; \kappa); \kappa)$. Placing more period-0 debt raises the period-1 debt stock, B_1, lowering the price q_0 received at issuance—the price-impact cost. The same schedule disciplines period-1 issuance and curbs dilution—the dilution-curbing force. These are the two sides of Equation (14), analyzed below.

[10]The role of default risk in curbing dilution is well-established in the literature (e.g., Chatterjee and Eyigungor, 2012; Hatchondo et al., 2016; Aguiar et al., 2019; Aguiar and Amador, 2020). The closest analog in a stylized three-period setting is Hatchondo et al. (2022b).

how a steeper demand schedule shapes these incentives.

Default-risk discipline

Let F admit a strictly positive, continuously differentiable density f, and define the inverse hazard rate $H(y) \equiv \frac{1-F(y)}{f(y)}$. Under perfectly elastic demand, $H(\phi B_1)$ weighs the proceeds from issuing an additional unit of debt against the marginal increase in default risk. We also define the *debt pass-through* $\Upsilon(B_0;\kappa) \equiv \partial B_1(B_0;\kappa)/\partial B_0$, the marginal response of total debt B_1 to inherited debt B_0, which governs the extent of dilution. Let $\{B_0^\star, B_1^\star\}$ be a solution under $\kappa = 0$ at which default risk is non-degenerate. The following proposition characterizes this solution.

Proposition 1. *Under $\kappa = 0$, suppose $0 < F(\phi B_1^\star) < 1$ and the period-1 second-order condition $1-(1-\beta)H_y(\phi B_1^\star) > 0$ holds, where $H_y(y) \equiv \partial H(y)/\partial y$. Then period-1 issuance is strictly positive and given by*

$$B_1^\star - B_0^\star = \frac{1-\beta}{\phi} H(\phi B_1^\star) > 0, \tag{11}$$

and period-0 borrowing satisfies

$$B_0^\star = \frac{1-\beta}{\phi} \frac{H(\phi B_1^\star)}{\Upsilon(B_0^\star;0)}, \qquad \Upsilon(B_0^\star;0) = \frac{1}{1-(1-\beta)H_y(\phi B_1^\star)} > 0. \tag{12}$$

Equation (11) characterizes period-1 issuance. Absent commitment, the government issues additional debt at $t=1$, diluting the long-term bonds sold at $t=0$. The amount issued is governed by the inverse hazard rate. Intuitively, the government borrows until the gain from front-loading consumption by selling bonds at price $1-F(\phi B_1)$ is offset by the marginal price decline from higher default risk, $\phi f(\phi B_1)$, on the bonds issued at $t=1$.

Lenders anticipate this dilution and discount the long-term bonds sold at $t=0$, which feeds back into period-0 borrowing. Equation (12) captures this and shows that $B_0^\star$ is larger when the debt pass-through $\Upsilon(B_0^\star;0)$ is smaller. A lower pass-through means period-1 borrowing responds little to inherited debt, curbing dilution, raising bond prices, and allowing the government to issue more at $t=0$. The pass-through, in turn, is smaller when $H_y(\phi B_1^\star) < 0$ is large in magnitude, because in this case an extra unit of debt carried into $t=1$ sharply worsens the issuance trade-off behind Equation (11), discouraging additional period-1 borrowing. At $\kappa=0$, this is the only force curbing dilution, which we call *default-risk discipline*.

Under commitment to a debt path, by contrast, optimal period-1 issuance is zero and the government attains a higher level of period-0 debt (Proposition A.1 in Appendix A.3). As we show next, imperfectly elastic demand ($\kappa > 0$) can move the no-commitment solution toward the commitment allocation. A larger price impact at $t = 1$ curbs dilution by making period-1 borrowing less attractive, reducing the debt pass-through and supporting a higher B_0.

Market discipline

We formalize the dilution-curbing force by locally perturbing the perfectly elastic solution, increasing κ from zero. To isolate the slope effect from level shifts in prices, we set reference demand to $\overline{\mathcal{B}} = B_1^\star$. This implies that $\mathcal{S}(B_1;\kappa)(B_1 - \overline{\mathcal{B}}) = 0$ at the perfectly elastic allocation, so a small increase in κ changes the slope of the price schedule without mechanically lowering the price level there. We return to an arbitrary $\overline{\mathcal{B}}$ below.

Proposition 2 shows that a steeper demand schedule reduces bond issuance at $t = 1$, mitigating dilution. Proposition 3 characterizes when this dilution-curbing force offsets the price impact of placing additional period-0 debt, leading to an increase in long-term debt B_0.

Proposition 2. *Under the assumptions of Proposition 1, and assuming the slope $\mathcal{S}_\kappa(\cdot)$ is twice continuously differentiable, hold inherited debt fixed at $B_0 = B_0^\star$ and set $\overline{\mathcal{B}} = B_1^\star$. For sufficiently small $\kappa > 0$, period-1 issuance is lower than in the perfectly elastic case:*

$$B_1(B_0^\star;\kappa) - B_0^\star \; < \; B_1^\star - B_0^\star. \tag{13}$$

Proposition 3. *Consider an interior solution $(B_0^\star, B_1^\star)$ under $\kappa = 0$. Assume that $B_1(B_0;\kappa)$ and $\Upsilon(B_0;\kappa)$ are locally differentiable in κ at $(B_0^\star, 0)$, and that the period-0 second-order condition holds, as characterized in Appendix A.3. Set $\overline{\mathcal{B}} = B_1^\star$. For a small $\kappa > 0$, B_0 increases relative to the $\kappa = 0$ case if and only if*

$$\underbrace{-\phi f(\phi B_1^\star)\,\frac{\Upsilon_\kappa(B_0^\star)}{\Upsilon(B_0^\star;0)}}_{\textit{Dilution-curbing force}} \; > \; \underbrace{\mathcal{S}_\kappa(B_1^\star)\;\Upsilon(B_0^\star;0)}_{\textit{Price-impact cost}}, \tag{14}$$

where $\Upsilon_\kappa(B_0^\star) \equiv \frac{\partial}{\partial\kappa}\Upsilon(B_0^\star;\kappa)\Big|_{\kappa=0}$ is the derivative of the debt pass-through with respect to κ.[11]

[11] Condition (14) is stated with strict inequality. In the knife-edge case in which it holds with equality, the sign of $B_0(\kappa) - B_0^\star$ is determined by higher-order terms, which we do not characterize. The same qualification applies to Theorem 1 and Corollaries A.1 and A.2.

The left-hand side of Equation (14) is the dilution-curbing force. When $\Upsilon_\kappa(B_0^\star) < 0$, a higher κ makes period-1 borrowing respond less to inherited debt, reducing dilution and improving the period-0 borrowing margin by $-\phi f(\phi B_1^\star)\Upsilon_\kappa(B_0^\star)/\Upsilon(B_0^\star;0) > 0$. The right-hand side captures the price-impact cost: a higher B_0 raises B_1 by $\Upsilon(B_0^\star;0)$, and the steeper schedule lowers the price by $\mathcal{S}_\kappa(B_1^\star)$ per additional unit. The government increases B_0 when the dilution-curbing force dominates.[12]

Theorem 1 restates this condition in terms of the model's primitives. In particular, the condition is easier to satisfy when default-risk discipline is weaker and when the demand slope itself increases with additional borrowing.

Theorem 1. *Under the assumptions of Propositions 2 and 3, suppose the inverse hazard rate is twice continuously differentiable in a neighborhood of the $\kappa = 0$ allocation. Define $\varepsilon_{\mathcal{S}}(B_1) \equiv \frac{\partial \log \mathcal{S}_\kappa(B_1)}{\partial B_1}$ as the semi-elasticity of the demand slope with respect to debt. For sufficiently small $\kappa > 0$, B_0 increases relative to the $\kappa = 0$ case if and only if*

$$2\left(1 + H_y(\phi B_1^\star)\right) \;+\; \Omega\, H_{yy}(\phi B_1^\star) \;+\; \frac{2H(\phi B_1^\star)}{\phi}\varepsilon_{\mathcal{S}}(B_1^\star) \;>\; 1, \tag{15}$$

where $\Omega \equiv (1-\beta)\, H(\phi B_1^\star)\, \Upsilon(B_0^\star;0) > 0$.

The right-hand side of Condition (15) is the normalized price-impact cost of a steeper demand schedule. The left-hand side collects the forces that determine how strongly the steeper schedule curbs dilution. The first two terms describe the discipline already imposed by default risk alone. When $H_y(\phi B_1^\star) < 0$ is large in magnitude, the incentive to issue at $t = 1$ deteriorates quickly as debt rises (Proposition 1), leaving the demand slope little room to further curb dilution.[13] When $H_{yy}(\phi B_1^\star) > 0$, this default-risk discipline weakens as debt accumulates, making market discipline more valuable. The last term captures an additional source of discipline, which depends on how the slope changes with borrowing. When $\varepsilon_{\mathcal{S}}(B_1^\star) > 0$, increasing debt steepens the demand schedule, strengthening the dilution-curbing force and making Condition (15) easier to satisfy.[14]

Extensions. Appendix A.3 considers two extensions. First, we allow period-0 bonds to mature partly at $t = 1$, with maturity governed by $\lambda \in [0, 1)$, where a higher λ

[12] Given $\overline{\mathcal{B}} = B_1^\star$, a small increase in κ rotates the price schedule without shifting its level at the perfectly elastic allocation. At the $\kappa > 0$ solution, this level shift is second order in κ (see Appendix A.3).

[13] To see this more clearly, note that $1 + H_y(y) = -H(y) f_y(y)/f(y)$, where $f_y(y) \equiv \partial f(y)/\partial y$. Thus, the $1 + H_y$ term in Condition (15) is governed by the slope of the density at the default cutoff.

[14] A positive $\varepsilon_{\mathcal{S}}$ also implies that additional debt raises the price-impact cost. Under the normalization $\overline{\mathcal{B}} = B_1^\star$, however, this effect is second order in κ. Corollary A.2 characterizes the general case with arbitrary $\overline{\mathcal{B}}$.

corresponds to shorter maturity and $\lambda = 0$ corresponds to the case analyzed above. The terms in Condition (15) are unchanged, but shorter maturity introduces an additional term, increasing in λ, that makes the condition harder to satisfy because less of the period-0 debt remains outstanding to be diluted at $t = 1$ (Corollary A.1). Second, we extend the analysis to any $\overline{\mathcal{B}}$, including the limiting case with $\overline{\mathcal{B}} = 0$. When $\overline{\mathcal{B}} < B_1^\star$, residual supply is positive, so increasing κ mechanically lowers the price at the perfectly elastic allocation by $\mathcal{S}(B_1^\star; \kappa)\,(B_1^\star - \overline{\mathcal{B}})$, as in Equation (10). Corollary A.2 characterizes the resulting condition, which adds a level-effect term proportional to $B_1^\star - \overline{\mathcal{B}}$, and shows when a lower $\overline{\mathcal{B}}$ makes it harder for the dilution-curbing force to dominate.

The analytical characterizations in this subsection are local around the perfectly elastic allocation. Lemma 1 in Appendix A.3 shows that, in a first-order approximation around a $\kappa > 0$ solution that holds the local default-risk and demand-price slopes fixed, the dilution-curbing force declines as κ increases. Intuitively, once the demand schedule is already steep, further steepening has less room to reduce the debt pass-through. Thus, the local debt-expansion result in Proposition 3 and Theorem 1 need not extend to larger κ. The numerical exercise below illustrates this possibility.

Numerical illustration. We solve the model numerically, varying κ, β, λ, and $\overline{\mathcal{B}}$ to analyze how B_0 responds. We assume $y_2 \sim U[0, 1]$, so that $H_y(\phi B_1) = -1$ everywhere and default-risk discipline is constant. We also set $\mathcal{S}(B_1; \kappa) = \kappa \mathbb{V}(\mathcal{R}(B_1, y_2))$, matching the demand-slope specification used in the quantitative section. Locally, for small κ and $\overline{\mathcal{B}} = B_1^\star$, Theorem 1 reduces to $\beta > 0.71$ (see Appendix A.3). That is, for small κ, the government must be sufficiently patient for the dilution-curbing force to dominate. The threshold reflects the fact that a patient government chooses lower default risk, placing it in a region where repayment variance, $\mathbb{V}(\mathcal{R}) = F(\phi B_1)\,[1 - F(\phi B_1)]$, rises with borrowing. Since the demand slope is proportional to repayment variance, additional debt steepens the schedule, $\varepsilon_{\mathcal{S}}(B_1^\star) > 0$, so the marginal price impact rises with borrowing, disciplining period-1 issuance and curbing dilution.

Figure 1 reports the parameter regions where B_0 rises relative to the perfectly elastic case. Panel (a) varies β and κ. Near $\kappa = 0$, the boundary matches the local condition above, with period-0 borrowing increasing when the government is sufficiently patient. As κ rises, however, B_0 eventually falls below $B_0^\star$, indicating that the price-impact cost comes to dominate in this numerical example. Panel (b) shows that shorter maturity

Figure 1: Numerical illustration: A steeper demand slope

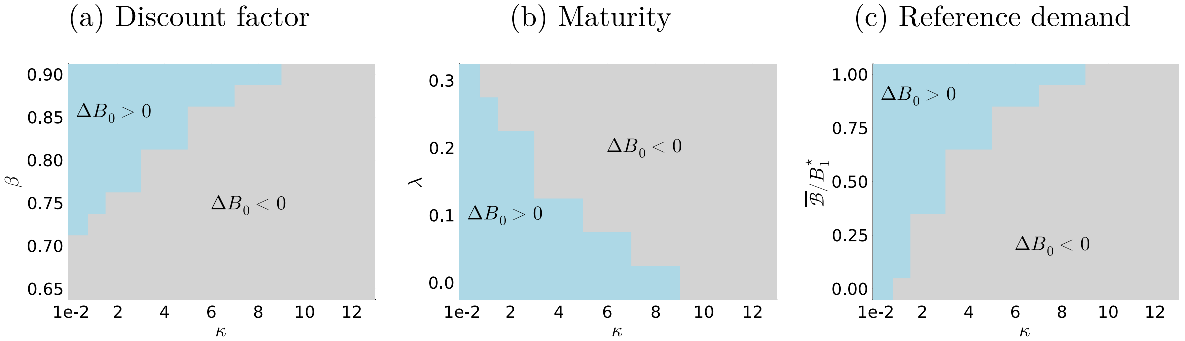

Note: Panels show parameter combinations for which B_0 increases with κ relative to the perfectly elastic solution. Baseline parameters are $y_2 \sim U[0,1]$, $\beta = 0.90$, $\phi = 1.5$, and $\overline{\mathcal{B}} = B_1^\star$. Panels vary (κ, β), (κ, λ), and $(\kappa, \overline{\mathcal{B}}/B_1^\star)$, respectively.

shrinks the region where B_0 rises, consistent with Corollary A.1. Panel (c) shows that the region also shrinks as reference demand falls below $B_1^\star$, consistent with the level effect in Corollary A.2.

Whether the dilution-curbing force dominates quantitatively, therefore, depends on the elasticity of investor demand. We now turn to the empirical variation that underlies our estimate of this elasticity.

3 Empirical Analysis

3.1 Shifts in Reference Demand and Link with the Model

Our empirical strategy uses shifts in the *passive*, benchmark-tracking component of reference demand to estimate reduced-form price responses, which we later map into the model's demand slope. When passive demand shifts—for instance, when an index rebalancing leads benchmark-tracking investors to hold more of a bond—residual supply changes even though the government's debt remains unchanged. The price then moves to clear the market, with larger responses when demand is less elastic.

To map this variation into the model, we decompose reference demand into two components:

$$\mathcal{B}(\tau, B') = \overline{\mathcal{A}} + \mathcal{T}(\tau, B') = \overline{\mathcal{A}} + \tau B', \tag{16}$$

where $\mathcal{T}(\tau, B')$ is passive benchmark-tracking demand and $\overline{\mathcal{A}}$ collects other price-insensitive holdings, such as buy-and-hold or preferred-habitat positions. We write $\mathcal{T}(\tau, B') = \tau B'$, so it scales with bonds outstanding. The share τ is determined by passive assets under

Figure 2: Shift in Reference Demand

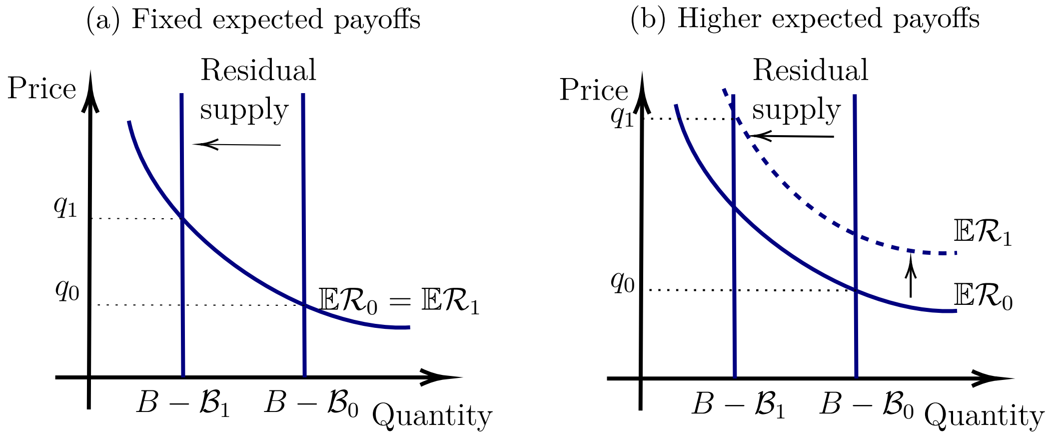

Note: The figure illustrates a decrease in residual supply induced by an increase in reference demand. In Panel (a), expected repayment is held fixed, so the price response reflects only movement along the demand schedule. In Panel (b), expected repayment also rises, so the price response reflects both movement along the demand schedule and an increase in expected repayment.

management (AUM) and index composition.[15] Index rebalancing events provide mechanical variation in τ because they change index composition and weights, while the debt stock B' is predetermined around the event.

With this notation in place, consider a within-period increase in passive demand, captured by a change from τ to τ' at fixed fundamentals $\{y, B'\}$. Here τ' denotes the next-period passive share, whose realization we assume becomes known within the current period. We formalize this timing in Section 4, where we extend the model with secondary markets. Reference demand rises by $\Delta\mathcal{B}' \equiv (\tau' - \tau)B'$ and residual supply falls by the same amount. Holding expected repayment fixed, the price response is governed by the demand slope $\mathcal{S}(y, \tau, B')$. Panel (a) of Figure 2 illustrates this case.

For risky long-term debt, however, the reduced-form price response need not identify $\mathcal{S}(y, \tau, B')$. A persistent shift in reference demand can alter the government's future incentives to borrow and default, changing dilution and expected repayment, which feed back into current prices. Panel (b) of Figure 2 illustrates this additional effect. Treating the full response as movement along the demand schedule would therefore confound the demand slope with changes in expected repayment. Our indirect-inference strategy addresses this by reproducing the same demand shift inside the model, where expected repayment is endogenous. We choose κ so that the model-implied total price response

[15] In bond units, passive demand equals AUM times the bond's index weight, divided by the price. Since the weight is value-weighted (proportional to qB'), the price cancels and passive demand becomes proportional to bonds outstanding, $\mathcal{T} = \tau B'$, with τ a quantity share rather than the index weight itself. Appendix A.1 provides the full derivation.

matches the reduced-form estimate.

3.2 Flows Induced by Index Rebalancing

We measure shifts in benchmark-tracking demand using end-of-month rebalancing in the J.P. Morgan Emerging Markets Bond Index Global Diversified (EMBIGD). The index tracks the performance of emerging market sovereign and quasi-sovereign U.S. dollar-denominated bonds issued in international markets.[16] It is the most widely used index for emerging-market hard-currency debt and served as a benchmark for funds with a combined AUM of around US$300 billion in 2018 (Appendix Figure B1).

A central feature of the EMBIGD, which we exploit in our identification strategy, is that, unlike traditional market-capitalization schemes, it applies a *cap rule* that limits the weights of countries with large eligible debt outstanding. Country c's index weight is

$$w_{c,t} = \frac{q_{c,t} B_{c,t} f_{c,t}}{MV_t}, \quad \text{with} \quad MV_t = \sum_{n \in \mathcal{I}} q_{n,t} B_{n,t} f_{n,t}, \tag{17}$$

where $q_{c,t}$ is the bond price, $B_{c,t}$ is the face amount included in the index, and MV_t is the aggregate market value of the index.[17] The term $f_{c,t} \in (0, 1]$ is a diversification coefficient implied by the EMBIGD cap rule. Each month, J.P. Morgan computes the cross-country average of the face amounts $B_{c,t}$ and labels it the *Index Country Average* (ICA). Countries with face amounts above this average receive $f_{c,t} < 1$, which scales down their weights, while countries below the ICA receive $f_{c,t} = 1$ (see Appendix B.1 for details). We refer to $f_{c,t} B_{c,t}$ as the *diversified face amount* (DFA).

Rebalancing occurs on the last business day of each month and is triggered by bond inclusions and exclusions.[18] We measure the resulting change in passive holdings using the *Flows Induced by Index Rebalancing* (FIR):

$$\text{FIR}_{c,t} \equiv \frac{(w_{c,t} - w_{c,t}^{BH}) \mathcal{W}_t}{q_{c,t-1} B_{c,t-1} - w_{c,t-1} \mathcal{W}_{t-1}}, \quad \text{with} \quad w_{c,t}^{BH} \equiv w_{c,t-1} \frac{q_{c,t}/q_{c,t-1}}{q_t^{\mathcal{I}}/q_{t-1}^{\mathcal{I}}}, \tag{18}$$

where $q_t^{\mathcal{I}}$ denotes the index price. The numerator captures the rebalancing-induced change in passive holdings, where $\mathcal{W}_t$ denotes the AUM of investors that track the EMBIGD, $w_{c,t}$ is the post-rebalancing index weight, and $w_{c,t}^{BH}$ is the *buy-and-hold* weight that would prevail absent index-composition changes, adjusting only for relative price movements. If

[16] To be considered for inclusion, a bond must have a maturity of at least 2.5 years and a minimum face amount outstanding of US$500 million.

[17] To simplify notation, we assume each country has a single bond that is fully included in the index. In practice, weights are constructed by aggregating all eligible bonds for a given country.

[18] J.P. Morgan announces these updates through a report detailing the updated index composition.

there are no inclusions or exclusions from the index, weights evolve only through relative price changes, so $w_{c,t}^{BH} = w_{c,t}$ and $\text{FIR}_{c,t} = 0$. The denominator normalizes flows by the market value of index-included bonds net of passive benchmark-tracking holdings. As defined, a 1 p.p. FIR corresponds to a passive flow equal to 1% of these non-passive holdings.

Because the index weight $w_{c,t}$ depends on current prices and a country's own issuance, the FIR is endogenous. We account for this in two ways. First, we restrict attention to countries whose outstanding eligible amounts in the index remain unchanged between two consecutive rebalancing events—that is, countries that neither issue new eligible bonds, repurchase existing ones, nor have bonds exit the index due to maturity or loss of eligibility. Second, we exploit the EMBIGD cap rule by instrumenting FIR with the relative change in a *synthetic index weight* based only on diversified face amounts, rather than actual bond prices:

$$\tilde{w}_{c,t} \equiv \frac{B_{c,t} f_{c,t}}{\sum_{n \in \mathcal{I}} B_{n,t} f_{n,t}} \quad \text{and} \quad Z_{c,t} \equiv \frac{\Delta \tilde{w}_{c,t}}{\tilde{w}_{c,t-1}}. \tag{19}$$

Variation in the instrument $Z_{c,t}$ comes from issuances, retirements, and eligibility changes of *other* countries in the index. When another country's bond enters or exits at a rebalancing event, all incumbent synthetic weights adjust mechanically. In a market-capitalization index, this adjustment would be proportional across countries. Under the cap rule, instead, such entry or exit can change the ICA, so the diversification coefficients $f_{c,t}$ are recomputed and synthetic weights adjust differently across countries, depending on how tightly each country's cap binds. We next present a simple example to illustrate these mechanisms.

Illustration of the cap rule. Consider five countries, $c = \{A, B, C, D, E\}$, each with one qualifying bond and no issuance or redemption during month t. Another country, F, issues an eligible bond that enters the index at the end-of-month rebalancing. Table 1 reports face amounts $FA_{c,t}$, diversified face amounts $DFA_{c,t}$—computed using Appendix Equation (B2)—and synthetic weights $\tilde{w}_{c,t}$ before and after rebalancing. In this example, countries D and E are capped in both $t-1$ and t, while country F becomes capped in t.

Table 1 illustrates the variation we exploit. Country F's inclusion reduces incumbent weights. The EMBIGD cap rule creates within-event variation across countries because F alters the distribution of eligible face amounts used to compute diversification coefficients.

Table 1: Illustration of the cap rule

Country	Before Rebalancing			After Rebalancing			$Z_{c,t}$
	$FA_{c,t-1}$	$DFA_{c,t-1}$	$\tilde{w}_{c,t-1}$	$FA_{c,t}$	$DFA_{c,t}$	$\tilde{w}_{c,t}$	
A	1,000	1,000	4.46%	1,000	1,000	3.19%	-28.54%
B	2,000	2,000	8.92%	2,000	2,000	6.37%	-28.54%
C	3,000	3,000	13.38%	3,000	3,000	9.56%	-28.54%
D	7,000	6,429	28.66%	7,000	6,769	21.57%	-24.75%
E	12,000	10,000	44.59%	12,000	11,000	35.05%	-21.39%
F	-	-	-	8,000	7,615	24.26%	-
ICA	5,000			5,500			

Note: The table illustrates the cap rule. Columns $FA_{c,t-1}$ and $FA_{c,t}$ list face amounts before and after rebalancing; $DFA_{c,t-1}$ and $DFA_{c,t}$ report diversified amounts. $\tilde{w}_{c,t-1}$ and $\tilde{w}_{c,t}$ are synthetic index weights, and $Z_{c,t} \equiv \Delta\tilde{w}_{c,t}/\tilde{w}_{c,t-1}$ is our instrument. ICA_t denotes the average face amount across index-eligible countries; when $FA_{c,t} > ICA_t$, the cap binds and $DFA_{c,t} < FA_{c,t}$.

In the example, the higher ICA relaxes the cap for previously constrained countries such as D and E, partly offsetting the reduction in their weights. Our instrument $Z_{c,t}$ captures this variation.

3.3 Data and Summary Statistics

We combine bond-level data from Datastream and J.P. Morgan with fund-level holdings from Morningstar. Our primary dataset covers the 2016–2018 period and includes daily observations for about 750 bonds from 68 countries. From Datastream, we obtain daily prices for all securities in the EMBIGD. From J.P. Morgan, we collect index weights, face amounts, and bond characteristics, including spreads, duration, average life, maturity, credit rating, and sovereign versus quasi-sovereign category, as well as the total AUM benchmarked to the EMBIGD. Appendix B.1 provides additional details.

Table 2 reports summary statistics for our main variables. The bonds in our sample are primarily medium- to long-term securities, with an average life of approximately 10 years and an average face amount of US$1.3 billion. The average stripped spread is 278 basis points.

We estimate $\mathcal{W}_t$, the AUM that mechanically tracks the EMBIGD, by combining Morningstar fund-level holdings with the total AUM benchmarked to the EMBIGD, as reported by J.P. Morgan. The AUM reported by J.P. Morgan does not distinguish between funds that passively track the benchmark and those that actively deviate from it.[19]

[19] Even if such data were available, many non-index funds may manage a large share of their portfolios in a semi-passive way, as emphasized by Pavlova and Sikorskaya (2022).

Table 2: Summary statistics

Variable	Mean	Std. Dev.	25th Pctl	75th Pctl	Min	Max
log(Price)	4.64	0.13	4.59	4.68	3.07	5.19
Stripped spread (bps)	278	288	128	356	0	4904
EIR duration (%)	6.36	3.92	3.48	7.71	-0.03	19.08
Average life (years)	9.6	8.9	4.0	9.9	1.0	99.8
Face amount (billion U.S. dollars)	1.3	0.8	0.7	1.6	0.5	7.0

Note: This table displays summary statistics for the main variables in the analysis. *Stripped spread* is the difference between a bond's yield-to-maturity and the corresponding point on the U.S. Treasury spot curve, with collateralized flows "stripped" from the bond. *Effective interest rate (EIR) duration* measures the percentage change in a bond's dirty price in response to a 100 basis point parallel shift in the U.S. Treasury yield curve. *Average life* is the weighted average time to principal repayment.

To address this, we proceed in two steps. First, using fund-level holdings, we compute each fund's *Passive Share* as $1 -$ *Active Share*, where *Active Share* is measured as in Cremers and Petajisto (2009) and captures the fraction of a fund's portfolio that deviates from benchmark weights. Aggregating across funds yields an AUM-weighted average passive share of 50%.[20,21] Second, we construct $\mathcal{W}_t$ by multiplying the total benchmarked AUM by this AUM-weighted passive share.

