

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

Germán Villegas-Bauer

International Monetary Fund

Tomás Williams

George Washington University

June 2, 2025

Abstract

We formulate a sovereign debt model with a rich yet flexible demand structure, featuring passive and active investors with asset-allocation mandates. In our framework, bond prices depend not only on government policies and default risk but also on investor composition and their price elasticity. We estimate this elasticity by combining our model with a novel strategy exploiting changes in the composition of the largest emerging-market bond index. We show that a downward-sloping demand acts as a disciplining device that mitigates debt dilution and lowers default risk. With greater market-driven discipline, fiscal rules have mild effects on spreads and default risk.

Keywords: sovereign debt, investor base, downward-sloping demand, debt dilution, 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 Jefferson Martinez Saavedra and Patricio Yunis for research assistance. The World Bank Chile Research and Development Center and Knowledge for Change Program (KCP) and George Washington University Facilitating Fund (GWU UFF) provided financial support for this paper. Lorenzo Pandolfi gratefully acknowledges financial support 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 or 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: gvillegasbauer@imf.org. Williams: tomaswilliams@email.gwu.edu.

1 Introduction

Governments in emerging markets (EM) rely heavily on bonds issued in international capital markets for financing. The composition of investors in these markets, and their responsiveness to changes in bond supply, are important determinants of borrowing costs, default risk, and debt management. Standard sovereign debt models typically abstract from these aspects, assuming perfectly elastic bond demand curves—implying that investors willingly absorb any borrowing at the risk-free rate plus a premium for default risk (Aguar and Gopinath, 2006; Arellano, 2008; Chatterjee and Eyigungor, 2012). This assumption contrasts with a growing demand-based asset pricing literature documenting downward-sloping demand curves across various asset markets (Kojen and Yogo, 2019; Gabaix and Kojen, 2021; Vayanos and Vila, 2021). It also conflicts with the rising presence of global institutional investors in EM debt—such as mutual funds and ETFs—whose demand is often perfectly inelastic, as they simply track benchmark indices.

In this paper, we formulate a sovereign debt model with a more realistic investor base, comprising both passive and active investors with asset-allocation mandates. In our framework, prices of long-term bonds depend not only on government policies and default risk, but also on the composition of investors and their price elasticity. We estimate this elasticity by combining the structural model with a novel identification strategy that exploits variation in the composition of the largest global index for EM government bonds. We find evidence of downward-sloping demand, which becomes less elastic as default risk increases. We then use the model, disciplined by this elasticity, to quantify the implications for bond prices, default risk, and optimal debt management.

The introduction of a richer investor structure reshapes key aspects of sovereign debt dynamics. Three results stand out. First, the presence of inelastic investors—broadly defined as those with downward-sloping demand curves—serves as a disciplining device. Each additional unit of debt issued lowers bond prices even when default risk remains constant, a cost that governments internalize by limiting borrowing. This force intensifies during periods of high default risk, as demand becomes less elastic. Second, this market-based pressure reshapes the classic debt-dilution channel faced by risky long-term bonds. When governments lack commitment over future issuance, investors anticipating additional borrowing require higher spreads to compensate for default risk. By curbing future borrowing, a downward-sloping demand mitigates dilution, lowering borrowing costs and allowing the government to sustain a *larger* debt stock. Third, our disciplining device directly influences issuance

and thus differs fundamentally from the standard default-cost mechanism, which operates through default incentives. This distinction carries important policy implications: fiscal rules such as debt ceilings have a more limited impact on default risk and bond spreads, as much of the discipline is already imposed by the market.

We begin the analysis by formulating a model in which a government issues long-term bonds in international markets, has limited commitment, and can endogenously default on its debt obligations. In the event of default, the government faces exogenous output losses and market exclusion. We depart from standard models by introducing a rich yet flexible investor demand structure with multiple investor types, explicitly distinguishing between active and passive investors. As in [Gabaix and Koijen \(2021\)](#), we assume that active investors operate under a mandate that dictates how they allocate funds across a finite set of bonds. For example, the mandate might require investors to adjust their holdings based on bonds' expected returns but imposes limits on the extent of these adjustments. This flexible specification generates a downward-sloping demand for active investors. On the other hand, passive investors are perfectly inelastic as they just replicate the composition of a benchmark index.

A key parameter in our model is the demand elasticity of active investors. While several studies estimate this elasticity across a range of asset markets using price reactions to supply-shifting shocks, identifying it in the context of risky long-term bonds is particularly challenging. This is because future borrowing decisions influence current default risk and, in turn, the pricing of outstanding bonds. We show that even when the supply-shifting shocks are orthogonal to *current* fundamentals, they can still affect the government's *future* borrowing decisions, bond payoffs, and ultimately current bond prices. As a result, the estimated price reactions may not reflect the true slope of the demand curve. We address this issue by combining novel reduced-form estimates with our structural model, which allows us to isolate endogenous changes in default risk and bond payoffs and to recover the underlying demand elasticity.

We first estimate a reduced-form demand elasticity for active investors using exogenous variation in passive demand. Since passive demand is perfectly inelastic, increases in it reduce the quantity of bonds available to active investors, effectively acting as supply-shifting shocks. We construct these shocks by exploiting monthly changes in the composition of the largest benchmark index for EM dollar-denominated sovereign bonds—the J.P. Morgan Emerging Markets Bond Index Global Diversified (EMBIGD). These routine rebalancings reflect the entry of newly qualifying bonds and exit of maturing ones, reallocate the portfolios of passive

funds tracking the EMBIGD, and shift the supply available to active investors.

We derive a measure of the flows implied by index rebalancings (FIR) by combining the assets passively tracking the EMBIGD with the index’s monthly changes. These flows are inherently endogenous, since index weights depend on bond prices, which in turn reflect country-specific fundamentals. To address this, we construct an instrument based on a diversification rule of the EMBIGD that caps country weights according to the *face value*—rather than the market value—of outstanding bonds. This rule allows us to compute *synthetic* weights that depend only on bond quantities, not prices. To ensure that these weights are not influenced by a country’s own bond issuances or retirements, we focus on countries with no changes in the face value of their index-eligible bonds during a given rebalancing. Variation in the instrument then comes from *other* countries’ bond issuances or retirements. Specifically, bond issuance or retirement by country i affects the *synthetic* weight of country j , and the magnitude of this effect depends on how tightly the cap binds for j . We combine this instrument with the known timing of index rebalancings—on the last business day of each month—to study bond price responses within narrow windows around these events.

According to our estimates, a 1 percentage point (p.p.) increase in the FIR—normalized by the size of active demand—raises bond prices by 0.30%, implying a decline in yield to maturity of over 5 basis points. This corresponds to an inverse price elasticity for active investors of approximately -0.30 . These price responses vary across countries with different levels of default risk. For riskier countries, a 1 p.p. increase in the FIR raises bond prices by up to 0.40%, while for safer countries the effects are small and statistically insignificant. Overall, these results suggest that active investors demand a premium for holding risky bonds—above and beyond their default risk—which we refer to as an “inconvenience” yield.

We discipline our model using these reduced-form estimates and recover a structural demand elasticity that accounts for endogenous changes in future policies and bond payoffs. Our approach is a form of indirect inference. First, we extend the baseline model to include secondary markets, allowing us to replicate our high-frequency empirical exercise. Next, we calibrate the demand-slope parameter—embedded in the asset-allocation mandates of active investors—so that the model matches the reduced-form estimate. Then, we decompose that estimate to recover the structural elasticity. We find that more than a third of the reduced-form elasticity reflects endogenous changes in bond payoffs, resulting in an average model-implied inverse elasticity of -0.20 . Overall, our results highlight the importance of accounting for issuers’ endogenous responses and changes in expected asset repayments when estimating demand elasticities, a dimension that has been widely overlooked in the literature.

Ignoring these differences would lead to an overestimation of the true slope of the demand curve.

Using the calibrated model, we then explore the implications of a downward-sloping demand curve for the optimal supply of risky sovereign bonds. First, we show that such demand acts as a disciplining device, particularly in periods of high default risk, when the demand is less elastic. The reason is that the government internalizes that each additional unit of debt lowers bond prices, even when default risk is held constant. Consequently, it finds it optimal to limit bond issuance, as large debt issuances become increasingly costly. The strength of this mechanism depends on the slope of the demand curve, underscoring the importance of carefully pinning down this parameter.

Second, we show that this market-based discipline significantly lowers default risk and bond spreads. This result hinges on the well-known debt-dilution mechanism associated with risky long-term bonds. In particular, the presence of inelastic investors reduces the perceived risk of future borrowing, placing far less upward pressure on current default probabilities. As a result, spreads decline substantially. We find that the resulting improvement in bond prices is strong enough that the government optimally issues, on average, *more* debt relative to the perfectly elastic case.

Third, in standard sovereign debt models with perfectly elastic investors, discipline arises solely from output losses and market exclusion triggered by default. In these models, higher default costs discourage the government from defaulting, which lowers default risk and borrowing costs, and indirectly supports higher debt levels. This contrasts with our market-based disciplining device, where the presence of inelastic investors directly curbs debt issuance. To analyze the implications of these mechanisms, we compare our baseline model with a recalibrated perfectly elastic benchmark that matches all targeted moments for bond spreads and debt. We find that market-based discipline gives rise to smaller default costs and different dynamics around default episodes. In particular, it leads to faster deleveraging and an extended decline in output prior to default.

This alternative source of discipline also carries important policy implications. We study the effects of a debt ceiling policy that sets an upper bound on the debt-to-output ratio. In our baseline model with inelastic investors, such a policy has only mild effects on default probabilities, bond spreads, and average debt levels. By contrast, in the perfectly elastic case, the same policy leads to a significantly larger drop in default probabilities and spreads, and induces the government to issue more debt. These results highlight that, with greater market-based discipline, fiscal rules have more limited effects on borrowing and sovereign risk.

Lastly, our model also allows us to quantify the effects of changes in the composition of the investor base. Motivated by the long-term rise of passive investors in emerging economies, we analyze the consequences of increased passive bond investing. We show that a larger passive share enables the government to sustain a higher debt-to-output ratio without raising borrowing costs. This is because passive investors are willing to absorb any amount of debt in order to replicate the composition of the index they track, generating a convenience yield that offsets the higher default risk.

Related Literature. Our findings contribute to several strands of literature. First, a growing body of work on inelastic financial markets emphasizes the role of the demand side in explaining asset prices across various markets (Koijen and Yogo, 2019; Gabaix and Koijen, 2021; Vayanos and Vila, 2021). Most of these studies focus on safe government bonds (e.g., U.S. Treasuries) and are largely silent on how a downward-sloping demand affects the provision of such assets (Krishnamurthy and Vissing-Jorgensen, 2012; Greenwood, Hanson, and Stein, 2015; Mian, Straub, and Sufi, 2022; Ray, Droste, and Gorodnichenko, 2024; Jiang, Lustig, Van Nieuwerburgh, and Xiaolan, 2024).¹ The closest paper is Choi, Kirpalani, and Perez (2024), who quantify the effects of downward-sloping demand on the underprovision of riskless U.S. Treasuries.² In contrast, we focus on the supply of risky bonds and show that, under default risk, a downward-sloping demand can lead to a *larger* bond supply.

Second, our paper is closely related to a large literature on quantitative sovereign debt models (Aguiar and Gopinath, 2006; Arellano, 2008; Chatterjee and Eyigungor, 2012). We extend the standard framework by introducing a more realistic investor base, with passive and active investors that follow asset-allocation mandates.³ We also contribute to the literature on debt dilution (Hatchondo, Martinez, and Sosa-Padilla, 2016; Aguiar, Amador, Hopenhayn, and Werning, 2019; Aguiar and Amador, 2020), emphasizing how investor demand elasticity shapes the government’s incentives to dilute the value of debt outstanding. In our framework, a less elastic demand curve acts as a disciplining device, limiting future borrowing and mitigating dilution concerns. In this regard, our paper relates to a larger literature on the use

¹Koijen and Yogo (2020), Gourinchas, Ray, and Vayanos (2022), Greenwood, Hanson, Stein, and Sunderam (2023), and Monteiro (2024) study international markets. Dathan and Davydenko (2020), Bretscher, Schmid, Sen, and Sharma (2022), Calomiris, Larrain, Schmukler, and Williams (2022), and Kubitzka (2023) focus on corporate bonds.

²Kaldorf and Rottger (2023) analyze how downward-sloping demand affects the pricing and issuance of European sovereign bonds. Like Choi et al. (2024), their framework features investors paying a convenience yield for collateral services.

³Similarly, recent work by Costain, Nuño, and Thomas (2022) introduces endogenous default risk into a Vayanos–Vila preferred habitat model to analyze the term structure of interest rates in the European Monetary Union.

of fiscal rules as commitment devices (Alfaro and Kanczuk, 2017; Dosis and Kirpalani, 2020; Hatchondo, Martinez, and Roch, 2022; Bianchi, Ottonello, and Presno, 2023). If the demand for bonds is downward sloping, the market itself creates incentives that discourage borrowing.

In our analysis, we remain agnostic about the underlying sources of investors' mandates that give rise to a downward-sloping demand. Borri and Verdelhan (2010), Lizarazo (2013), Pouzo and Presno (2016), and Arellano, Bai, and Lizarazo (2017) study sovereign debt models with risk-averse investors, where price elasticity arises because investors require compensation for each additional unit of risky debt. Other mechanisms can also generate downward-sloping demand, such as regulatory constraints (e.g., Value-at-Risk limits) (Gabaix and Maggiori, 2015; Miranda-Agrippino and Rey, 2020), liquidity considerations (He and Milbradt, 2014; Moretti, 2020; Passadore and Xu, 2022; Bigio, Nuño, and Passadore, 2023; Chaumont, 2024), buy-and-hold investment strategies, or inertia. Our framework adopts a flexible demand structure that can accommodate any of these features. Our goal is not to identify the source of inelasticity, but to study its aggregate consequences.

Third, our work relates to recent empirical studies examining how changes in the investor base influence the pricing of government and corporate bonds (Faia, Lewis, and Zhou, 2024; Faia, Salomao, and Veghazy, 2024; Jansen, Li, and Schmid, 2024; Zhou, 2024; Fang, Hardy, and Lewis, 2025). While we also analyze differences across investors (active versus passive), our approach differs by combining bond-level price reactions with a structural model that endogenizes issuer policies and bond payoffs. This framework allows us to move beyond reduced-form patterns to study how investor composition shapes bond yields, default risk, and optimal debt issuance.