3.4 Estimation Strategy and Results

We exploit variation in FIR to estimate reduced-form price responses to shifts in passive demand. Specifically, we estimate daily bond-level price responses within symmetric five-day windows centered on each rebalancing date. FIR varies at the country-event level, while prices are observed at the bond-event-day level. Our baseline specification is

$$\log(q_{i,t,h}) = \theta_{c(i),t} + \theta_{b(i),t} + \gamma \mathbf{1}_{h \in Post} + \delta(\text{FIR}_{c(i),t} \times \mathbf{1}_{h \in Post}) + \mathbf{X}'_{i,t}\alpha + \varepsilon_{i,t,h}, \quad (20)$$

where $\text{FIR}_{c(i),t}$ denotes the flows induced by rebalancing into country c at event t. The variable $q_{i,t,h}$ is the price of bond i at rebalancing event t, observed h trading days before or after the event. For instance, $h = 1$ corresponds to the first trading day after J.P. Morgan releases the EMBIGD's new composition, typically during trading hours on the

[20]We compute *Active Share* at the country level, which is the level at which the FIR is constructed. Using Morningstar holdings, we measure each fund's country weights among EMBIGD bonds and compare them with the corresponding benchmark weights. Funds whose country allocations more closely track the benchmark have lower *Active Share*. For this calculation, we include funds benchmarked to the EMBIGD and the EMBI Global, which have the same underlying bond constituents. Restricting the calculation to funds explicitly benchmarked to the EMBIGD yields an *Active Share* of 48%.

[21]At the bond level, we obtain an *Active Share* of 71%, in line with Cremers and Petajisto (2009), who report an *Active Share* between 55% and 80%. The difference between the country- and bond-level passive shares might reflect funds that closely track benchmark country weights while deviating from the benchmark at the bond level. The country-level passive share can thus be interpreted as an upper bound on the AUM that mechanically replicates the benchmark at the security level.

last business day of each month. The indicator $\mathbf{1}_{h\in Post}$ equals 1 in the days following the rebalancing and 0 in the days preceding it. We estimate Equation (20) by 2SLS, instrumenting $\text{FIR}_{c(i),t} \times \mathbf{1}_{h\in Post}$ with $Z_{c(i),t} \times \mathbf{1}_{h\in Post}$.[22] We include country-month fixed effects $\theta_{c(i),t}$ and bond-characteristic-month fixed effects $\theta_{b(i),t}$, which account for maturity, credit rating, and whether the bond is sovereign or quasi-sovereign. The vector $\mathbf{X}_{i,t}$ includes bond-level controls such as face value and beginning-of-month spread. We also estimate a specification that replaces $\theta_{c(i),t}$, $\theta_{b(i),t}$, and $\mathbf{X}_{i,t}$ with bond-month fixed effects $\theta_{i,t}$. Our coefficient of interest is δ, the price response to FIR.

The previous specification exploits variation both within and across rebalancing events. To better isolate within-event variation, we leverage the heterogeneity induced by the cap rule and estimate an alternative specification that includes month-by-$\mathbf{1}_{h\in Post}$ fixed effects. These absorb average differences between the pre- and post-windows for each rebalancing event—for instance, those driven by global shocks around the rebalancing date, such as changes in risk-free rates or risk premia. The coefficient δ is then identified from cross-country differences in cap-rule-induced FIRs within the same event.

Table 3 reports the results of our baseline estimation using a five-day window around each rebalancing event (i.e., $h \in [-5, 5]$). Across all specifications, our coefficient of interest, δ, is positive and statistically significant. The estimate in column (4), which includes bond-month and month-by-$\mathbf{1}_{h\in Post}$ fixed effects, implies that a 1 p.p. increase in the FIR raises bond prices by approximately 0.30%.[23] In yield terms, this corresponds to a decline in yield to maturity of about 5.5 basis points—see Appendix Table B3.

The documented effects are heterogeneous across bonds with varying default risk levels. In Table 4, we divide the sample into terciles based on bond spreads and estimate Equation (20) for each subsample. We find that prices of high-spread bonds respond more strongly to FIR, with a 1 p.p. increase in FIR associated with a nearly 0.40% increase in bond prices. In contrast, the effect for low-spread bonds is smaller and not statistically significant. In yield terms, the response is roughly three times larger for high- than for low-spread bonds (Appendix Table B4).

[22] In specifications where the level of FIR is not absorbed by fixed effects, we also instrument $\text{FIR}_{c(i),t}$ with $Z_{c(i),t}$.

[23] Appendix Table B1 presents OLS estimates that do not instrument the FIR. The coefficients are slightly smaller than in the 2SLS specification, consistent with a potential downward bias from the endogeneity of FIR to price movements around the rebalancing. Our findings are also robust to alternative event windows (Appendix Table B2).

Table 3: Effects of FIR on bond prices

Dependent Variable: Log Price	Symmetric window: $[-5, 5]$				Excl. $h = -1$	
FIR × Post	0.232	0.233	0.232	0.301	0.264	0.320
	(0.099)	(0.101)	(0.099)	(0.135)	(0.098)	(0.135)
Bond FE	Yes	Yes	No	No	No	No
Month FE	Yes	No	No	No	No	No
Bond Characteristics × Month FE	No	Yes	No	No	No	No
Country × Month FE	No	Yes	No	No	No	No
Bond × Month FE	No	No	Yes	Yes	Yes	Yes
Month × Post FE	No	No	No	Yes	No	Yes
Bond Controls	No	Yes	No	No	No	No
Observations	105,548	105,508	105,548	105,548	84,433	84,433
N. of Bonds	738	738	738	738	738	738
N. of Countries	68	68	68	68	68	68
N. of Clusters	1,576	1,575	1,576	1,576	1,576	1,576
SW F (FIR × Post)	1,863	1,618	1,667	477	1,672	477

Note: The table presents 2SLS estimates of log bond prices on FIR interacted with a post-rebalancing indicator around rebalancing events. Post is an indicator equal to 1 for the five trading days following each rebalancing, and 0 otherwise. Bond-characteristic fixed effects refer to interactions of maturity bins, credit rating grades, and sovereign or quasi-sovereign category indicators. Maturity bins classify bonds into four groups: short (less than 5 years), medium (5–10 years), long (10–20 years), and very long (over 20 years). Rating bins are based on Moody's categorical grades. Bond controls indicate whether the regression includes the log face amount and the beginning-of-month log stripped spread. The baseline specification uses a symmetric five-day window around each rebalancing; the last two columns exclude the trading day before each event. The FIR and Post regressors are included when not absorbed by the fixed effects, but their estimated coefficients are omitted from the table. Standard errors are clustered at the country-month level. The reported first-stage statistic is the Sanderson-Windmeijer (SW) F-statistic for FIR × Post. The sample period is 2016–2018.

Overall, the estimates point to larger reduced-form price responses for riskier bonds. This pattern motivates the functional form we adopt in the quantitative model, where the demand slope scales with the conditional variance of repayment (Equation (4)).

One potential concern with the previous analysis is that countries with larger FIRs may have been on distinct price trends before the rebalancing. To assess this possibility, we estimate a leads-and-lags specification in which the instrumented FIR is interacted with trading-day dummies around each rebalancing event. This allows us to trace the dynamic response of bond prices. Specifically, we estimate the following regression by 2SLS:

$$\log(q_{i,t,h}) = \theta_{i,t} + \sum_{\ell \neq -2} \gamma_\ell \mathbf{1}_{h=\ell} + \sum_{\ell \neq -2} \delta_\ell \left(\text{FIR}_{c(i),t} \times \mathbf{1}_{h=\ell}\right) + \varepsilon_{i,t,h}, \tag{21}$$

where $\theta_{i,t}$ are bond-month fixed effects and $\mathbf{1}_{h=\ell}$ are indicators for each trading day in the window. The estimated δ_ℓ coefficients are plotted in Figure 3. The day $\ell = -2$ is the

Table 4: Effects of FIR on bond prices: The role of default risk

Dependent Variable: Log Price						
	High Spread		Medium Spread		Low Spread	
FIR × Post	0.381	0.382	0.326	0.323	0.087	0.087
	(0.167)	(0.165)	(0.152)	(0.151)	(0.099)	(0.098)
Bond FE	Yes	No	Yes	No	Yes	No
Month FE	Yes	No	Yes	No	Yes	No
Bond × Month FE	No	Yes	No	Yes	No	Yes
Observations	28,105	28,104	28,055	28,053	28,276	28,276
N. of Bonds	381	381	453	453	375	375
N. of Countries	58	58	51	51	43	43
N. of Clusters	975	975	837	837	634	634
SW F (FIR × Post)	2,309	2,344	813	720	870	883

Note: The table presents 2SLS estimates of log bond prices on FIR interacted with a post-rebalancing indicator around rebalancing events. The sample is split into high-spread bonds (above the 66th percentile of stripped spreads), medium-spread bonds, and low-spread bonds (below the 33rd percentile). The sample period and estimation procedure follow those in Table 3, excluding the trading day before each rebalancing. Standard errors are clustered at the country-month level.

omitted reference period, so each δ_ℓ measures the price response relative to two trading days before the rebalancing.

In the days leading up to the rebalancing, FIR is not systematically related to bond price movements. After the rebalancing, the δ_ℓ coefficients turn positive and statistically significant, stabilizing around 0.30 by the end of the window. We observe a modest anticipation effect on the day before the rebalancing, which is not uncommon in this type of setting.[24] In the last two columns of Table 3, we report estimates from our preferred specification of Equation (20), excluding the trading day before the rebalancing. These range from 0.26 to 0.32 and serve as our baseline, as they exclude potential anticipation effects in the pre-period.

A related concern is the potential for anticipation earlier in the month. Between mid- and end-month, J.P. Morgan releases preliminary bond weight estimates, which could allow investors to adjust their portfolios ahead of the rebalancing. However, our data show no evidence of such behavior. We find no correlation between FIR and bond returns in the week leading up to the event, except on the day before. We cannot examine earlier periods, as they overlap with the prior rebalancing cycle. If some trades occur in advance, our FIR measure would overstate rebalancing-day inflows, implying that the estimates in Table 3 should be interpreted as a lower bound.

[24]This is consistent with the portfolio rebalancing patterns documented by Escobar et al. (2026), who show that institutional investors often begin trading one day prior to the official index rebalancing.

Figure 3: Effects of FIR on bond prices: Leads and lags coefficients

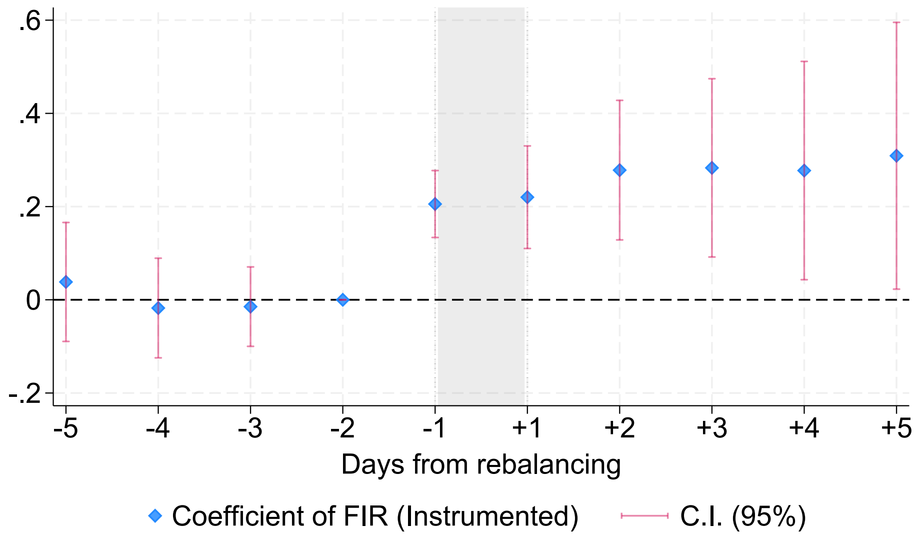

Note: The figure presents lead and lag coefficients from a 2SLS estimation of log bond prices on FIR interacted with trading-day dummies around each rebalancing event, following Equation (21). The specification includes bond-month fixed effects and uses the same procedure as in Table 3. Coefficients are measured relative to $h = -2$, the omitted trading day. The shaded area indicates the rebalancing day, and $h = +1$ is the first post-rebalancing trading day. Vertical lines represent 95% confidence intervals. Standard errors are clustered at the country-month level.

Interpreting the estimates. The estimates imply an inverse demand elasticity of about -0.30, corresponding to a yield semi-elasticity of 5.5 basis points. These values are larger in magnitude than estimates for advanced-economy sovereign bonds (Krishnamurthy and Vissing-Jorgensen, 2012; Greenwood et al., 2015; Mian et al., 2025; Jiang et al., 2026), underscoring the relative inelasticity of emerging-market bond markets. They are, however, comparable to estimates for advanced-economy corporate bonds (Bretscher et al., 2026). Within emerging-market government debt, our estimates imply more elastic demand than Fang et al. (2025), consistent with our focus on large and liquid EMBIGD bonds.[25]

These reduced-form price responses do not necessarily map one-for-one into the demand slope $\mathcal{S}$ in the model, because expected payoffs may also adjust. Consistent with this mechanism, we find stronger price responses for longer-maturity bonds, whose valuations are more sensitive to future payoffs (Appendix Table B5). We also find that CDS spreads decline with FIR (Appendix Table B6), suggesting that perceived credit risk adjusts. These patterns motivate our structural model, which disentangles pure demand

[25]The bond sample in Fang et al. (2025) includes a broader set of countries, loans, and local-currency bonds, assets typically associated with less elastic demand because they are concentrated among domestic investors (Pandolfi and Williams, 2019).

pressure from endogenous payoff responses.

A further challenge is the frequency mismatch. The empirical strategy isolates price responses in narrow daily windows around rebalancing events, effectively holding fundamentals fixed. The model operates at a quarterly frequency, with bond prices and fundamentals adjusting concurrently. To bridge this gap, the next section adds secondary markets, allowing us to capture intra-period trading and replicate the high-frequency empirical exercise.

4 Quantitative Analysis

4.1 Calibration

We calibrate the model at a quarterly frequency using data on Argentina. The calibration follows a two-step procedure. We first fix a subset of parameters to standard values in the literature or based on historical Argentine data. We then internally calibrate the remaining parameters to match moments of external debt and bond spreads, and our reduced-form estimates from Section 3.

In terms of functional forms, preferences are CRRA: $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, where γ denotes the coefficient of risk aversion. Log output follows an AR(1) process: $\log(y') = \rho_y \log(y) + \epsilon'_y$, with $\epsilon'_y \sim N(0, \sigma_y^2)$. In the event of default, output costs are governed by a quadratic loss function: $\Phi(y) = \max\{d_0 y + d_1 y^2, 0\}$, where $d_0 < 0$ and $d_1 > 0$. This implies that output costs are zero whenever $0 \leq y \leq -\frac{d_0}{d_1}$, and rise more than proportionally with y beyond that threshold. This specification, drawn from Chatterjee and Eyigungor (2012), allows us to match the observed mean and volatility of sovereign spreads.

Table 5 lists the parameters. For the subset of fixed parameters (Panel a), we set $\gamma = 2$, a standard value for risk aversion in the literature. The quarterly gross risk-free rate is set to $r^f = 1.01$, consistent with the average real rate observed in the United States. The probability of re-entry into international markets is set to $\chi = 0.0385$, implying an average exclusion duration of 6.5 years. We choose $\lambda = 0.05$ to target a debt maturity of five years, and $\nu = 0.03$ to match Argentina's average debt service. The parameters governing the endowment process, ρ_y and σ_y, are based on log-linearly detrended quarterly real GDP data for Argentina. All these parameters are taken from Morelli and Moretti (2023).

Following Equation (16), we let the passive component of reference demand be pro-

Table 5: Model calibration

Panel a: Fixed Parameters			Panel b: Calibrated Parameters		
Param.	Description	Value	Param.	Description	Value
γ	Risk aversion	2.00	β	Discount factor	0.948
r^f	Gross risk-free rate	1.01	d_0	Default cost—level	−0.224
λ	Debt maturity	0.05	d_1	Default cost—curvature	0.27
ν	Debt services	0.03	κ	Slope parameter	70.0
χ	Reentry probability	0.0385	$\bar{\mathcal{A}}$	Reference demand term	0.4493
ρ_y	Output, autocorrelation	0.93			
σ_y	Output, shock volatility	0.02			
$\tau^\star$	Share of passive demand	0.123			
ρ_τ	Passive share, autocor.	0.66			
σ_τ	Passive share, volatility	0.02			

Note: The table reports the set of fixed parameters (Panel a) and calibrated parameters (Panel b).

portional to end-of-period bonds outstanding: $\mathcal{T}(\tau, B') = \tau B'$. The passive share τ is stochastic and follows $\tau' = (1-\rho_\tau)\tau^\star + \rho_\tau \tau + \epsilon'_\tau$, where $\epsilon'_\tau \sim N(0, \sigma_\tau^2)$. We parameterize this process using the passive share implied by Argentina's EMBIGD weights and passive benchmark-tracking AUM. Estimating the process at a quarterly frequency yields $\rho_\tau = 0.66$ and $\sigma_\tau = 0.02$. Appendix C.3 details the construction, its mapping to the model, and alternative estimates, which deliver similar persistence values.

We internally calibrate the remaining parameters (Table 5, Panel b). We jointly set the default-cost parameters $\{d_0, d_1\}$ and the government's discount factor β to match Argentina's average external debt-to-GDP ratio, mean sovereign spread, and spread volatility.[26] We calibrate κ, the parameter governing the demand slope, to match the reduced-form inverse demand elasticity implied by price responses around rebalancing events. To this end, we first extend the baseline framework to include secondary-market trading within the period. The timing is as follows:

1. *The endowment y is realized. The initial states are $\{y, \tau, B\}$.*
2. *The government chooses $d(y, \tau, B)$ and $B'(y, \tau, B)$.*
3. *The primary and secondary markets open; $q^{SM,0}(y, \tau, B')$ denotes the opening price.*
4. *Next-period passive share τ' is realized. Bond prices are updated.*

[26]The debt-to-output target is based on Argentina's "external debt stocks, public and publicly guaranteed," as reported by the World Bank. Since the model abstracts from recovery in default, we target only the unsecured portion of debt, approximately 70%, following Chatterjee and Eyigungor (2012). For spreads, we focus on the 2016–2019 period, when a pro-market administration was in power. This timeframe includes the 2016–2018 window used in our empirical analysis. Spreads are higher outside this period, under governments led by a different political party with a history of favoring default. Since our model does not capture these reputational factors, which account for a substantial share of sovereign spreads (Morelli and Moretti, 2023), we exclude these periods from the calibration.

5. *The secondary market closes;* $q^{SM,1}(y, \tau', B')$ *denotes the closing price.*

This extension allows us to reproduce the empirical experiment by computing price responses $q^{\mathrm{SM},1} - q^{\mathrm{SM},0}$ to changes in the passive share, holding the government's debt and endowment fixed. Appendix C.1 provides the pricing functions for this extension. In the absence of secondary markets, the timing remains identical to the baseline framework, so the extension nests our baseline model.

The resulting κ governs prices at issuance, but our mapping does not impose the same overall price response to rebalancing and issuance. Rebalancing changes residual supply through shifts in passive demand while leaving the government's debt position unchanged. Issuance, by contrast, alters the debt stock and therefore also affects default risk and expected repayment. The model accounts for these channels endogenously through $\mathbb{E}\mathcal{R}(y', \tau', B')$.[27]

For each state $\{y, \tau, B'\}$ and realized τ', the reduced-form inverse elasticity is

$$\eta^{\mathrm{RF}}(y, \tau, B', \tau') = -\frac{\Delta \log q}{\Delta\mathcal{T}' / (B' - \mathcal{T}(\tau, B'))}, \tag{22}$$

where $\Delta\mathcal{T}' \equiv \mathcal{T}(\tau', B') - \mathcal{T}(\tau, B')$ is the mechanical change in passive demand and $\Delta \log q \equiv \log q^{\mathrm{SM},1}(y, \tau', B') - \log q^{\mathrm{SM},0}(y, \tau, B')$ is the corresponding log price change. As constructed, the denominator $\Delta\mathcal{T}' / (B' - \mathcal{T}(\tau, B'))$ mirrors the empirical FIR in Equation (18). We simulate the model and average η^{RF} over the simulated joint distribution of (y, τ, B', τ') conditional on no default and on a nonzero change in the passive share.

In our baseline calibration, we normalize the level of reference demand by setting $\overline{\mathcal{A}}$ such that the average *inelasticity premium*, $\vartheta(\cdot)$, is zero. This premium is the difference between the spread under imperfectly elastic demand and the counterfactual spread under perfectly elastic demand, holding the baseline model's borrowing and default policies fixed.[28] Imperfectly elastic demand therefore operates primarily through a steeper demand slope rather than through a level effect. We relax this normalization in the

[27] The mapping relies on two assumptions. First, newly issued and outstanding bonds are close substitutes, so that secondary-market price responses are informative about primary-market prices. This is a reasonable approximation for the large, liquid hard-currency bonds in the EMBIGD segment (Appendix C.1). Second, the demand elasticity estimated on the EMBIGD segment applies to the government's entire external debt stock (Appendix C.3). To the extent that demand for other debt is less elastic (Fang et al., 2025), price impact would be even stronger for that portion of the market. Figure 6 reports model outcomes across a wide range of demand elasticities.

[28] Formally, $\vartheta(y, \tau, B') \equiv SP(y, \tau, B') - SP_{\kappa=0}(y, \tau, B')$, where $SP_{\kappa=0}(\cdot)$ is computed from bond prices under $\kappa = 0$ and the policies implied by our baseline model. Annualized spreads are $SP(y, \tau, B') \equiv \left(\frac{1+i(y,\tau,B')}{r^f}\right)^4 - 1$, where the internal quarterly return $i(y, \tau, B')$ solves $q(y, \tau, B') = \frac{\lambda + (1-\lambda)(\nu + q(y,\tau,B'))}{1+i(y,\tau,B')}$.

Table 6: Targeted moments

Target	Description	Data	Model	Alternative slope parameters $\kappa = 0$	$\kappa = 30$	$\kappa = 50$	$\kappa = 90$
$\bar{\mathbb{E}}(B/y)$	Debt to output (%)	53.0	53.1	51.0	52.4	52.8	53.2
$\bar{\mathbb{E}}(SP)$	Bond spreads (%)	5.08	5.06	8.59	6.46	5.61	4.68
$\bar{\sigma}(SP)$	Volatility of spreads (%)	1.43	1.56	4.49	2.62	1.95	1.34
$\bar{\mathbb{E}}(\eta^{\mathrm{RF}})$	Reduced-form elasticity	-0.29	-0.28	-	-0.17	-0.23	-0.31

Note: The table reports the moments targeted in the calibration and their model counterparts. The targeted reduced-form elasticity corresponds to the midpoint of the last two columns of Table 3. The last four columns report moments for alternative values of κ; in each case, we adjust $\overline{\mathcal{A}}$ so that the inelasticity premium is zero on average, as in the baseline calibration. The operators $\bar{\mathbb{E}}(\cdot)$ and $\bar{\sigma}(\cdot)$ denote the sample mean and standard deviation, respectively.

Figure 4: Decomposing the inverse demand elasticity

(a) Reduced-form inverse elasticity (b) Reduced-form vs model-implied

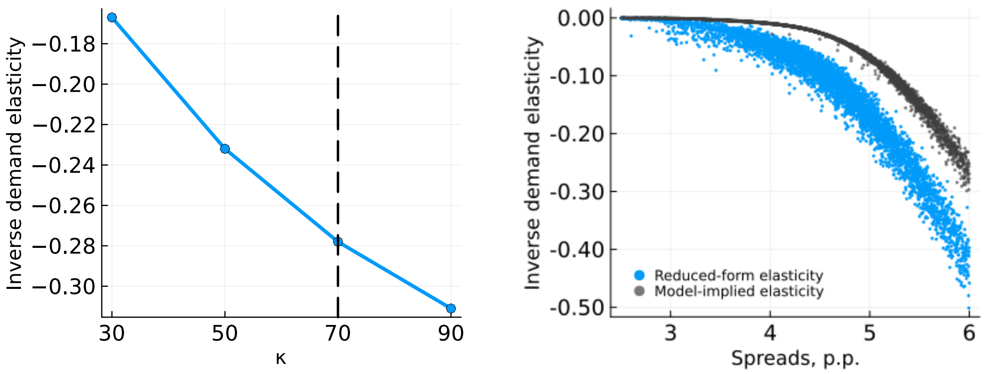

Note: Panel (a) displays the average reduced-form inverse demand elasticity, η^{RF}, as a function of κ. Panel (b) shows a binscatter of the simulated reduced-form inverse elasticity η^{RF} (blue dots) and the model-implied η (black dots) against annualized bond spreads.

sensitivity exercise of Table 8, where we consider positive premia.

Table 6 reports the set of targeted moments. The model closely matches all four targets. Figure 4, Panel (a) shows that the parameter κ is well identified, as it generates a monotonic relation with the reduced-form elasticity.

With the calibrated κ, we recover the demand slope, $\mathcal{S}(y,\tau,B') = \frac{\kappa}{r^f}\mathbb{V}_{s'|s}\mathcal{R}(y',\tau',B')$, and compute the associated *model-implied* elasticity:

$$\eta\left(y,\tau,B'\right) \equiv -\,\mathcal{S}\left(y,\tau,B'\right)\frac{B' - \mathcal{T}\left(\tau,B'\right)}{q\left(y,\tau,B'\right)}, \tag{23}$$

where we normalize by $B' - \mathcal{T}(\tau, B')$, as in the reduced-form elasticity. Figure 4, Panel (b), plots the simulated reduced-form inverse elasticity η^{RF} alongside the model-implied η as a binscatter against bond spreads. The magnitudes of both increase with spreads.

The difference between the two elasticities reflects the part of the price response driven by endogenous changes in repayment, through both its expected value and its

Table 7: Persistence of passive demand and estimated elasticities

Moment	Description	Baseline model	Alternative persistence values		
			Low	Medium	High
$\bar{\mathbb{E}}(\eta^{\mathrm{RF}})$	Reduced-form elasticity	-0.28	-0.23	-0.25	-0.37
$\bar{\mathbb{E}}(\eta)$	Model-implied elasticity	-0.18	-0.18	-0.18	-0.18
$1-\bar{\mathbb{E}}(\eta)/\bar{\mathbb{E}}(\eta^{\mathrm{RF}})$	Relative difference	35%	23%	29%	51%

Note: The table compares the reduced-form inverse demand elasticity η^{RF} with its model-implied counterpart η. The "Baseline" column uses the baseline calibration. The "Low," "Medium," and "High persistence" columns set ρ_τ to 0.30, 0.50, and 0.90, respectively, holding all other parameters fixed.

variance. On average, η is about two-thirds as large as η^{RF} (first column of Table 7). The magnitude of this difference depends on the persistence of the τ process. The last three columns of Table 7 compare these elasticities across different ρ_τ values. When the process is more (less) persistent, the model-implied η accounts for a smaller (larger) share of η^{RF} in absolute value. Greater persistence in the τ process makes expected repayment more responsive to a given change in the passive share, both because the shift in residual supply lasts longer and because it induces stronger adjustments in subsequent issuance decisions. These channels generate larger price responses, increasing the magnitude of the reduced-form inverse elasticity. Appendix C.7 provides a more detailed analysis.

4.2 Market Discipline in the Dynamic Model

The three-period analysis in Section 2.4 shows that less elastic demand has an ambiguous effect on long-term debt. Whether the dilution-curbing force outweighs the price-impact cost at the calibrated κ in the full dynamic model is, therefore, a quantitative question. We assess this by comparing the baseline calibration with the perfectly elastic case, $\kappa = 0$, holding the rest of the calibration fixed. We then consider counterfactual economies with alternative values of κ.

Panel (a) of Figure 5 plots the bond policy $B'(y,\tau,B)$ under $\kappa > 0$ and $\kappa = 0$. Under imperfectly elastic demand, an additional unit of B' lowers bond prices through both higher default risk and the price impact of the larger residual supply. This second channel intensifies in states with higher debt and larger repayment variance, where the demand schedule is steeper. The government internalizes this steeper schedule and, when inherited debt is large, limits borrowing relative to the perfectly elastic case.[29]

[29] Figure C2 in the appendix reports a phase diagram comparing the long-run target-debt locus across endowment levels.

Figure 5: Comparison with the perfectly elastic case, $\kappa = 0$

(a) Bond policy (b) Bond price (c) Default probability

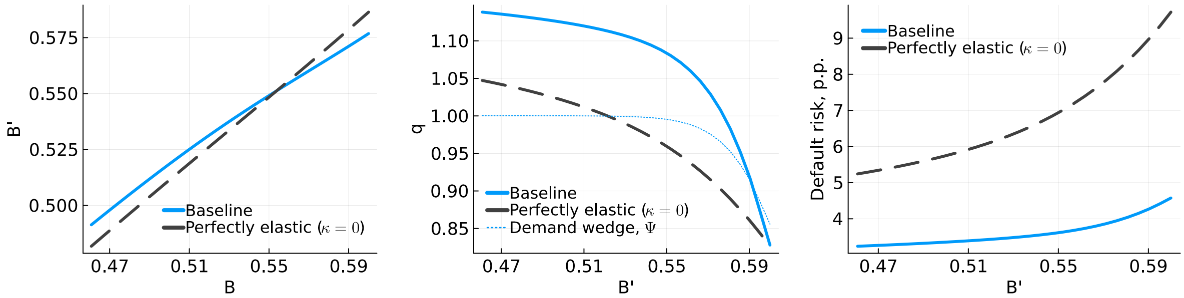

Note: The figure compares debt policies, bond prices, and default risk under the baseline model (blue) and a perfectly elastic demand case (black), in which we set $\kappa = 0$ while holding all other parameters fixed. Panel (a) plots next-period debt, B', as a function of inherited debt, B. Panel (b) shows bond prices and Ψ, defined as $\Psi \equiv q/(\frac{1}{r^f}\mathbb{E}\mathcal{R}')$. Panel (c) reports the annualized default probability over the bond's expected maturity. All functions are evaluated at the mean values of y and τ.

This restraint carries over to bond prices. Panel (b) shows that for low values of B', prices are *higher* when $\kappa > 0$. This is not mechanically driven by a contemporaneous quantity-based price adjustment that raises current prices above discounted expected repayment: the same panel shows that $\Psi \equiv q/(\frac{1}{r^f}\mathbb{E}\mathcal{R}')$ does not exceed one over the plotted range. The higher prices instead reflect larger expected repayment, due in part to lower default risk relative to the $\kappa = 0$ case, as shown in Panel (c).

Taken together, the panels in Figure 5 illustrate the dynamic counterpart of the dilution-curbing force in Theorem 1. By restraining issuance in high-debt states, imperfectly elastic demand lowers default risk and curbs expected dilution. Lenders price this in, raising the value of long-term bonds and supporting higher issuance in states with low debt, where price impact is weaker (Panel (a)). Section 4.4 quantifies the contribution of dilution to spreads.

Table 6 summarizes the main quantitative implications. Relative to the perfectly elastic case, the baseline economy features spreads roughly 40% lower (a reduction of about 350 basis points) and a debt-to-output ratio that is 2 p.p. higher. Thus, at the empirically calibrated value of κ, the dilution-curbing force is strong enough for less elastic demand to lower borrowing costs while supporting more debt. These results are not tied to a single calibrated κ, as shown in the last columns of Table 6.

Figure 6 further traces debt policies across a wider range of κ values, recalibrating $\overline{\mathcal{A}}$ in each case so that the average inelasticity premium remains zero. Panel (a) plots the average debt-to-output ratio against average $|\eta|$, so moving right corresponds to less

Figure 6: Debt and borrowing policies as demand becomes less elastic

(a) Average debt-to-output (b) Borrowing differential

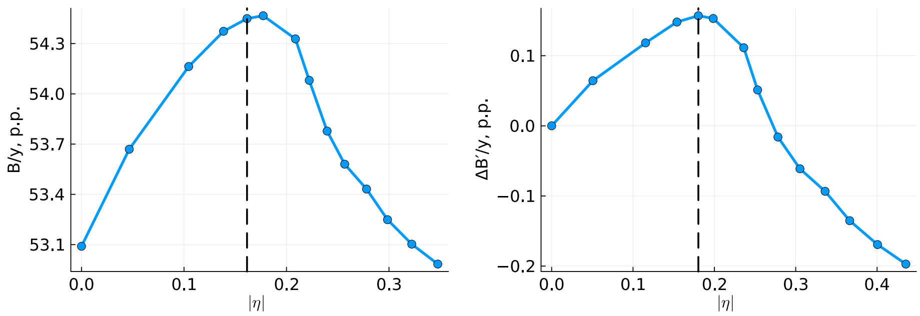

Note: Panel (a) reports average debt-to-output in the model with imperfectly elastic demand for different values of κ. Panel (b) reports $\Delta B'(y,\tau,B) \equiv B'(y,\tau,B) - B'_{\kappa=0}(y,\tau,B)$, scaled by output. That is, the average differential between the borrowing policy under imperfectly elastic and perfectly elastic demand, evaluated at the same simulated repayment states of the imperfectly elastic economy. The horizontal axis is the absolute value of the average model-implied inverse elasticity, $|\eta|$, obtained by varying κ while holding the remaining calibration fixed, except that $\overline{\mathcal{A}}$ is adjusted so that the average inelasticity premium is zero.

elastic demand. The debt ratio is hump-shaped: it initially rises above the perfectly elastic level and grows with $|\eta|$, but the price-impact cost eventually dominates, and debt falls below that level once $|\eta|$ exceeds about 0.30.[30] Panel (b) isolates differences in borrowing relative to the perfectly elastic case by evaluating the $\kappa > 0$ and $\kappa = 0$ debt policies at the *same* states (y,τ,B), drawn from simulations of the corresponding $\kappa > 0$ economy. It then plots the average borrowing differential, $\Delta B'(y,\tau,B) \equiv B'(y,\tau,B) - B'_{\kappa=0}(y,\tau,B)$, scaled by output, where $B'_{\kappa=0}(\cdot)$ denotes the debt policy in the $\kappa = 0$ case. This exercise removes the composition effect that would otherwise arise because the perfectly elastic economy spends more time near default and re-entry, when debt is low. Consistent with the three-period model, the average $\Delta B'(y,\tau,B)$ is positive for moderate $|\eta|$ but turns negative once the price-impact cost prevails.

We next allow the average inelasticity premium to be positive. Table 8 reports alternative calibrations of $\overline{\mathcal{A}}$ that target average premia of 0.50, 1.00, and 1.50 p.p., with κ adjusted slightly in each case to match the targeted reduced-form elasticity. As the premium rises, the government reduces debt, lowering default risk and offsetting the higher premium, so spreads barely change relative to the baseline. Even at 1.50 p.p., when the premium accounts for almost one-third of spreads, average debt remains above

[30]The figure panel excludes three quarters before each default and three after re-entry, so that differences across economies mainly reflect borrowing policy rather than debt swings around default. Absent this filter, average B/y is still hump-shaped in $|\eta|$ but stays above the perfectly elastic case, since the government defaults less often under imperfectly elastic demand.

Table 8: Alternative inelasticity premia

Moment	Description	Baseline model	Perfectly elastic ($\kappa = 0$)	Alternative inelasticity premia		
				High	Higher	Highest
$\bar{\mathbb{E}}(B/y)$	Debt to output (%)	53.1	51.0	52.7	52.1	51.5
$\bar{\mathbb{E}}(SP)$	Bond spreads (%)	5.06	8.59	4.96	4.98	4.91
$\bar{\sigma}(SP)$	Volatility of spreads (%)	1.56	4.49	1.48	1.47	1.41
$\bar{\mathbb{E}}(\vartheta)$	Inel. premium (%)	0.00	0.00	0.50	1.00	1.52
$\bar{\mathbb{E}}(d_\lambda)$	Default probability (%)	4.01	6.57	3.61	3.28	2.86
$\bar{\mathbb{E}}(\eta^{\mathrm{RF}})$	Reduced-form elasticity	-0.28	-	-0.28	-0.27	-0.27
$\bar{\mathbb{E}}(\eta)$	Model-implied elasticity	-0.18	-	-0.18	-0.17	-0.16
$\bar{\mathbb{E}}(\mathrm{CEC})$	Cons. equivalent (%)	0.32	-	0.31	0.29	0.27

Note: The table reports moments for our baseline calibration and alternative scenarios with higher average inelasticity premia, obtained by adjusting the $\overline{\mathcal{A}}$ parameter. In these scenarios, κ is also slightly adjusted to match the targeted reduced-form elasticity. The "High" column targets an inelasticity premium of 0.50 p.p.; "Higher" targets 1.00 p.p.; and "Highest" targets 1.50 p.p. The perfectly elastic column shows the results under $\kappa = 0$, holding all other parameters fixed at their baseline values. The last row reports the certainty-equivalent consumption (CEC) gain relative to the corresponding perfectly elastic counterfactual, in percent of permanent consumption.

the perfectly elastic case. Thus, the dilution-curbing force continues to outweigh the price-impact cost throughout the range considered.

The last row of Table 8 reports the welfare effect of imperfectly elastic demand. Viewed in isolation, the baseline 2 p.p. increase in debt-to-output may appear modest. The key quantitative result, however, is that reduced expected dilution leads the government to sustain *more* debt while facing substantially *lower* spreads and default risk. We summarize these combined effects in welfare units using a certainty-equivalent consumption (CEC) measure. It is the permanent proportional change in consumption that makes the government indifferent between the baseline economy and the perfectly elastic case.[31] The average implied gain is 0.32% in the baseline and 0.27% in the highest-premium case we consider.[32]

Additional results. Appendix C.5 reports several sensitivity tests. First, consistent with the three-period model, we show that the dilution-curbing force weakens as maturity shortens. Second, we consider environments with greater debt capacity by raising output costs of default, which lowers default risk. The government still sustains more debt

[31] The CEC is the value of $\tilde{x}$ such that $\sum_{t=0}^{\infty}\beta^t\mathbb{E}_0 u(c_t) = \sum_{t=0}^{\infty}\beta^t\mathbb{E}_0 u\left((1+\tilde{x})\tilde{c}_t\right)$, where c_t and $\tilde{c}_t$ denote consumption in the baseline and counterfactual scenarios, respectively. Under power utility, this simplifies to $\tilde{x} = \left[V(y,\tau,B)/\tilde{V}(y,\tau,B)\right]^{1/(1-\gamma)} - 1$, where $V(\cdot)$ and $\tilde{V}(\cdot)$ are the corresponding value functions. A positive $\tilde{x}$ indicates that the government prefers the baseline scenario.

[32] For comparison, eliminating the dilution remaining in our baseline model, through a covenant we introduce in Section 4.4, yields a CEC gain of 0.25% relative to that baseline (Table 10).

than under perfectly elastic demand, by a relative margin similar to that in the baseline. Third, we document suggestive comovement between the reduced-form elasticity and global financial conditions, consistent with the frictions captured by κ tightening in stress periods. Motivated by this pattern, we extend the model by allowing κ to rise in low-output states and find that the dilution-curbing force strengthens, with spreads declining further and the government sustaining higher debt.[33]

Appendices C.6 and C.7 report three additional exercises on reference demand and its passive component. First, we show that reference demand is what allows the model to generate a steep demand schedule, since it separates marginal price responses from the average inelasticity premium. When $\mathcal{B} = 0$, even values of κ large enough for the premium to account for most of spreads imply demand far more elastic than observed in the data; the resulting η^{RF} is only about one-sixth of its empirical counterpart. Second, we quantify the contribution of passive-demand shock volatility to the model's moments. Shutting down innovations to τ lowers spread variance by only about 4% and leaves the other targeted moments largely unchanged.[34] Third, our results are robust to the level of the mean passive share, $\tau^\star$. Solving the model for $\tau^\star$ between 3% and 18%, adjusting $\overline{\mathcal{A}}$ in each case so that the average inelasticity premium remains zero, leaves the main moments and elasticities essentially unchanged.

4.3 Role of Default Costs

The previous subsection compared the imperfectly elastic economy with a perfectly elastic (PE) case, holding all parameters fixed. Here, we ask what default costs are needed under perfectly elastic demand to account for Argentina's spreads. To answer this question, we adjust the default-cost parameters $\{d_0, d_1\}$ to match the average and standard deviation of bond spreads, leaving the remaining parameters unchanged. We refer to the resulting economy as the *recalibrated PE model* and compare it with our baseline.

Column $\{d_0, d_1\}$ of Table 9 reports the results. Matching the same spread moments requires larger output costs of default than in the baseline at low output levels, precisely

[33]Even with constant κ, the demand slope is state dependent because repayment variance changes with the state. This exercise adds state dependence in κ itself.

[34]This does not mean that τ shocks are irrelevant. In a counterfactual in which the government's borrowing does not respond to them, spread variance is about 17.5% higher than in the fixed-τ economy; the government's issuance response offsets most of this effect.

Table 9: Disciplining devices: Default costs vs. downward-sloping demand

Moment	Description	Baseline model	Recalibrated PE $\{d_0, d_1\}$	Recalibrated PE $\{d_0, d_1, \beta\}$
(a) Targeted				
$\bar{\mathbb{E}}(B/y)$	Debt to output (%)	53.1	56.2	53.0
$\bar{\mathbb{E}}(SP)$	Bond spreads (%)	5.06	5.06	5.08
$\bar{\sigma}(SP)$	Std of spreads (%)	1.56	1.52	1.61
(b) Untargeted				
$\bar{\sigma}(c)/\bar{\sigma}(y)$	Std of consumption	1.28	1.28	1.26
$\text{cor}(tb/y, y)$	Corr. trade balance and output	-0.65	-0.67	-0.68
$\bar{\sigma}(tb/y)/\bar{\sigma}(y)$	Std of trade balance	0.38	0.36	0.34
(c) Default episodes				
$\bar{\mathbb{E}}(\Delta y_t/y_{t-1} \mid d_t = 1)$	Drop in output at default (%)	-4.71	-6.30	-6.00
$\bar{\mathbb{E}}(y_{t-1} - \bar{y} \mid d_t = 1)/\bar{\sigma}(y)$	Output before default (std)	-1.16	-0.63	-0.66
$\bar{\mathbb{E}}(b_{t-1}/y_{t-1} \mid d_t = 1)$	Debt ratio before default (%)	54.48	59.72	56.25
$\bar{\mathbb{E}}(SP_{t-1} \mid d_t = 1)$	Spreads before default (%)	8.48	7.64	7.89

Note: The table compares targeted and untargeted model moments, including outcomes around default episodes, across our baseline model and two recalibrated PE counterfactuals. The $\{d_0, d_1\}$ column sets $\kappa = 0$ and adjusts the default-cost parameters to match the targeted mean and standard deviation of bond spreads. In that column, the debt-to-output ratio is an untargeted outcome. The $\{d_0, d_1, \beta\}$ column jointly adjusts all three parameters to match the targeted spread moments and debt-to-output ratio.

the states in which defaults typically occur (Figure 7, Panel (a)). In these states, repayment variance is high and the demand schedule is steep, so the price impact in our baseline model already restrains current and future borrowing. In good times, the ranking reverses: repayment variance is lower and demand is more elastic, so the restraint from price impact weakens and the baseline requires higher default costs to generate similar spreads. Because defaults in good times are rare, realized costs of default are, on average, higher in the recalibrated PE model, shifting the entire distribution to the right (Panel (b)).

Despite targeting similar spreads, the two economies sustain different debt levels: the government issues more debt in the recalibrated PE model, with a debt-to-output ratio near 56% (Table 9, Panel (a)). In that model, the recalibrated default costs make default less attractive, lowering default risk and borrowing costs at a given debt stock. The government optimally responds by borrowing more. In our baseline, default risk also shapes borrowing costs, but the price impact of issuance places an additional restraint on borrowing. Panel (c) of Figure 7 shows that the recalibrated PE model's entire debt distribution lies to the right of the baseline distribution. Consistent with the higher debt it sustains, the government prefers the recalibrated PE economy, with a CEC gain of

Figure 7: Default costs and debt: Baseline vs. recalibrated PE model

(a) Default costs, function of y (b) Default costs (c) Debt-to-output

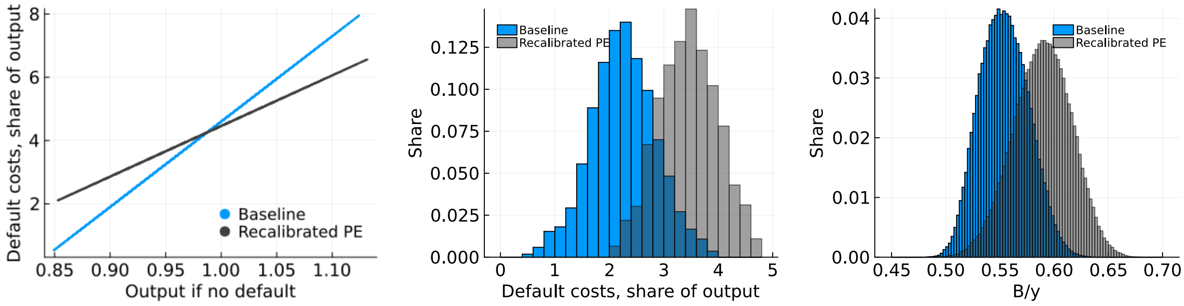

Note: The figure compares model-implied default costs and debt-to-output distributions under our baseline model and the recalibrated PE model, in which the parameters $\{d_0, d_1\}$ are adjusted to match the targeted mean and standard deviation of bond spreads. Panel (a) plots default costs as a share of output, $\Phi(y)/y$ (in percentage points), against output. Panel (b) shows the distribution of $\Phi(y)/y$ at the time of default. Panel (c) shows the distribution of debt-to-output ratios, excluding observations within three years of a default event.

approximately 0.09%.

In the last column of Table 9, we jointly recalibrate $\{d_0, d_1, \beta\}$ to match the targeted mean and standard deviation of bond spreads and the debt-to-output ratio.[35] We refer to this version as the *fully recalibrated PE model*. Under this calibration, the untargeted moments in Panel (b)—consumption and trade-balance volatility, and the trade-balance-output correlation—remain similar to those in the baseline and broadly consistent with empirical patterns in emerging economies. There are, however, two important differences between the two models that reflect their distinct sources of discipline.

The first difference concerns the dynamics around default episodes, summarized in Panel (c) of Table 9. The drop in y at default is smaller in the baseline (4.7% versus 6.0%), reflecting differences not only in calibrated default costs but also in the timing of default. In the period before default, output in our baseline is already 1.16 standard deviations below average—compared to 0.66 in the fully recalibrated PE model. Figure 8 offers a more complete picture. Panel (a) shows the average output path around default: in the baseline, output declines gradually over four years, whereas in the fully recalibrated PE model, it remains slightly above average until about a year before default. The smoother decline in the baseline is consistent with empirical studies such as Levy Yeyati and Panizza (2011), who document that output contractions typically begin more than three years before a sovereign default, while the drop during the default period itself is

[35]The calibration slightly increases β, which reduces borrowing incentives. The implied default costs are nearly identical to those in Figure 7, and we report them in Appendix C.4.

Figure 8: Periods around default episodes

(a) Endowment (b) Spreads (c) Debt

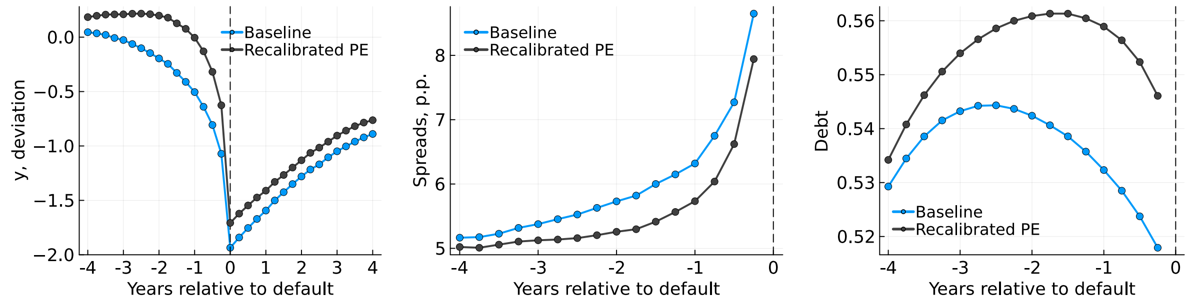

Note: The figure shows the average paths for endowment y, spreads, and debt around sovereign default episodes under the baseline model (blue) and the fully recalibrated PE model (black), which adjusts $\{d_0, d_1, \beta\}$. Time is centered on the default quarter ($t = 0$). Simulations in which the government was already in default during the four years preceding a new default episode are excluded. Panel (a) shows deviations from the average endowment y, expressed in standard deviations. The x-axis is in years.

relatively small.

The remaining panels of Figure 8 trace the mechanism behind these default dynamics. Panel (b) shows that spreads in the baseline rise earlier and more gradually than in the fully recalibrated PE model. This pattern reflects both rising default risk and the demand slope, which steepens as fundamentals deteriorate. Panel (c) displays the corresponding debt dynamics. In the baseline, the price impact of issuance leads the government to accumulate less debt and begin deleveraging sooner. This early adjustment delays default, so defaults occur after a longer, smoother decline in the endowment. In contrast, the government in the fully recalibrated PE model faces no such restraint and continues borrowing as fundamentals weaken. Debt is therefore larger, and default follows a shorter and more abrupt decline in output.

The second difference lies in the dilution component embedded in bond prices. In our baseline, the demand slope restrains future issuance, so lenders expect less dilution and price the bonds accordingly. This force is absent in the perfectly elastic model. The next subsection quantifies dilution in each economy and analyzes its implications for borrowing policies and fiscal rules.

4.4 Quantifying Debt Dilution and the Role of Fiscal Rules

Following Hatchondo et al. (2016), we introduce a debt covenant under which the government compensates existing bondholders for dilution caused by the increase in default risk associated with new issuance. The covenant is defined with respect to the perfectly elastic price. That is, it covers price declines arising from higher default risk, but not the direct price impact of issuance under imperfectly elastic demand. This facilitates comparison with the PE model and isolates the default-risk component of dilution. We solve this counterfactual for both our baseline and the fully recalibrated PE model. Tildes denote prices and policies in the corresponding covenant economy.

Formally, the debt covenant is

$$\mathcal{C}\left(y, \tau, B, B'\right) = \max\left\{0, \widetilde{q}_{\kappa=0}\left(y, \tau, (1-\lambda)B\right) - \widetilde{q}_{\kappa=0}\left(y, \tau, B'\right)\right\}, \tag{24}$$

where $\widetilde{q}_{\kappa=0}(\cdot)$ is the bond price function under perfectly elastic demand, evaluated at the policies of the covenant economy. When the covenant is applied to our baseline model, $\widetilde{q}_{\kappa=0}(\cdot)$ is an auxiliary pricing object that values the repayment stream as if demand were perfectly elastic. When the covenant is applied to the PE model, it coincides with the covenant economy's equilibrium bond price. In both cases, it is defined recursively as $\widetilde{q}_{\kappa=0}\left(y, \tau, B'\right) \equiv \frac{1}{r^f}\mathbb{E}_{s'|s}\widetilde{\mathcal{R}}_{\kappa=0}\left(y', \tau', B'\right)$, where the repayment function is

$$\widetilde{\mathcal{R}}_{\kappa=0}\left(y', \tau', B'\right) = \left[1 - \widetilde{d}\left(y', \tau', B'\right)\right]\left[\lambda + (1-\lambda)\left(\nu + \widetilde{q}'_{\kappa=0} + \mathcal{C}\left(y', \tau', B', \widetilde{B}''\right)\right)\right],$$

where $\widetilde{q}'_{\kappa=0} \equiv \widetilde{q}_{\kappa=0}\left(y', \tau', \widetilde{B}''\right)$, and $\widetilde{B}'' = \widetilde{B}'(y', \tau', B')$ is the next-period debt choice in the covenant economy.

With the covenant in place, the bond price function faced by the government is

$$\widetilde{q}\left(y, \tau, B'\right) = \frac{1}{r^f}\mathbb{E}_{s'|s}\widetilde{\mathcal{R}}\left(y', \tau', B'\right) - \frac{\kappa}{r^f}\mathbb{V}_{s'|s}\widetilde{\mathcal{R}}\left(y', \tau', B'\right)\left(B' - \mathcal{B}(\tau, B')\right), \tag{25}$$

where the repayment function reflects the covenant defined in (24):

$$\widetilde{\mathcal{R}}\left(y', \tau', B'\right) = \left[1 - \widetilde{d}\left(y', \tau', B'\right)\right]\left[\lambda + (1-\lambda)\left(\nu + \widetilde{q}\left(y', \tau', \widetilde{B}''\right) + \mathcal{C}\left(y', \tau', B', \widetilde{B}''\right)\right)\right].$$

Notice that in the $\kappa = 0$ case, this repayment function is identical to $\widetilde{\mathcal{R}}_{\kappa=0}(\cdot)$.

When not in default, and for a given B', the government's budget constraint becomes

$$\widetilde{c}\left(y, \tau, B, B'\right) = y + \widetilde{q}\left(y, \tau, B'\right)\left(B' - (1-\lambda)B\right) - \left[\lambda + (1-\lambda)\left(\nu + \mathcal{C}\left(y, \tau, B, B'\right)\right)\right]B. \tag{26}$$

All other equilibrium conditions and definitions follow analogously from the baseline model.

Table 10 compares the economies with dilution—our baseline and the fully recalibrated PE model—to their *no-dilution* counterfactuals. In both economies, default probabilities and bond spreads decline sharply, by more than 80%, highlighting the central role of debt dilution in the pricing of long-term risky bonds.[36] Two differences emerge in the no-

[36]The magnitudes are comparable to those reported in Hatchondo et al. (2016).

Table 10: Quantifying debt dilution and fiscal policy

Moment	Description	Imperfectly elastic case			Fully recalibrated PE		
		Baseline economy	No dilution	Debt ceiling	Baseline economy	No dilution	Debt ceiling
$\bar{\mathbb{E}}(B/y)$	Debt to output	53.1	51.2	52.7	53.0	55.5	53.8
$\bar{\mathbb{E}}(SP)$	Bond spreads	5.06	0.71	3.84	5.08	0.32	2.22
$\bar{\sigma}(SP)$	Volatility of spreads	1.56	0.40	1.71	1.61	0.18	1.54
$\bar{\mathbb{E}}(d_\lambda)$	Default probability	4.01	0.64	3.60	4.32	0.29	2.02
$\bar{\mathbb{E}}(\vartheta)$	Inelasticity premium	0.00	0.02	-0.64	0.00	0.00	0.00
$\bar{\mathbb{E}}(\text{CEC})$	Cons. equivalent	-	0.25	0.10	-	0.54	0.30

Note: The table reports selected moments for our baseline model with imperfectly elastic demand and the fully recalibrated perfectly elastic (PE) case. Within each panel, the first column reports the economy with no intervention. The second introduces the covenant of Equation (24), under which the government compensates existing bondholders for price declines caused by the higher default risk associated with new issuance. The third imposes a debt-to-output ceiling, $B'/y \leq 0.55$. Debt-to-output, spreads, default probabilities, and inelasticity premia are reported in percent. The last row reports the certainty-equivalent consumption (CEC) gain from each intervention relative to the corresponding economy with no intervention, in percent of permanent consumption.

dilution counterfactual. First, the covenant moves debt in opposite directions: relative to the calibrated target of 53%, the average debt-to-output ratio falls in our baseline (51.2%) but rises in the PE case (55.5%). Second, despite its lower debt, the baseline features default risk and spreads about twice as high as under perfectly elastic demand.

To understand these results, consider the two opposing forces through which the covenant affects borrowing. By protecting existing bondholders from dilution, the covenant raises bond prices, which supports more debt. At the same time, the government must compensate bondholders whenever new issuance raises default risk, which makes issuance more costly. In the PE model, even with the recalibrated default costs, expected dilution absent the covenant is large enough that the price improvement due to the covenant dominates its cost, and debt *rises*. In our baseline, instead, the demand slope already restrains future issuance, leaving less expected dilution for the covenant to eliminate. The compensation cost then dominates, and debt *falls*. These opposite debt responses echo the three-period analysis: once the covenant neutralizes the default-risk component of dilution, the dilution-curbing benefit of downward-sloping demand largely disappears, while the price-impact cost remains.

This logic extends to fiscal rules, such as a debt ceiling. Like the covenant, a ceiling curbs dilution, but it does so by limiting future borrowing rather than by compensating existing bondholders. Its effects should therefore be largest when dilution is high. To quantify this, we impose a debt-to-output ceiling of 55% in both economies. Results are reported in the debt-ceiling columns of Table 10. In the fully recalibrated PE model, the ceiling has large effects. Default risk and spreads fall by roughly half, while average debt *rises* slightly. The effects are smaller in our baseline model with imperfectly elastic demand. Default risk barely changes, and average debt falls slightly. Spreads also fall, but the decline is largely mechanical, as the lower residual supply raises bond prices and turns the average inelasticity premium negative.

The last row of Table 10 reports the average welfare effect of each intervention, measured as the CEC gain relative to the corresponding economy with no intervention. Both interventions raise welfare in the two economies, but the gains are smaller under imperfectly elastic demand: 0.25% versus 0.54% for the covenant, and 0.10% versus 0.30% for the debt ceiling. The smaller gains reflect that the demand slope already curbs dilution, leaving less for either intervention to remove.

Taken together, the two experiments show that the effects of policies that limit dilution depend on the source of discipline. In our calibration, the covenant and the debt ceiling have smaller welfare effects under imperfectly elastic demand and support more debt only in the fully recalibrated PE model.

5 Conclusion

For riskless or short-term debt, downward-sloping demand is typically understood to reduce borrowing, since additional issuance lowers the price the issuer receives for its bonds. We show that for risky long-term debt, the effect can run in the opposite direction. The same price impact that discourages current borrowing also restrains future issuance, reducing expected dilution and raising the value of debt issued today. When this dilution-curbing force outweighs the price-impact cost, less elastic demand can support *more* debt. After estimating a reduced-form demand elasticity from index-rebalancing flows and mapping it into the model through indirect inference, we find that the dilution-curbing force dominates: spreads and default risk fall, and the government sustains a larger debt stock. Because the demand slope itself restrains borrowing, observed spreads are consistent with lower realized default costs, and fiscal rules such as debt ceilings have more limited effects.

More broadly, our analysis brings the demand side into the study of sovereign debt supply. How much a government can borrow, and at what price, depends not only on its fundamentals and the penalties it would face in default, but also on who holds its bonds and how elastic investor demand is. We model these forces parsimoniously, leaving room for richer treatments of investor heterogeneity and institutional frictions that underlie them. As benchmark-driven and passive investing continue to grow in emerging markets, the demand side of the sovereign bond market is changing, with implications for debt pricing, default risk, and sovereign borrowing capacity.

References

Aguiar, M. and M. Amador (2020). Self-fulfilling debt dilution: Maturity and multiplicity in debt models. *American Economic Review 110*(9), 2783–2818.