More generally, our analysis contributes to a long-standing literature that uses index rebalancings as supply-shifting shocks to estimate asset price reactions and demand elasticities (Harris and Gurel, 1986; Shleifer, 1986; Greenwood, 2005; Hau, Massa, and Peress, 2010; Chang, Hong, and Liskovich, 2014; Raddatz, Schmukler, and Williams, 2017; Pandolfi and Williams, 2019; Pavlova and Sikorskaya, 2022; Beltran and Chang, 2024). A key contribution of our work is to show that part of the observed price movement reflects endogenous changes in expected payoffs—not just an inelastic demand. This decomposition is absent from existing studies, yet it has important quantitative implications, particularly for long-term securities. Furthermore, our indirect inference approach is general and can be applied to other assets and markets.

The paper is structured as follows. Section 2 presents a sovereign debt model with endogenous default and price-elastic investors. Section 3 outlines the empirical strategy and

main estimates. Section 4 reports the quantitative results. Section 5 concludes.

2 A Sovereign Debt Model with Price-elastic Investors

We formulate a quantitative sovereign debt model with price-elastic investors. On the issuer side, the setup follows a standard framework, featuring an endowment economy in which the government issues long-term debt in international markets. The government lacks commitment and can default on its debt, triggering temporary exclusion from financial markets and output losses. On the investor side, we introduce a rich demand structure with a downward-sloping curve and heterogeneous investors subject to asset-allocation mandates.

2.1 Model Setup

We consider a small open economy with incomplete markets in which an infinitely lived government issues long-term bonds in international markets. Output in this economy, y , is exogenous and follows a continuous Markov process with transition function $f_y(y_{t+1} | y_t)$. The government chooses debt to maximize the utility of a risk-averse representative household, which is given by

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

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

Let B_{t-1} denote the beginning-of-period stock of government bonds. Each bond unit matures in the next period with probability λ . If the bond does not mature, it pays a coupon ν . The government lacks commitment and can default on its debt. Let $d_t = \{0, 1\}$ denote the default policy, where $d = 1$ indicates a default. Default leads to a temporary exclusion from international debt markets and an exogenous output loss, $\phi(y_t)$. The government is benevolent and chooses $\{d_t, B_t\}$ to maximize Equation (1), subject to the economy's resource constraint.

If the government is currently not in default and given a choice of debt, the resource constraint of the economy is

$$c_t = y_t + q_t \times (B_t - (1 - \lambda)B_{t-1}) - (\lambda + (1 - \lambda)\nu) B_{t-1}, \quad (2)$$

where q_t denotes the bond price, which is endogenous and depends on the government's policies and the underlying assumptions on the demand side. The term $B_t - (1 - \lambda)B_{t-1}$ captures new bond issuance, while $(\lambda + (1 - \lambda)\nu) B_{t-1}$ represents current debt services. If the government

is in default, the resource constraint simplifies to $c_t = y_t - \phi(y_t)$.

Bonds are priced by global investors operating in international markets. We assume that the sovereign bonds of our small open economy are part of a benchmark index \mathcal{I} , which includes bonds from other countries. We consider two types of investors, active and passive, who differ on how they allocate funds towards these bonds. Active investors are price elastic: for a given expected payoff, their demand decreases with the quantity of bonds issued. In contrast, passive investors are perfectly inelastic, as they simply replicate the composition of the index \mathcal{I} . We next describe the investor block in detail.

2.2 Investor Block

Let $j \in \{1, \dots, J\}$ denote an individual investor. We define $x_{j,t}^i = \frac{q_t^i B_{j,t}^i}{W_{j,t}}$ as the share of wealth that investor j allocates to bond (or equivalently, country) i at time t , where q_t^i is the unit price of bond i , $B_{j,t}^i$ the investor's holdings, and $W_{j,t}$ their total wealth. We assume that bond i is included in the benchmark bond index \mathcal{I} , and let $\mathbf{w}_t = \{w_t^1, \dots, w_t^N\}$ denote the vector of time-varying index weights across all N constituent bonds.

Following [Gabaix and Koijen \(2021\)](#), we assume that $x_{j,t}^i$ is determined by the following mandate:

$$x_{j,t}^i = \theta_{j,t} \xi_{j,t}^i e^{\Lambda_j \pi_t^i} + (1 - \theta_{j,t}) w_t^i, \quad (3)$$

where $\theta_{j,t}$ governs the degree of activeness. Purely passive investors are characterized by $\theta_{j,t} = 0$, meaning their portfolios replicate the benchmark index \mathcal{I} . Conversely, $\theta_{j,t} \in (0, 1]$ corresponds to active or semi-active investors. Within the active share, investors allocate a fraction $\xi_{j,t}^i$ of their wealth to bond i , adjusted by a price-elastic component $e^{\Lambda_j \pi_t^i}$, where π_t^i is a function of the expected excess return r_{t+1}^i . For example, if $\pi_t^i = \mathbb{E}_t(r_{t+1}^i)$, investors tilt their portfolios toward bonds with higher expected excess returns. The parameter $\Lambda_j > 0$ captures how sensitive investor j is to changes in π_t^i , and thus governs their demand elasticity.

The simple mandate in Equation (3) allows us to introduce an aggregate demand elasticity for bond i that can be parameterized by $\mathbf{\Lambda} \equiv \{\Lambda_1, \dots, \Lambda_J\}$. While this mandate can have different microfoundations (as shown in [Appendix A](#)), we take it as given for our analysis.

Let $\mathcal{W}_t = \sum_j W_{j,t}(1 - \theta_{j,t})$ denote the assets under management of all passive investors tracking the index \mathcal{I} . The passive demand for country i is then given by $\mathcal{T}_t^i \equiv \frac{1}{q_t^i} \mathcal{W}_t w_t^i$. 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 $w_t^i = \frac{\alpha_t^i q_t^i B_t^i}{MV_t}$. Combining these expressions, the passive demand can be written as $\mathcal{T}_t^i = \tau_t^i B_t^i$, where $\tau_t^i \equiv \frac{\mathcal{W}_t}{MV_t} \alpha_t^i$ is

a time-varying share that determines how much of each unit of country i 's debt is held by passive investors.

As for active demand, we let $\xi_{j,t}^i = \xi_j^i q_t^i$, so that ξ_j^i captures the fixed component of investor j 's mandate. This specification yields the following active demand for investor j : $\mathcal{A}_{j,t}^i = \overline{\mathcal{A}}_{j,t}^i e^{\Lambda_j^i \pi_t^i}$, where $\overline{\mathcal{A}}_{j,t}^i \equiv W_{j,t} \theta_{j,t} \xi_j^i$ represents investor j 's wealth allocated to country i , as implied by their activeness and mandate. Let $\overline{\mathcal{A}}_t^i \equiv \sum_j \overline{\mathcal{A}}_{j,t}^i$ denote total targeted wealth for country i .

To derive a closed-form expression for the bond price, we proceed in two steps. First, we linearize each investor j 's active demand around $\pi_t^i = 0$ and a reference allocation $\overline{\mathcal{A}}_{j,t}^i = \overline{\mathcal{A}}_j^{i*}$, which yields $\mathcal{A}_t^i \approx \overline{\mathcal{A}}_t^i + \Lambda^i \pi_t^i$. Here, Λ^i denotes the sensitivity of demand for bond i with respect to changes in π_t^i . It is given by $\Lambda^i \equiv \overline{\mathcal{A}}^{i*} \sum_j s_j^{i*} \Lambda_j^i$, where $s_j^{i*} \equiv \overline{\mathcal{A}}_j^{i*} / \overline{\mathcal{A}}^{i*}$ is the share of investor j in the reference allocation, with $\overline{\mathcal{A}}^{i*} \equiv \sum_j \overline{\mathcal{A}}_j^{i*}$.

Second, we assume the functional form $\pi_t^i(r_{t+1}^i) = \frac{\mathbb{E}_t(r_{t+1}^i)}{\mathbb{V}_t(r_{t+1}^i)}$ so that active investors respond to Sharpe ratios, allocating a larger share of their wealth to bonds with higher expected excess returns relative to risk.⁴ This assumption is for tractability and does not rule out consideration of cross-bond covariances, which we interpret as embedded in the fixed mandates ξ_j^i —for instance, higher baseline allocations to bonds with better hedging properties.⁵

Given these assumptions, it is straightforward to derive the following pricing equation (see Appendix A for details):

$$q_t^i = \frac{\mathbb{E}_t(\mathcal{R}_{t+1}^i)}{r^f} - \frac{\kappa^i}{r^f} \mathbb{V}_t(\mathcal{R}_{t+1}^i) (B_t^i - \mathcal{T}_t^i - \overline{\mathcal{A}}_t^i), \quad (4)$$

where $\kappa^i \equiv 1/\Lambda^i$ and \mathcal{R}_{t+1}^i denotes the next-period repayment per unit of the bond. That is, $r_{t+1}^i \equiv \frac{\mathcal{R}_{t+1}^i}{q_t^i} - r^f$, where r^f denotes the gross risk-free rate at which investors discount payoffs. For now, we treat \mathcal{R}_{t+1}^i as given—we endogenize it in the next section.

The first term in Equation (4) captures bond i 's discounted expected repayment, which would correspond to the bond price under a perfectly elastic demand. The second term follows from the asset-allocation mandate and captures both investor composition and the downward-sloping component of active demand. Specifically, $B_t^i - \mathcal{T}_t^i$ is the bond supply available to active investors, while $\overline{\mathcal{A}}_t^i$ reflects the inelastic portion of their demand. The difference $B_t^i - \mathcal{T}_t^i - \overline{\mathcal{A}}_t^i$ thus represents the price-elastic segment. The slope of the demand is given by $\mathcal{S}_t^i \equiv -\frac{\kappa^i}{r^f} \mathbb{V}_t(\mathcal{R}_{t+1}^i)$. When $\kappa^i = 0$, demand is perfectly elastic and the bond price reflects only expected repayment. When $\kappa^i > 0$, the demand curve slopes downward,

⁴This specification follows Gabaix and Koijen (2021), where the $\pi_t^i(\cdot)$ function depends on expected excess returns and perceived risk.

⁵Appendix A provides a microfoundation that captures this.

and its slope becomes steeper as the variance of repayments increases. For risk-free bonds, Equation (4) implies a perfectly elastic demand. In Section 3, we show that this specification is consistent with our empirical findings.

2.3 Recursive Formulation

We focus on a Recursive Markov Equilibrium (RME) and represent the infinite-horizon decision problem of the government as a recursive dynamic programming problem (see Appendix A.3 for the equilibrium definition).

To obtain a recursive formulation, we let $\mathcal{T}' = \mathcal{T}(\tau, B')$, where B' denotes the end-of-period stock of government bonds and τ is the share of bonds included in the benchmark index. For the remainder of the analysis, we omit the superscript i . We assume that τ is time-varying and follows a continuous Markov process with transition function $f_\tau(\tau' | \tau)$. By market clearing, the end-of-period active demand is given by $\mathcal{A}' = \mathcal{A}(\tau, B') = B' - \mathcal{T}(\tau, B')$. We further assume that investors' wealth and the ξ component of their active mandates are constant, implying a fixed $\overline{\mathcal{A}}$.

The state space can be summarized by the n -tuple (h, B, s) , where h captures the government's current default status, B is the beginning-of-period stock of debt, and $s = (y, \tau)$ are the exogenous states. For a given default status h and choice of B' , the resource constraint of the economy can be written as

$$\begin{aligned} c(h = 0, B, y, \tau; B') &= y + q(y, \tau, B') (B' - (1 - \lambda)B) - (\lambda + (1 - \lambda)\nu) B, \\ c(h = 1, y) &= y - \phi(y), \end{aligned} \quad (5)$$

where $q(y, \tau, B')$ denotes the price of a unit of debt, $B' - (1 - \lambda)B$ are new bond issuances, and $(\lambda + (1 - \lambda)\nu) B$ are current debt services.

If the government is not in default, its value function is given by

$$V(y, \tau, B) = \max_{d \in \{0, 1\}} \left\{ V^r(y, \tau, B), V^d(y) \right\}, \quad (6)$$

where $V^r(\cdot)$ denotes the value function in case of repayment and $V^d(\cdot)$ denotes the default value. If the government chooses to repay, its value function is given by the following Bellman equation:

$$\begin{aligned} V^r(y, \tau, B) &= \max_{B'} u(c) + \beta \mathbb{E}_{s'|s} V(y', \tau', B'), \\ \text{subject to } c &= y + q(y, \tau, B') (B' - (1 - \lambda)B) - (\lambda + (1 - \lambda)\nu) B. \end{aligned} \quad (7)$$

While in default, the country is excluded from debt markets and cannot issue new debt.

The government exits default with probability θ , with no recovery value. The value function in case of default is given by

$$V^d(y) = u(y - \phi(y)) + \beta \mathbb{E}_{s'|s} \left[\theta V(y', \tau', 0) + (1 - \theta) V^d(y') \right]. \quad (8)$$

Given these assumptions and the demand structure in Section 2.2, we can rewrite the bond price in Equation (4) recursively as a function of B' and the exogenous states $\{y, \tau\}$:

$$q(y, \tau, B') = \frac{1}{r^f} \mathbb{E}_{s'|s} \mathcal{R}(y', \tau', B') + \mathcal{S}(y, \tau, B') (B' - \mathcal{T}(\tau, B') - \bar{\mathcal{A}}), \quad (9)$$

where $\mathcal{R}(y', \tau', B')$ denotes the next-period repayment function, and the last term in the equation, $\mathcal{S}(y, \tau, B') (B' - \mathcal{T}(\tau, B') - \bar{\mathcal{A}})$, captures the downward-sloping component of active demand.