Aguiar, M., M. Amador, G. Gopinath, and E. Farhi (2013). Crisis and commitment: Inflation credibility and the vulnerability to sovereign debt crises. NBER Working Paper 19516.

Aguiar, M., M. Amador, H. Hopenhayn, and I. Werning (2019). Take the short route: Equilibrium default and debt maturity. *Econometrica 87*(2), 423–462.

Aguiar, M., S. Chatterjee, H. Cole, and Z. Stangebye (2016). Quantitative models of sovereign debt crises. Volume 2 of *Handbook of Macroeconomics*, Chapter 21, pp. 1697–1755.

Aguiar, M. and G. Gopinath (2006). Defaultable debt, interest rates, and the current account. *Journal of International Economics 69*(1), 64–83.

Alfaro, L. and F. Kanczuk (2017). Fiscal rules and sovereign default. NBER Working Paper 23370.

Amador, M., I. Werning, and G.-M. Angeletos (2006). Commitment vs. flexibility. *Econometrica 74*(2), 365–396.

Arellano, C. (2008). Default risk and income fluctuations in emerging economies. *American Economic Review 98*(3), 690–712.

Arellano, C., Y. Bai, and S. Lizarazo (2017). Sovereign risk contagion. NBER Working Paper 24031.

Beltran, P. and C. He (2025). Inelastic financial markets and foreign exchange interventions. Working paper.

Bianchi, J., P. Ottonello, and I. Presno (2023). Fiscal stimulus under sovereign risk. *Journal of Political Economy 131*(9), 2328–2369.

Bigio, S., G. Nuño, and J. Passadore (2023). Debt-maturity management with liquidity costs. *Journal of Political Economy Macroeconomics 1*(1), 119–190.

Borri, N. and A. Verdelhan (2010). Sovereign risk premia. AFA 2010 Atlanta Meetings Paper.

Bretscher, L., L. Schmid, I. Sen, and V. Sharma (2026). Institutional corporate bond pricing. *The Review of Financial Studies 39*(3), 605–660.

Broner, F. A., G. Lorenzoni, and S. L. Schmukler (2013). Why do emerging economies borrow short term? *Journal of the European Economic Association 11*, 67–100.

Chang, Y., H. Hong, and I. Liskovich (2014). Regression discontinuity and the price effects of stock market indexing. *The Review of Financial Studies 28*(1), 212–246.

Chari, A., K. Dilts Stedman, and C. Lundblad (2022). Global fund flows and emerging market tail risk. NBER Working Paper 30577.

Chatterjee, S. and B. Eyigungor (2012). Maturity, indebtedness, and default risk. *American Economic Review 102*(6), 2674–2699.

Chaumont, G. (2025). Sovereign debt, default risk, and the liquidity of government bonds. Working paper.

Choi, J., R. Kirpalani, and D. J. Perez (2026). US public debt and safe asset market power. *Journal of Political Economy 134*(5), 1506–1560.

Cole, H. L., D. Neuhann, and G. Ordoñez (2025). Information spillovers and sovereign debt: Theory meets the eurozone crisis. *The Review of Economic Studies 92*(1), 197–237.

Costain, J., G. Nuño, and C. Thomas (2025). The term structure of interest rates in a heterogeneous monetary union. *Journal of Finance 80*(4), 2389–2434.

Cremers, M. and A. Petajisto (2009). How active is your fund manager? A new measure that predicts performance. *The Review of Financial Studies 22*(9), 3329–3365.

DeMarzo, P., Z. He, and F. Tourre (2023). Sovereign debt ratchets and welfare destruction. *Journal of Political Economy 131*(10), 2825–2892.

Dovis, A. and R. Kirpalani (2020). Fiscal rules, bailouts, and reputation in federal governments. *American Economic Review 110*(3), 860–888.

Escobar, M., L. Pandolfi, A. Pedraza, and T. Williams (2026). Who trades index rebalancings? Evidence on benchmarking and inelastic demand. CEPR Discussion Paper 21526.

Faia, E., J. Salomao, and A. V. Veghazy (2025). Market segmentation and international bond prices: The role of ECB asset purchases. *American Economic Journal: Macroeconomics 17*(4), 391–421.

Fang, X., B. Hardy, and K. Lewis (2025). Who holds sovereign debt and why it matters. *The Review of Financial Studies 38*(8), 2326–2361.

Gabaix, X. and R. Koijen (2021). In search of the origins of financial fluctuations: The inelastic markets hypothesis. NBER Working Paper 28967.

Gabaix, X. and M. Maggiori (2015). International liquidity and exchange rate dynamics.

Quarterly Journal of Economics 130(3), 1369–1420.

Greenwood, R. (2005). Short- and long-term demand curves for stocks: Theory and evidence on the dynamics of arbitrage. *Journal of Financial Economics 75*(3), 607–649.

Greenwood, R., S. Hanson, and J. Stein (2015). A comparative-advantage approach to government debt maturity. *Journal of Finance 70*(4), 1683–1722.

Greenwood, R., S. Hanson, and D. Vayanos (2024). Supply and demand and the term structure of interest rates. *Annual Review of Financial Economics 16*, 115–151.

Halac, M. and P. Yared (2021). Instrument-based versus target-based rules. *The Review of Economic Studies 89*(1), 312–345.

Halac, M. and P. Yared (2022). Fiscal rules and discretion under limited enforcement. *Econometrica 90*(5), 2093–2127.

Harris, L. and E. Gurel (1986). Price and volume effects associated with changes in the S&P 500 list: New evidence for the existence of price pressures. *Journal of Finance 41*(4), 815–829.

Hatchondo, J. C., L. Martinez, and F. Roch (2022a). Fiscal rules and the sovereign default premium. *American Economic Journal: Macroeconomics 14*(4), 244–273.

Hatchondo, J. C., L. Martinez, and F. Roch (2022b). Numerical fiscal rules for economic unions: The role of sovereign spreads. *Economics Letters 210*, 110168.

Hatchondo, J. C., L. Martinez, and C. Sosa-Padilla (2016). Debt dilution and sovereign default risk. *Journal of Political Economy 124*(5), 1383–1422.

Hau, H., M. Massa, and J. Peress (2010). Do demand curves for currencies slope down? Evidence from the MSCI Global Index change. *The Review of Financial Studies 23*(4), 1681–1717.

He, Z., B. Kelly, and A. Manela (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics 126*(1), 1–35.

Jiang, Z., H. Lustig, S. Van Nieuwerburgh, and M. Z. Xiaolan (2024). The U.S. public debt valuation puzzle. *Econometrica 92*(4), 1309–1347.

Jiang, Z., H. Lustig, S. Van Nieuwerburgh, and M. Z. Xiaolan (2026). Bond convenience yields in the eurozone currency union. *The Review of Financial Studies*.

Kaldorf, M. and J. Röttger (2025). Convenient but risky government bonds. *European Economic Review 180*, 105152.

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

Krishnamurthy, A. and A. Vissing-Jorgensen (2012). The aggregate demand for Treasury debt. *Journal of Political Economy 120*(2), 233–267.

Levy Yeyati, E. and U. Panizza (2011). The elusive costs of sovereign defaults. *Journal of Development Economics 94*(1), 95–105.

Lizarazo, S. (2013). Default risk and risk averse international investors. *Journal of International Economics 89*(2), 317–330.

Longstaff, F. A., J. Pan, L. H. Pedersen, and K. J. Singleton (2011). How sovereign is sovereign credit risk? *American Economic Journal: Macroeconomics 3*(2), 75–103.

Mian, A., L. Straub, and A. Sufi (2025). A goldilocks theory of fiscal deficits. *American Economic Review 115*(12), 4253–4291.

Monteiro, R. A. (2024). A debt crisis with strategic investors: Changes in demand and the role of market power. Working paper.

Morelli, J. and M. Moretti (2023). Information frictions, reputation, and sovereign spreads. *Journal of Political Economy 131*(11), 3066–3102.

Moretti, M. (2020). Financial innovation and liquidity premia in sovereign markets: The case of GDP-linked bonds. Working paper.

Pandolfi, L. and T. Williams (2019). Capital flows and sovereign debt markets: Evidence from index rebalancings. *Journal of Financial Economics 132*(2), 384–403.

Passadore, J. and Y. Xu (2022). Illiquidity in sovereign debt markets. *Journal of International Economics 137*, 103618.

Pavlova, A. and T. Sikorskaya (2022). Benchmarking intensity. *The Review of Financial Studies 36*(3), 859–903.

Petajisto, A. (2009). Why do demand curves for stocks slope down? *Journal of Financial and Quantitative Analysis 44*(5), 1013–1044.

Raddatz, C., S. Schmukler, and T. Williams (2017). International asset allocations and capital flows: The benchmark effect. *Journal of International Economics 108*, 413–430.

Ray, W., M. Droste, and Y. Gorodnichenko (2024). Unbundling quantitative easing: Taking a cue from treasury auctions. *Journal of Political Economy 132*(9), 3115–3172.

Shleifer, A. (1986). Do demand curves for stocks slope down? *Journal of Finance 41*(3), 579–590.

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

Zhou, H. (2024). The fickle and the stable: Global financial cycle transmission via heterogeneous investors. Working paper.

A Model Appendix

A.1 Investor Block

This appendix derives the pricing schedule in the main text under four alternative specifications: (i) institutional mandates, (ii) risk-averse arbitrageurs with preferred-habitat investors, (iii) Value-at-Risk constraints, and (iv) non-pecuniary benefits. All four deliver a pricing equation analogous to Equation (3).

Institutional mandates

We first detail the derivation of the pricing schedule starting from the Gabaix and Koijen (2021) investor mandates described in the main text. Let $j \in \{1, \ldots, J\}$ denote an individual investor. We define $x^i_{j,t} \equiv \frac{q^i_t B^i_{j,t}}{W_{j,t}}$ as the portfolio weight that investor j assigns to bond i at time t, where q^i_t is the unit price, $B^i_{j,t}$ represents physical holdings, and $W_{j,t}$ is investor j's wealth. We allow for a benchmark index $\mathcal{I}$ that includes bond i, with time-varying index weights $\boldsymbol{w}_t = \{w^1_t, \ldots, w^N_t\}$ across all N constituent bonds. Throughout, bond i corresponds to a single country.

Given a function of bond i's expected excess return, $\pi^i_t(r^i_{t+1})$, we specify investors' mandates as:

$$x^i_{j,t}\left(\pi^i_t\right) = \theta_{j,t}\, \xi^i_{j,t}\, e^{\Lambda_j \pi^i_t} + (1 - \theta_{j,t})\; w^i_t, \tag{A1}$$

where $\theta_{j,t} \in [0, 1]$ captures an active share. The mandate comprises two distinct motives. The first term is a return-sensitive (active) target: a reference exposure, $\theta_{j,t}\xi^i_{j,t}$, scaled by a tilt $e^{\Lambda_j \pi^i_t}$ that shifts this exposure with expected (risk-adjusted) excess returns, where $\Lambda_j > 0$ governs the strength of the response. The second term, $(1 - \theta_{j,t})\, w^i_t$, is a benchmark-tracking (passive) component that mechanically replicates index weights.

We assume that π^i_t is given by a mean-variance–style risk-adjusted excess return:

$$\pi^i_t\left(r^i_{t+1}\right) = \frac{\mathbb{E}_t\left(r^i_{t+1}\right)}{\mathbb{V}_t\left(r^i_{t+1}\right)} = q^i_t \frac{\mathbb{E}_t\left(\mathcal{R}^i_{t+1}\right) - q^i_t\, r^f}{\mathbb{V}_t\left(\mathcal{R}^i_{t+1}\right)}, \tag{A2}$$

where r^i_{t+1} denotes next-period excess returns, defined as $\mathcal{R}^i_{t+1}/q^i_t - r^f$, with $\mathcal{R}_{t+1}$ denoting next-period repayment per bond unit.

We first characterize the return-sensitive component of the mandate, which we refer to as the active position. We define $\Gamma^i_{j,t} \equiv W_{j,t}\theta_{j,t}\xi^i_{j,t}$ as investor j's time-varying reference exposure to bond i in dollars (i.e., the active dollar position when $\pi^i_t = 0$). The total active position is the sum of these exposures scaled by their return sensitivity: $\sum_j \Gamma^i_{j,t} e^{\Lambda_j \pi^i_t}$. We linearize this object around zero expected excess returns ($\pi_t^i = 0$) and around a reference cross-sectional vector $\{\Gamma_j^{i,\star}\}_{\forall j}$. This yields:

$$\sum_j \Gamma_{j,t}^i e^{\Lambda_j \pi_t^i} \approx \sum_j \Gamma_j^{i,\star} + \sum_j \left(\Gamma_{j,t}^i - \Gamma_j^{i,\star}\right) + \sum_j \Gamma_j^{i,\star} \Lambda_j \left(\pi_t^i - 0\right)$$
$$= \Gamma_t^i + \Lambda^i \pi_t^i,$$

where $\Gamma_t^i \equiv \sum_j \Gamma_{j,t}^i$ is the aggregate active dollar exposure, and the aggregate sensitivity Λ^i is defined as the exposure-weighted sum of individual sensitivities:

$$\Lambda^i \equiv \Gamma^{i,\star} \sum_j s_j^{i,\star} \Lambda_j, \text{ with } \quad \Gamma^{i,\star} \equiv \sum_j \Gamma_j^{i,\star} \quad \text{and} \quad s_j^{i,\star} \equiv \frac{\Gamma_j^{i,\star}}{\Gamma^{i,\star}}.$$

We let $\mathcal{W}_t \equiv \sum_j W_{j,t}(1-\theta_{j,t})$ denote the assets under management (AUM) of passive, benchmark-tracking investors. The total passive position for bond i (in dollars) is then $\mathcal{W}_t w_t^i$. We assume that a fraction α_t^i of country i's bonds are included in the index, so that the index's market value is $MV_t \equiv \sum_{c\in\mathcal{I}} \alpha_t^c q_t^c B_t^c$, and index weights are $w_t^i = \frac{\alpha_t^i q_t^i B_t^i}{MV_t}$.[1]

We next convert these dollar positions into physical bond units. We assume that $\xi_{j,t}^i = \xi_j^i q_t^i$, so that ξ_j^i captures a fixed component of investor j's mandate. This normalization ensures that the baseline allocation does not mechanically scale up in quantity terms as prices fall. Under this assumption, the active demand in units of the bond, denoted by $\mathcal{A}_t^i$, can be written as:

$$\mathcal{A}_t^i \approx \overline{\mathcal{A}_t^i} + \Lambda^i \frac{\pi_t^i}{q_t^i}, \tag{A3}$$

where $\overline{\mathcal{A}_t^i} \equiv \sum_j W_{j,t}\theta_{j,t}\xi_j^i$ represents a price-insensitive component of demand (a preferred-habitat demand as in Vayanos and Vila, 2021).

We denote passive demand in units of bond i by $\mathcal{T}_t^i \equiv \mathcal{W}_t w_t^i / q_t^i$. Using the fact that w_t^i is proportional to the price of bond i, we can write this as:

$$\mathcal{T}_t^i = \tau_t^i B_t^i, \quad \text{with} \quad \tau_t^i \equiv \frac{\mathcal{W}_t}{MV_t}\alpha_t^i. \tag{A4}$$

Here, τ_t^i is a time-varying share that moves with passive AUM and index composition. This implies that passive demand is price-inelastic in terms of quantities.[2]

We denote by $\mathcal{B}_t^i$ the total reference demand, which we define as bond i's demand

[1]For simplicity, this derivation uses a pure market-capitalization index. For the EMBI Global Diversified used in the empirical application, the formulae are analogous, but weights also depend on the country's diversification coefficient, f_t^i, and the market value of the index is computed over the adjusted amounts—see Equation (17).

[2]Changes in q_t^i still influence $\mathcal{T}_t^i$ through their effects on MV_t. Under the small open economy assumption, we treat the index-wide market value MV_t as exogenous to local price fluctuations.

when expected excess returns are zero. From the analysis above, this is given by

$$\mathcal{B}_t^i \equiv \sum_j \frac{W_{j,t} x_{j,t}^i (0)}{q_t^i} = \overline{\mathcal{A}_t^i} + \mathcal{T}_t^i. \tag{A5}$$

Combining the previous elements and imposing market clearing:

$$B_t^i = \mathcal{B}_t^i + \Lambda^i \frac{\pi_t^i}{q_t^i}.$$

Substituting π_t^i in Equation (A2) yields the following pricing equation:

$$q_t^i = \frac{\mathbb{E}_t \left(\mathcal{R}_{t+1}^i \right)}{r^f} - \frac{\kappa^i}{r^f} \mathbb{V}_t \left(\mathcal{R}_{t+1}^i \right) \left(B_t^i - \mathcal{B}_t^i \right), \tag{A6}$$

where $\kappa^i \equiv 1/\Lambda^i$. It is also useful to write this pricing schedule as

$$q_t^i = \frac{\mathbb{E}_t \left(\mathcal{R}_{t+1}^i \right)}{r^f} \Psi_t^i, \quad \text{with} \quad \Psi_t^i \equiv 1 - \kappa^i \frac{\mathbb{V}_t \left(\mathcal{R}_{t+1}^i \right)}{\mathbb{E}_t \left(\mathcal{R}_{t+1}^i \right)} \left(B_t^i - \mathcal{B}_t^i \right), \tag{A7}$$

where Ψ_t^i captures a proportional wedge relative to the risk-neutral price, induced by the downward-sloping component of demand.

Risk-averse arbitrageurs and preferred-habitat investors

We next consider a preferred-habitat framework. There are two types of investors: (i) risk-averse arbitrageurs, whose demand responds to expected excess returns, and (ii) "preferred-habitat" investors whose demand is inelastic with respect to expected returns. We depart slightly from Vayanos and Vila (2021) by allowing the arbitrageurs' compensation to depend on both absolute returns and performance relative to a benchmark index. This modification introduces the notion of passive demand into the arbitrageurs' problem, penalizing them for deviating from the index weights.

Consider risk-averse arbitrageurs $j \in \{1, ..., J\}$ with mean-variance preferences over both total portfolio returns and performance relative to a benchmark index $\mathcal{I}$. Let $i \in \{1, ..., N\}$ denote the set of constituent bonds with supplies $\boldsymbol{B}_t = \left\{ B_t^1, ..., B_t^N \right\}$ and index weights $\boldsymbol{w}_t = \left\{ w_t^1, ..., w_t^N \right\}$. Let $\boldsymbol{r}_{t+1} = \left\{ r_{t+1}^1, ..., r_{t+1}^N \right\}$ denote the vector of next-period excess returns.

For arbitrageur j, compensation is a convex combination of portfolio returns and returns relative to the benchmark $\mathcal{I}$. Let $\boldsymbol{x}_{j,t} = \{x_{j,t}^1, \ldots, x_{j,t}^N\}$ denote portfolio weights. Total compensation (TC) is

$$TC_{j,t} = \theta_j \left(\boldsymbol{x}_{j,t} \right)' \cdot \boldsymbol{r}_{t+1} + (1 - \theta_j) \left(\boldsymbol{x}_{j,t} - \boldsymbol{w}_t \right)' \cdot \boldsymbol{r}_{t+1} = \left(\boldsymbol{x}_{j,t} - (1 - \theta_j) \, \boldsymbol{w}_t \right)' \cdot \boldsymbol{r}_{t+1},$$

where $1 - \theta_j$ governs the importance of relative performance. For instance, when $\theta_j = 1$, the investor is fully active and ignores the benchmark.

Each investor chooses portfolio weights $\boldsymbol{x}_{j,t}$ to maximize $\mathbb{E}_t\left(TC_{j,t}\right) - \frac{\sigma_j}{2}\mathbb{V}_t\left(TC_{j,t}\right)$, where σ_j captures the investor's risk aversion. Their problem is as follows:

$$\max_{\boldsymbol{x}_{j,t}} \left[\boldsymbol{x}_{j,t} - \left(1-\theta_j\right)\boldsymbol{w}_t\right]' \boldsymbol{\mu}_t - \frac{\sigma_j}{2}\left[\boldsymbol{x}_{j,t} - \left(1-\theta_j\right)\boldsymbol{w}_t\right]' \boldsymbol{\Sigma}_t \left[\boldsymbol{x}_{j,t} - \left(1-\theta_j\right)\boldsymbol{w}_t\right],$$

where $\boldsymbol{\mu}_t \equiv \mathbb{E}_t\left(\boldsymbol{r}_{t+1}\right)$ denotes expected excess returns and $\boldsymbol{\Sigma}_t \equiv \mathbb{V}_t\left(\boldsymbol{r}_{t+1}\right)$ is the variance-covariance matrix. The optimal portfolio for investor j is given by

$$\boldsymbol{x}_{j,t} = \frac{1}{\sigma_j}\boldsymbol{\Sigma}_t^{-1}\boldsymbol{\mu}_t + \left(1-\theta_j\right)\boldsymbol{w}_t. \tag{A8}$$

The first term in Equation (A8) is the standard mean-variance portfolio. The second term captures a reluctance to deviate from the benchmark weights $\boldsymbol{w}_t$, generating an inelastic demand component. This component is independent of expected returns and risk, and depends only on the penalty for benchmark deviations. Purely passive investors (i.e., $\theta_j = 0$ and $\sigma_j \to \infty$) never deviate from $\boldsymbol{w}_t$ and therefore exhibit perfectly inelastic demand. Let $W_{j,t}$ denote investor j's wealth so their holdings of bond i are $B^i_{j,t} = W_{j,t}x^i_{j,t}/q^i_t$.

In addition to arbitrageurs, we introduce preferred-habitat investors, who have preferences over specific bonds. We consider the simplest case in which this demand is perfectly inelastic, and denote it with $\overline{\mathcal{A}^i_t}$.

Combining the two investor types, market clearing requires $q^i_t\left(B^i_t - \overline{\mathcal{A}^i_t}\right) = \sum_j W_{j,t}x^i_{j,t}$. Replacing with the arbitrageurs' optimal portfolios, this can be written in vector form as

$$\boldsymbol{q}_t \circ \left(\boldsymbol{B}_t - \overline{\boldsymbol{\mathcal{A}_t}}\right) = \sum_j W_{j,t}\left[\frac{1}{\sigma_j}\boldsymbol{\Sigma}_t^{-1}\boldsymbol{\mu}_t + \left(1-\theta_j\right)\boldsymbol{w}_t\right]. \tag{A9}$$

To simplify, we analyze the demand for bond i conditional on the allocation of the remaining $N-1$ bonds, which we treat as a composite portfolio z. Let $\sigma_{z,j,t}$ and $\mu_{z,j,t}$ denote the volatility and excess return of this background portfolio for investor j. We index these moments by j because the weights within the rest of the portfolio may vary across investor types. Finally, let $\rho_{iz,j,t}$ denote the correlation between bond i and portfolio z. From Equation (A8), the optimal share for bond i is:

$$x^i_{j,t} = \frac{1}{\sigma_j}\frac{\mu_{i,t}}{\sigma^2_{i,t}}\zeta^i_{j,t} + (1-\theta_j)w^i_t, \quad \text{where} \quad \zeta^i_{j,t} \equiv \frac{1}{1-\rho^2_{iz,j,t}}\left[1 - \frac{\mu_{z,j,t}}{\mu_{i,t}}\frac{\sigma_{i,t}}{\sigma_{z,j,t}}\rho_{iz,j,t}\right]. \tag{A10}$$

The $\zeta^i_{j,t}$ term implies that investor j allocates more wealth to bond i when it is a good hedge ($\rho_{iz,j,t} < 0$) or has lower relative volatility.

Replacing Equation (A10) in (A9), the market-clearing condition for bond i becomes

$$B^i_t = \frac{1}{q^i_t}\sum_j \frac{1}{\sigma_j}\left[\frac{\mathbb{E}_t\left(r^i_{t+1}\right)}{\mathbb{V}_t\left(r^i_{t+1}\right)}\zeta^i_{j,t}\right]W_{j,t} + \mathcal{T}^i_t + \overline{\mathcal{A}^i_t},$$

where $\mathcal{T}_t^i \equiv \frac{1}{q_t^i} \sum_j W_{j,t} (1-\theta_j)\, w_t^i$ denotes the benchmark-tracking passive demand which, as shown in (A4), is perfectly inelastic and can be written as a time-varying fraction of B_t^i. We then define $\mathcal{B}_t^i = \overline{\mathcal{A}_t^i} + \mathcal{T}_t^i$ as the reference demand. Replacing with $r_{t+1}^i = \mathcal{R}_{t+1}^i / q_t^i - r^f$ and rearranging, we get an analogous pricing schedule to that of Equation (3):

$$q_t^i = \frac{1}{r^f} \mathbb{E}_t \left(\mathcal{R}_{t+1}^i \right) - \frac{\widetilde{\kappa_t}^i}{r^f} \mathbb{V}_t \left(\mathcal{R}_{t+1}^i \right) \left(B_t^i - \mathcal{B}_t^i \right), \tag{A11}$$

where $\widetilde{\kappa_t}^i \equiv 1/\left(\sum_j \frac{1}{\sigma_j} W_{j,t} \zeta_{j,t}^i \right)$ is the inverse of aggregate effective risk-bearing capacity (the wealth-weighted sum of investors' risk tolerances $1/\sigma_j$, adjusted for hedging via $\zeta_{j,t}^i$).

The primary distinction between this approach and the mandate-based framework is that the slope $\widetilde{\kappa_t}^i$ is explicitly linked to the arbitrageurs' aggregate effective risk-bearing capacity, rather than a mandate-driven sensitivity parameter. However, both specifications result in a pricing schedule where the slope scales with the conditional variance of repayment.[3]

Value-at-Risk constraints

An analogous pricing equation arises for risk-neutral arbitrageurs subject to a Value-at-Risk (VaR) constraint (as in Gabaix and Maggiori, 2015). Specifically, investors maximize the expected value of the compensation structure defined in the previous subsection, subject to a bound on its variance. Investor j chooses portfolio weights $\boldsymbol{x}_{j,t}$ to solve:

$$\begin{aligned} \max_{\boldsymbol{x}_{j,t}} \quad & \left[\boldsymbol{x}_{j,t} - (1-\theta_j)\, \boldsymbol{w}_t \right]' \boldsymbol{\mu}_t \\ \text{subject to} \quad & \frac{\mathit{\Phi}_t}{2} \left[\boldsymbol{x}_{j,t} - (1-\theta_j)\, \boldsymbol{w}_t \right]' \boldsymbol{\Sigma}_t \left[\boldsymbol{x}_{j,t} - (1-\theta_j)\, \boldsymbol{w}_t \right] \leq 1, \end{aligned}$$

where $\mathit{\Phi}_t > 0$ governs the tightness of the risk constraint and can be interpreted as a regulatory limit on the amount of risk investors are allowed to take. Let $\varrho_{j,t}$ denote the Lagrange multiplier on the VaR constraint. The optimal portfolio is then:

$$\boldsymbol{x}_{j,t} = \frac{1}{\varrho_{j,t} \mathit{\Phi}_t} \boldsymbol{\Sigma}_t^{-1} \boldsymbol{\mu}_t + (1-\theta_j)\, \boldsymbol{w}_t. \tag{A12}$$

Conditional on a binding VaR constraint, this optimal portfolio takes the same form as Equation (A8), with the primitive risk-aversion parameter σ_j replaced by the endogenous product $\varrho_{j,t} \mathit{\Phi}_t$. Following the same steps as before, we can then derive an analogous pricing equation.

[3] Unlike the baseline specification based on mandates, this expression is not fully closed-form because $\zeta_{j,t}^i$ depends on bond prices.

Non-pecuniary benefits

Lastly, we introduce an alternative case in which investors derive non-pecuniary benefits (or costs) from holding assets (as in Greenwood et al., 2024; Choi et al., 2026, and many others). To keep the analysis simple, we assume a representative, risk-neutral, deep-pocketed investor—it is straightforward to extend the framework to include multiple investors and passive demand. The investor chooses b_t^i to maximize:

$$\max_{b_t^i} \frac{\mathbb{E}_t\left(\mathcal{R}_{t+1}^i\right)}{r^f} b_t^i \; - \; q_t^i b_t^i \; + \; \left(a_t^i b_t^i - \tfrac{\varphi_t^i}{2}\left(b_t^i\right)^2\right). \tag{A13}$$

The quadratic term $a_t^i b_t^i - \frac{\varphi_t^i}{2}\left(b_t^i\right)^2$ can be viewed as capturing a non-pecuniary benefit from holding b_t^i units of bond i, such as liquidity or collateral services.

Using the first-order condition for b_t^i and imposing market clearing yields:

$$q_t^i = \frac{1}{r^f}\mathbb{E}_t \mathcal{R}_{t+1}^i - \varphi_t^i\left(B_t^i - \mathcal{B}_t^i\right), \tag{A14}$$

with $\mathcal{B}_t^i = a_t^i/\varphi_t^i$. The pricing schedule in the main text is analogous to Equation (A14), provided the curvature of the non-pecuniary benefit function φ_t^i scales with the volatility of repayment, i.e., $\varphi_t^i \propto \mathbb{V}_t(\mathcal{R}_{t+1}^i)$.

A.2 Definition of Equilibrium

A Recursive Markov Equilibrium is a collection of value functions $\left\{V(\cdot), V^r(\cdot), V^d(\cdot)\right\}$; policy functions $\{B'(\cdot), d(\cdot)\}$; and bond prices $q(\cdot)$ such that:

1. Taking as given the bond price function $q(\cdot)$, the government's policy functions $B'(\cdot)$ and $d(\cdot)$ solve the optimization problem in Equations (5), (6), and (7), and $V(\cdot)$, $V^r(\cdot)$, and $V^d(\cdot)$ are the associated value functions.
2. Given $B'(\cdot)$ and $d(\cdot)$, the repayment function $\mathcal{R}(\cdot)$ satisfies Equation (8).
3. Taking the repayment function as given, bond prices $q(\cdot)$ are consistent with Equation (9).