The next-period repayment function, in turn, is given by

$$\mathcal{R}(y', \tau', B') = [1 - d(y', \tau', B')] [\lambda + (1 - \lambda) (\nu + q(y', \tau', B''))], \quad (10)$$

where $d(y', \tau', B')$ is the next-period default choice and $q(y', \tau', B'')$ denotes the next-period bond price, which is a function of next-period exogenous states, $\{y', \tau'\}$, and the next-period debt, $B'' \equiv B'(y', \tau', B')$. Lastly, the term characterizing the downward-sloping component of the demand is

$$\mathcal{S}(y, \tau, B'; \kappa) \equiv -\frac{1}{r^f} \kappa \mathbb{V}_{s'|s} \mathcal{R}(y', \tau', B'), \quad (11)$$

which increases with the volatility of payments. For the decomposition in our quantitative analysis, it is useful to express the bond price as $q(\cdot) = \frac{1}{r^f} \mathbb{E}_{s'|s} \mathcal{R}(y', \tau, B') \Psi(\cdot)$, where

$$\Psi(y, \tau, B') \equiv 1 - \kappa \frac{\mathbb{V}_{s'|s} \mathcal{R}(y', \tau', B')}{\mathbb{E}_{s'|s} \mathcal{R}(y', \tau', B')} (B' - \mathcal{T}(\tau, B') - \bar{\mathcal{A}}). \quad (12)$$

A value of $\Psi(y, \tau, B') < 1$ indicates that investors price the bond below its expected repayment value.

Based on this recursive formulation, the inverse elasticity for active investors is

$$\eta(y, \tau, B'; \kappa) = \mathcal{S}(y, \tau, B'; \kappa) \frac{\mathcal{A}(\tau, B')}{q(y, \tau, B')}. \quad (13)$$

As in standard sovereign debt models, Equations (9)-(10) show that the bond price declines with the expected default probability. In particular, for a given y , a higher B' (weakly) raises default risk, implying that $q(y, \tau, B')$ is (weakly) decreasing in B' . The term $\mathcal{S}(y, \tau, B')$ introduces an additional channel through which the bond price falls with B' : the price-elastic component of active investors. Changes in investor composition or increases in passive demand—as captured by τ —affect bond prices only if the active demand is not

perfectly elastic. If it were, shifts in the passive demand would leave bond prices, and thus government policy, unchanged.

2.4 Illustration: The Effects of a Downward-sloping Demand

When choosing its optimal debt policy, the government internalizes the effect of B' on the bond price $q(y, \tau, B')$, accounting for both changes in expected repayment and the downward-sloping demand component. Let $\epsilon(y, \tau, B'; \kappa) \equiv \frac{\partial \log q(y, \tau, B')}{\partial \log B'}$ denote the (inverse) supply elasticity. To simplify the analysis in this section, we impose two assumptions. First, additional debt does not affect passive demand, so that $\frac{\partial \mathcal{T}(\tau, B')}{\partial B'} = 0$. Second, we consider the case where $B' - \mathcal{T}(\tau, B') - \bar{A} \approx 0$, which implies that changes in κ do not directly influence bond prices. Under these assumptions, the supply elasticity can be written as:

$$\epsilon(y, \tau, B'; \kappa) \approx \left[\frac{1}{rf} \frac{\partial \mathbb{E}_{s'|s} \mathcal{R}(y', \tau', B')}{\partial B'} + \mathcal{S}(y, \tau, B'; \kappa) \right] \frac{B'}{q(y, \tau, B')} < 0, \quad (14)$$

The first term on the right-hand side captures how expected repayment changes with bond supply. This elasticity is (weakly) negative: a higher B' increases default risk, reducing expected repayments. The second term—also weakly negative—captures the additional price decline due to the slope of the demand curve.

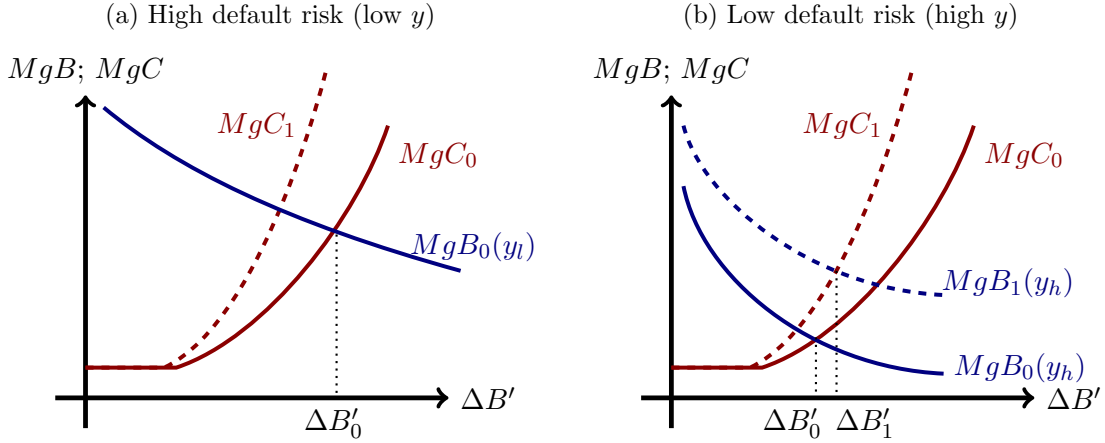
Under further simplifying assumptions (see Appendix A.4 for a full derivation), the first-order condition for B' can be expressed as:

$$u_c(c) q(y, \tau, B') = \beta \mathbb{E}_{s'|s} [u_c(c') \mathcal{R}'(\cdot)] - u_c(c) \left[\frac{1}{rf} \frac{\partial \mathbb{E}_{s'|s} \mathcal{R}'(\cdot)}{\partial B'} + \mathcal{S}(\cdot; \kappa) \right] \Delta B', \quad (15)$$

where $u_c(\cdot) \equiv \frac{\partial u(\cdot)}{\partial c}$ is the marginal utility of consumption, $\mathcal{S}(\cdot; \kappa) \equiv \mathcal{S}(y, \tau, B'; \kappa)$ is the slope of active demand, and $\Delta B' \equiv B' - (1 - \lambda)B$ denotes new issuance. The left-hand side of Equation (15) represents the marginal benefit of an additional unit of debt. It depends on the proceeds from issuing one more bond, $q(\cdot)$, and on how much the household values the resulting increase in consumption. The right-hand side is the marginal cost, consisting of two components. The first captures the value of the additional repayment the government must make next period. The second term reflects how an additional issuance affects bond prices, through both expected repayment and the slope of the demand curve.

Figure 1 illustrates how changes in κ influence debt issuance. Panel (a) shows a high-default-risk (low-endowment) case, and Panel (b) a low-risk (high-endowment) case. In both cases, a higher κ implies a steeper demand curve and thus higher marginal costs. For any given marginal benefit, this results in lower issuance. The effect is stronger when default risk is high, as the demand is less elastic. We refer to this mechanism as a “disciplining device,”

Figure 1: Effects of a lower demand elasticity



Note: The figure illustrates the effects of a higher slope parameter κ on the government's marginal benefits and costs of issuing an additional unit of debt. Panel (a) shows a case in which default risk is high (low endowment) and thus the demand is less elastic. Panel (b) depicts a scenario with low default risk (high endowment). In Panel (a), the $MgB_0(y_l)$ curve lies above the $MgB_0(y_h)$ of Panel (b) because marginal utility is higher when endowment is low.

since it restrains borrowing precisely when repayment risk is high.

By reducing future bond issuance—especially in high-risk periods—a higher κ raises current bond prices. The mechanism operates through the classic debt dilution channel: fewer expected future issuances lower future default risk, which in turn raises prices today. As shown in Panel (b), this leads to an upward shift in the marginal benefit curve.⁶ If the price response is strong enough, the government may even optimally issue *more* debt despite facing a steeper demand curve (from $\Delta B'_0$ to $\Delta B'_1$ in the graph). Whether this occurs is a quantitative question, which we return to in Section 4.

3 Empirical Analysis

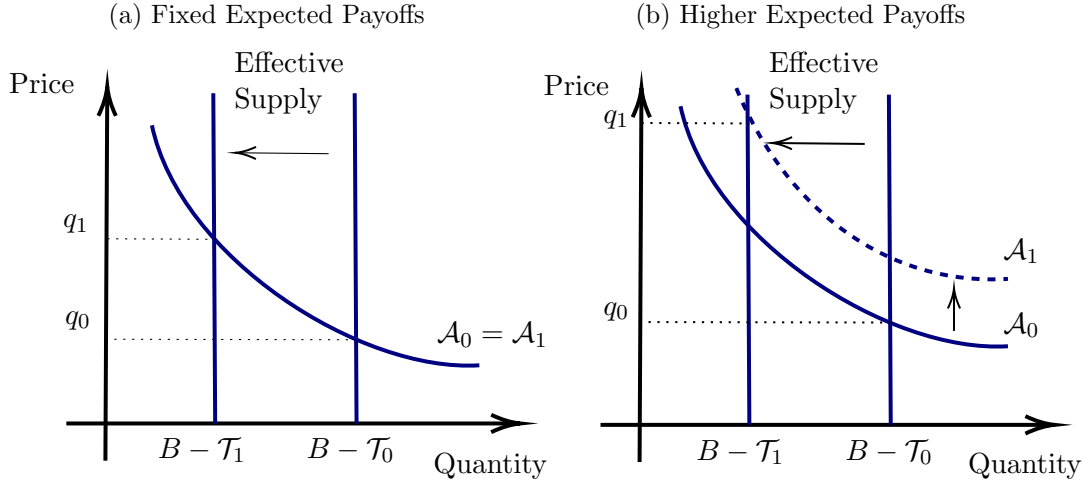
3.1 Passive Demand Shocks and the Slope of Active Demand

To solve the model and quantify the implications of a downward-sloping demand curve, we first need an estimate of κ —the structural parameter governing the elasticity of active demand.

Based on the model pricing equations—Equations (9)–(11)—and holding $\{y, B'\}$ and repayment fixed, bond price reactions to shocks in passive demand, $\Delta T' \equiv T(\tau', B') - T(\tau, B')$, are informative about the slope of active investors' demand. Panel (a) of Figure 2 illustrates this point. For a given total supply, an increase in passive demand reduces the amount of bonds available to active investors—a leftward shift in the “effective” supply they face. These

⁶A similar shift occurs in Panel (a), omitted for simplicity.

Figure 2: Index rebalancing and the demand elasticity



Note: The figure depicts a decrease in the effective supply available to active investors driven by an increase in the passive demand \mathcal{T} . Panel (a) considers a case in which expected payoffs do not change as a consequence of the lower effective supply. Panel (b) shows a case in which the expected payoffs increase.

shifts thus provide a natural basis for identifying the slope parameter κ .

This strategy—while common in the literature—is not sufficient to identify κ and the structural elasticity $\eta(y, \tau, B'; \kappa)$ in Equation (13), as it ignores the potential effect of the $\Delta \mathcal{T}'$ shock on bond payoffs, $\mathcal{R}(y', \tau', B')$.⁷ To illustrate this more clearly, suppose we compute the following reduced-form inverse elasticity based on bond price changes in response to exogenous shifts in passive demand:

$$\hat{\eta}(y, \tau, B'; \kappa) \equiv - \frac{\Delta \log q(y, \tau, B')}{\Delta \mathcal{T}' / (B' - \mathcal{T}(\tau, B'))}, \quad (16)$$

where the change in passive demand is normalized by the size of active demand.

After some algebra and under mild assumptions—see Appendix A.5 for details—, one can decompose the previous expression in two components, the structural elasticity and changes in bond prices driven by changes in expected future payoffs. That is,

$$\eta(y, \tau, B'; \kappa) \approx \hat{\eta}(y, \tau, B'; \kappa) \times \left(1 - \frac{\frac{1}{r_f} \Delta \mathbb{E}_{s'|s} \mathcal{R}'(\cdot)}{\Delta q(\cdot)} \right), \quad (17)$$

where $\Delta \mathbb{E}_{s'|s} \mathcal{R}'(\cdot)$ and $\Delta q(\cdot)$ are the changes in expected-repayment and bond prices for a given $\Delta \mathcal{T}'$. If a larger passive demand leads to higher expected payoffs and bond prices, then the reduced-form estimate $\hat{\eta}(y, \tau, B'; \kappa)$ would overstate the magnitude of the structural elasticity, $\eta(y, \tau, B'; \kappa)$. Panel (b) of Figure 2 illustrates this case.

⁷There is a large literature that uses index rebalancings as supply-shifting shocks to estimate asset price reactions and demand elasticities; see, for example, 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 Chang (2024).

The rest of our analysis proceeds as follows. In this section, we begin by constructing an instrument for $\Delta\mathcal{T}'$ using index rebalancings. We then use it to estimate a reduced-form elasticity, $\hat{\eta}$, based on high-frequency bond price reactions around these events. In Section 4, we extend our baseline model to replicate this same exercise within the model framework and identify $\eta(\cdot)$ by indirect inference. That is, we calibrate κ so that the model matches the estimated reduced-form elasticity, $\hat{\eta}(\cdot)$, and use the resulting structure to isolate the component driven by the unobserved term $\Delta\mathbb{E}_{s'|s}\mathcal{R}'(\cdot)$.

3.2 Flows Implied by Index Rebalancings

We construct a measure of shocks to the passive demand based on changes in the composition of a benchmark bond index. To this end, we exploit end-of-month rebalancings in the J.P. Morgan EMBIGD to identify exogenous shifts in the supply of bonds available to active investors. The EMBIGD tracks the performance of emerging market sovereign and quasi-sovereign U.S. dollar-denominated bonds issued in international markets.⁸ Among bond indexes for emerging economies, the EMBIGD is the most widely used benchmark and was tracked by funds with combined assets under management (AUM) of around US\$300 billion in 2018 (Appendix Figure B1). Unlike most indices that use a traditional market capitalization-based weighting scheme, the EMBIGD limits the weights of countries with above-average debt outstanding (relative to other countries in the index) by including only a fraction of their face amount in the index—referred to as the *diversified face amount*. Its goal is to achieve greater diversification by lowering the index weight of large countries, and we refer to it as a cap rule.⁹

Rebalancings in the EMBIGD index occur on the last business day of each month in the United States and are triggered by bond inclusions and exclusions. J.P. Morgan announces these updates through a report detailing the updated index composition. Upon the announcement, passive investors mimicking the index composition need to adjust their portfolios by buying or selling bonds to match the new index weights.

We construct a measure of flows implied by index rebalancings (FIR) for each country c on each rebalancing date. The FIR quantifies the amount of funds that, on a given rebalancing date, enter or leave a country's bonds due to the rebalancing in the portfolio of passive

⁸To 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.