A.3 Additional derivations for the three-period model

This appendix collects the main derivations for the three-period model in Section 2.4. We maintain $t \in \{0, 1, 2\}$, $(y_0, y_1) = (0, 0)$, and y_2 supported on $(\underline{y}, \overline{y})$ with CDF $F(\cdot)$ and density $f(\cdot)$. Throughout, we impose linear preferences $u(c) = c$ and default costs $\Phi(y_2) = y_2/\phi$, with $\phi \geq 1$.

Government problem

Period 2. Given B_1, the government defaults if $y_2 < \phi B_1$. We can write the period-2 expected continuation payoff as

$$\mathbb{E}[c_2(B_1, y_2)] = \int_{\phi B_1}^{\overline{y}} (y_2 - B_1)\, f(y_2)\, dy_2 + \int_{\underline{y}}^{\phi B_1} y_2\Big(1 - \frac{1}{\phi}\Big)\, f(y_2)\, dy_2. \tag{A15}$$

Applying Leibniz's rule:

$$\frac{d}{dB_1}\mathbb{E}[c_2(B_1, y_2)] = -\int_{\phi B_1}^{\overline{y}} f(y_2)\, dy_2 = -\Big[1 - F(\phi B_1)\Big]. \tag{A16}$$

Bond Prices. The default rule implies that $\Pr(\text{default} \mid B_1) = F(\phi B_1)$. There is no recovery and thus repayment per unit of promised payoff is $\mathcal{R}(B_1, y_2) = 1 - \mathbf{1}\{y_2 < \phi B_1\}$. Following the pricing kernel of Section 2.2, the period-1 bond price is

$$q_1(B_1; \kappa) = 1 - F(\phi B_1) - \mathcal{S}(B_1; \kappa)\,(B_1 - \overline{\mathcal{B}}), \tag{A17}$$

where $\mathcal{S}(B_1; \kappa)$ is the demand slope. Since $\mathcal{S}(B_1; 0) = 0$ and $\mathcal{S}(\cdot)$ is differentiable in κ, we work throughout with the first-order expansion $\mathcal{S}(B_1; \kappa) = \kappa\, \mathcal{S}_\kappa(B_1) + o(\kappa)$, with $\mathcal{S}_\kappa(B_1) \equiv \frac{\partial \mathcal{S}(B_1;\kappa)}{\partial \kappa}|_{\kappa=0}$, and drop higher-order terms in what follows. Since there are no shocks between $t = 0$ and $t = 1$, $q_0(B_0; \kappa) = q_1(B_1(B_0; \kappa); \kappa)$.

Period 1. At $t = 1$ the government inherits B_0 and issues one-period bonds in the amount $B_1 - B_0$. The period-1 budget constraint is $c_1 = q_1(B_1; \kappa)\,(B_1 - B_0) \geq 0$. We can then write the period-1 problem as

$$V_1(B_0; \kappa) = \max_{B_1 \geq B_0}\; q_1(B_1; \kappa)\,(B_1 - B_0) + \beta\, \mathbb{E}_{y_2}\Big[c_2(B_1, y_2)\Big], \tag{A18}$$

where $c_2(B_1, y_2)$ is the consumption implied by the cutoff default rule above. Let $\mu_1 \geq 0$ denote the multiplier on the $B_1 \geq B_0$ constraint. Using (A16), the Kuhn–Tucker conditions for the period-1 choice are

$$\Psi_1(B_1; B_0, \kappa) \equiv q_1(B_1; \kappa) + q_{1,B_1}(B_1; \kappa)(B_1 - B_0) - \beta\Big[1 - F(\phi B_1)\Big] + \mu_1 = 0 \tag{A19}$$
$$\text{with}\;\; \mu_1 \geq 0,\; B_1 - B_0 \geq 0,\; \mu_1(B_1 - B_0) = 0.$$

Here $q_{1,B_1}(B_1; \kappa)$ is the derivative of the price with respect to B_1:

$$q_{1,B_1}(B_1; \kappa) = -\phi f(\phi B_1) - \kappa \mathcal{S}_{\kappa,B_1}(B_1)\Big(B_1 - \overline{\mathcal{B}}\Big) - \kappa \mathcal{S}_\kappa(B_1)\,, \tag{A20}$$

with $\mathcal{S}_{\kappa,B_1}(B_1) \equiv \partial \mathcal{S}_\kappa(B_1)\,/\partial B_1$.

Period 0. At $t = 0$, the government starts with no debt and chooses B_0. Its budget constraint is $c_0 = q_0\left(B_0;\kappa\right)B_0 \geq 0$. The period-0 value function is:

$$V_0\left(\kappa\right) = \max_{B_0 \geq 0}\; q_1\left(B_1\left(B_0;\kappa\right);\kappa\right)B_0 + \beta V_1(B_0;\kappa). \tag{A21}$$

From the envelope condition for the period-1 problem,

$$\frac{\partial V_1(B_0;\kappa)}{\partial B_0} = -q_1(B_1;\kappa) - \mu_1. \tag{A22}$$

Let $\Upsilon(B_0;\kappa) \equiv \frac{\partial B_1(B_0;\kappa)}{\partial B_0}$ denote the derivative of the period-1 debt policy with respect to B_0 (when it exists). Let $\mu_0 \geq 0$ denote the multiplier on $B_0 \geq 0$. Using (A22), the optimality conditions are

$$\Psi_0(B_0;\kappa) \equiv (1-\beta)\,q_1(B_1;\kappa) + q_{1,B_1}(B_1;\kappa)\,\Upsilon(B_0;\kappa)\,B_0 - \beta\,\mu_1 + \mu_0 = 0 \tag{A23}$$
$$\text{with}\;\; \mu_0 \geq 0,\; B_0 \geq 0,\; \mu_0 B_0 = 0,$$

where $B_1 \equiv B_1\left(B_0;\kappa\right)$ denotes the period-1 policy. The pricing function in (A17) and Equations (A19) and (A23) characterize the solution of the three-period model.

Perfectly elastic case: $\kappa = 0$

We characterize the solution $(B_0^\star, B_1^\star)$ under the $\kappa = 0$ case. In this case, bond prices simplify to $q_1(B_1;0) = 1 - F(\phi B_1)$ and $q_{1,B_1}(B_1;0) = -\phi f(\phi B_1)$. We focus on solutions with non-degenerate default risk, $0 < F(\phi B_1^\star) < 1$. The constraint $B_1 \geq B_0$ is then slack: if $B_1 = B_0$, Equation (A19) would require $\mu_1 = -(1-\beta)\Big[1 - F(\phi B_1)\Big] < 0$, a contradiction. With $\mu_1 = 0$, the period-1 optimality condition simplifies to

$$\Psi_1(B_1;B_0,0) = (1-\beta)\Big[1 - F(\phi B_1)\Big] - \phi f(\phi B_1)(B_1 - B_0) = 0. \tag{A24}$$

Defining the inverse hazard rate $H(x) \equiv \frac{1-F(x)}{f(x)}$ and dividing (A24) by $\phi f(\phi B_1)$ yields

$$B_1 - B_0 - \frac{1-\beta}{\phi}H(\phi B_1) = 0. \tag{A25}$$

This proves the first part of Proposition 1. Since $H(\phi B_1^\star) > 0$ under non-degenerate default risk, the condition implies $B_1^\star > B_0^\star$. A sufficient local second-order condition is:

$$(1-\beta)\,H_y\left(\phi B_1^\star\right) - 1 < 0, \tag{A26}$$

where $H_y\left(y\right) \equiv \frac{\partial H(y)}{\partial y}$. Note that if $H_y\left(y\right) \leq 0$ globally, the solution $B_1(B_0;0)$ is unique. Implicitly differentiating (A25) yields the pass-through:

$$\Upsilon(B_0;0) = \frac{1}{1-(1-\beta)H_y(\phi B_1)}, \tag{A27}$$

which corresponds to Equation (12) in the main text.

Turning to period 0, Equation (A23) implies $\mu_0^\star = 0$. To see this, suppose $B_0^\star = 0$.

The condition becomes $(1-\beta)\Big(1-F(\phi B_1^\star)\Big)+\mu_0^\star=0$. Since $F(\phi B_1^\star)<1$, this requires $\mu_0^\star<0$, a contradiction. Thus $B_0^\star>0$ and $\mu_0^\star=0$, reducing the optimality condition to

$$\Psi_0(B_0;0)=(1-\beta)\Big[1-F(\phi B_1)\Big]-\phi f(\phi B_1)\,\Upsilon(B_0;0)\,B_0=0. \tag{A28}$$

Dividing by $f(\phi B_1)$ and using the definition of $H(\cdot)$, we obtain

$$B_0=\frac{1}{\phi}(1-\beta)\,\frac{H(\phi B_1)}{\Upsilon(B_0;0)}. \tag{A29}$$

Evaluating at $(B_0^\star,B_1^\star)$ yields the expression in Proposition 1.

Commitment benchmark

The following proposition compares the $\kappa=0$ case to a commitment scenario in which the government chooses the entire debt path at $t=0$ and cannot revise it at $t=1$, while it still retains the option to repay or default at $t=2$ after observing y_2.

Proposition A.1. Commitment benchmark. *Suppose that at $t=0$ the government chooses $\{B_0,B_1\}$ and commits to that path. On $(\underline{y},\overline{y})$, assume that $H(\cdot)$ is twice continuously differentiable, satisfies $H_{yy}(y)\geq 0$, and that the function $m(y)\equiv y/H(y)$ is strictly increasing. Then, under $\kappa=0$, the commitment allocation satisfies $B_1^c-B_0^c=0$ and $B_0^c>B_0^\star$ for any no-commitment equilibrium $(B_0^\star,B_1^\star)$ with non-degenerate default risk.*

Proof. The default cutoff rule is identical to the baseline: default occurs if and only if $y_2<\phi B_1$. Under $\kappa=0$ and no shocks between $t=0$ and $t=1$, the bond price equals $q_1(B_1;0)=1-F(\phi B_1)$. The $t=0$ commitment problem is

$$\max_{B_0\geq 0,\,B_1\geq B_0}\; q_1(B_1;0)\,B_0+\beta\,q_1(B_1;0)\,(B_1-B_0)+\beta^2\,\mathbb{E}\Big[\max\Big\{y_2-B_1,\;y_2\Big(1-\frac{1}{\phi}\Big)\Big\}\Big].$$

Holding B_1 fixed, the only term that varies with B_0 is $(1-\beta)q_1(B_1;0)B_0$, which is strictly increasing in B_0 whenever $q_1(B_1;0)>0$. At any optimum, $q_1(B_1;0)>0$ because an allocation with $F(\phi B_1)=1$ raises no revenue and is strictly dominated by $B_0=B_1=0$. Hence, the constraint $B_0\leq B_1$ binds and $B_0^c=B_1^c\equiv B^c$. The commitment problem therefore reduces to

$$\max_{B\geq 0}\; q_1(B;0)\,B+\beta^2\,\mathbb{E}\Big[\max\Big\{y_2-B,\;y_2\Big(1-\frac{1}{\phi}\Big)\Big\}\Big].$$

The Kuhn–Tucker conditions are

$$\begin{aligned}&\Psi_c(B)\equiv q_1(B;0)+q_{1,B}(B;0)\,B-\beta^2\Big[1-F(\phi B)\Big]+\hat{\mu}=0\\ &\quad\text{with }\;\hat{\mu}\geq 0,\quad B\geq 0,\quad \hat{\mu}B=0,\end{aligned} \tag{A30}$$

where $\hat{\mu} \geq 0$ denotes the Lagrange multiplier. Substituting for $q_1(B;0)$ and $q_{1,B}(B;0)$, (A30) becomes

$$(1-\beta^2)\Big[1-F(\phi B)\Big] - \phi f(\phi B)\, B + \hat{\mu} = 0, \qquad \hat{\mu} \geq 0,\ B \geq 0,\ \hat{\mu} B = 0. \tag{A31}$$

We first rule out the corner $B = 0$. If $B = 0$, (A31) implies $\hat{\mu} = -(1-\beta^2)\big[1-F(0)\big]$. Since $q_1(0;0) = 1 - F(0) > 0$, this yields $\hat{\mu} < 0$, contradicting $\hat{\mu} \geq 0$. Hence $B^c > 0$, so $\hat{\mu} = 0$ and the interior first-order condition is

$$(1-\beta^2)\Big[1-F(\phi B)\Big] - \phi f(\phi B)\, B = 0.$$

The previous equation can be written as

$$m(\phi B^c) = \frac{\phi B^c}{H(\phi B^c)} = 1 - \beta^2. \tag{A32}$$

Since $m(\cdot)$ is strictly increasing on $(\underline{y}, \overline{y})$, this commitment solution is unique.

We next compare B^c with a no-commitment equilibrium $(B_0^\star, B_1^\star)$ under $\kappa = 0$ with non-degenerate default risk. From the period-1 condition (A25),

$$\phi B_1^\star - \phi B_0^\star = (1-\beta)\, H(\phi B_1^\star). \tag{A33}$$

From the period-0 condition (A29) and the $\kappa = 0$ identity (A27),

$$\phi B_0^\star = (1-\beta)\, H(\phi B_1^\star)\Big[1-(1-\beta)H_y(\phi B_1^\star)\Big] = (1-\beta)\,\frac{H(\phi B_1^\star)}{\Upsilon(B_0^\star;0)}. \tag{A34}$$

The assumption $H_{yy} \geq 0$ implies the tangent inequality at $\phi B_1^\star$:

$$H(\phi B_0^\star) \geq H(\phi B_1^\star) + H_y(\phi B_1^\star)\Big(\phi B_0^\star - \phi B_1^\star\Big).$$

Substituting Equation (A33), we get

$$H(\phi B_0^\star) \geq H(\phi B_1^\star)\Big[1-(1-\beta)H_y(\phi B_1^\star)\Big] = \frac{H(\phi B_1^\star)}{\Upsilon(B_0^\star;0)}. \tag{A35}$$

Multiplying both sides by $1-\beta$ and combining with Equation (A34) gives

$$\phi B_0^\star \leq (1-\beta)\, H(\phi B_0^\star) \quad \Rightarrow \quad m(\phi B_0^\star) = \frac{\phi B_0^\star}{H(\phi B_0^\star)} \leq 1-\beta.$$

Since $\beta \in (0,1)$, we have $1-\beta < 1-\beta^2$. Under commitment, (A32) implies $m(\phi B^c) = 1-\beta^2$, so $m(\phi B_0^\star) < m(\phi B^c)$. Because $m(\cdot)$ is strictly increasing, this implies $B_0^\star < B^c$. □

Downward-sloping demand: $\kappa > 0$

Proof of Proposition 2. Let $B^\star = \{B_0^\star, B_1^\star\}$ denote the solution under $\kappa = 0$, where $B_1^\star \equiv B_1(B_0^\star; 0)$. Holding inherited debt fixed at $B_0^\star$, consider a sufficiently small $\kappa > 0$. By continuity of $B_1(B_0^\star;\kappa)$ at $\kappa = 0$, the strict inequality $B_1^\star > B_0^\star$ continues to hold locally. The constraint therefore remains slack, so $\mu_1 = 0$ and the period-1 first-order condition is

$$\Psi_1\left(B_1;B_0,\kappa\right)=(1-\beta)\left[1-F\left(\phi B_1\right)\right]-\phi f\left(\phi B_1\right)\left(B_1-B_0\right)-\kappa\mathcal{S}_\kappa(B_1)\left(B_1-\overline{\mathcal{B}}\right)$$
$$-\kappa\left[\mathcal{S}_\kappa(B_1)+\mathcal{S}_{\kappa,B_1}(B_1)\left(B_1-\overline{\mathcal{B}}\right)\right]\left(B_1-B_0\right)=0. \tag{A36}$$

Evaluating at the $\kappa=0$ solution, $B^\star=\{B_0^\star,B_1^\star\}$, and setting $\overline{\mathcal{B}}=B_1^\star$ gives

$$\Psi_1\left(B_1^\star;B_0^\star,\kappa\right)=-\kappa\,\mathcal{S}_\kappa(B_1^\star)\left(B_1^\star-B_0^\star\right)<0.$$

By the period-1 second-order condition, $\Psi_{1,B_1}<0$ locally, so $\Psi_1(\cdot;B_0^\star,\kappa)$ is strictly decreasing in B_1 around $B_1^\star$. Since $\Psi_1(B_1^\star;B_0^\star,\kappa)<0$ and $\Psi_1(B_1(B_0^\star;\kappa);B_0^\star,\kappa)=0$, it follows that $B_1(B_0^\star;\kappa)<B_1^\star$. □

Proof of Proposition 3. Let $B_0(\kappa)$ denote the local solution to the period-0 problem, with $B_0(0)=B_0^\star$. Since $B_0^\star>0$, continuity implies that $B_0(\kappa)>0$ for sufficiently small $\kappa>0$, so the period-0 constraint remains slack. The first-order condition is then

$$\Psi_0\left(B_0;\kappa\right)=(1-\beta)\left[1-F\left(\phi B_1\left(B_0;\kappa\right)\right)-\kappa\,\mathcal{S}_\kappa(B_1\left(B_0;\kappa\right))\left(B_1\left(B_0;\kappa\right)-\overline{\mathcal{B}}\right)\right] \tag{A37}$$
$$-\left[\phi f\left(\phi B_1\left(B_0;\kappa\right)\right)+\kappa\left[\mathcal{S}_{\kappa,B_1}(B_1\left(B_0;\kappa\right))\left(B_1\left(B_0;\kappa\right)-\overline{\mathcal{B}}\right)+\mathcal{S}_\kappa(B_1\left(B_0;\kappa\right))\right]\right]\Upsilon\left(B_0;\kappa\right)B_0$$
$$=0$$

To simplify notation, we define $\mathring{B}_1\equiv B_1(B_0^\star;\kappa)$. Evaluating the first-order condition at $\{B_0^\star,\mathring{B}_1\}$, we get:

$$\Psi_0\left(B_0^\star;\kappa\right)=(1-\beta)\left[1-F\left(\phi\mathring{B}_1\right)\right]-\left[\phi f\left(\phi\mathring{B}_1\right)+\kappa\,\mathcal{S}_\kappa\left(\mathring{B}_1\right)\right]\Upsilon\left(B_0^\star;\kappa\right)B_0^\star$$
$$-\kappa\left[\mathring{B}_1-\overline{\mathcal{B}}\right]\left[(1-\beta)\,\mathcal{S}_\kappa\left(\mathring{B}_1\right)+\mathcal{S}_{\kappa,B_1}\left(\mathring{B}_1\right)\Upsilon\left(B_0^\star;\kappa\right)B_0^\star\right].$$

To determine whether long-term issuance rises with κ, we apply the implicit function theorem to $\Psi_0(B_0;\kappa)=0$, which gives

$$\left.\frac{dB_0(\kappa)}{d\kappa}\right|_{\kappa=0}=-\frac{\Psi_{0,\kappa}(B_0^\star;0)}{\Psi_{0,B_0}(B_0^\star;0)}.$$

Under the period-0 second-order condition, $\Psi_{0,B_0}(B_0^\star;0)<0$, which can be written in primitive terms as

$$1+\frac{(1-\beta)^2\,H(\phi B_1^\star)\,H_{yy}(\phi B_1^\star)}{\left[\,1-(1-\beta)H_y(\phi B_1^\star)\,\right]^2}>0,$$

with the simple sufficient condition $H_{yy}(\phi B_1^\star)\geq 0$. Hence, $B_0(\kappa)>B_0^\star$ for all sufficiently small $\kappa>0$ if and only if

$$\Psi_{0,\kappa}(B_0^\star;0)\equiv\left.\frac{\partial\Psi_0(B_0^\star;\kappa)}{\partial\kappa}\right|_{\kappa=0}>0.$$

This derivative is given by

$$\Psi_{0,\kappa}(B_0^\star;0) = \phi f\left(\phi B_1^\star\right)(1-\beta)\, H_y\left(\phi B_1^\star\right)\, B_{1,\kappa}\left(B_0^\star\right) \tag{A38}$$
$$-\phi f\left(\phi B_1^\star\right) B_0^\star\, \Upsilon_\kappa\left(B_0^\star\right) - \mathcal{S}_\kappa(B_1^\star)\,\Upsilon\left(B_0^\star;0\right) B_0^\star,$$

where

$$\Upsilon_\kappa(B_0^\star) \equiv \frac{\partial\Upsilon(B_0^\star;\kappa)}{\partial\kappa}\Big|_{\kappa=0}, \qquad B_{1,\kappa}(B_0^\star) \equiv \frac{\partial B_1(B_0^\star;\kappa)}{\partial\kappa}\Big|_{\kappa=0},$$

and, by the implicit function theorem applied to the period-1 optimality condition, using $B_1^\star - B_0^\star = \Upsilon(B_0^\star;0)\, B_0^\star$ from Equations (A25) and (A29),

$$B_{1,\kappa}(B_0^\star) = -\frac{\mathcal{S}_\kappa(B_1^\star)\,\Upsilon(B_0^\star;0)^2\, B_0^\star}{\phi f(\phi B_1^\star)}. \tag{A39}$$

Substituting Equation (A39) in (A38) yields the following necessary and sufficient local condition for $B_0(\kappa) > B_0^\star$, for a small $\kappa > 0$:

$$-\phi f\left(\phi B_1^\star\right)\frac{\Upsilon_\kappa\left(B_0^\star\right)}{\Upsilon\left(B_0^\star;0\right)} > \mathcal{S}_\kappa(B_1^\star)\,\Upsilon\left(B_0^\star;0\right), \tag{A40}$$

which is equivalent to Equation (14) in the main text. □

Proof of Theorem 1. We start from Equation (A40) and we compute $\Upsilon_\kappa\left(B_0^\star\right)$. First, from the implicit function theorem, we have that

$$\Upsilon\left(B_0^\star;\kappa\right) \equiv \frac{\partial B_1\left(B_0;\kappa\right)}{\partial B_0}|_{B_0^\star} = -\frac{\Psi_{1,B_0}(B_1;B_0^\star,\kappa)}{\Psi_{1,B_1}(B_1;B_0^\star,\kappa)}.$$

We work with a first-order expansion in κ around $\kappa = 0$ with $\overline{\mathcal{B}} = B_1^\star$. Since the period-1 policy is differentiable in κ at $(B_0^\star, 0)$, we have $B_1(B_0^\star;\kappa) = B_1^\star + \kappa\, B_{1,\kappa}(B_0^\star) + o(\kappa)$. Since $\overline{\mathcal{B}} = B_1^\star$, multiplying by κ yields $\kappa\Big(B_1(B_0^\star;\kappa) - \overline{\mathcal{B}}\Big) = O(\kappa^2)$. Hence, terms proportional to $\kappa(B_1 - \overline{\mathcal{B}})$ are of second order (or higher) in κ and therefore vanish when computing the derivative with respect to κ at $\kappa = 0$. From here, one can show that

$$\Upsilon\left(B_0^\star;\kappa\right) = \frac{\phi f\left(\phi\mathring{B}_1\right) + \mathcal{S}_\kappa\left(\mathring{B}_1\right)\kappa}{(2-\beta)\,\phi f\left(\phi\mathring{B}_1\right) + \phi^2 f_y\left(\phi\mathring{B}_1\right)\left(\mathring{B}_1 - B_0^\star\right) + 2\kappa\,\mathcal{S}_\kappa\left(\mathring{B}_1\right) + 2\kappa\,\mathcal{S}_{\kappa,B_1}\left(\mathring{B}_1\right)\left(\mathring{B}_1 - B_0^\star\right)}$$
$$+\, o\left(\kappa\right).$$

From this expression, we can now compute $\Upsilon_\kappa\left(B_0^\star\right)$. Differentiating at $\kappa = 0$ yields

$$\Upsilon_\kappa\left(B_0^\star\right) = \frac{\left[\Upsilon\left(B_0^\star;0\right)\right]^2}{\phi f\left(\phi B_1^\star\right)}\mathcal{S}_\kappa(B_1^\star)\left(-\beta + \Theta\frac{\left[\Upsilon\left(B_0^\star;0\right)\right]^2 B_0^\star}{\phi f\left(\phi B_1^\star\right)} - 2\frac{\mathcal{S}_{\kappa,B_1}(B_1^\star)}{\mathcal{S}_\kappa(B_1^\star)}\left(B_1^\star - B_0^\star\right)\right), \tag{A41}$$

where $\Theta \equiv (3-\beta)\,\phi^2 f_y\left(\phi B_1^\star\right) + \phi^3 f_{yy}\left(\phi B_1^\star\right)\left(B_1^\star - B_0^\star\right)$.

After replacing Equation (A41) in Equation (A40), and using Equations (11)-(12) to replace $B_1^\star - B_0^\star$ and $B_0^\star$, one obtains the following necessary and sufficient condition for B_0 to increase with κ:

$$\frac{2}{\phi}\frac{\mathcal{S}_{\kappa,B_1}(B_1^\star)}{\mathcal{S}_\kappa(B_1^\star)}H\left(\phi B_1^\star\right) > 1 + \Theta\frac{\Upsilon\left(B_0^\star;0\right)H\left(\phi B_1^\star\right)}{\phi^2 f\left(\phi B_1^\star\right)}. \tag{A42}$$

Rewriting f_y and f_{yy} in Θ in terms of the inverse hazard and its derivatives, and collecting terms, we obtain the condition in the main text, written in terms of the semi-elasticity $\varepsilon_{\mathcal{S}}(B_1^\star) = \mathcal{S}_{\kappa,B_1}(B_1^\star)/\mathcal{S}_\kappa(B_1^\star)$:

$$2\left(1 + H_y(\phi B_1^\star)\right) \;+\; \Omega\, H_{yy}(\phi B_1^\star) \;+\; \frac{2H(\phi B_1^\star)}{\phi}\varepsilon_{\mathcal{S}}(B_1^\star) \;>\; 1, \tag{A43}$$

where $\Omega \equiv (1-\beta)\, H(\phi B_1^\star)\, \Upsilon(B_0^\star;0) > 0$.

Uniform example. Consider the case in which $\mathcal{S}(B_1;\kappa) = \kappa\,\mathbb{V}(\mathcal{R}(B_1,y_2))$. Since $\mathcal{R}(B_1,y_2) \sim$ Bernoulli $(1-F(\phi B_1))$, $\mathcal{S}_\kappa(B_1^\star) = F\left(\phi B_1^\star\right)\left(1-F\left(\phi B_1^\star\right)\right)$ and $\mathcal{S}_{\kappa,B_1}(B_1^\star) = \phi f\left(\phi B_1^\star\right)\left[1-2F\left(\phi B_1^\star\right)\right]$. Suppose further that $y_2 \sim U[0,1]$. Since $H_y(y) = -1$ and $H_{yy}(y) = 0$, the condition reduces to $F\left(\phi B_1^\star\right) < 0.40$. After replacing with the $B_1^\star$ solution, $B_1^\star = \frac{(1-\beta)(3-\beta)}{(2-\beta)^2\phi}$, the previous condition is equivalent to $\beta > 0.71$. That is, the government must be sufficiently patient for the dilution-curbing force to dominate. □

The role of maturity

We extend our baseline model to consider a more general maturity structure in which a fraction $\lambda \in [0,1)$ of bonds issued at $t=0$ matures at $t=1$, and the remainder matures at $t=2$. This nests the baseline case $\lambda = 0$. The knife-edge case $\lambda = 1$ corresponds to purely short-term period-0 debt, with no remaining long-term claim to dilute, and is therefore not covered by the local maturity condition below. Consumption levels are $c_0 = q_0(B_0;\kappa)B_0 \geq 0$ and $c_1 = q_1(B_1;\kappa)\Big(B_1 - (1-\lambda)B_0\Big) - \lambda B_0 \geq 0$. The following corollary characterizes when B_0 increases as we steepen the price schedule.[4]

Corollary A.1. *Consider the same assumptions as in Theorem 1, extended to allow for $\lambda \in [0,1)$. Set $\overline{\mathcal{B}} = B_1^\star$. Suppose the $\kappa = 0$ allocation satisfies $c_0 > 0$ and $c_1 > 0$. For sufficiently small $\kappa > 0$, B_0 increases relative to the $\kappa = 0$ case if and only if*

$$2\left(1 + H_y(\phi B_1^\star)\right) + \Omega H_{yy}(\phi B_1^\star) + \frac{2H(\phi B_1^\star)}{\phi}\varepsilon_{\mathcal{S}}(B_1^\star) > 1 + \frac{\lambda}{(1-\lambda)\left(1-F(\phi B_1^\star)\right)+\lambda}, \tag{A44}$$

where $\Omega \equiv \frac{(1-\beta)\,H\left(\phi B_1^\star\right)}{1-(1-\beta)H_y\left(\phi B_1^\star\right)} > 0$.

The left-hand side has the same form as Condition (15) in the main text. The additional

[4]Unlike the baseline model, non-degenerate default risk does not necessarily imply $c_1 > 0$ when λ is large, which is why we impose the assumption explicitly.

term on the right-hand side is nonnegative and, holding the default cutoff fixed, increasing in λ. Shorter maturity therefore directly makes the condition harder to satisfy. Intuitively, when more debt matures at $t = 1$, the government has less scope to dilute long-term creditors, which reduces the value of market discipline as a deterrent. As a result, a steeper demand schedule is less effective at expanding debt capacity.