⁹In comparison, the J.P. Morgan Emerging Markets Bond Index Global (EMBIG) has the same bond inclusion criteria as the EMBIGD but it adopts a standard market-capitalization weighting scheme, without caps.

investors tracking the EMBIGD index. It is defined as:

$$\text{FIR}_{c,t} \equiv \frac{\Delta \tilde{\mathcal{T}}_{c,t}}{q_{c,t-1}B_{c,t-1} - w_{c,t-1}\mathcal{W}_{t-1}}. \quad (18)$$

The numerator, $\Delta \tilde{\mathcal{T}}_{c,t} \equiv (w_{c,t} - w_{c,t}^{BH})\mathcal{W}_t$, captures the shift in passive demand resulting from the rebalancing. Here, \mathcal{W}_t denotes the AUM of funds that passively track the EMBIGD, $w_{c,t}$ is the new index weight for country c after the rebalancing, and $w_{c,t}^{BH}$ denotes the “buy-and-hold” weight. The latter reflects the end-of-month weight of country c in passive fund portfolios just before the rebalancing. The denominator normalizes the flow by the market value of bonds available to active investors, computed as $q_{c,t-1}B_{c,t-1} - w_{c,t-1}\mathcal{W}_{t-1}$.

Index weights are given by $w_{c,t} \equiv \frac{q_{c,t}B_{c,t}f_{c,t}}{MV_t}$, where $q_{c,t}$ is the bond price, $B_{c,t}$ is the face amount, and $f_{c,t} \in (0, 1]$ is a diversification coefficient determined by the EMBIGD’s cap rule.¹⁰ The product $f_{c,t}B_{c,t}$ defines the diversified face amount (DFA). The term $MV_t \equiv q_t^I I_t^I$ denotes the market value of the EMBIGD, where q_t^I denotes its price and I_t^I the number of units. Finally, the “buy-and-hold” weight is calculated as the lagged weight adjusted for relative price movements: $w_{c,t}^{BH} \equiv w_{c,t-1} \frac{q_{c,t}/q_{c,t-1}}{q_t^I/q_{t-1}^I}$. If there are no inclusions or exclusions from the index, weights adjust only due to price changes. In that case, $w_{c,t}^{BH} = w_{c,t}$, and thus $\text{FIR}_{c,t} = 0$.

As defined, a 1 p.p. FIR corresponds to a 1% reduction in the supply of a country’s bonds available to active investors at the time of the rebalancing. Since it reflects changes in index composition, however, the FIR is endogenous and may correlate with country fundamentals. First, it is affected by bond issuances: when a country issues new bonds that enter the index—or redeems existing ones—its index weight shifts, altering the FIR. Second, because index weights depend on bond prices, the FIR is mechanically linked to price movements.

Our empirical strategy addresses these endogeneity concerns in two ways. First, we focus on countries whose outstanding bond 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 a specific feature of the EMBIGD’s weighting scheme: a cap rule that introduces a diversification coefficient, $f_{c,t}$, which limits a country’s index weight based on the face value of its bonds, independently of market prices.

Each month, J.P. Morgan calculates the *Index Country Average* (ICA), defined as the

¹⁰To simplify our notation, we assume that each country has a single bond in the index. In practice, J.P. Morgan aggregates the market values of all eligible bonds issued by a given country and applies an analogous formula to that in Equation (18). What matters for the analysis is that weights are ultimately constructed at the country level.

average face amount of bonds across all index-eligible countries. For countries with a face amount below the ICA, $f_{c,t} = 1$. For those above the ICA, $f_{c,t} < 1$, with the exact value depending on how far the country is from the ICA—see Equation (B1) in Appendix B for the specific formula used by J.P. Morgan. Based on this coefficient, the diversified face amount is defined as $DFA_{c,t} \equiv f_{c,t}B_{c,t}$.

Using this cap rule, we construct a synthetic index in which country weights are only a function of the diversified face amount, $\tilde{w}_{c,t} \equiv \frac{f_{c,t}B_{c,t}}{\sum_{n \in \mathcal{I}} f_{n,t}B_{n,t}}$. We then instrument the FIR based on the relative changes in these synthetic weights, $Z_{c,t} \equiv \frac{\Delta \tilde{w}_{c,t}}{\tilde{w}_{c,t-1}}$. Given our focus on countries whose debt outstanding in the index remains unchanged ($B_{c,t} = B_{c,t-1}$), this instrument simplifies to:

$$Z_{c,t} \equiv \left(\frac{f_{c,t}}{\sum_{n \in \mathcal{I}} f_{n,t}B_{n,t}} - \frac{f_{c,t-1}}{\sum_{n \in \mathcal{I}} f_{n,t-1}B_{n,t-1}} \right) / \left(\frac{f_{c,t-1}}{\sum_{n \in \mathcal{I}} f_{n,t-1}B_{n,t-1}} \right). \quad (19)$$

The $Z_{c,t}$ instrument captures two sources of variation. First, it accounts for fluctuations *across* rebalancing events arising from changes in the diversified face amount of *other* countries in the index. For instance, when a newly issued bond from country n is added to the index, $Z_{c,t}$ decreases proportionally for every country c . Second, the EMBIGD cap rule allows us to exploit variation across countries *within* each rebalancing event. This variation arises because, for any issuance or redemption from some country n , the $f_{c,t}$ coefficient adjusts differently across countries, depending on whether the face value of country c 's bonds is above or below the average face value in the index. We consider next a simple example to illustrate these mechanisms.

The Cap Rule: A Simple Example

To illustrate how rebalancings and the cap rule operate, we consider an example with five countries, $c = \{A, B, C, D, E\}$. Each of these countries has only one qualifying bond and does not issue or redeem bonds during month t . We assume that another country, F , issues an eligible bond during month t , which is included in the index at the rebalancing date at the end of that month. Table 1 shows the assumed face amounts FA_c , the corresponding diversified face amounts $DFA_{c,t}$ —computed based on Equation (B1)—and the synthetic weights $\tilde{w}_{c,t}$, both before and after rebalancing in period t . In this example, countries D and E are capped in both $t-1$ and t (since their face values exceed the ICA), while country F becomes capped in t .

Table 1 illustrates the sources of variation we exploit in our analysis. The issuance of a new bond by country F and its inclusion in the index reduce the weights of all other countries.

Table 1: The cap rule

Country	Before Rebalancing			After Rebalancing			Instrument $Z_{c,t} \equiv \frac{\Delta \tilde{w}_{c,t}}{\tilde{w}_{c,t-1}}$
	$FA_{c,t-1}$	$DFA_{c,t-1}$	$\tilde{w}_{c,t-1}$	$FA_{c,t}$	$DFA_{c,t}$	$\tilde{w}_{c,t}$	
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%	-
<i>ICA</i>	5,000			5,500			
<i>FA_{max}</i>	12,000			12,000			

Note: The table presents a simple example to illustrate the cap rule mechanism. Columns FA_c list each country's face amount before and after rebalancing, while DFA_c reports the diversified (capped) amount. \tilde{w}_c are the resulting synthetic index weights and Z_c is our instrument. *ICA* denotes the average face amount across index-eligible countries; when $FA_c > ICA$, the cap binds and $DFA_c < FA_c$. FA_{\max} is the largest face amount in the sample.

Absent the cap rule, this reduction would be proportional across all remaining countries. The cap rule, however, introduces heterogeneity in these relative changes. This additional variation arises because diversified face amounts are capped for countries with above-average outstanding bonds, and the cap itself adjusts following the inclusion of country *F*. In particular, the resulting increase in the *ICA* relaxes the cap for countries *D* and *E*, dampening the relative decrease in their index weights after the introduction of *F*.

3.3 Data and Summary Statistics

We collect data from three main sources: Datastream, J.P. Morgan, and Morningstar. Datastream provides daily bond price data for securities included in the EMBIGD. From J.P. Morgan, we obtain index weights, bond face amounts, and characteristics such as maturity and duration. Morningstar supplies asset holdings of funds benchmarked against the EMBIGD, which we use to construct our measure of AUM that passively track the index. The final dataset contains 131,820 bond-time observations, covering 751 bonds from 68 countries over the 2016–2018 period.

While J.P. Morgan reports the total AUM benchmarked against its indexes, it does not distinguish between passive and active funds.¹¹ To compute the AUM that passively track the EMBIGD, \mathcal{W}_t , we proceed as follows. First, we collect data from Morningstar on the holdings of funds benchmarked to the EMBIGD and the EMBI Global Core.¹² Second,

¹¹Even if such data were available, many active funds may manage a large share of their portfolios in a semi-passive way, as emphasized by Pavlova and Sikorskaya (2022).

¹²We include the EMBI Global Core due to its similarity to the EMBIGD. It follows the same diversification methodology and applies a comparable bond inclusion criterion. The main difference is that bonds must have a minimum face amount of US\$1 billion to be eligible.

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
Instrumented FIR (%)	-0.15	0.20	-0.32	0.00	-0.66	0.23

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. *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.

we compute each fund’s *Passive Share* = $100 - \text{Active Share}$, where *Active Share* is the measure developed by [Cremers and Petajisto \(2009\)](#) that quantifies how much a fund’s country-level holdings deviate from its benchmark’s weights. This allows us to identify, even for active funds, the portion of assets that behave passively.¹³ Third, we compute an aggregate *Passive Share* by weighting each fund’s passive share by its AUM, yielding a value of 50%.¹⁴ Lastly, we compute \mathcal{W}_t by multiplying the aggregate passive share by the total AUM benchmarked against the EMBIGD, as reported by J.P. Morgan.

Table 2 displays summary statistics for our main variables. The bonds in our sample have an average stripped spread of 278 basis points, and an average face amount of US\$1.3 billion. In terms of maturity, most of these bonds are medium and long-term bonds—the average maturity is around 10 years. Appendix Figure B5 shows a strong positive relationship between the residualized FIR and our instrument Z , with an R^2 of 86%.

3.4 Estimation Strategy and Results

We use our instrumented FIR and the schedule of rebalancing events to estimate a demand elasticity for active investors. Specifically, we examine daily bond-level price responses within 5-day symmetric windows centered on each rebalancing date.¹⁵

We implement an instrumented difference-in-differences design and estimate the following specification using two-stage least squares (2SLS):

$$\log(q_{i,t,h}) = \theta_{c(i),t} + \theta_{b(i),t} + \gamma \mathbb{1}_{h \in Post} + \beta(\widehat{FIR}_{c(i),t} \times \mathbb{1}_{h \in Post}) + \mathbf{X}_{i,t} + \varepsilon_{i,t,h}, \quad (20)$$

¹³We focus on country-level holdings, as this is the level at which the FIR is computed. Bonds are assigned to a country only if they are included in the EMBIGD.

¹⁴We also construct a bond-level active share, obtaining a value-weighted average of 72%. As a comparison, [Cremers and Petajisto \(2009\)](#) report an *Active Share* between 55% and 80%.

¹⁵Appendix Figure B6 presents a visual representation of a rebalancing event and the specific days included in the estimation window.

where $\widehat{FIR}_{c(i),t}$ denotes the flows implied by rebalancing into country c at event t , instrumented with $Z_{c,t}$ in the first stage. 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 last business day of each month. The indicator $\mathbb{1}_{h \in Post}$ equals 1 in the h days following the rebalancing and 0 in the days preceding it. We include country-month fixed effects $\theta_{c(i),t}$ and bond-type-month fixed effects $\theta_{b(i),t}$ to 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 β , which captures the price impact of changes in the supply of a country’s index-eligible bonds.

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- $\mathbb{1}_{h \in Post}$ fixed effects. These absorb all average differences between the pre- and post-windows for each rebalancing event—for instance, those driven by global shocks occurring around the rebalancing date, such as changes in risk-free rates or risk aversion. As a result, β is identified from cross-country variation in FIRs within the same event window.

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, β , is positive and statistically significant. The estimate in column (4), which includes bond-month and month-by- $\mathbb{1}_{h \in Post}$ fixed effects, implies that a 1 p.p. increase in the FIR raises bond prices by approximately 0.30%.¹⁶ In yield terms, this corresponds to a decline in yield to maturity of about 5.5 basis points—see Appendix Table B7.

The documented effects are heterogeneous across bonds with varying levels of default risk. 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 are more sensitive to rebalancing shocks, with a 1 p.p. increase in FIR associated with an almost 0.40% increase in bond prices (about a 10 basis point decline in yield to maturity; see Appendix Table B8).

¹⁶Appendix B 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. We also consider an alternative FIR measure that holds prices fixed at their levels from the previous rebalancing, yielding similar results. Our findings are robust to alternative event windows and to excluding quasi-sovereign bonds from the sample.

Table 3: Effects of FIR on bond prices

Dependent Variable: Log Price						
	Symmetric window: [-5:+5]				Excl. h=-1	
FIR X Post	0.231** (0.099)	0.232** (0.100)	0.231** (0.099)	0.300** (0.134)	0.263*** (0.098)	0.319** (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
F (FS)	654	1,616	1,666	476	1,670	476

Note: The table presents 2SLS estimates of log bond prices on the instrumented FIR measure around rebalancing events. *Post* is an indicator equal to 1 for the five trading days following each rebalancing, and 0 otherwise. *Bond Characteristics* refer to fixed effects formed by interacting maturity bins, credit rating grades, and bond type 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, and bond type indicators distinguish between sovereign and quasi-sovereign bonds. *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. Coefficients for *Post* and FIR are included in the regression but omitted from the table for brevity. Standard errors are clustered at the country-month level. The sample covers the 2016–2018 period. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

In contrast, the effect for low-spread bonds is smaller and not statistically significant. Overall, these findings suggest that investors demand a premium for holding riskier bonds, which we refer to as an *inconvenience yield*. These results align with our assumption that investors penalize return variance and are consistent with an elasticity that varies with default risk—as shown in Equations (11) and (13).

One potential concern with the previous analysis is that bonds with higher or lower FIRs may already be on different price trends before the rebalancing. To assess this possibility, we estimate a leads-and-lags regression where the instrumented FIR is interacted with trading-day dummies around each rebalancing event. This approach allows us to trace the dynamic response of bond prices. Specifically, we estimate the following regression via 2SLS:

$$\log(q_{i,t,h}) = \theta_{c(i),t} + \theta_{b(i),t} + \sum_{h \notin -2} \gamma_h \mathbb{1}_h + \sum_{h \notin -2} \beta_h (\widehat{FIR}_{c(i),t} \times \mathbb{1}_h) + \mathbf{X}_{i,t} + \varepsilon_{i,t,h}, \quad (21)$$

where $\mathbb{1}_h$ are indicators for each trading day in the $[-5, 5]$ window. The estimated β_h

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

Dependent Variable: Log Price						
	High Spread		Median Spread		Low Spread	
FIR X Post	0.380**	0.381**	0.325**	0.322**	0.087	0.087
	(0.166)	(0.165)	(0.152)	(0.151)	(0.098)	(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
F (FS)	501	2,342	436	720	0	882

Note: The table presents 2SLS estimates of log bond prices on the instrumented FIR measure 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. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

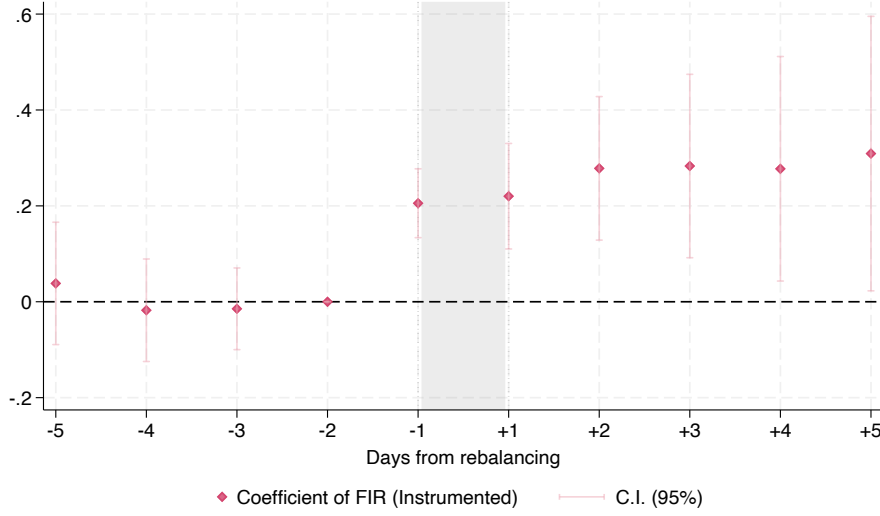
coefficients are plotted in Figure 3.

In the days leading up to the rebalancing, FIR changes are not systematically related to bond price movements. After the rebalancing, the coefficients turn positive and statistically significant, stabilizing just below 0.35 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.¹⁷ 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.

¹⁷This is consistent with the portfolio rebalancing patterns documented by Escobar, Pandolfi, Pedraza, and Williams (2021), 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: This figure presents leads and lags coefficients from a 2SLS estimation of log bond prices on trading-day dummies around each rebalancing event, using the same 2SLS procedure as in Table 3. The specification includes bond characteristics-month fixed effects (maturity, rating, and bond type). The shaded gray area indicates the rebalancing day, and $h = +1$ captures the price change on that day. Vertical red lines represent 95% confidence intervals. Standard errors are clustered at the country-month level.

Discussion of Results

As emphasized at the beginning of the section, the previous estimates do not necessarily capture a structural demand elasticity, since the supply-shifting shock may also affect bonds' expected future payoffs—and this, in turn, is reflected in current prices. This is consistent with our finding that price responses are stronger for longer-term bonds, whose prices are more sensitive to changes in future payoffs (Appendix Table B5). Additionally, we find that credit default swap spreads decline with the instrumented FIR (Appendix Table B6), further supporting the idea that expected bond payoffs may respond to the supply shock. While these results are suggestive, they underscore the need for a model to disentangle the different forces affecting prices—a task we take up in the next section.

Nonetheless, we can still map these estimates to a *reduced-form* demand elasticity for active investors, assuming that bond expected payoffs remain fixed. Based on our FIR definition, the inverse elasticity in Equation (16) can be written as $\hat{\eta}^i \equiv (-) \frac{\Delta \log(q_t^i)}{FIR_{c,t}}$, which corresponds to the β coefficient in Equation (20). Our estimates thus imply an inverse price elasticity of demand for active investors of approximately -0.30 (corresponding to a yield semi-elasticity of about 5.5 basis points).

Our reduced-form estimates are higher in magnitude than those for sovereign bonds in advanced economies (Jiang, Lustig, Van Nieuwerburgh, and Xiaolan, 2021; Krishnamurthy

and Vissing-Jorgensen, 2012; Greenwood et al., 2015; Mian et al., 2022), highlighting the relatively inelastic nature of EM bond markets. The estimates are similar to those for corporate bonds in advanced economies (Bretscher et al., 2022), which typically carry higher risk. These findings are consistent with our result that price responses to the instrumented FIR are larger for countries with higher default risk. Within EM government debt, our estimates imply a more elastic demand than Fang et al. (2025), which is expected given our focus on large and liquid EM bonds.¹⁸

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 relevant moments for external debt-to-output and spreads. In particular, we set the structural parameter behind the active demand elasticity, κ , to match our reduced-form estimates of Section 3.

In terms of functional forms and stochastic processes, we assume CRRA preferences: $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, where γ denotes the coefficient of risk aversion. Log output follows an AR(1) process: $\log(y') = \rho_y \log(y) + \epsilon'_y$, with $\epsilon'_y \sim N(0, \sigma_y)$. 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. Lastly, the passive demand is proportional to the (end-of-period) amount of bonds outstanding: $\mathcal{T}(\tau, B') = \tau B'$, where τ is stochastic and follows an AR(1) process, $\tau' = (1 - \rho_\tau)\tau^* + \rho_\tau \tau + \epsilon'_\tau$, where $\epsilon'_\tau \sim N(0, \sigma_\tau)$.

Table 5 lists the calibrated parameters. For the subset of fixed parameters (Panel a), we set $\gamma = 2$, a standard value for risk aversion in the literature. The quarterly risk-free rate is set to $r^f = 0.01$, consistent with the average real rate observed in the United States. The probability of re-entry into international markets is set to $\theta = 0.0385$, implying an average exclusion duration of 6.5 years. We choose $\lambda = 0.05$ to target a debt maturity of five years,

¹⁸The bond sample in Fang et al. (2025) includes a broader set of countries, loans, and local-currency bonds—assets typically less elastic due to their concentration among domestic investors (see Pandolfi and Williams, 2019).

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 rate	0.948
r^f	Risk-free interest rate	0.01	\bar{d}_0	Default cost—level	-0.224
λ	Debt maturity	0.05	\bar{d}_1	Default cost—curvature	0.27
ν	Debt services	0.03	κ	Slope parameter	70.0
θ	Reentry probability	0.0385	$\bar{\mathcal{A}}$	Active investors demand	0.4495
ρ_y	Output, autocorrelation	0.93			
σ_y	Output, shock volatility	0.02			
τ^*	Share of passive demand	0.123			
ρ_τ	FIR, autocorrelation	0.3			
σ_τ	FIR, shock volatility	0.02			

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

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 \(2023\)](#). Finally, we set τ^* to match the average share of bonds held by passive investors, and calibrate ρ_τ and σ_τ to reproduce the persistence and volatility of passive demand.¹⁹

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.²⁰

Regarding the parameters on the demand side, we calibrate κ to match the estimated (inverse) reduced-form demand elasticity, $\hat{\eta}(\cdot)$, derived from high-frequency price responses within a short window around rebalancing events. To capture the high-frequency nature of the empirical analysis, we extend our baseline model to include secondary markets, allowing for within-period trading. We assume two rounds of trading in secondary markets within each period, with the following timing:

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. Let $q^{SM,0}(y, \tau, B')$ denote the opening price.*

¹⁹We assume that the entire stock of debt is included in the \mathcal{I} index. This simplifies the analysis by reducing the dimensionality of the state space. For Argentina, approximately 80% of its external government debt is included in a bond index.

²⁰Our 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 encompasses the 2016–2018 window used in our empirical analysis. Spreads are significantly 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 (see [Morelli and Moretti, 2023](#))—we exclude them from the calibration.

Table 6: Targeted moments

Target	Description	Data	Model	Alternative elasticity values			
				$\kappa = 0$	$\kappa = 30$	$\kappa = 50$	$\kappa = 90$
$\bar{\mathbb{E}}(B/y)$	Debt to output (%)	53.0	53.1	51.0	52.3	52.8	53.2
$\bar{\mathbb{E}}(SP)$	Bond spreads (%)	5.08	5.06	8.59	6.47	5.61	4.67
$\bar{\sigma}(SP)$	Volatility of spreads (%)	1.43	1.56	4.49	2.65	1.96	1.35
$\bar{\mathbb{E}}(\hat{\eta})$	Reduced-form elasticity	-0.29	-0.28	-	-0.17	-0.23	-0.32

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 operators $\bar{\mathbb{E}}(\cdot)$ and $\bar{\sigma}(\cdot)$ denote the sample mean and standard deviation, respectively.

4. *Next-period index weights τ' are realized. Bond prices are updated.*
5. *The secondary market closes. Let $q^{SM,1}(y, \tau', B')$ denote the closing price.*

Appendix C provides the pricing functions for this extension. Importantly, in the absence of secondary markets, the timing remains identical to that of the baseline model. As a result, the proposed extension nests the baseline framework introduced in Section 2.

With this extension, we can compute bond price reactions to exogenous changes in index weights while holding debt supply and endowment fixed—mirroring the identification strategy used in our empirical analysis. The only difference between $q^{SM,1}$ and $q^{SM,0}$ arises from the update in the passive share τ . We simulate the model, generate a path for $\{y, \tau, B'\}$, and compute the same reduced-form inverse elasticity as in our empirical analysis:

$$\hat{\eta}(y, \tau, B'; \kappa) = - \frac{\Delta \log q}{\Delta \mathcal{T}' / (B' - \mathcal{T}(\tau, B'))}, \quad (22)$$

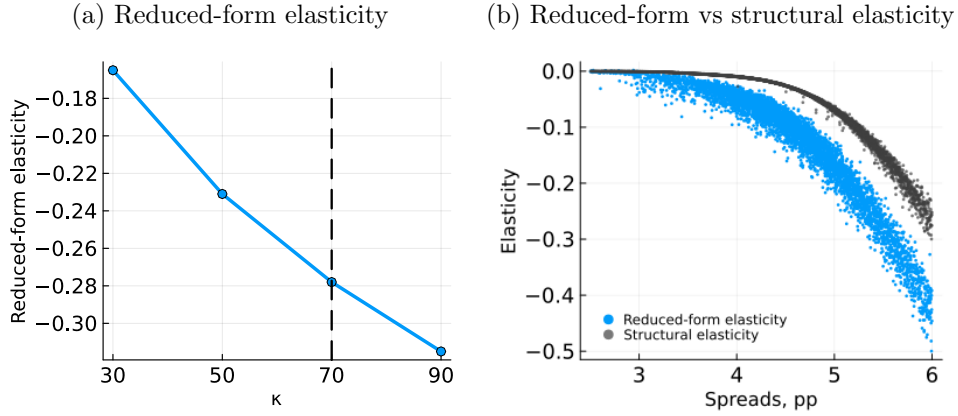
where $\Delta \mathcal{T}' \equiv \mathcal{T}(\tau', B') - \mathcal{T}(\tau, B')$ is the change in the passive demand and $\Delta \log q \equiv \log q^{SM,1}(y, \tau', B') - \log q^{SM,0}(y, \tau, B')$ is the log change in the bond price before and after the change in the passive share. Lastly, we normalize $\bar{\mathcal{A}}$ so that the average inconvenience yield, $\vartheta(\cdot)$, is zero.²¹ The introduction of inelastic investors thus affects only the sensitivity—but not the level—of the pricing kernel to changes in B' in the neighborhood of the average $\{y, \tau, B'\}$. In the next section, we analyze how our results change with different values of $\bar{\mathcal{A}}$.

Table 6 reports the set of targeted moments.²² The model successfully matches all four targeted moments. Panel (a) of Figure 4 shows that the parameter governing the slope of the demand curve, κ , is well identified, as it generates a monotonic relationship with the

²¹We compute the inconvenience yield, $\vartheta(\cdot)$, as follows. First, we calculate bond prices under a counterfactual scenario with perfectly elastic demand ($\kappa = 0$), holding fixed the policies from our baseline model. We then define $\vartheta(y, \tau, B')$ as the difference between the bond spread implied by the baseline model and that of the counterfactual: $\vartheta(y, \tau, B') \equiv SP(y, \tau, B') - SP(y, \tau, B'; \kappa = 0)$. A positive value indicates that, for a given level of default risk, investors require a higher spread to hold the bond.

²²Annualized spreads are computed as $SP(y, \tau, B') \equiv \left(\frac{1+i(y, \tau, B')}{1+r_f} \right)^4 - 1$, where $i(y, \tau, B')$ is the internal quarterly return rate, defined as the value that solves $q(y, \tau, B') = \frac{\lambda+(1-\lambda)(\nu+q(y, \tau, B'))}{\lambda+i(y, \tau, B')}$.

Figure 4: Decomposing the demand elasticity



Note: Panel (a) displays the average reduced-form inverse demand elasticity, $\hat{\eta}$, as a function of κ . Panel (b) shows a binscatter plot of annualized bond spreads against both the reduced-form elasticity (blue dots) and the model-implied elasticity (black dots).

Table 7: Role of shock persistence

Moment	Description	Baseline model	Alternative persistence values		
			Lower	Low	High
$\bar{\mathbb{E}}(\hat{\eta})$	Reduced-form elasticity	-0.28	-0.23	-0.25	-0.37
$\bar{\mathbb{E}}(\eta)$	Structural elasticity	-0.18	-0.18	-0.18	-0.18
$1 - \bar{\mathbb{E}}(\eta)/\bar{\mathbb{E}}(\hat{\eta})$	Bias (percent)	35%	23%	29%	51%

Note: The table compares the reduced-form inverse demand elasticity $\hat{\eta}$ with the model-implied counterpart η . The “Baseline” column reports elasticities under our baseline calibration. In the “Lower persistence” case, we set the persistence of the $\{\tau\}$ process to $\rho_\tau = 0.30$. The “Low persistence” column corresponds to $\rho_\tau = 0.50$, and the “Higher persistence” column to $\rho_\tau = 0.90$.

reduced-form elasticity.

Given the calibrated value of κ and the endogenous repayment function $\mathcal{R}(y', \tau', B')$, we can decompose the components underlying the reduced-form inverse elasticity $\hat{\eta}(y, \tau, B'; \kappa)$ and recover the structural demand elasticity $\eta(y, \tau, B'; \kappa)$, following Equation (17). Panel (b) of Figure 4 presents the results of this decomposition, showing a binscatter plot of simulated reduced-form and the model-implied elasticities as a function of bond spreads. The vertical distance between the two series captures the portion of the reduced-form elasticity attributable to endogenous variation in the repayment function. On average, we find that the model-implied elasticity accounts for about two-thirds of the reduced-form elasticity—first column of Table 7.