Proof of Corollary A.1. The proof follows the same local expansion as in the proof of Theorem 1, so we briefly describe the steps that differ. In this maturity extension, the period-0 price is $q_0 = \lambda + (1-\lambda)q_1$ and the period-1 budget constraint nets out the principal maturing at $t = 1$. The period-1 first-order condition implies $\Delta B_\lambda = \frac{1-\beta}{\phi} H(\phi B_1)$, as in Equation (A25), where $\Delta B_\lambda \equiv B_1 - (1-\lambda)B_0$. The analogue of the debt pass-through in Equation (A27) is $\Upsilon(B_0; 0) = \frac{1-\lambda}{1-(1-\beta)H_y(\phi B_1)}$. Lastly, at $\kappa = 0$, the analogue of Equation (A29), which pins down $B_0^\star$, is:

$$\phi \Upsilon(B_0^\star; 0) B_0^\star = (1-\beta)\left[H(\phi B_1^\star) + \frac{\lambda}{(1-\lambda) f(\phi B_1^\star)}\right].$$

The one substantive difference with respect to our baseline model derivation is that, for $\lambda > 0$, $\Delta B_\lambda \neq \Upsilon(B_0^\star; 0) B_0^\star$. Thus, the baseline substitution between ΔB_λ and the pass-through term, used after Equations (A40)–(A41), is no longer possible. Carrying the two terms separately introduces the coefficient $(1-\beta)\phi\left[f H_y - \frac{f_y}{f}\frac{\lambda}{1-\lambda}\right]$ on $B_{1,\kappa}$ in the analogue of Equation (A38). Substituting the expressions for $B_{1,\kappa}$ and Υ_κ implied by the maturity extension into its local sign condition (the counterpart of Equation (A40)), the inverse-hazard terms simplify exactly as in the proof of Theorem 1, since ΔB_λ retains the form $\frac{1-\beta}{\phi} H(\phi B_1^\star)$. Collecting terms yields Condition (A44). □

A generalized local condition when $\overline{\mathcal{B}} \neq B_1^\star$

We allow the reference demand $\overline{\mathcal{B}}$ to be arbitrary. Relative to the normalization $\overline{\mathcal{B}} = B_1^\star$ in the main text, changing κ affects bond prices through two channels: the *slope* of the price schedule and the price *level*. Corollary A.2 characterizes the necessary and sufficient condition for B_0 to increase with κ.

Corollary A.2. *Consider the assumptions of Theorem 1. For sufficiently small $\kappa > 0$ and any reference demand $\overline{\mathcal{B}}$, period-0 debt increases relative to the $\kappa = 0$ case if and only if*

$$2\left(1 + H_y(\phi B_1^\star)\right) + \Omega H_{yy}(\phi B_1^\star) + \frac{2H(\phi B_1^\star)}{\phi}\varepsilon_{\mathcal{S}}(B_1^\star) \; > \; 1 + \Xi\left(B_1^\star - \overline{\mathcal{B}}\right), \tag{A45}$$

where

$$\Xi \equiv -\phi\,\frac{H_y(\phi B_1^\star)}{H(\phi B_1^\star)} - 2\left(\tfrac{1}{2} + H_y(\phi B_1^\star)\right)\varepsilon_{\mathcal{S}}(B_1^\star) - \frac{H(\phi B_1^\star)}{\phi}\,\frac{\mathcal{S}_{\kappa,B_1,B_1}(B_1^\star)}{\mathcal{S}_\kappa(B_1^\star)}$$
$$- H_{yy}(\phi B_1^\star)\,\Upsilon(B_0^\star;0)\left(\phi + (1-\beta)H(\phi B_1^\star)\,\varepsilon_{\mathcal{S}}(B_1^\star)\right),$$

with $\mathcal{S}_{\kappa,B_1,B_1}(B_1) \equiv \frac{\partial^2 \mathcal{S}_\kappa(B_1)}{\partial B_1^2}$ *and* $\Omega \equiv \frac{(1-\beta)\,H\left(\phi B_1^\star\right)}{1-(1-\beta)H_y\left(\phi B_1^\star\right)} > 0.$

The left-hand side is identical to that of our baseline Condition (15). For $\overline{\mathcal{B}} < B_1^\star$, Condition (A45) is harder to satisfy than the baseline condition, and becomes increasingly restrictive as $\overline{\mathcal{B}}$ falls, if and only if $\Xi > 0$. For the uniform example, $\Xi > 0$ in the region where the baseline condition holds, $F(\phi B_1^\star) < 0.4$. There, $H_y(\phi B_1^\star) = -1$, $H_{yy}(\phi B_1^\star) = 0$, and $\mathcal{S}_\kappa(B_1) = \mathbb{V}(\mathcal{R}(B_1, y_2)) = F(\phi B_1)\,[1 - F(\phi B_1)]$. Since $F(\phi B_1^\star) < 0.4$, $\mathcal{S}_\kappa(B_1)$ is thus increasing and concave around $B_1^\star$, so $\mathcal{S}_{\kappa,B_1}(B_1^\star) \geq 0$ and $\mathcal{S}_{\kappa,B_1,B_1}(B_1^\star) \leq 0$.

Proof of Corollary A.2. The proof follows the same steps as the proof of Theorem 1, so we briefly highlight the differences. The main difference from the baseline derivation is that, when $\overline{\mathcal{B}} \neq B_1^\star$, the demand term $\mathcal{S}_\kappa(B_1)(B_1 - \overline{\mathcal{B}})$ does not vanish at $B_1^\star$, so terms proportional to $B_1^\star - \overline{\mathcal{B}}$ survive at first order. Relative to Equations (A39) and (A38), the period-1 response $B_{1,\kappa}$ and the derivative $\Psi_{0,\kappa}$ therefore contain additional terms in $B_1^\star - \overline{\mathcal{B}}$. The pass-through response Υ_κ also acquires a term involving the curvature of the demand slope, $\mathcal{S}_{\kappa,B_1,B_1}$, relative to Equation (A41). This term is absent from the baseline because it is multiplied by $B_1^\star - \overline{\mathcal{B}}$. Carrying these additional terms through the local sign condition $\Psi_{0,\kappa}(B_0^\star;0) > 0$ and collecting terms yields Condition (A45), with Ξ summarizing their combined coefficient. □

Diminishing dilution-curbing force

The local characterizations above are derived around the perfectly elastic allocation. Lemma 1 considers a first-order approximation around a solution for the $\kappa > 0$ economy and shows that, within this approximation, the dilution-curbing force is decreasing in κ. Thus, the local expansion result in Proposition 3 and Theorem 1 need not extend to larger κ.

Fix some $\kappa > 0$ and let $\hat{B} = \{\hat{B}_0, \hat{B}_1\}$ denote a solution with $\hat{B}_0 > 0$, $\hat{B}_1 > \hat{B}_0$, and non-degenerate default risk, where $\hat{B}_1 \equiv B_1(\hat{B}_0;\kappa)$. We approximate the period-1 optimality condition around this solution, holding the local default-risk and demand-price slopes at $\hat{B}_1$ fixed as κ varies, while still allowing period-1 debt to respond marginally to inherited debt.

Lemma 1 (Diminishing dilution-curbing force)**.** *Consider a first-order approximation to the period-1 price schedule around $\hat{B}_1$, with the local default-risk and demand-price slopes evaluated at $\hat{B}_1$, while allowing B_1 to respond marginally to B_0. Let $\widetilde{B}_1(B_0;\kappa)$ denote the resulting period-1 policy and define the associated debt pass-through as*

$$\widetilde{\Upsilon}(\kappa) \equiv \frac{\partial \widetilde{B}_1(B_0;\kappa)}{\partial B_0}\bigg|_{B_0=\hat{B}_0}.$$

Assume that $1+\varepsilon_{\mathcal{S}}(\hat{B}_1)(\hat{B}_1-\overline{\mathcal{B}})>0$, so that the demand-price adjustment steepens the price schedule locally. Then,

$$-\phi f(\phi\hat{B}_1)\frac{\widetilde{\Upsilon}_\kappa(\kappa)}{\widetilde{\Upsilon}(\kappa)}>0 \quad \textit{and} \quad \frac{\partial}{\partial\kappa}\left[-\phi f(\phi\hat{B}_1)\frac{\widetilde{\Upsilon}_\kappa(\kappa)}{\widetilde{\Upsilon}(\kappa)}\right]<0,$$

where $\widetilde{\Upsilon}_\kappa(\kappa)\equiv\partial\widetilde{\Upsilon}(\kappa)/\partial\kappa$.

In this approximation, a higher κ lowers the debt pass-through, but each additional increase reduces it by less. The associated dilution-curbing force, $-\phi f(\phi\hat{B}_1)\,\widetilde{\Upsilon}_\kappa(\kappa)/\widetilde{\Upsilon}(\kappa)$, therefore declines with κ. This gives a local reason why period-0 debt expansion may become less likely at larger κ, a possibility we evaluate numerically in the full nonlinear model.

Proof of Lemma 1. Maintain the solution $\hat{B}=\{\hat{B}_0,\hat{B}_1\}$ defined above. Around $\hat{B}_1$, define

$$\mathcal{D}(\hat{B}_1)\equiv\mathcal{S}_\kappa(\hat{B}_1)+\mathcal{S}_{\kappa,B_1}(\hat{B}_1)(\hat{B}_1-\overline{\mathcal{B}})=\mathcal{S}_\kappa(\hat{B}_1)\left[1+\varepsilon_{\mathcal{S}}(\hat{B}_1)(\hat{B}_1-\overline{\mathcal{B}})\right].$$

This object is the derivative of $\mathcal{S}_\kappa(B_1)(B_1-\overline{\mathcal{B}})$ with respect to B_1, evaluated at $\hat{B}_1$. The assumption in the lemma, together with $\mathcal{S}_\kappa(\hat{B}_1)>0$, implies $\mathcal{D}(\hat{B}_1)>0$.

From the analysis around Equation (A17), the bond price is $q_1(B_1;\kappa)=1-F(\phi B_1)-\kappa\mathcal{S}_\kappa(B_1)(B_1-\overline{\mathcal{B}})$. In what follows, we use a first-order Taylor approximation around $\hat{B}_1$ for the objects that enter the period-1 first-order condition in Equation (A19). For the price schedule, this gives

$$\widetilde{q}_1(B_1;\kappa)=1-F(\phi\hat{B}_1)-\phi f(\phi\hat{B}_1)(B_1-\hat{B}_1)-\kappa\mathcal{S}_\kappa(\hat{B}_1)(\hat{B}_1-\overline{\mathcal{B}})-\kappa\mathcal{D}(\hat{B}_1)(B_1-\hat{B}_1). \quad \text{(A46)}$$

Thus,

$$\widetilde{q}_{1,B_1}(B_1;\kappa)=-\phi f(\phi\hat{B}_1)-\kappa\mathcal{D}(\hat{B}_1). \quad \text{(A47)}$$

Because $\hat{B}_1>\hat{B}_0$, the issuance constraint is locally slack and $\mu_1=0$. Substituting the approximated objects into Equation (A19), we obtain

$$\begin{aligned}\widetilde{\Psi}_1(B_1;B_0,\kappa)=&(1-\beta)\left[1-F(\phi\hat{B}_1)-\phi f(\phi\hat{B}_1)(B_1-\hat{B}_1)\right]-\phi f(\phi\hat{B}_1)(B_1-B_0)\\ &-\kappa\left[\mathcal{S}_\kappa(\hat{B}_1)(\hat{B}_1-\overline{\mathcal{B}})+\mathcal{D}(\hat{B}_1)(B_1-\hat{B}_1)\right]-\kappa\mathcal{D}(\hat{B}_1)(B_1-B_0)=0.\end{aligned}$$

This approximation holds fixed the default-risk slope $\phi f(\phi\hat{B}_1)$ and the demand-price derivative $\mathcal{D}(\hat{B}_1)$ at the point of approximation, letting the repayment probability and the price vary linearly. It therefore omits higher-order curvature terms, such as $f_y(\phi\hat{B}_1)$ and $\mathcal{S}_{\kappa,B_1,B_1}(\hat{B}_1)$.

By the implicit function theorem, the pass-through in the approximated problem is

$$\widetilde{\Upsilon}(\kappa) = -\frac{\widetilde{\Psi}_{1,B_0}}{\widetilde{\Psi}_{1,B_1}},$$

where the partial derivatives are:

$$\widetilde{\Psi}_{1,B_0} = \phi f(\phi\hat{B}_1) + \kappa\mathcal{D}(\hat{B}_1) \quad \text{and} \quad -\widetilde{\Psi}_{1,B_1} = (2-\beta)\phi f(\phi\hat{B}_1) + 2\kappa\mathcal{D}(\hat{B}_1).$$

Therefore,

$$\widetilde{\Upsilon}(\kappa) = \frac{\phi f(\phi\hat{B}_1) + \kappa\mathcal{D}(\hat{B}_1)}{(2-\beta)\phi f(\phi\hat{B}_1) + 2\kappa\mathcal{D}(\hat{B}_1)}.$$

Differentiating this object with respect to κ, holding the approximation point and the local slopes fixed,

$$\begin{aligned}\widetilde{\Upsilon}_\kappa(\kappa) =& \frac{\mathcal{D}(\hat{B}_1)\left[(2-\beta)\phi f(\phi\hat{B}_1) + 2\kappa\mathcal{D}(\hat{B}_1)\right] - 2\mathcal{D}(\hat{B}_1)\left[\phi f(\phi\hat{B}_1) + \kappa\mathcal{D}(\hat{B}_1)\right]}{\left[(2-\beta)\phi f(\phi\hat{B}_1) + 2\kappa\mathcal{D}(\hat{B}_1)\right]^2} \\ =& -\frac{\beta\phi f(\phi\hat{B}_1)\mathcal{D}(\hat{B}_1)}{\left[(2-\beta)\phi f(\phi\hat{B}_1) + 2\kappa\mathcal{D}(\hat{B}_1)\right]^2} < 0.\end{aligned}$$

Thus, a higher κ lowers the pass-through in the approximated problem. The associated dilution-curbing force is

$$-\phi f(\phi\hat{B}_1)\frac{\widetilde{\Upsilon}_\kappa(\kappa)}{\widetilde{\Upsilon}(\kappa)} = \frac{\beta\left[\phi f(\phi\hat{B}_1)\right]^2\mathcal{D}(\hat{B}_1)}{\left[(2-\beta)\phi f(\phi\hat{B}_1) + 2\kappa\mathcal{D}(\hat{B}_1)\right]\left[\phi f(\phi\hat{B}_1) + \kappa\mathcal{D}(\hat{B}_1)\right]} > 0.$$

Moreover, since both terms in the denominator are positive and increasing in κ, while the numerator is fixed, we get

$$\frac{\partial}{\partial\kappa}\left[-\phi f(\phi\hat{B}_1)\frac{\widetilde{\Upsilon}_\kappa(\kappa)}{\widetilde{\Upsilon}(\kappa)}\right] < 0.$$

Thus, additional increases in κ continue to reduce the debt pass-through, but the associated marginal dilution-curbing force is decreasing as the schedule becomes steeper.

□

B Empirical Appendix

B.1 Additional information on the EMBIGD and the FIR

EMBIGD. Figure B1 shows the assets under management (AUM) benchmarked to the EMBIGD and other major benchmark indexes. Among emerging-market bond indexes, the EMBIGD is the most widely used and, as of 2018, served as a benchmark for funds managing approximately US$300 billion. To put this in perspective, the share of U.S. dollar-denominated sovereign debt included in the EMBIGD relative to each country's total general government debt issued in international markets exceeds 60% for more than half the countries in our sample—highlighting the index's central role in external sovereign debt markets.

Appendix Figure B1: AUM benchmarked to emerging market bond indexes

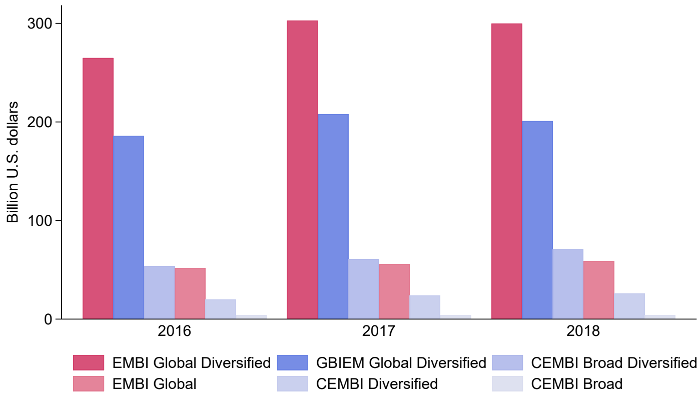

Note: The figure shows AUM (in billions of U.S. dollars) benchmarked to major J.P. Morgan emerging-market bond indexes during 2016–2018. The EMBI Global covers U.S. dollar-denominated bonds issued by emerging-market sovereign and quasi-sovereign entities. The EMBI Global Diversified (EMBIGD) covers the same eligible securities but applies diversification caps to countries with larger debt stocks. The GBIEM Global Diversified tracks local-currency government bonds. The CEMBI indexes cover U.S. dollar-denominated corporate bonds; "Broad" versions apply wider eligibility criteria, while "Diversified" versions cap country weights.

Diversification Methodology. Relative to a market-capitalization-weighted index, the EMBIGD employs a diversification methodology designed to achieve a more balanced distribution of country weights: it limits the index weight of countries with above-average debt by including only a portion of their outstanding bonds.

The methodology centers on the average country-level face value of bonds included in the index, referred to as the Index Country Average (ICA). It is defined as:

$$ICA_t \equiv \frac{1}{|\mathcal{I}|} \sum_{c \in \mathcal{I}} B_{c,t}, \tag{B1}$$

where $B_{c,t}$ denotes the total face value of bonds from country c included in the index at time t, and $|\mathcal{I}|$ is the number of countries represented in the index. Using this average, the *Diversified Face Amount* for each country, denoted $DFA_{c,t} \equiv f_{c,t}B_{c,t}$, is defined as:

$$f_{c,t}B_{c,t} = \begin{cases} ICA_t + \frac{ICA_t}{\max_{\forall c \in \mathcal{I}}\{B_{c,t}\} - ICA_t}\left(B_{c,t} - ICA_t\right) & \text{if } B_{c,t} > ICA_t \\ B_{c,t} & \text{if } B_{c,t} \leq ICA_t. \end{cases} \tag{B2}$$

This formula implies that $f_{c,t} = 1$ for countries whose face amount is below the index average. For the largest issuer, $f_{c,t} < 1$ precisely when $\max_{\forall c \in \mathcal{I}}\{B_{c,t}\} > 2 \times ICA_t$. This inequality holds in our sample period, reflecting the presence of very large issuers (e.g., China and Mexico) relative to the cross-country average. As a result, all countries with $B_{c,t} > ICA_t$ are mechanically scaled down by a factor $f_{c,t} < 1$, implying a proportional reduction applied uniformly across all of their bonds in the index.[5] The index's country weights are then computed as $\frac{q_{c,t}B_{c,t}f_{c,t}}{\sum_{n \in \mathcal{I}} q_{n,t}B_{n,t}f_{n,t}}$.

FIR. Panel (a) of Figure B2 shows a scatter plot of our FIR measure against the Z instrument, after both variables have been residualized using rebalancing-month and country fixed effects. The two variables exhibit a strong positive relation, with an R^2 of around 80%. Panel (b) displays the distribution of the instrumented FIR, which ranges from -0.7% to 0.25% and is skewed toward negative values. Since we focus on countries with a constant face amount and inclusions from other countries predominate over exclusions, their index weights generally fall.[6]

B.2 Additional Results and Robustness

Alternative specifications and event windows. We present a set of robustness checks that complement our main analysis. Table B1 reports estimates from an OLS regression in which the FIR is not instrumented. Our findings are also robust to alternative event windows around rebalancing dates (Table B2).

Yield to maturity. One can map the $\hat{\delta}$ estimates from Table 3 into a semi-elasticity of bond yields to FIR. Using $\frac{\Delta Y_t}{FIR_t} \approx -\frac{\hat{\delta}}{D_t}$ with the average sample duration gives approximately 5 basis points. Table B3 re-estimates the baseline specification using yield

[5]In addition to this rule, country weights are capped at 10%. Any excess weight above this cap is redistributed pro rata to smaller countries below the cap, across all bonds from countries not capped at 10%.

[6]For our main analysis, we apply the following data cleaning procedures. We exclude extreme values of daily returns, stripped spreads, and the instrument $Z_{c,t}$. Specifically, we drop observations with stripped spreads below 0 or above 5,000 basis points, as well as those below the 5th or above the 95th percentile of the $Z_{c,t}$ distribution. The rationale is that extreme values of $Z_{c,t}$ may reflect large, pre-announced changes to the EMBIGD and are therefore not suitable for our identification strategy, which assumes that most information becomes available on the last business day of the month. Lastly, we exclude bond-month observations with daily returns below the 1st or above the 99th percentile.

Appendix Figure B2: Flows induced by rebalancing (FIR)

(a) Relation between FIR and Z (b) Distribution of instrumented FIR values

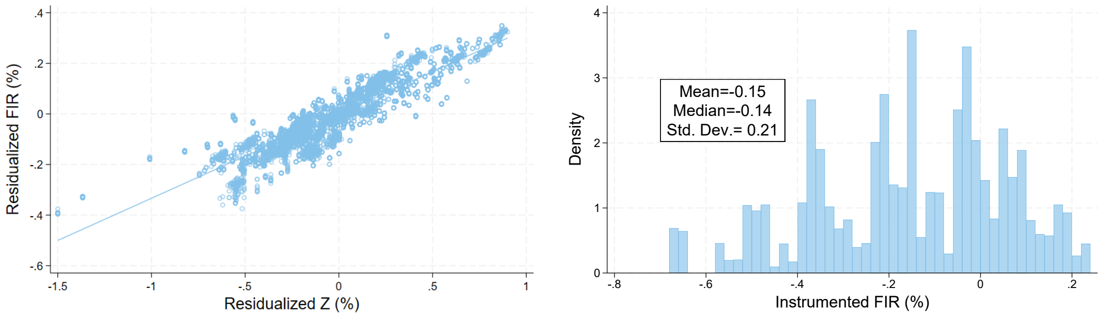

Notes: Panel (a) presents a scatter plot of the FIR and the Z instrument. Both variables are residualized based on a regression with rebalancing-month and country fixed effects. Panel (b) shows a histogram of the instrumented FIR. The sample period is 2016–2018.

Appendix Table B1: OLS regression

Dependent Variable: Log Price	$[-5, 5]$				Excl. $h = -1$	
FIR × Post	0.197	0.199	0.198	0.134	0.237	0.168
	(0.076)	(0.077)	(0.076)	(0.087)	(0.077)	(0.093)
Bond FE	Yes	Yes	No	No	No	No
Month FE	Yes	No	No	No	No	No
Bond Characteristics × Month FE	No	Yes	No	No	No	No
Country × Month FE	No	Yes	No	No	No	No
Bond × Month FE	No	No	Yes	Yes	Yes	Yes
Month × Post FE	No	No	No	Yes	No	Yes
Bond Controls	No	Yes	No	No	No	No
Observations	105,548	105,508	105,548	105,548	84,433	84,433
N. of Bonds	738	738	738	738	738	738
N. of Countries	68	68	68	68	68	68
N. of Clusters	1,576	1,575	1,576	1,576	1,576	1,576

Note: The table presents OLS estimates of log bond prices on FIR interacted with a post-rebalancing indicator around rebalancing events. Standard errors are clustered at the country-month level, and the sample covers the 2016–2018 period. See the note in Table 3 for additional details.

to maturity as the dependent variable. The results closely align with the back-of-the-envelope calculation above. Lastly, Table B4 shows that yield responses are also larger for countries with higher spreads.

Bond maturity. Table B5 examines heterogeneity in price responses across three maturity buckets—short (below 4.8 years), medium, and long (above 8.8 years)—based on the 33rd and 66th percentiles of the maturity distribution. Price responses to FIR increase sharply with maturity: short-term bonds show a modest, statistically insignificant effect, while long-term bonds exhibit a larger response of around 0.41%. These results are consistent with Broner et al. (2013), who show that short-term bonds are more tightly

Appendix Table B2: Log price and FIR: Different event windows

Panel A. Dependent Variable: Log Price	$[-2,2]$	$[-3,3]$	$[-4,4]$	$[-5,5]$
FIR × Post	0.147	0.198	0.222	0.232
	(0.053)	(0.072)	(0.086)	(0.099)
Bond × Month FE	Yes	Yes	Yes	Yes
Observations	42,217	63,327	84,435	105,548
N. of Bonds	738	738	738	738
N. of Countries	68	68	68	68
N. of Clusters	1,576	1,576	1,576	1,576
SW F (FIR × Post)	1,661	1,663	1,665	1,667

Panel B. Dependent Variable: Log Price (Excluding $h=-1$)	$[-2,1]$	$[-3,2]$	$[-4,3]$	$[-5,4]$
FIR × Post	0.221	0.257	0.272	0.264
	(0.056)	(0.075)	(0.088)	(0.098)
Bond × Month FE	Yes	Yes	Yes	Yes
Observations	21,106	42,216	63,325	84,433
N. of Bonds	738	738	738	738
N. of Countries	68	68	68	68
N. of Clusters	1,576	1,576	1,576	1,576
SW F (FIR × Post)	1,668	1,668	1,670	1,672

Note: The table presents 2SLS estimates of log bond prices on FIR interacted with a post-rebalancing indicator around rebalancing events. Each column reports estimates for a different symmetric h-day window. The sample period and estimation procedure follow those in Table 3. Standard errors are clustered at the country-month level.

Appendix Table B3: Effects of FIR on yield to maturity

Dependent Variable: Yield to Maturity	$[-5,5]$				Excl. $h=-1$	
FIR × Post	-3.523	-3.554	-3.556	-5.505	-4.049	-5.790
	(1.677)	(1.699)	(1.674)	(2.613)	(1.673)	(2.598)
Bond FE	Yes	Yes	No	No	No	No
Month FE	Yes	No	No	No	No	No
Bond Characteristics × Month FE	No	Yes	No	No	No	No
Country × Month FE	No	Yes	No	No	No	No
Bond × Month FE	No	No	Yes	Yes	Yes	Yes
Month × Post FE	No	No	No	Yes	No	Yes
Bond Controls	No	Yes	No	No	No	No
Observations	103,726	103,686	103,726	103,726	82,976	82,976
N. of Bonds	717	717	717	717	717	717
N. of Countries	66	66	66	66	66	66
N. of Clusters	1,546	1,545	1,546	1,546	1,546	1,546
SW F (FIR × Post)	1,895	1,641	1,689	470	1,691	470

Note: The table presents 2SLS estimates of yield to maturity (in basis points) on FIR interacted with a post-rebalancing indicator around rebalancing events. The sample period and estimation procedure follow those in Table 3. Standard errors are clustered at the country-month level.

anchored to policy rates and less sensitive to broader shifts in investor flows.

Credit default swaps. Table B6 asks whether the rebalancing events affect credit risk by regressing 5-year log CDS spreads on the FIR, using the same 2SLS setup and five-day window as in the main text. An increase in the FIR is associated with a decline in CDS spreads (significant at the 10% level), suggesting that benchmark-tracking shifts affect not only bond prices but also credit risk.

Appendix Table B4: Effects of FIR on yield to maturity: The role of default risk

Dependent Variable: Yield to Maturity	High Spread		Medium Spread		Low Spread	
FIR × Post	-6.401	-6.427	-3.096	-3.127	-2.242	-2.242
	(3.276)	(3.252)	(2.119)	(2.102)	(1.573)	(1.561)
Bond FE	Yes	No	Yes	No	Yes	No
Month FE	Yes	No	Yes	No	Yes	No
Bond × Month FE	No	Yes	No	Yes	No	Yes
Observations	27,697	27,696	27,520	27,520	27,760	27,760
N. of Bonds	370	370	445	445	366	366
N. of Countries	57	57	51	51	43	43
N. of Clusters	954	954	834	834	626	626
SW F (FIR × Post)	2,320	2,355	830	732	866	879

Note: The table presents 2SLS estimates of yield to maturity (in basis points) on FIR interacted with a post-rebalancing indicator around rebalancing events. The sample is split into high-spread bonds (above the 66th percentile of stripped spreads), medium-spread bonds, and low-spread bonds (below the 33rd percentile). The sample period and estimation procedure follow those in Table 3, excluding the trading day before each rebalancing. Standard errors are clustered at the country-month level.

Appendix Table B5: Effects of FIR on bond prices across maturities

Dependent Variable: Log Price	Long Maturity		Medium Maturity		Short Maturity	
FIR × Post	0.413	0.414	0.304	0.304	0.095	0.094
	(0.199)	(0.197)	(0.094)	(0.093)	(0.061)	(0.061)
Bond FE	Yes	No	Yes	No	Yes	No
Month FE	Yes	No	Yes	No	Yes	No
Bond × Month FE	No	Yes	No	Yes	No	Yes
Observations	28,199	28,198	28,145	28,144	28,092	28,091
N. of Bonds	333	333	295	295	324	324
N. of Countries	55	55	65	65	57	57
N. of Clusters	937	937	1,364	1,364	1,135	1,135
SW F (FIR × Post)	1,153	990	2,582	2,571	1,300	1,275

Note: The table presents 2SLS estimates of log bond prices on FIR interacted with a post-rebalancing indicator around rebalancing events. The sample is divided into three maturity buckets (based on the 33rd and 66th percentiles): "short" (< 4.8 years), "medium", and "long" (> 8.8 years). The sample period and the 2SLS procedure are identical to those described in Table 3. The estimation excludes the trading day before each rebalancing day. Standard errors are clustered at the country-month level.

Appendix Table B6: Log CDS and FIR

Dependent variable: Log CDS			
FIR × Post	-0.799	-0.799	-0.799
	(0.461)	(0.461)	(0.460)
Country FE	Yes	Yes	No
Month FE	No	Yes	No
Country × Month FE	No	No	Yes
Observations	10,160	10,160	10,160
N. of Countries	44	44	44
N. of Clusters	1,016	1,016	1,016

Note: The table presents 2SLS estimates of log 5-year CDS spreads on FIR interacted with a post-rebalancing indicator around rebalancing events. Standard errors are clustered at the country-month level, and the sample period is 2016–2018.

C Quantitative Appendix

In this appendix, we provide additional details for our quantitative analysis in Section 4.

C.1 Secondary Markets

Our reduced-form elasticity (Section 3) exploits high-frequency variation in passive demand around EMBIGD rebalancing announcements. The baseline model already includes passive demand shares, τ, but is calibrated at quarterly frequency, so it cannot directly speak to high-frequency price responses to shifts in passive demand. To bridge this gap, we add secondary markets that capture the high-frequency nature of our empirical elasticity. Within each period, trading occurs twice: before and after the realization of the new passive share, τ'.

The first trading instance (SM_0) occurs at the beginning of the period, immediately after the government announces its default and debt choices. The bond price is

$$q^{\mathrm{SM},0}(y,\tau,B') = \frac{1}{r^f}\mathbb{E}_{y',\tau'|y,\tau}\mathcal{R}\left(y',\tau',B'\right) - \frac{\kappa}{r^f}\mathbb{V}_{y',\tau'|y,\tau}\mathcal{R}\left(y',\tau',B'\right)\left(B' - \mathcal{B}(\tau,B')\right), \quad \text{(C1)}$$

where the first term is the expected next-period repayment of the bond, conditional on the information available when the secondary market opens; the second is a price adjustment proportional to residual supply, $B' - \mathcal{B}(\tau, B')$. Notice that $q^{\mathrm{SM},0}\left(y,\tau,B'\right)$ coincides with the price in the primary market $q\left(y,\tau,B'\right)$, which is the price relevant to the government.

The second trading instance (SM_1) occurs at the end of the period, when the secondary market closes and after τ' is realized. The price is

$$q^{\mathrm{SM},1}\left(y,\tau',B'\right) = \frac{1}{r^f}\mathbb{E}_{y'|y}\mathcal{R}\left(y',\tau',B'\right) - \frac{\kappa}{r^f}\mathbb{V}_{y'|y}\mathcal{R}\left(y',\tau',B'\right)\left(B' - \mathcal{B}(\tau',B')\right), \quad \text{(C2)}$$

where both terms now condition on the realized passive share τ'. The only difference between $q^{\mathrm{SM},1}$ and $q^{\mathrm{SM},0}$ arises from the update of τ', since both the endowment and the stock of debt remain fixed while the secondary market is open. Moreover, in the absence of secondary markets, the timing assumption is exactly the same as in the baseline model.

Using secondary-market responses to discipline the demand slope relevant for primary issuance requires newly issued and outstanding bonds to be close substitutes and issuance prices to be anchored to secondary-market valuations. This approximation is plausible in our setting because the EMBIGD comprises large, liquid hard-currency bonds, with a minimum face value of US$500 million and an average of about US$1.3 billion. Our sample period also largely falls outside episodes of widespread sovereign default or severe global market stress, and primary and secondary prices are therefore more likely to remain closely aligned. For instance, Cole et al. (2025) find primary–secondary yield spreads near zero for major European sovereigns outside crisis periods, while Bigio et al. (2023) document auction–secondary price gaps of only a few basis points for Spain.

C.2 Solution Method

We employ a global solution method to solve the model. We discretize the output process y and the process for the passive demand share τ using Tauchen's method. We select 35 gridpoints for y and 9 for τ. As for B, we construct a grid of 251 equally spaced points between $B_{\min} = 0$ and $B_{\max} = 1.2$. We ensure that $B_{\max}$ is sufficiently large so that it never binds in our simulations. The steps of the algorithm are as follows:

1. We start with initial guesses for the value functions $V^r(y, \tau, B)$ and $V^d(y, \tau)$, as well as for the bond price function $q(y, \tau, B')$.
2. Based on these guesses, we solve for the optimal bond policy $B'(y, \tau, B)$, as described in Equation (6). To this end, we use an optimizing routine based on Brent's method and employ cubic splines to interpolate the value functions and bond prices when evaluating off-grid points. Given $B'(y, \tau, B)$, we then update $V^r(y, \tau, B)$.
3. We compute the value function for the case in which the government defaults in the current period, $V^d(y, \tau)$, as given by Equation (7).
4. We solve for the government's optimal default choice, as shown in Equation (5). As is standard in the literature, we convexify the default decision to achieve convergence. We assume that in each period, the government's value function $V^d(y, \tau)$ is subject to an i.i.d. shock $\epsilon_v \sim \mathcal{N}(1, \sigma_v^2)$ so that the government defaults if $V^r(\cdot) < V^d(\cdot) \times \epsilon_v$. We choose σ_v^2 to be small enough ($\sigma_v^2 = 2.25 \times 10^{-6}$) so that the convexified solution does not significantly differ from the "true" solution of the model. Let $d(y, \tau, B)$ denote the optimal default choice.
5. Taking the policy functions $B'(y, \tau, B)$ and $d(y, \tau, B)$ as given, we update $q(y, \tau, B')$ according to Equation (9). We use cubic splines to evaluate the right-hand side of the pricing equation at $B'' \equiv B'(y', \tau', B')$.[7]
6. We iterate over the previous steps until convergence of the value functions and the bond price function.

[7] In the numerical algorithm, we restrict bond prices to the interval $[0, q^f]$, where $q^f \equiv \frac{\lambda+(1-\lambda)\nu}{r^f-(1-\lambda)}$ is the price of a risk-free bond with the same maturity structure.

C.3 Estimation of the Passive-Share Process

This appendix details how we discipline the persistence and volatility of the passive share, ρ_τ and σ_τ. We measure the passive share in the EMBIGD-included bond segment, which is the benchmark universe used in the empirical estimation. Let $B^i_{\mathcal{I},t}$ denote the face value of country i's bonds included in the EMBIGD. The country weight is

$$w^i_t = \frac{f^i_t\, q^i_t B^i_{\mathcal{I},t}}{MV_t}, \qquad MV_t = \sum_{c\in\mathcal{I}} f^c_t\, q^c_t B^c_{\mathcal{I},t}, \tag{C3}$$

where f^i_t is the country's diversification coefficient. Let $\mathcal{W}_t$ denote the assets under management of passive investors tracking the EMBIGD. Passive holdings of country i's EMBIGD bonds are $w^i_t\mathcal{W}_t/q^i_t$. We therefore measure the passive share as

$$\tau^i_t \equiv \frac{w^i_t\mathcal{W}_t}{q^i_t B^i_{\mathcal{I},t}}. \tag{C4}$$

This expression is the one we take to the data: the index weight, the assets under management of passive investors, the bond price, and the face value of EMBIGD-included bonds are all observable. Substituting Equation (C3) into Equation (C4) gives $\tau^i_t = f^i_t\mathcal{W}_t/MV_t$. Thus, changes in $B^i_{\mathcal{I},t}$ do not mechanically affect the measured passive share, except through the aggregate index value MV_t and the diversification coefficient f^i_t.

The quantitative model abstracts from index eligibility to avoid introducing an additional state variable, so we use this EMBIGD-segment passive share as the empirical counterpart of the model state τ. This mapping preserves the normalization relevant for the reduced-form elasticity. In both the data and the model, what matters is the passive flow relative to non-passive holdings. With $\mathcal{T}(\tau, B') = \tau B'$ and fixed B' around rebalancing,

$$\frac{\Delta\mathcal{T}}{B' - \mathcal{T}} = \frac{\tau' - \tau}{1-\tau},$$

so the scale of the bond segment cancels. The level of the passive share still matters, as it determines the non-passive share $1-\tau$. Because the empirical FIR is also normalized by non-passive holdings within the *same* EMBIGD bond segment, the model-implied elasticity calculation uses the same normalization. In the quantitative analysis of Section 4, the average level of reference demand is disciplined separately by $\overline{\mathcal{A}}$. Appendix C.7 shows that recalibrating the model with alternative levels of the average passive share, $\tau^\star$, adjusting $\overline{\mathcal{A}}$ so that the average inelasticity premium remains zero, leaves the main moments and the model-implied elasticities essentially unchanged.

Appendix Table C1: Persistence and volatility of the passive share

Parameter	Baseline	Monthly estimation	All Countries
ρ_τ	0.660	0.600	0.636
σ_τ	0.020	0.028	0.007

Note: The table reports estimates of the persistence and innovation volatility of the passive-share process, τ. The "Baseline" column estimates a quarterly AR(1) using end-of-quarter observations from Argentina's monthly passive-share series over 2016–2018; these are the values used in the calibration. "Monthly" estimates the AR(1) on the same Argentina series at the monthly frequency and converts the coefficients to quarterly frequency. "All Countries" estimates a quarterly AR(1) on country-level passive-share series with country fixed effects.

We compute τ_t^{ARG} for Argentina at a monthly frequency over 2016–2018. To match the quarterly frequency of the model, our baseline calibration keeps end-of-quarter observations and estimates an AR(1) process for the passive share. This yields $\rho_\tau = 0.66$ and $\sigma_\tau = 0.02$, the values used in the baseline calibration. Because the sample period is short, we complement these estimates in two ways, reported in Table C1. First, we estimate the AR(1) on the monthly (m) Argentina series and convert the coefficients to a quarterly frequency, using $\rho_\tau = \hat{\rho}_m^3$ and $\sigma_\tau^2 = \hat{\sigma}_m^2(1 + \hat{\rho}_m^2 + \hat{\rho}_m^4)$. Second, we construct τ_t^c for every country c in the index, keep end-of-quarter observations, and estimate an AR(1) with country fixed effects. The persistence estimates are similar across the three approaches, ranging from 0.60 to 0.66. The volatility estimates differ more, especially in the country panel, so we use the panel mainly as a check on persistence.

Table 7 and Table C7 show how the persistence of the passive share process affects the decomposition of the reduced-form elasticity and the main unconditional moments.

C.4 Additional Figures and Tables

Default risk and prices. Figure C1 displays the default probability and the bond price $q(y, \tau, B')$ across different values of B' and y. Panel (a) shows that default risk rises with B' and falls with y. Accordingly, Panel (b) shows that the bond price decreases with B' (blue line). The figure also highlights how downward-sloping demand affects bond prices. The solid black line in Panel (b) plots $q(\cdot)/\Psi(\cdot)$, where

$$\Psi\left(y, \tau, B'\right) \equiv \frac{q\left(y, \tau, B'\right)}{\mathbb{E}_{s'|s}\mathcal{R}\left(y', \tau', B'\right)/r^f}$$

denotes the ratio of the bond price to discounted expected repayment. At low B' levels, default risk and repayment volatility are small, so $\Psi(\cdot)$ is close to one. As B' increases, repayment volatility rises, $\Psi(\cdot)$ falls, and the bond price moves below discounted expected

Appendix Figure C1: Default risk and bond prices

(a) Default probability (b) Bond price

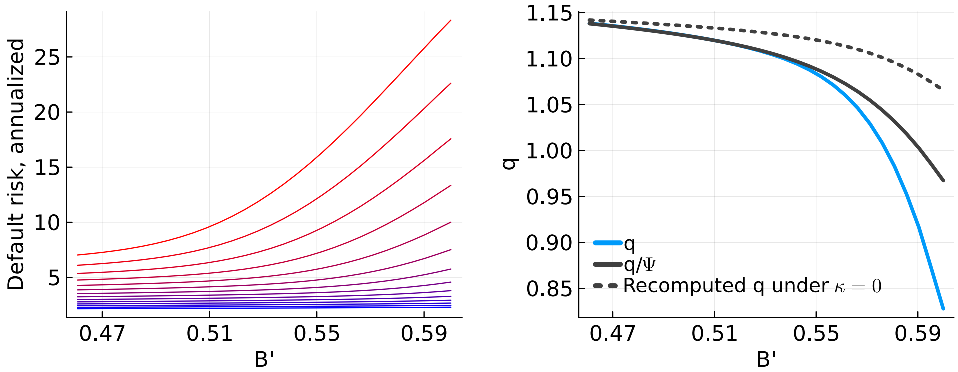

Note: Panel (a) displays the $1/\lambda$-period-ahead (annualized) default risk across combinations of B' and y. Color intensity ranges from red (low output) to blue (high output). Panel (b) plots bond prices as a function of B', evaluated at the mean values of y and τ. The solid blue line shows the bond pricing schedule $q(y, \tau, B')$. The solid black line plots discounted expected repayment, $q(\cdot)/\Psi(\cdot)$. The dotted line shows bond prices recomputed at $\kappa = 0$ while holding the baseline policies $\{B'(\cdot), d(\cdot)\}$ fixed.

repayment. The dotted line keeps the baseline policies $\{B'(\cdot), d(\cdot)\}$ fixed but recomputes prices under $\kappa = 0$. Its lower curvature relative to q/Ψ reflects that imperfectly elastic demand affects prices both contemporaneously, through the quantity-based price adjustment, and dynamically, through expected future prices embedded in repayment.

Phase diagram. Figure C2 plots the target-debt locus $B^\star(y)$, defined as the fixed point of the policy function $B'(B^\star, y) = B^\star$ for each y. The arrows show quarterly transition dynamics from different initial pairs (y, B), holding y fixed. The shaded area displays the default region. In the baseline model, all objects are evaluated at the mean value of τ.

Appendix Figure C2: Debt dynamics

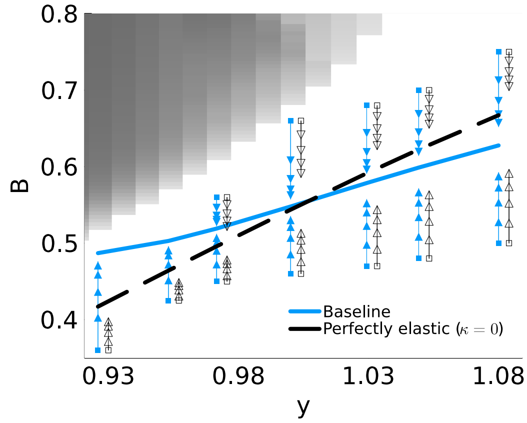

Notes. Lines show the target-debt locus $B^\star(y)$, the fixed point of the policy function for each y. Arrows illustrate quarterly transitions from different initial combinations of (y, B), holding y fixed, with squares marking the starting points. Solid blue lines and arrows correspond to the baseline model ($\kappa > 0$); dashed black lines and open arrows correspond to the perfectly elastic case ($\kappa = 0$). The shaded area is the default region in the $\kappa > 0$ economy. Because the default decision is subject to a smoothness shock, the default policy takes values in $[0, 1]$ rather than following a binary $\{0, 1\}$ rule (see Appendix C.2).

Appendix Figure C3: Default costs: Fully recalibrated perfectly elastic model

(a) Default costs, function of y (b) Default costs (c) Debt-to-output

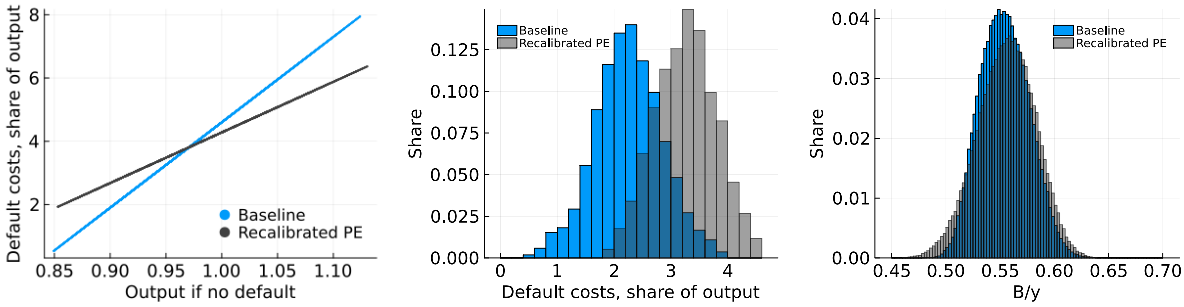

Note: The figure compares model-implied default costs and debt-to-output distributions under our baseline model and a counterfactual with perfectly elastic demand. In the counterfactual, the parameters $\{d_0, d_1\}$ and β are recalibrated to match the targeted mean and standard deviation of bond spreads, as well as the average debt-to-output ratio. Panel (a) plots default costs, $\Phi(y)/y$, as a function of output. Panel (b) shows the distribution of output costs at the time of default. Panel (c) shows the distribution of debt-to-output ratios, excluding observations within three years of a default event.

Relative to the perfectly elastic case, the baseline target-debt locus is flatter, and adjustment toward it is faster, particularly when current debt lies above $B^\star(y)$. At these points, our baseline economy deleverages more quickly, since the larger price impact strengthens the incentive to reduce debt when outstanding obligations are large.

Default costs. Figure 7 in the main text compares default costs and the debt-to-output distribution between our baseline model and the recalibrated PE model, in which $\{d_0, d_1\}$ are adjusted to match the average and standard deviation of bond spreads. Figure C3 presents the analogous comparison for the fully recalibrated PE model, which additionally adjusts β to match the targeted debt-to-output ratio. The implied default cost parameters (Panels a and b) are similar to those in Figure 7. The calibration slightly increases β (from 0.948 to 0.95), allowing the fully recalibrated PE model to generate a debt-to-output distribution comparable to that of the baseline (Panel c).

C.5 Sensitivity Exercises

We examine the robustness of our findings to three alternative parameter choices: bond maturity, safer government environments, and a state-dependent demand elasticity.

Role of maturity

In our three-period environment, we show that the dilution-curbing force is tightly linked to debt maturity, weakening as maturity shortens. We revisit this link quantitatively by varying λ around the baseline calibration, reducing expected maturity from

Appendix Table C2: Role of maturity

λ	Maturity	Baseline model		Perfectly elastic		Difference			
		B/y	SP	B/y	SP	$\Delta B/y$	$\frac{\Delta B/y}{(B/y)_{PE}}$	ΔSP	$\frac{\Delta SP}{SP_{PE}}$
0.05	20.00	53.1	5.06	51.0	8.59	2.1	4.1%	-3.53	-41.1%
0.07	14.29	55.5	3.72	53.1	6.08	2.4	4.5%	-2.36	-38.8%
0.09	11.11	57.4	2.92	55.2	4.57	2.2	4.0%	-1.66	-36.2%
0.11	9.09	59.0	2.39	57.0	3.62	1.9	3.3%	-1.23	-33.9%
0.13	7.69	60.1	2.02	58.5	2.97	1.6	2.7%	-0.95	-32.1%
0.15	6.67	61.1	1.74	59.8	2.52	1.3	2.2%	-0.78	-30.8%
0.17	5.88	61.8	1.53	60.7	2.18	1.1	1.8%	-0.65	-29.7%

Note: The table reports simulated moments for economies with different maturity parameters, λ, with expected maturity given by $1/\lambda$ quarters. The columns labeled Baseline model report economies with the baseline calibrated value of κ, recalibrating only $\overline{\mathcal{A}}$ so that the average inelasticity premium is zero, $\bar{\mathbb{E}}(\vartheta) = 0$. The columns labeled Perfectly elastic report the corresponding economies with $\kappa = 0$. The difference columns report the baseline value minus the perfectly elastic value; percentage differences divide by the corresponding perfectly elastic moment. B/y denotes the average debt-to-output ratio, reported in percentage points. SP denotes annualized bond spreads, also in percentage points.

20 to about 6 quarters. For each λ, we compare the imperfectly elastic model with its perfectly elastic counterpart, $\kappa = 0$. In the imperfectly elastic case, we keep κ fixed, recalibrate only $\overline{\mathcal{A}}$ to keep the average inelasticity premium at zero, $\bar{\mathbb{E}}(\vartheta) = 0$, and leave all other parameters at their baseline values.

As maturity shortens, Table C2 shows that spreads fall in both economies, because less outstanding debt remains exposed to future dilution. As a result, the government optimally increases its average debt. Comparing the two economies, debt remains higher under imperfectly elastic demand throughout, although the gap narrows as maturity falls, consistent with Corollary A.1. The spread reduction associated with imperfectly elastic demand follows the same pattern and is largest at the baseline maturity. This decline has two sources. The reduced exposure to dilution leaves the dilution-curbing force with less to shield. Shorter maturity also lowers default risk and repayment variance, which, in our quantitative specification, flattens the demand schedule even when κ is held fixed.

Safer government environments

Our calibration targets Argentina, a high-spread emerging-market sovereign. In this appendix, we ask whether our mechanism remains relevant in environments closer to advanced sovereigns, where spreads are lower and sustainable debt is higher (e.g., Italy or Spain). To keep the analysis transparent, we do not perform a full country-specific recalibration. Instead, we move the economy progressively toward a safer environment by lowering $|d_0|$, which raises default costs in the relevant region and makes default less

Appendix Table C3: Safer government environment

Moment	Description	Safer		Safer +		Safer ++	
		$\kappa > 0$	$\kappa = 0$	$\kappa > 0$	$\kappa = 0$	$\kappa > 0$	$\kappa = 0$
$\bar{\mathbb{E}}(B/y)$	Debt to output	62.7	59.9	73.1	69.4	83.9	79.7
$\bar{\mathbb{E}}(SP)$	Bond spreads	4.26	7.09	3.59	5.92	3.08	5.00
$\bar{\sigma}(SP)$	Volatility of spreads	1.15	3.24	0.91	2.42	0.72	1.86
$\bar{\mathbb{E}}(d_\lambda)$	Default probability	3.47	5.67	3.03	4.90	2.62	4.26
$\bar{\mathbb{E}}(\eta)$	Model-implied elasticity	-0.17	-	-0.17	-	-0.16	-
$\bar{\mathbb{E}}(\vartheta)$	Inel. premium	0.03	0.00	0.00	0.00	0.03	0.00

Note: The table compares model moments as the government environment becomes progressively safer (*Safer* $\rightarrow$ *Safer* + $\rightarrow$ *Safer* ++), implemented by lowering $|d_0|$ so that default costs are higher and default is less likely. Within each pair, the first column is the baseline economy ($\kappa > 0$) and the second is the perfectly elastic case ($\kappa = 0$). In the $\kappa > 0$ economies, $\overline{\mathcal{A}}$ is recalibrated so that the average inelasticity premium is approximately zero. Debt-to-output, spreads, default probabilities, and inelasticity premia are reported in percentage points.

likely. We keep all other parameters fixed and, in the $\kappa > 0$ economies, recalibrate only $\overline{\mathcal{A}}$ so that the average inelasticity premium is approximately zero.

The exercise confirms that our results do not rely on Argentina being a high-spread sovereign. The economy with imperfectly elastic demand continues to sustain more debt while facing lower spreads and default risk than the perfectly elastic case. As debt capacity increases, the debt-to-output gap widens in absolute terms, from 2.8 p.p. in the *Safer* case to 4.2 p.p. in the *Safer* ++ case. The proportional debt gap remains similar to that in the baseline.

State-dependent elasticity

Throughout our analysis, κ, the parameter governing the demand slope, is constant across economic states. Because κ summarizes underlying market frictions—regulatory constraints, limited risk-bearing capacity and the price of risk it generates, liquidity, market segmentation, and investor composition—a natural extension is to let κ *itself* vary with the state. These frictions are likely to tighten in periods of stress, when risk-bearing capacity is lower and constraints bind more sharply.

Given our focus on international bonds included in the EMBIGD, we first examine how our reduced-form estimates vary with global factors that proxy for the risk-bearing capacity and portfolio constraints of global institutional investors: the VIX, US GDP growth, the US yield-curve slope (10-year minus 2-year yields), and financial-intermediary wealth shocks. We standardize each factor and denote it by G_t. We then augment our baseline regression, Equation (20), with an additional interaction term and estimate the

Appendix Figure C4: Reduced-form price responses and global financial conditions

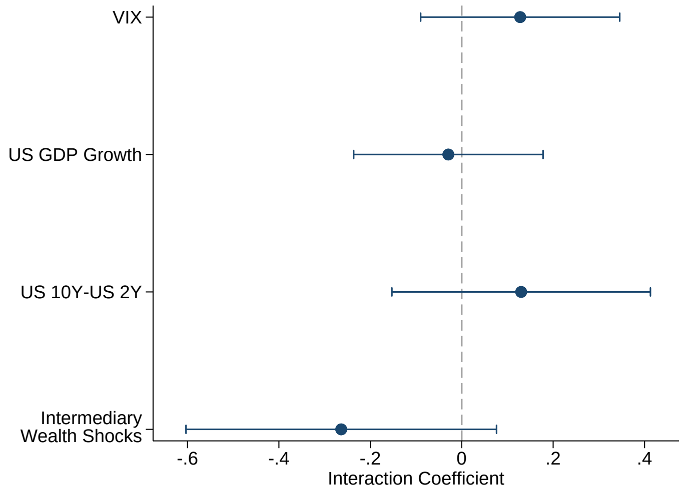

Note: The figure reports the interaction coefficient $\widehat{\delta}_G$, estimated separately for each standardized global factor. A positive coefficient indicates that the price response to FIR is larger when the factor is high. Each regression follows the specification in column (6) of Table 3, including bond-month and month-by-$\mathbf{1}_{h\in Post}$ fixed effects and excluding $h = -1$. The endogenous variables are instrumented with $Z_{c(i),t} \times \mathbf{1}_{h\in Post}$ and $Z_{c(i),t} \times \mathbf{1}_{h\in Post} \times G_t$. US GDP Growth is the quarterly growth rate of seasonally adjusted real US GDP. US 10Y–US 2Y is the US 10-year constant maturity yield minus the US 2-year constant maturity yield. Intermediary Wealth Shocks are unexpected shocks to financial intermediaries' wealth from He et al. (2017) and are available through 2018Q3. Horizontal bars represent 95% confidence intervals. Standard errors are clustered at the country-month level.

following pooled 2SLS specification:

$$\log(q_{i,t,h}) = \theta_{i,t} + \gamma_t \mathbf{1}_{h\in Post} + \delta \left(\mathrm{FIR}_{c(i),t} \times \mathbf{1}_{h\in Post}\right) + \delta_G \left(\mathrm{FIR}_{c(i),t} \times \mathbf{1}_{h\in Post} \times G_t\right) + \varepsilon_{i,t,h},$$

where $\theta_{i,t}$ denotes bond-month fixed effects and $\gamma_t \mathbf{1}_{h\in Post}$ denotes month-by-$\mathbf{1}_{h\in Post}$ fixed effects. We instrument $\mathrm{FIR}_{c(i),t} \times \mathbf{1}_{h\in Post}$ and its interaction with G_t using $Z_{c(i),t} \times \mathbf{1}_{h\in Post}$ and $Z_{c(i),t} \times \mathbf{1}_{h\in Post} \times G_t$, respectively. The coefficient of interest, δ_G, measures how the price response to FIR varies with a one-standard-deviation increase in the corresponding factor. Each regression follows the specification in column (6) of Table 3, using the same event window and excluding the trading day immediately before rebalancing.

Our empirical sample (2016–2018) covers a relatively tranquil period, without any major global crises, and therefore does not allow us to cleanly identify how the parameter κ itself varies with global financial conditions. Figure C4 nevertheless documents how the estimated price responses comove with these global factors. Price responses are larger when the VIX is high, US GDP growth is low, the yield curve slope is steep, and intermediaries experience negative wealth shocks, consistent with tighter constraints

Appendix Table C4: State-dependent elasticity

Target	Description	Baseline model	State-dependent $\mathcal{K}(y)$			
			$\varpi = 1.15$	$\varpi = 1.25$	$\varpi = 1.55$	$\varpi = 2.0$
$\bar{\mathbb{E}}(SP)$	Bond spreads (%)	5.06	4.87	4.76	4.48	4.17
$\bar{\sigma}(SP)$	Volatility of spreads (%)	1.56	1.46	1.40	1.30	1.34
$\bar{\mathbb{E}}(B/y)$	Debt to output (%)	53.1	53.2	53.3	53.5	53.7
$\bar{\mathbb{E}}(\eta^{\mathrm{RF}})$	Reduced-form elasticity	-0.28	-0.30	-0.31	-0.34	-0.38

Note: The table reports simulated moments for the baseline model with constant κ and for counterfactual economies in which the parameter governing the demand slope is state dependent, $\mathcal{K}(y)$, as defined in Equation (C5).

during periods of financial stress.[8] Given the short sample, these estimates are noisy, so we view this evidence as suggestive. We therefore use it to motivate a quantitative sensitivity exercise in which κ is higher in bad states.

One caveat is worth emphasizing. Our estimates capture reduced-form price responses, not κ itself. In our quantitative specification, the demand slope, $\mathcal{S}(y, \tau, B') = \frac{\kappa}{r^f}\mathbb{V}_{s'|s}\mathcal{R}\left(y', \tau', B'\right)$, varies across states even when κ is constant: in bad states, default risk and repayment variance rise, and the price impact of issuance grows. Larger price responses in periods of stress are thus consistent with our baseline specification. The exercise below asks whether our conclusions change when demand becomes even less elastic in bad states, beyond this baseline state dependence.