The difference between the structural and reduced-form elasticity 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 elasticity accounts for a smaller (larger) share of the reduced-form elasticity. In other words, greater persistence

in the τ process increases the portion of the total price response driven by endogenous changes in expected repayment. This is because more persistent shocks generate larger shifts in default risk and induce stronger adjustments in the government’s current and future issuance decisions. Appendix C.4 provides a more detailed analysis of these differences.

Overall, our analysis highlights the importance of accounting for issuers’ endogenous responses to exogenous supply shocks and the resulting changes in expected bond payoffs. Ignoring these factors can lead to substantial biases in estimated demand elasticities, particularly when the underlying supply shock is persistent.

4.2 Implications of a Downward-sloping Demand

Next, we analyze and quantify the effects of a downward-sloping demand on bond prices, default risk, and government policies. To do so, we compare our baseline model with alternative calibrations that vary the κ parameter, holding all other parameters constant.

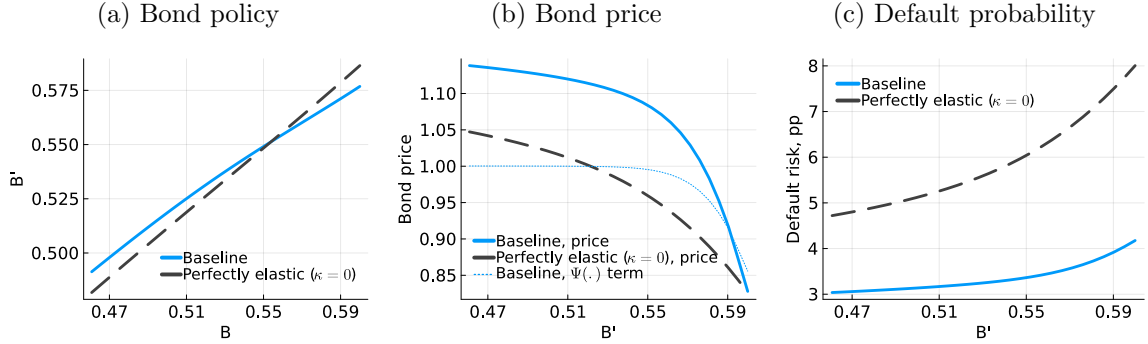
Disciplining device

The last four columns of Table 6 report the targeted moments under alternative values of κ . A striking result is that average spreads are *lower* when demand is less elastic (i.e., when κ is larger), while the average debt-to-output ratio is *higher*.

There are two interconnected mechanisms behind this result. First, when facing inelastic investors, each additional unit of B' lowers bond prices both by increasing default risk and through the price sensitivity implied by the downward-sloping demand. This latter effect intensifies as B increases, since the variance of repayment rises and the demand becomes less elastic. The government internalizes these forces and thus has stronger incentives to limit bond issuance when κ is higher. Panel (a) of Figure 5 illustrates this mechanism by comparing the optimal debt policy $B'(y, \tau, B)$ under our baseline calibration with the $\kappa = 0$ case—the perfectly elastic benchmark. In the baseline, the government issues less debt at high levels of B to avoid sharp price declines.

Second, these differences in issuance have important effects on default risk and bond prices. Panel (b) shows that for low values of B' , bond prices are actually *higher* when $\kappa > 0$. Since $\Psi(\cdot)$ is below one, the higher price is not mechanically driven by a convenience component—that is, a premium investors are willing to pay to hold the bond. Instead, it reflects a lower default probability, shown in Panel (c), which results directly from the government’s stronger incentives to limit bond issuance. In our baseline model, the average default probability is about 2 p.p. lower than in the perfectly elastic case.

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



Note: The figure compares policies and bond prices between 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 bond issuance policies as a function of B . Panel (b) shows bond prices, and Panel (c) reports the annualized $1/\lambda$ -period-ahead default probability, where $1/\lambda$ denotes the bonds' expected time to maturity. All functions are evaluated at the mean values of output y and index weight τ .

Table 8: Alternative inconvenience yields

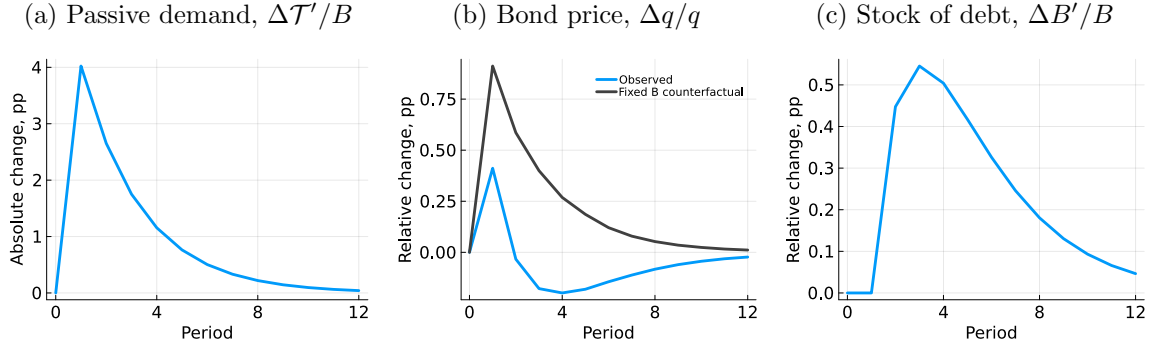
Moment	Description	Baseline model	Perfectly elastic ($\kappa = 0$)	Alternative inconvenience yields		
				High	Higher	Highest
$\bar{\mathbb{E}}(B/y)$	Debt to output (%)	53.1	51.0	52.7	52.1	51.5
$\bar{\mathbb{E}}(SP)$	Bond spreads (%)	5.06	8.59	4.96	4.98	4.91
$\bar{\sigma}(SP)$	Volatility of spreads (%)	1.56	4.49	1.48	1.47	1.41
$\bar{\mathbb{E}}(\vartheta)$	Incov. yield (%)	0.00	0.00	0.50	1.00	1.52
$\bar{\mathbb{E}}(d_\lambda)$	Default probability (%)	3.60	5.70	3.27	3.00	2.64
$\bar{\mathbb{E}}(\hat{\eta})$	Reduced-form elasticity	-0.28	-	-0.28	-0.27	-0.27
$\bar{\mathbb{E}}(\eta)$	Structural elasticity	-0.18	-	-0.18	-0.17	-0.16

Note: The table reports moments for our baseline calibration and alternative scenarios with higher average inconvenience yields, obtained by adjusting the \bar{A} parameter. In these scenarios, κ is also slightly adjusted to match the targeted reduced-form elasticity. The 'High' column targets an inconvenience yield 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.

Overall, these results imply that less elastic demand acts as a market-based disciplining device that curbs bond issuance, especially when bond supply is already large and default risk is high. This lowers default risk and borrowing costs, allowing the government to sustain a higher debt-to-output ratio. The effects are sizable: spreads are roughly 40% lower, and the debt-to-output ratio is 2 p.p. higher relative to the perfectly elastic benchmark.

In our baseline analysis, we assume an average inconvenience yield of zero. The disciplining mechanism we document—and its effects on spreads and debt issuance—still holds when the inconvenience yield is positive. Table 8 compares our model to alternative calibrations in which the average inconvenience yield is set to 0.50, 1.00, and 1.50 p.p., respectively. In each case, we slightly adjust κ to match our targeted reduced-form elasticity. Bond spreads remain almost unaffected relative to our baseline, despite now incorporating a premium beyond

Figure 6: Transitory increase in passive demand



Note: The figure shows impulse responses to a temporary increase in passive demand (blue lines). The black line reports the evolution of bond prices in a counterfactual scenario where the stock of debt is held constant.

default risk. This is because the government responds by lowering the average debt-to-output ratio, which reduces default risk. In this sense, market discipline is even stronger than in our baseline calibration. Still, even when we target a relatively large inconvenience yield of 1.50 p.p., the average debt-to-output ratio remains slightly higher than in the perfectly elastic case.

Changes in investor base

Unlike the standard sovereign debt model with perfectly elastic investors, in our framework, changes in the composition of investors have important implications for default risk, the pricing of bonds, and optimal debt policies.

We begin by analyzing the impulse responses to a temporary two-standard-deviation increase in the passive share—Figure 6. The assumed ϵ'_τ shock implies a 4 p.p. on-impact increase in T/B , as shown in Panel (a). Despite the larger demand, bond prices initially rise and then fall—see the blue line in Panel (b). This pattern arises because the government optimally responds by expanding the bond supply—Panel (c)—which raises default risk. For comparison, the black line in Panel (b) shows bond prices in a counterfactual scenario where supply is held fixed. In that case, prices rise by almost 1% on impact, reflecting the higher bond demand—see Appendix C.4 for additional details.

This simple exercise highlights a key departure from the standard model, where bond supply responds only to changes in output. In our framework, the composition of investors is a direct and quantitatively important force shaping both bond prices and optimal issuance decisions.

We now turn to a longer-term scenario and quantify the effects of a permanent increase in the passive share, τ^* . This exercise is aimed at capturing the long-term trend behind the

Table 9: Permanent increase in passive demand

Moment	Description	Baseline model	Higher passive demand
$\bar{\mathbb{E}}(B/y)$	Debt to output	53.06	54.41
$\bar{\mathbb{E}}(SP)$	Bond spreads	5.06	4.95
$\bar{\sigma}(SP)$	Std of spreads	1.6	1.6
$\bar{\mathbb{E}}(d_\lambda)$	Default probability	3.60	4.54
$\bar{\mathbb{E}}(\vartheta)$	Inconvenience yield	0.00	-2.05
$\bar{\mathbb{E}}(CEC)$	Certainty equivalent consumption	-	0.10

Note: The table reports the effects of a permanent increase in passive demand. It compares key moments from the baseline model to those from an alternative calibration in which τ^* is raised by 50%, holding all other parameters constant. All results are expressed in percentage points.

rise of passive investors in emerging economies.²³ To this end, we solve our baseline model under an alternative calibration in which τ^* is increased by 50%, keeping all other parameters unchanged.²⁴

From Equation (9), a permanent increase in the passive share shifts the bond price function upward by lowering the inconvenience yield required for any given B' . However, when we simulate the model, bond spreads and their volatility remain largely unchanged—Table 9. The reason is that the government responds by raising its average debt-to-output ratio by nearly 1.5 p.p., increasing default risk by about 1 p.p. This rise in risk offsets much of the 2 p.p. *convenience* yield associated with the larger passive share.

Since the (relatively impatient) government is able to increase the external debt without raising financing costs, it prefers the counterfactual scenario with higher passive demand over the baseline model. To quantify the welfare implications, we compute the certainty equivalent consumption (CEC), defined as the proportional increase in consumption that makes the household indifferent between the two scenarios.²⁵ As reported in the last row of Table 9, the average CEC is 0.10%, indicating welfare gains from a higher passive demand.

4.3 Role of Default Costs and Debt Dilution

In a standard sovereign debt model with perfectly elastic investors, the curvature of the pricing function arises from exogenous output losses triggered by default. Higher default costs discourage the government from defaulting, lower borrowing costs, and indirectly *increase*

²³Chari, Dilts Stedman, and Lundblad (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 driven by ETFs.

²⁴Despite its relative size, the shock raises the passive share τ^* by only 6 percentage points.

²⁵Formally, the CEC is defined as the value of \tilde{x} such that $\sum_{t=0}^{\infty} \beta^t \mathbb{E}_t u(\tilde{c}_t) = \sum_{t=0}^{\infty} \beta^t \mathbb{E}_t u((1 + \tilde{x})c_t)$, where \tilde{c}_t denotes consumption in the counterfactual scenario and c_t in the baseline. Under power utility, this simplifies to: $\tilde{x} = \left[\frac{\tilde{V}(y, \tau, B)}{V(y, \tau, B)} \right]^{\frac{1}{1-\gamma}} - 1$, where $\tilde{V}(\cdot)$ is the government's value function in the counterfactual.

Table 10: Disciplining devices: Exogenous output costs vs. inelastic investors

Moment	Description	Baseline model	Perfectly elastic	
			$\Delta\{d_0, d_1\}$	$\Delta\{d_0, d_1, \beta\}$
<i>A. Targeted</i>				
$\bar{\mathbb{E}}(B/y)$	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
<i>B. Untargeted</i>				
$\bar{\sigma}(c)/\bar{\sigma}(y)$	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
<i>C. Default episodes</i>				
$\bar{\mathbb{E}}(\Delta y_t/y_{t-1} \mid d_t = 1)$	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

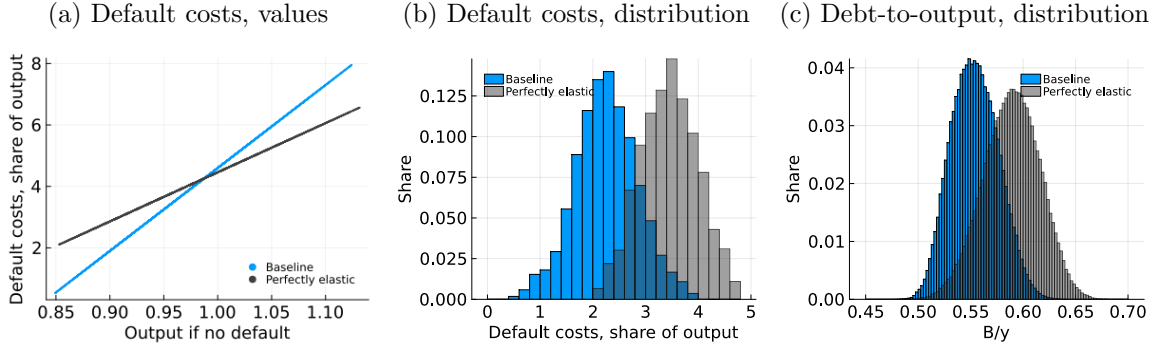
Note: The table compares targeted moments and default risk across our baseline model (column i) and various counterfactual scenarios. The $\Delta\{d_0, d_1\}$ columns presents a case with perfectly elastic investors ($\kappa = 0$) but adjusts default cost parameters $\{d_0, d_1\}$ to match the average and volatility of bond spreads. The last column shows a fully recalibrated version with $\{d_0, d_1, \beta\}$ set to match all three targeted moments.

debt issuance. This contrasts with our disciplining mechanism, in which a less elastic demand directly *reduces* debt issuance. In this section, we compare the implications of these two mechanisms.