To this end, we consider a minimal departure from our baseline specification and let the parameter governing the demand slope vary with the state:

$$\mathcal{S}(y, \tau, B') = \frac{\mathcal{K}(y)}{r^f}\mathbb{V}_{s'|s}\mathcal{R}(y', \tau', B') \quad \text{with} \quad \mathcal{K}\left(y\right) = \begin{cases} \kappa & \text{if } y \geq \bar{y} \\ \varpi\kappa & \text{if } y < \bar{y}, \end{cases} \tag{C5}$$

where $\bar{y}$ denotes mean output and $\varpi > 1$. This simple specification allows demand to become less elastic in low-output states. While we do not introduce an additional state variable for global conditions, we use low domestic output as a parsimonious proxy for stress states in which global financial constraints may also tighten.[9]

Table C4 reports the implied changes in the targeted moments. We find that allowing κ to rise in low-y states strengthens the dilution-curbing force: as ϖ increases, equilibrium

[8]These results are in line with Monteiro (2024), who studies primary auctions for Portuguese bonds over a longer horizon (2004–2020) and finds that demand is less elastic in periods of stress, in particular during the European debt crisis. Similarly, Fang et al. (2025) show that marginal holders of sovereign debt change during periods of stress, which directly influences the aggregate demand elasticity.

[9]This assumption is similar in spirit to the lenders' stochastic discount factor in Arellano (2008), which loads on the borrower's endowment shocks, so that lenders value payoffs more in the borrower's bad states.

spreads decline and the government sustains a slightly higher stock of debt.

C.6 Reference Demand and Model-Implied Elasticities

This appendix clarifies why reference demand, $\mathcal{B}(\tau, B')$, matters quantitatively. We proceed in three steps. First, we derive a general expression linking the inelasticity premium to the expected path of demand elasticities and residual supply. Second, using this expression in a simplified stationary environment, we show that absent reference demand, even a large inelasticity premium implies a demand schedule far more elastic than what we estimate. Third, we re-solve the full model with $\mathcal{B}(\cdot) = 0$ and confirm that even values of κ for which the premium accounts for most of the average spread imply highly elastic demand.

Let $SP(y, \tau, B')$ denote annualized spreads, as a function of the current state. The inelasticity premium is defined as

$$\vartheta(y, \tau, B') \equiv SP(y, \tau, B') - SP^{\kappa=0}(y, \tau, B'),$$

where $SP^{\kappa=0}(\cdot)$ are annualized spreads under perfectly elastic demand, evaluated at the *same* borrowing and default policies as in the baseline, imperfectly elastic, economy. For what follows, we define residual supply relative to debt net of passive demand,

$$\psi(y, \tau, B') \equiv \frac{B' - \mathcal{B}(\tau, B')}{B' - \mathcal{T}(\tau, B')}.$$

This scaling matches the normalization used to define the inverse demand elasticity, η. For the rest of the analysis, we use time subscripts to simplify notation.

The following lemma expresses the inelasticity premium as a weighted present value of current and future elasticities and residual supply, $\eta_{t+j}\psi_{t+j}$. Future terms enter because bonds are long term. A lower future bond price reduces their value today, so the current premium reflects discounts along the entire expected repayment path.

Lemma 2. *The quarterly inelasticity premium is*

$$\vartheta_t^{quarterly} = -J_t \, \mathbb{E}_t \left[\sum_{j=0}^{\infty} \left(\tfrac{1-\lambda}{r^f} \right)^j \left(\prod_{h=1}^{j} (1 - d_{t+h}) \right) \frac{q_{t+j}}{q_t} \, \eta_{t+j} \, \psi_{t+j} \right], \tag{C6}$$

where $J_t \equiv [\lambda + (1-\lambda)\nu]/(r^f q_t^{\kappa=0})$*. The annualized premium is* $\vartheta_t \approx 4 \, \vartheta_t^{quarterly}$.

The date-$t+j$ term $\eta_{t+j}\psi_{t+j}$ is the price adjustment in period $t+j$ implied by the demand schedule. The inverse demand elasticity η_{t+j} measures the marginal price response per unit of debt net of passive demand, while ψ_{t+j} scales it by residual supply relative to that same quantity. This term affects today's price only if the bond is still outstanding and has not defaulted. The factor $\big((1-\lambda)/r^f\big)^j$ discounts the surviving fraction j periods ahead, $\prod_{h=1}^{j}(1-d_{t+h})$ keeps only no-default histories, and q_{t+j}/q_t converts future price adjustments into units of today's bond price.

Consider now the following example. Assume that the economy is stationary with constant default risk, prices, and elasticities. Suppose also that there is no reference demand, $\mathcal{B}=0$, so $\psi=1$. The sum in Equation (C6) then collapses to a geometric series with quarterly survival probability $s \equiv \mathbb{E}[1-d]$, giving

$$\eta \approx -\frac{\vartheta}{4J}\left(1-\frac{1-\lambda}{r^f}s\right) \quad \text{with} \quad J=\frac{1}{s}\left(1-\frac{1-\lambda}{r^f}s\right), \tag{C7}$$

where the last equality uses the fact that $q^{\kappa=0}=\frac{s(\lambda+(1-\lambda)\nu)}{r^f-s(1-\lambda)}$. Take an annual default probability of 5%, so $s=(0.95)^{1/4}$, and suppose the inelasticity premium accounts for 30% of spreads. Under perfectly elastic demand the spread is $(0.95)^{-1}-1=5.26$ p.p. A 30% premium implies spreads under imperfectly elastic demand of $5.26/(1-0.30)=7.52$ p.p. and $\vartheta=2.26$ p.p. Using Equation (C7) with $\lambda=0.05$, $\nu=0.03$, and $r^f=1.01$ gives $\eta=-0.0056$, which is thirty times smaller in magnitude than the -0.18 we recover in the quantitative section.

Absent reference demand, the example above shows that even a large premium implies demand far more elastic than we estimate. Importantly, the calculation imposes no functional form on $\mathcal{S}(\cdot)$ and is agnostic about its drivers—risk aversion, liquidity, or others—though it relies on the additive structure of Equation (3).

The calculation above holds government policies and future prices fixed. We now re-solve the full model with $\mathcal{B}=0$. As in the rest of the quantitative analysis, the slope scales with repayment risk, $\mathcal{S}_t=(\kappa/r^f)\mathbb{V}_t(\mathcal{R}_{t+1})$, so the bond price becomes

$$q_t=\frac{1}{r^f}\mathbb{E}_t\mathcal{R}_{t+1}-\frac{\kappa}{r^f}\,\mathbb{V}_t(\mathcal{R}_{t+1})\,B_t'. \tag{C8}$$

This limiting case is instructive because Equation (C8) is the price a mean-variance risk-averse investor would set.

Table C5 reports the results, which are in line with those of the example above. With $\kappa=2$, the premium accounts for more than one-third of spreads, above the risk-premium estimates of Longstaff et al. (2011) for sovereign bonds. Yet the implied η^{RF} and η are about twenty times smaller in magnitude than in the baseline. Raising κ until the premium is close to the full spread (last columns) still leaves η about ten times smaller

Appendix Table C5: No-reference-demand counterfactuals

Moment	Description	Baseline model	No reference demand				
			$\kappa = 0$ (PE)	$\kappa = 2$	$\kappa = 2$ bounded $\mathcal{S}$	$\kappa = 10$	$\kappa = 25$
$\bar{\mathbb{E}}(B/y)$	Debt to output	53.1	51.0	50.5	50.5	47.2	42.7
$\bar{\mathbb{E}}(SP)$	Bond spreads	5.06	8.59	8.01	8.02	7.26	6.30
$\bar{\sigma}(SP)$	Volatility of spreads	1.56	4.49	3.46	3.50	2.52	1.78
$\bar{\mathbb{E}}(d_\lambda)$	Default probability	4.01	6.57	3.93	4.07	1.12	0.32
$\bar{\mathbb{E}}(\eta^{\mathrm{RF}})$	Reduced-form elasticity	-0.278	-	-0.014	-0.002	-0.048	-0.042
$\bar{\mathbb{E}}(\eta)$	Model-implied elasticity	-0.180	-	-0.008	-0.008	-0.016	-0.016
$\bar{\mathbb{E}}(\vartheta)$	Inelasticity premium	0.003	-	3.433	3.240	6.088	5.972

Note: The table reports the baseline economy and the no-reference-demand case ($\mathcal{B} = 0$) for different values of κ, holding all other parameters fixed. The column $\kappa = 0$ corresponds to the perfectly elastic (PE) case. The column labeled bounded $\mathcal{S}$ caps the conditional repayment variance at $\mathbb{V}(\mathcal{R}) \leq 0.04$. Moments are in percentage points, except elasticities.

than in the baseline economy, and η^{RF} at about one-sixth of its empirical counterpart. The slope rises only modestly with κ because the government deleverages, which lowers default risk and repayment variance.

The first row quantifies this deleveraging. With $\mathcal{B} = 0$, the parameter that steepens the demand schedule also raises the premium on the entire debt stock, which ranges from 3 to 6 p.p. across the columns, and the government holds less debt in response. This reflects the level effect characterized in Corollary A.2. These results do not imply that a positive inelasticity premium necessarily leads to less debt than under perfectly elastic demand. In Table 8 in the main text, we adjust reference demand to generate positive average premia and show that, even in calibrations in which the premium reaches almost one-third of spreads, average debt remains above the perfectly elastic case.

Figure C5 shows that these elasticity differences hold across the state space. Panel (a) plots η against spreads for the baseline economy and for the $\mathcal{B} = 0$, $\kappa = 2$ case. Over the range visited by the baseline, the no-reference-demand elasticity stays close to zero. Panel (b) extends to the far right tail of the $\kappa = 2$ economy's own spread distribution, states it rarely reaches. Only there does $|\eta|$ rise, up to 0.09. One might worry that these rare states still discipline future issuance because the schedule is steeper there. They do not. Capping $\mathcal{S}$ in this tail leaves debt, spreads, and elasticities essentially unchanged (the bounded-$\mathcal{S}$ column of Table C5).

Taken together, these exercises show that, when price impact applies to the entire debt stock rather than to residual supply, even a high κ does not generate the steep slope

Appendix Figure C5: Demand elasticity in the no-reference-demand counterfactual

(a) Baseline vs. no-reference-demand (b) High-spread tail

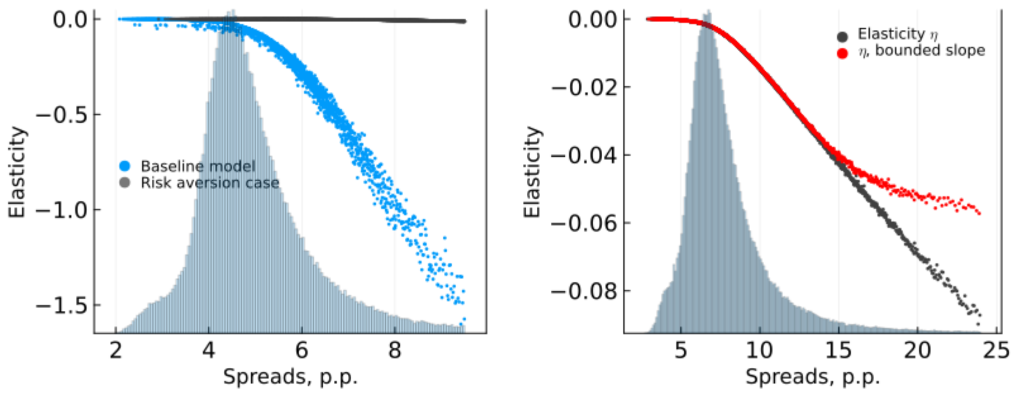

Note: Panel (a) plots the model-implied inverse demand elasticity η against bond spreads for the baseline and the no-reference-demand case ($\mathcal{B} = 0$, $\kappa = 2$); the histogram shows the baseline spread distribution. Panel (b) restricts attention to the $\kappa = 2$ economy, with the histogram showing its spread distribution; it compares the implied elasticity (black) with the elasticity obtained by holding $\mathcal{S}$ constant in the far right tail (red). For the latter, we cap the variance at $\mathbb{V}(\mathcal{R}) \leq 0.04$. The cutoff is chosen so that the demand slope is mostly constant once spreads exceed 15%.

that disciplines issuance in our baseline. These results echo Petajisto (2009) and Gabaix and Koijen (2021), who show that frictionless models, including those with risk aversion, imply far more elastic demand than the data reveal. Anchoring demand around reference positions is what lets the model match the empirical elasticity.

Proof of Lemma 2. In both the $\kappa > 0$ and $\kappa = 0$ cases, prices are evaluated at the same borrowing and default policies, those of the $\kappa > 0$ economy. Let

$$q_t = \tfrac{1}{r^f}\,\mathbb{E}_t\Big[(1-d_{t+1})\Big(\lambda + (1-\lambda)\nu + (1-\lambda)q_{t+1}\Big)\Big] - \mathcal{S}_t(B_t' - \mathcal{B}_t),$$
$$q_t^{\kappa=0} = \tfrac{1}{r^f}\,\mathbb{E}_t\Big[(1-d_{t+1})\Big(\lambda + (1-\lambda)\nu + (1-\lambda)q_{t+1}^{\kappa=0}\Big)\Big].$$

Subtracting the first equation from the second gives the recursion

$$\Delta q_t = \mathcal{S}_t(B_t' - \mathcal{B}_t) + \tfrac{1-\lambda}{r^f}\mathbb{E}_t\Big[(1-d_{t+1})\Delta q_{t+1}\Big],$$

where $\Delta q_t \equiv q_t^{\kappa=0} - q_t$. Iterating forward N times via the law of iterated expectations,

$$\Delta q_t = \mathbb{E}_t\Bigg[\sum_{j=0}^{N}\Big(\tfrac{1-\lambda}{r^f}\Big)^j\Big(\textstyle\prod_{h=1}^{j}(1-d_{t+h})\Big)\,\mathcal{S}_{t+j}(B_{t+j}' - \mathcal{B}_{t+j})\Bigg] + R_{t,N},$$
$$R_{t,N} \equiv \mathbb{E}_t\Bigg[\Big(\tfrac{1-\lambda}{r^f}\Big)^{N+1}\Big(\textstyle\prod_{h=1}^{N+1}(1-d_{t+h})\Big)\Delta q_{t+N+1}\Bigg].$$

Bond prices lie in $[0, q^f]$ with $q^f = [\lambda + (1-\lambda)\nu]/[r^f - (1-\lambda)]$, so Δq is bounded; the survival products lie in $[0,1]$ and $(1-\lambda)/r^f < 1$, so $|R_{t,N}| \leq \Big((1-\lambda)/r^f\Big)^{N+1} q^f \to 0$.

Therefore,

$$\Delta q_t = \mathbb{E}_t\Big[\sum_{j=0}^{\infty}\left(\tfrac{1-\lambda}{r^f}\right)^j\left(\prod_{h=1}^{j}(1-d_{t+h})\right)\mathcal{S}_{t+j}(B'_{t+j}-\mathcal{B}_{t+j})\Big].$$

Using $\mathcal{S}_{t+j}(B'_{t+j}-\mathcal{B}_{t+j}) = -q_{t+j}\eta_{t+j}\psi_{t+j}$ and dividing by q_t yields the bracketed sum in Equation (C6). To recover J_t, note that the quarterly internal rate of return satisfies $q_t(1+i_t) = \lambda + (1-\lambda)(\nu + q_t)$, so $i_t = [\lambda + (1-\lambda)\nu]/q_t - \lambda$ and

$$i_t - i_t^{\kappa=0} = [\lambda + (1-\lambda)\nu]\left(\frac{1}{q_t} - \frac{1}{q_t^{\kappa=0}}\right) = \frac{\lambda + (1-\lambda)\nu}{q_t^{\kappa=0}}\,\frac{\Delta q_t}{q_t}.$$

Dividing by r^f gives $\vartheta_t^{\text{quarterly}} = J_t\,\Delta q_t/q_t$, which is Equation (C6). The approximate annualized premium is then $\vartheta_t = 4 \times \vartheta_t^{\text{quarterly}}$. □

C.7 Passive Demand Shifts, Spreads, and Borrowing

We next analyze how variation in the passive component of reference demand, τ, affects bond spreads and borrowing. First, we quantify the contribution of passive-demand shocks to spread volatility. Second, we study how these shocks propagate to prices, expected repayment, and borrowing, and how persistence shapes these responses. Third, we examine the role of the level of passive demand, considering alternative values of the mean share $\tau^\star$.

Passive demand and spread volatility

We begin by asking how much of the model's spread volatility is driven by passive-demand shocks. Some of the differences between the baseline and the perfectly elastic case could in principle reflect the volatility induced by τ shocks rather than the demand slope. To gauge this, we shut down innovations to τ ($\sigma_\tau = 0$), keeping $\kappa > 0$ and all other calibrated parameters fixed. Spread volatility falls only from 1.56 to 1.53 percentage points—roughly 4% of spread variance—while average spreads remain nearly unchanged. Passive-demand shocks therefore play only a small role in spreads. The differences between our baseline model and the perfectly elastic case—lower spreads, lower default risk, and higher sustainable debt—are driven primarily by the demand slope and the government's issuance response.

This comparison already incorporates the government's policy response to changes in τ. To isolate how much this response dampens spread volatility, we evaluate the baseline spread function $SP(y, \tau, B')$ at the fixed-τ borrowing policy, $B'_{\tau^\star}(y, \tau^\star, B)$. For

Appendix Table C6: Passive demand and spread volatility

	Fixed-τ model	Baseline model	Counter-factual
$\bar{\mathbb{E}}(SP)$, pp	5.052	5.064	5.059
$\bar{\sigma}(SP)$, pp	1.533	1.563	1.662
Variance ratio (vs fixed-τ model)	1.0	1.039	1.175

Note: The table compares the distribution of bond spreads across three economies. The *Fixed-τ* model holds passive demand fixed at $\tau = \tau^\star$. The *Baseline* model is our calibrated economy. The *Counterfactual* evaluates the baseline spread function using the borrowing policy from the fixed-τ economy, $SP_t^{\text{ctfl}} \equiv SP(y_t, \tau_t, B'_{\tau^\star}(y_t, \tau^\star, B_t))$, so that prices respond to τ as in the baseline but borrowing does not. $\bar{\mathbb{E}}(SP)$ and $\bar{\sigma}(SP)$ denote the mean and standard deviation of spreads. The variance-ratio row expresses spread variance relative to the fixed-τ model.

Appendix Table C7: Passive-demand persistence and model moments

Moment	Description	Baseline model	Alternative persistence values: Low	Medium	High	Fixed passive demand
$\bar{\mathbb{E}}(B/y)$	Debt to output (%)	53.1	53.1	53.1	53.0	53.0
$\bar{\mathbb{E}}(SP)$	Bond spreads (%)	5.06	5.05	5.05	5.11	5.05
$\bar{\sigma}(SP)$	Volatility of spreads (%)	1.56	1.55	1.55	1.64	1.53
$\bar{\mathbb{E}}(d_\lambda)$	Default probability (%)	4.01	3.98	3.99	4.12	3.94
$\bar{\mathbb{E}}(\eta^{\text{RF}})$	Reduced-form elasticity	-0.28	-0.23	-0.25	-0.37	-
$\bar{\mathbb{E}}(\eta)$	Model-implied elasticity	-0.18	-0.18	-0.18	-0.18	-0.18

Note: The table reports model-implied moments under alternative assumptions about passive demand. The "Baseline" column uses the calibrated process for τ. The "Low," "Medium," and "High" columns set ρ_τ to 0.30, 0.50, and 0.90, respectively, holding all other parameters fixed. The last column fixes passive demand at $\tau = \tau^\star$. The reduced-form inverse elasticity is not reported in the fixed-passive-demand case because there is no rebalancing shock.

this counterfactual, we compute spreads as

$$SP_t^{\text{ctfl}} \equiv SP(y_t, \tau_t, B'_{\tau^\star}(y_t, \tau^\star, B_t)) .$$

With the borrowing policy held fixed, spread variance is about 17.5% higher than in the fixed-τ economy (Table C6). Thus, aggregate spread volatility changes little in the baseline because the sovereign adjusts issuance in response to these shocks, dampening their impact on spreads.

Passive-demand persistence and price responses

In the main text, we show that the gap between the reduced-form and model-implied inverse elasticities depends on the persistence of passive demand, ρ_τ. Table C7 extends this analysis by reporting debt-to-output, spreads, default probabilities, and inverse elasticities for alternative values of ρ_τ. Although the reduced-form inverse elasticity varies with ρ_τ, all other unconditional moments change little across the values considered.

Figure C6 illustrates this result by tracing how passive-demand changes affect prices,

Appendix Figure C6: Effects of changes in demand on prices and policies

(a) $\Delta q(\cdot)/q(\cdot)$ (b) $\Delta \mathbb{E}\mathcal{R}'(\cdot)/\mathbb{E}\mathcal{R}'(\cdot)$ (c) $\Delta B'(\cdot)/B'(\cdot)$

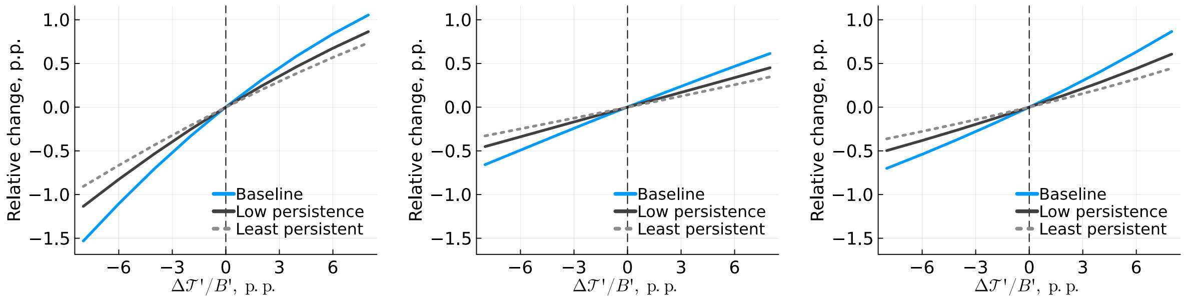

Note: The figure shows how changes in passive demand affect bond prices, expected repayment, and borrowing. The blue lines show the baseline calibration. The gray lines show parameterizations with lower persistence of the passive-share process, $\rho_\mathcal{T} = 0.50$ and $\rho_\mathcal{T} = 0.30$.

expected repayment, and borrowing. Panels (a) and (b) report within-period changes in bond prices and expected repayment following a change in passive demand. We evaluate the opening price, $q^{\mathrm{SM},0}(y, \tau, B')$, and expected repayment at the mean state $\{y, \tau, B'\}$, and then recompute both objects for different realizations of τ', holding $\{y, B'\}$ fixed. Both bond prices and expected repayment increase monotonically with passive demand. The expected-repayment response is large relative to the total price response and strengthens as the τ process becomes more persistent. Together with changes in repayment variance, this endogenous response accounts for the gap between the reduced-form and model-implied inverse elasticities and reflects adjustments in default risk and future issuance. Panel (c) illustrates the latter channel: the government's borrowing response to an increase in τ' is positive but muted, with a 5 p.p. increase (decrease) in the passive share raising (lowering) borrowing by less than 1%. Investors anticipate this response, which feeds back into current bond prices.

Figure C7 presents the impulse response to a positive shock to τ. The top panel shows the joint dynamics: debt rises and remains elevated, while bond prices initially increase by about 0.40%. As τ reverts, prices decline because the larger outstanding supply raises default risk. The bottom panel reports a counterfactual that keeps debt on its no-shock path. The initial price increase then more than doubles, reaching 0.90%, because the government's debt expansion is absent. Decomposing the response into the pricing ratio, $\Psi(\cdot)$, and expected repayment shows that the latter accounts for more than a third of the initial price response.

Appendix Figure C7: Impulse response to an increase in passive demand

(a) Effects on debt and bond prices

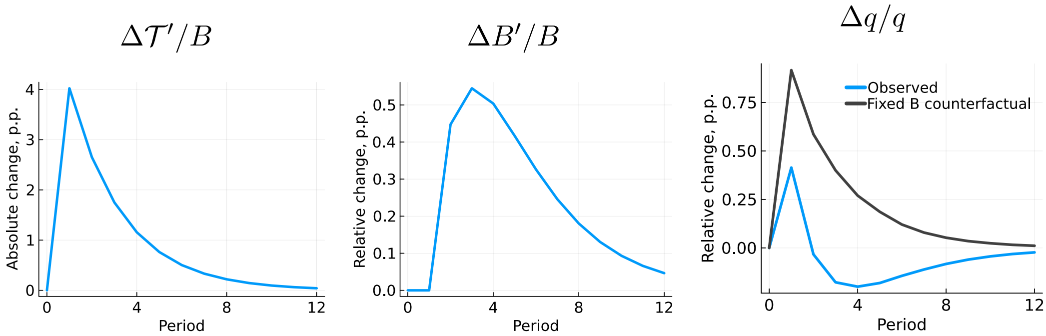

(b) Decomposition: Counterfactual with fixed debt

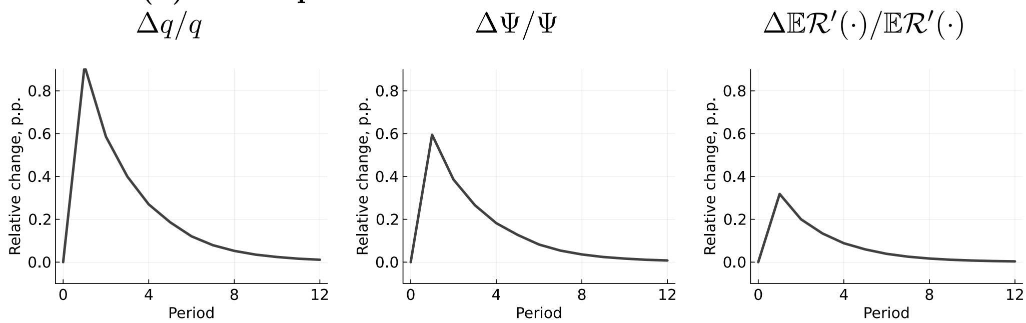

Note: The top panel displays impulse responses to an increase in passive demand. The bottom panel decomposes the bond-price response in a counterfactual in which debt follows its no-shock path.

The level of passive demand

Lastly, we examine the role of the mean passive share, $\tau^\star$. The baseline value is measured within the EMBIGD bond segment and applied to the entire debt stock in the model. This choice expresses the passive share in the same units as the empirical FIR, keeping the model's elasticity experiments comparable to the data (Appendix C.3). Alternative mappings would imply different values of $\tau^\star$, which could affect our analysis. To assess this possibility, we solve the model at $\tau^\star = 3\%$, 6%, and 18%, in each case adjusting $\overline{\mathcal{A}}$ so that the average inelasticity premium remains zero and holding κ and all remaining parameters fixed. Across this range, the reduced-form and model-implied inverse elasticities remain nearly unchanged, as do debt, spreads, spread volatility, and default risk (Table C8). Thus, the main quantitative results are robust to alternative values of the mean passive share once the average inelasticity premium is held at zero.

We then illustrate the effects of a higher mean passive share, $\tau^\star$, aimed at capturing the long-term rise of passive investors in emerging economies. Chari et al. (2022), for instance, document that assets under management in global funds investing in emerging markets have grown nearly twentyfold over the past two decades, with most of this growth

Appendix Table C8: The role of the passive share

Moment	Description	Baseline $\tau^\star = 12\%$	Alternative passive share $\tau^\star$			$\tau^\star = 18\%$, fixed $\overline{\mathcal{A}}$
			3%	6%	18%	
$\bar{\mathbb{E}}(B/y)$	Debt to output (%)	53.1	53.1	53.1	53.0	54.3
$\bar{\mathbb{E}}(SP)$	Bond spreads (%)	5.06	4.90	4.95	5.17	4.97
$\bar{\sigma}(SP)$	Volatility of spreads (%)	1.56	1.47	1.50	1.64	1.58
$\bar{\mathbb{E}}(d_\lambda)$	Default probability (%)	4.01	3.89	3.93	4.09	5.09
$\bar{\mathbb{E}}(\eta^{\mathrm{RF}})$	Reduced-form elasticity	-0.28	-0.29	-0.29	-0.27	-0.31
$\bar{\mathbb{E}}(\eta)$	Model-implied elasticity	-0.18	-0.19	-0.19	-0.17	-0.21
$\bar{\mathbb{E}}(\vartheta)$	Inelasticity premium	0.00	0.01	0.01	0.00	-1.87

Note: The first three alternative columns set $\tau^\star$ to 3%, 6%, and 18%, in each case recalibrating $\overline{\mathcal{A}}$ so that the average inelasticity premium is zero, as in the baseline. The last column sets $\tau^\star = 18\%$ but holds $\overline{\mathcal{A}}$ and all other model parameters fixed at their baseline values.

driven by ETFs. The last column of Table C8 sets $\tau^\star = 18\%$, a 50% increase over its baseline value of 12%, while, unlike the exercise above, holding $\overline{\mathcal{A}}$ fixed. The higher passive share then raises reference demand, lowering the average inelasticity premium to -1.87 p.p. The government responds by borrowing more, with the average debt-to-output ratio rising by about 1.2 p.p. and default risk by about 1.1 p.p. As a result, average spreads remain close to their baseline level.