We begin by comparing our baseline model to a recalibrated perfectly elastic benchmark, in which we adjust the default cost parameters $\{d_0, d_1\}$ to match the average and standard deviation of bond spreads, keeping all other parameters fixed. These results are reported in column $\Delta\{d_0, d_1\}$ of Table 10. Matching these moments under perfectly elastic demand requires larger default costs when output is low (Panel a of Figure 7)—precisely the periods in which defaults typically occur. In contrast, our baseline model relies on lower output costs in bad times—when demand is less elastic—to offset the stronger market discipline imposed by downward-sloping demand. In good times, when default risk is low and demand becomes more elastic, the disciplining effect weakens, and the model requires higher output costs to generate similar bond prices. However, defaults are rare in these periods. As a result, average default costs are higher in the perfectly elastic case, and the entire distribution of realized default costs shifts to the right, as shown in Panel b.

Despite targeting similar spreads, the government issues less debt in our baseline model because discipline operates directly through penalizing additional borrowing, rather than through changes in default incentives. The average debt level is substantially higher in the perfectly elastic case, with a debt-to-output ratio near 56%. As shown in Panel (c)

Figure 7: Source of discipline: Role of default costs



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 $\{d_0, d_1\}$ are recalibrated 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.

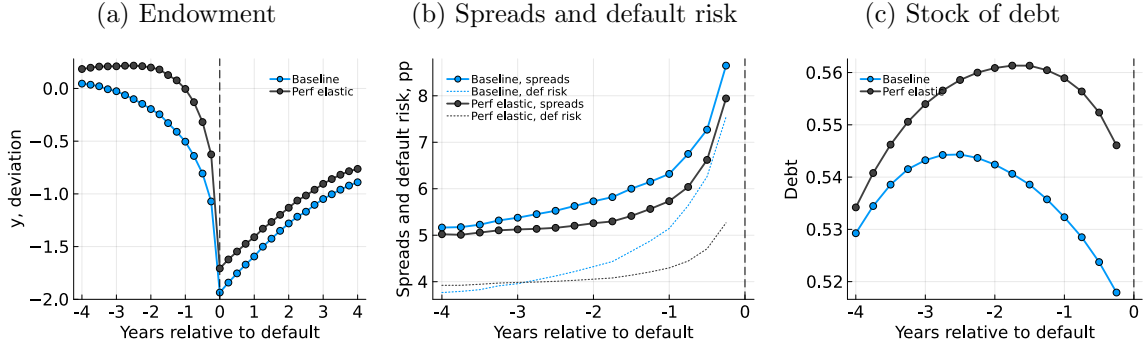
of Figure 7, the entire debt distribution shifts to the right. Given these differences, the source of discipline—whether driven by default costs or price-elastic investors—has important welfare implications. Overall, we find that the government prefers an economy with perfectly elastic investors, where discipline comes solely through default costs. On average, the CEC is approximately 0.10%.

In the last column of Table 10, we recalibrate the default cost parameters $\{d_0, d_1\}$ and the government's discount factor β to jointly match our targeted moments for bond spreads and the debt-to-output ratio.²⁶ Under this calibration, the set of untargeted moments reported in Panel B—volatility of consumption and trade balance, 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, important differences across the two models—beyond the implied default costs—that reflect their distinct sources of discipline.

First, as shown in the bottom panel of Table 10, the dynamics around default episodes differ. The observed drop in y at default is smaller in the baseline (4.7% versus 6%), reflecting not only differences in the calibrated default cost parameters but also in the timing of default. In our baseline, output is already 1.16 standard deviations below average in the period before default—nearly twice as low as in the perfectly elastic case. 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. In the perfectly elastic case, it remains flat until just before default, with output still above trend one year prior. The smoother decline in the baseline is

²⁶The calibration slightly increases β to reduce borrowing incentives. The implied default costs in this case are nearly identical to those in Figure 7, and we report them in Appendix C.3.

Figure 8: Periods around default episodes



Note: The figure shows the average paths of output, spreads, default risk, and debt around sovereign default episodes under the baseline model (blue) and the fully recalibrated perfectly elastic benchmark (black), which adjusts $\{d_0, d_1, \beta\}$. Time is centered at the quarter of default ($t = 0$). 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.

consistent with empirical studies such as [Levy Yeyati and Panizza \(2011\)](#), who document that (i) output contractions typically begin more than three years ahead of sovereign defaults, and (ii) the drop in output during the default period itself is relatively small.

The previous patterns reflect the different sources of fiscal discipline. Panel (b) shows that in the baseline, spreads rise earlier and more gradually, as they respond not only to higher default risk but also to the price-elastic component of demand (as fundamentals deteriorate, demand becomes less elastic). Given the worse prices faced by the government, default risk also rises earlier and more sharply in our baseline. Panel (c) displays debt dynamics. In the baseline, the government accumulates less debt and begins deleveraging earlier due to higher marginal borrowing costs. This early adjustment delays default and explains the longer, smoother decline in the endowment path. In contrast, with perfectly elastic demand, the absence of market-based discipline allows the government to continue borrowing even as fundamentals weaken. This leads to a larger stock of debt, and default can occur even when the endowment is just 0.66 standard deviations below average.

A second key difference lies in the role of debt dilution in shaping bond prices. The disciplining effect of a downward-sloping demand curve operates through the well-known dilution problem associated with long-term risky bonds. When governments cannot commit to future issuance plans, lenders anticipating additional borrowing demand lower prices to compensate for higher default risk. In the following analysis, we show that, by curbing incentives for future issuance, a downward-sloping demand curve weakens the dilution channel and raises bond prices, which explains the need for smaller default costs to attain similar spreads.

Quantifying debt dilution

We follow [Hatchondo et al. \(2016\)](#) and introduce a debt covenant that requires the government to compensate existing bondholders for any decline in bond prices resulting from an increase in default risk caused by new issuances. The covenant is defined as

$$\mathcal{C}(y, \tau, B, B') = \max \left\{ 0, \tilde{q}_{\kappa=0}(y, \tau, (1 - \lambda)B) - \tilde{q}_{\kappa=0}(y, \tau, B') \right\}, \quad (23)$$

where $\tilde{q}_{\kappa=0}(y, \tau, B') \equiv \mathbb{E}_{s'|s} \tilde{\mathcal{R}}(y', \tau', B')$ denotes the bond price under perfectly elastic demand, and

$$\tilde{\mathcal{R}}(y', \tau', B') = \left[1 - \tilde{d}'(\cdot) \right] \left[\lambda + (1 - \lambda) (\nu + \tilde{q}'(\cdot) + \mathcal{C}'(\cdot)) \right], \quad (24)$$

with $\tilde{d}'(\cdot) = \tilde{d}(y', \tau', B')$, $\tilde{q}'(\cdot) = \tilde{q}(y', \tau', \tilde{B}'')$, $\mathcal{C}'(\cdot) = \mathcal{C}(y', \tau', B', \tilde{B}'')$, and $\tilde{B}'' = \tilde{B}'(y', \tau', B')$. We use tildes to denote prices and policies in the counterfactual economy subject to the debt covenant.

Importantly, the covenant is defined with respect to the perfectly elastic price. That is, the government compensates bondholders for increases in default risk, but not for price changes arising from a downward-sloping demand. This approach facilitates comparison with the perfectly elastic benchmark and helps isolate the portion of debt dilution attributable purely to default risk.

Based on the covenant defined above, the bond price faced by the government is

$$\tilde{q}(y, \tau, B') = \frac{1}{r_f} \mathbb{E}_{s'|s} \tilde{\mathcal{R}}(y', \tau', B') - \frac{1}{r_f} \kappa \mathbb{V}_{s'|s} \tilde{\mathcal{R}}(y', \tau', B') (B' - \mathcal{T}(\tau, B') - \bar{\mathcal{A}}). \quad (25)$$

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

$$\tilde{c}(h = 0, B, y, \tau; B') = y + \tilde{q}(\cdot) (B' - (1 - \lambda)B) - (\lambda + (1 - \lambda) (\nu + \mathcal{C}(\cdot))) B. \quad (26)$$

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

Table 11 compares the model with dilution—our baseline and perfectly elastic cases—to the *no dilution* counterfactuals. In these counterfactuals, 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.²⁷

Two key differences emerge between the baseline and perfectly elastic cases. First, in the absence of dilution, the average debt-to-output ratio is significantly higher when investors are perfectly elastic (55.5% versus 51.6%). Second, despite the lower debt, default risk and spreads are *higher* in our baseline. In fact, the default probability more than doubles relative

²⁷The magnitudes are comparable to those reported in [Hatchondo et al. \(2016\)](#).

Table 11: Quantifying debt dilution

Moment	Description	With dilution			Without dilution		
		Baseline model	Higher inc. yield	Perfectly elastic	Baseline model	Higher inc. yield	Perfectly elastic
$\bar{\mathbb{E}}(B/y)$	Debt to output	53.1	52.1	53.0	51.6	51.0	55.5
$\bar{\mathbb{E}}(SP)$	Bond spreads	5.06	4.98	5.08	0.44	0.65	0.32
$\bar{\sigma}(SP)$	Std of spreads	1.56	1.47	1.61	0.46	0.53	0.18
$\bar{\mathbb{E}}(d_\lambda)$	Default prob.	3.60	3.00	3.95	0.61	0.57	0.29
$\bar{\mathbb{E}}(\vartheta)$	Inconv. yield	0.00	1.00	0.00	-0.23	0.03	0.00

Note: The table compares selected moments across our baseline model and various counterfactual scenarios, highlighting the role of debt dilution. The left panel reports results for economies with dilution: column (i) shows the baseline model; column (ii) introduces a higher average inconvenience yield of 1 p.p.; and column (iii) presents the fully recalibrated perfectly elastic case (adjusting $\{d_0, d_1, \beta\}$). The right panel reports results for the same scenarios without dilution. All results are expressed in percentage points.

to the perfectly elastic case. Spreads are also larger in the baseline, despite the presence of a small *convenience* yield—driven by the lower bond supply. In the *Higher inc. yield* column, we recalibrate $\bar{\mathcal{A}}$ so that the model delivers an average inconvenience yield of zero. Once the convenience yield is shut down, the average spread rises to 0.65 p.p., which is more than twice the spread observed under perfectly elastic demand.

Overall, debt dilution remains an important driver of bond prices, though its effects are dampened under downward-sloping demand. This is because market-based discipline curbs future bond issuance—especially in periods of high default risk—which mitigates dilution. As we show next, this distinction has important implications for the design of fiscal rules aimed at reducing default risk.

4.4 Policy Implications: Fiscal Rules

We examine the implications of fiscal rules for the pricing of sovereign bonds, comparing our baseline model to the perfectly elastic demand benchmark. For the latter, we focus on the fully recalibrated version—that is, the case in which we recalibrate $\{d_0, d_1, \beta\}$. This analysis further illustrates the distinct role played by the disciplining mechanism generated by a downward-sloping demand.

Table 12 presents the effects of a debt ceiling policy that limits the government’s debt-to-output ratio. For this analysis, we assume a debt-to-output ceiling of 55%. In our baseline model, the policy has mild effects on default probabilities. While bond spreads decline, this is driven almost entirely by a positive convenience yield resulting from the reduced supply of bonds. By contrast, under perfectly elastic demand, the same policy cuts both default probabilities and bond spreads by roughly half, and even prompts the government to issue more debt.

Table 12: Fiscal policy: A debt-to-output ceiling

Moment	Description	Baseline model		Perfectly elastic	
		No policy	Debt ceiling	No policy	Debt ceiling
$\mathbb{E}(B/y)$	Debt to output	53.1	52.7	53.0	53.8
$\mathbb{E}(SP)$	Bond spreads	5.06	3.83	5.08	2.22
$\bar{\sigma}(SP)$	Std of spreads	1.56	1.71	1.61	1.54
$\mathbb{E}(d_\lambda)$	Default probability	3.60	3.25	3.95	1.91
$\mathbb{E}(\vartheta)$	Inconvenience yield	0.00	-0.66	0.00	0.00

Note: The table compares the effects of a debt ceiling policy, defined by $B/y \leq \bar{b}$, under the baseline model and a perfectly elastic demand counterfactual. In the perfectly-elastic case, d_0, d_1, β are jointly recalibrated to match the three targeted moments of the baseline calibration. Results are shown for $\bar{b} = 0.55$.

These differences stem from the debt dilution mechanism. In the perfectly elastic case, the debt ceiling acts as a strong disciplining device: it directly limits future debt issuance, reducing default risk and spreads. In our baseline, much of the discipline is already provided by the market—through a downward-sloping demand curve—so the added effect of a statutory ceiling is limited. Overall, the implication for policy is clear: debt ceilings have surprisingly limited impact when inelastic investors already price in fiscal slack.

5 Conclusion

Our analysis shows that introducing a realistic investor base—passive funds that track benchmark indices and active investors with downward-sloping demand—fundamentally reshapes sovereign debt dynamics. Using variation from rebalancings of the largest global index for sovereign bonds, we estimate demand elasticities and show that demand becomes less elastic as default risk rises. We feed these estimates into a standard sovereign debt model with long-term debt, augmented with a rich demand block.

We show that the presence of inelastic investors acts as a market-based disciplining device: they limit borrowing precisely when default risk is high, dampen the effects of debt dilution, and lower spreads. The disciplining effect is strong enough that the government ultimately issues *more* debt when facing a less elastic demand. These mechanisms carry first-order policy implications. Since the market already penalizes excessive borrowing, debt ceilings or similar fiscal rules have only modest effects on spreads and default risk.

A natural next step is to allow the government to choose not only how much to borrow, but also at which maturities. Understanding how inelastic demand shapes the optimal maturity composition of risky bonds—and whether it tilts issuance toward longer maturities given its effects on dilution—remains an open question.

References

- Adrian, T. and H. Shin (2014). Procyclical leverage and value-at-risk. *Review of Financial Studies* 27, 373–403.
- 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, H. Hopenhayn, and I. Werning (2019). Take the short route: Equilibrium default and debt maturity. *Econometrica* 87(2), 423–462.
- 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 Papers 23370, National Bureau of Economic Research, Inc.
- 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 H. Chang (2024). 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 (2022). Institutional corporate bond pricing. Swiss Finance Institute Research Paper 21-07.
- 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.
- Calomiris, C., M. Larrain, S. Schmukler, and T. Williams (2022). Large international corporate bonds: Investor behavior and firm responses. *Journal of International Economics* 137(C).
- Chang, Y., H. Hong, and I. Liskovich (2014). Regression discontinuity and the price effects of stock market indexing. *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. Working Paper 30577, National Bureau of Economic Research.
- Chatterjee, S. and B. Eyigungor (2012). Maturity, indebtedness, and default risk. *American Economic Review* 102(6), 2674–2699.
- Chaumont, G. (2024). Sovereign debt, default risk, and the liquidity of government bonds. Working paper.
- Choi, J., R. Kirpalani, and D. Perez (2024). U.S. public debt and safe asset market power. Working paper.
- Costain, J., G. Nuño, and C. Thomas (2022). The term structure of interest rates in a heterogeneous monetary union. Banco de Espana Working Paper 2223.
- Cremers, M. and A. Petajisto (2009). How active is your fund manager? A new measure that predicts performance. *Review of Financial Studies* 22(9), 3329–3365.
- Dathan, M. and S. Davydenko (2020). Debt issuance in the era of passive investment. University of Toronto Working Paper.
- 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 (2021). The anatomy of index

- rebalancings: Evidence from transaction data. CSEF Working Paper 621.
- Faia, E., K. K. Lewis, and H. Zhou (2024). Do investor differences impact monetary policy spillovers to emerging markets? NBER Working Paper 32986.
- Faia, E., J. Salomao, and A. V. Veghazy (2024). Market segmentation and international bond prices: The role of ECB asset purchases. Working paper.
- Fang, X., B. Hardy, and K. Lewis (2025). Who holds sovereign debt and why it matters. *Review of Financial Studies Forthcoming*.
- 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.
- Gourinchas, P., W. Ray, and D. Vayanos (2022). A preferred-habitat model of term premia, exchange rates, and monetary policy spillovers. NBER Working Paper 29875.
- 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, J. Stein, and A. Sunderam (2023). A quantity-driven theory of term premia and exchange rates. *Quarterly Journal of Economics* 138(4), 2327–2389.
- Greenwood, R., S. Hanson, and D. Vayanos (2023, November). Supply and demand and the term structure of interest rates. Working Paper 31879, National Bureau of Economic Research.
- 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), 815–29.
- Hatchondo, J. C., L. Martinez, and F. Roch (2022). Fiscal rules and the sovereign default premium. *American Economic Journal: Macroeconomics* 14(4), 244–73.
- 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. *Review of Financial Studies* 23(4), 1681–1717.
- He, Z. and K. Milbradt (2014). Endogenous liquidity and defaultable bonds. *Econometrica* 82(4), 1443–1508.
- Jansen, K. A., W. Li, and L. Schmid (2024, December). Granular treasury demand with arbitrageurs. NBER Working Papers 33243, National Bureau of Economic Research, Inc.
- Jiang, Z., H. Lustig, S. Van Nieuwerburgh, and M. Xiaolan (2021). Bond convenience yields in the eurozone currency union. Stanford University Graduate School of Business Research Paper 3976.
- Jiang, Z., H. Lustig, S. Van Nieuwerburgh, and M. Z. Xiaolan (2024). The U.S. public debt valuation puzzle. *Econometrica* 92(4), 1309–1347.
- Kaldorf, M. and J. Rottger (2023). Convenient but risky government bonds. Deutsche Bundesbank discussion paper 15/2023.
- Koijen, R. and M. Yogo (2019). A demand system approach to asset pricing. *Journal of Political Economy* 127(4), 1475–1515.
- Koijen, R. and M. Yogo (2020). Exchange rates and asset prices in a global demand system. NBER Working Paper 27342.
- Krishnamurthy, A. and A. Vissing-Jorgensen (2012). The aggregate demand for Treasury

- debt. *Journal of Political Economy* 120(2), 233–267.
- Kubitza, C. (2023). Investor-driven corporate finance: Evidence from insurance markets. ECB Working Paper 2816.
- 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.
- Mian, A., L. Straub, and A. Sufi (2022). A Goldilocks theory of fiscal deficits. NBER Working Paper 29707.
- Miranda-Agrippino, S. and H. Rey (2020). U.S. monetary policy and the global financial cycle. *Review of Economic Studies* 87, 2754–2776.
- 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. *Review of Financial Studies* 36(3), 859–903.
- Pouzo, D. and I. Presno (2016). Sovereign default risk and uncertainty premia. *American Economic Journal: Macroeconomics* 8(3), 230–266.
- Raddatz, C., S. Schmukler, and T. Williams (2017). International asset allocations and capital flows: The benchmark effect. *Journal of International Economics* 108(C), 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), 579–90.
- 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.

Appendix for Inelastic Demand Meets Optimal Supply of Risky Sovereign Bonds

A Model Appendix

In this appendix, we first provide additional derivations for the investors' block—Section A.1. In Section A.2, we present alternative microfoundations for the assumed demand structure. Section A.3 defines the equilibrium of the model, and Section A.4 provides the derivation for the government's optimal debt issuance.

A.1 Investors' Block: Additional Derivations

Following the notation used in the main text, let $x_{j,t}^i = \frac{q_t^i B_{j,t}^i}{W_{j,t}}$ be the share of its wealth that investor j invests in bond i at time t . We have assumed that:

$$x_{j,t}^i = \theta_{j,t} \xi_{j,t}^i e^{\Lambda_j \pi_t^i(r_{t+1}^i)} + (1 - \theta_{j,t}) w_t^i,$$

where r_{t+1}^i is the next-period excess return. This is defined as $r_{t+1}^i \equiv \frac{\mathcal{R}_{t+1}^i}{q_t^i} - r^f$, where r^f denotes the gross risk-free rate and \mathcal{R}_{t+1}^i is the next-period repayment, per unit of bond.

Purely passive funds are those with $\theta_j = 0$ and they just replicate the weight of country i in the index \mathcal{I} , w_t^i . Active funds are those with $\theta_j \in (0, 1]$. They have an asset-allocation mandate to invest a share $\xi_{j,t}^i$ of their wealth in bond i . They may deviate from that target, depending on the $\pi_t^i(\cdot)$ function. In our analysis, this function is given by the Sharpe ratio:

$$\pi_t^i(r_{t+1}^i) = \frac{\mathbb{E}_t(r_{t+1}^i)}{\mathbb{V}_t(r_{t+1}^i)} = q_t^i \frac{\mathbb{E}_t(\mathcal{R}_{t+1}^i) - q_t^i r^f}{\mathbb{V}_t(\mathcal{R}_{t+1}^i)}. \quad (\text{A1})$$

The market clearing condition for bond i is $B_t^i = \frac{1}{q_t^i} \sum_j W_{j,t} x_{j,t}^i$, which can be decomposed into an active, \mathcal{A}_t^i , and passive, \mathcal{T}_t^i , component:

$$B_t^i = \frac{1}{q_t^i} \sum_j W_{j,t} \theta_{j,t} \xi_{j,t}^i e^{\Lambda_j \pi_t^i} + \frac{1}{q_t^i} \sum_j W_{j,t} (1 - \theta_{j,t}) w_t^i.$$

Regarding the passive demand, we let $\mathcal{W}_t = \sum_j W_{j,t} (1 - \theta_{j,t})$ denote the assets under management of passive investors tracking the index \mathcal{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 $w_t^i = \frac{\alpha_t^i q_t^i B_t^i}{MV_t}$. Combining these expressions, the passive demand can be written as $\mathcal{T}_t^i = \tau_t^i B_t^i$, where $\tau_t^i \equiv \frac{\mathcal{W}_t}{MV_t} \alpha_t^i$ is a time-varying share that determines how much of each unit of country i 's debt is held by passive investors.

As for the active demand, we let $\xi_{j,t}^i = \xi_j^i q_t^i$, so that ξ_j^i captures a fixed component of

investors' mandates. This specification leads to the following active demand for investor j : $\mathcal{A}_{j,t}^i = \overline{\mathcal{A}_{j,t}^i} e^{\Lambda_j^i \pi_t^i}$, where $\overline{\mathcal{A}_{j,t}^i} \equiv W_{j,t} \theta_{j,t} \xi_j^i$ captures investor j 's wealth targeted towards country i , as implied by their degree of activeness and mandates. Let $\overline{\mathcal{A}_t^i} \equiv \sum_j \overline{\mathcal{A}_{j,t}^i}$.

To derive a closed-form expression for the bond price, we linearize each investor j 's active demand around $\pi_t^i = 0$ and around a reference allocation $\overline{\mathcal{A}_{j,t}^i} = \overline{\mathcal{A}_j^{i*}}$. We can then write the active demand as:

$$\mathcal{A}_t^i \approx \overline{\mathcal{A}_t^i} + \Lambda^i \pi_t^i. \quad (\text{A2})$$

The term Λ^i denotes the aggregate demand elasticity for active investors, and it is given by a weighted average of investors' elasticities, where the weights are given by each investor's relative importance in the allocation to bond i around the reference point. That is, $\Lambda^i \equiv \overline{\mathcal{A}^{i*}} \sum_j s_j^{i*} \Lambda_j$, with $s_j^{i*} \equiv \overline{\mathcal{A}_j^{i*}} / \overline{\mathcal{A}^{i*}}$ and $\overline{\mathcal{A}^{i*}} \equiv \sum_j \overline{\mathcal{A}_j^{i*}}$.

Replacing with all these expressions, the market clearing condition boils down to:

$$B_t^i = \left(\overline{\mathcal{A}_t^i} + \Lambda^i \pi_t^i \right) + \mathcal{T}_t^i. \quad (\text{A3})$$

From the previous expression, after replacing with the $\pi_t^i(\cdot)$ function in Equation (A1) and solving for the bond price, we get our baseline pricing function:

$$q_t^i = \frac{1}{rf} \mathbb{E}_t(\mathcal{R}_{t+1}^i) - \frac{\kappa^i}{rf} \frac{\mathbb{V}_t(\mathcal{R}_{t+1}^i)}{\mathbb{E}_t(\mathcal{R}_{t+1}^i)} \left(B_t^i - \mathcal{T}_t^i - \overline{\mathcal{A}_t^i} \right), \quad (\text{A4})$$

where $\kappa^i \equiv 1/\Lambda^i$ parameterizes the inverse demand elasticity for active investors. It is also useful to write this pricing equation as

$$q_t^i = \frac{1}{rf} \mathbb{E}_t(\mathcal{R}_{t+1}^i) \Psi_t^i, \quad (\text{A5})$$

where

$$\Psi_t^i \equiv 1 - \kappa^i \frac{\mathbb{V}_t(\mathcal{R}_{t+1}^i)}{\mathbb{E}_t(\mathcal{R}_{t+1}^i)} \left(B_t^i - \mathcal{T}_t^i - \overline{\mathcal{A}_t^i} \right). \quad (\text{A6})$$

A.2 Microfoundations for Investors' Mandates

In this section, we provide microfoundations for investors' mandates. We show that one can obtain a pricing kernel analogous to the one in Equation (A4) under three alternative setups: (i) risk-averse investors, (ii) risk-neutral investors subject to a standard value-at-risk constraint, and (iii) investors who derive non-pecuniary benefits (or costs) from holding the bond.

Risk-averse Investors

Consider first a case where investors are risk averse and have mean-variance preferences (as in [Vayanos and Vila, 2021](#)). They care about both the total return of their portfolio and their return relative to a benchmark index \mathcal{I} they track. Investors differ in their degree of risk aversion and in how their compensation depends on total versus relative returns.

Following the same notation as in the main text, let $j = \{1, \dots, J\}$ denote the investor type. Let $i = \{1, \dots, N\}$ denote the set of bonds that are part of the \mathcal{I} index, and let $\mathbf{w}_t = \{w_t^1, \dots, w_t^N\}$ be the vector of index weights for each constituent bond. The vector $\mathbf{r}_{t+1} = \{r_{t+1}^1, \dots, r_{t+1}^N\}$ denotes the next-period (gross) excess returns. Last, let $\mathbf{B}_t = \{B_t^1, \dots, B_t^N\}$ denote the bond supply of each bond in the index.

For an investor j , their total compensation is a convex combination of the return of their portfolio and the return relative to the index \mathcal{I} . Let $\mathbf{x}_{j,t} = \{x_{j,t}^1, \dots, x_{j,t}^N\}$ be investor j 's vector of portfolio weights. The investor's total compensation (TC) is:

$$\begin{aligned} TC_{j,t} &= \theta_j (\mathbf{x}_{j,t})' \cdot \mathbf{r}_{t+1} + (1 - \theta_j) (\mathbf{x}_{j,t} - \mathbf{w}_t)' \cdot \mathbf{r}_{t+1} \\ &= [\mathbf{x}_{j,t} - (1 - \theta_j) \mathbf{w}_t]' \cdot \mathbf{r}_{t+1}, \end{aligned}$$

where $1 - \theta_j$ captures the weight of relative returns in the investor's compensation. For instance, if $\theta = 1$, the investor is fully active, as their compensation is not tied to the relative performance of their portfolio.

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

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

where $\boldsymbol{\mu}_t \equiv \mathbb{E}_t(\mathbf{r}_{t+1})$ denotes expected excess returns and $\boldsymbol{\Sigma}_t \equiv \mathbb{V}_t(\mathbf{r}_{t+1})$ is the variance-covariance matrix. It is straightforward to show that the optimal portfolio for investor j is:

$$\mathbf{x}_{j,t} = \frac{1}{\sigma_j} \boldsymbol{\Sigma}_t^{-1} \boldsymbol{\mu}_t + (1 - \theta_j) \mathbf{w}_t. \quad (\text{A7})$$

The first term on the right-hand side of Equation (A7) is the standard mean-variance portfolio. The second term reflects reluctance to deviate from the benchmark weights, \mathbf{w}_t , implying a perfectly inelastic demand component. It does not depend on expected returns or risk, but solely on the penalty for deviating from the benchmark. Purely passive investors (i.e., $\theta_j = 0$ and $\sigma_j \rightarrow \infty$) never deviate and have perfectly inelastic demand.

Let $W_{j,t}$ denote the wealth of investor j . Then $B_{j,t}^i = \frac{W_{j,t} x_{j,t}^i}{q_t^i}$ are investor j 's holdings

of bond i , where q_t^i is the bond price. The market-clearing condition is $q_t^i B_t^i = \sum_j W_{j,t} x_{j,t}^i$. Substituting the optimal portfolios, the market-clearing condition becomes:

$$\begin{bmatrix} q_t^1 B_t^1 \\ \vdots \\ q_t^N B_t^N \end{bmatrix} = \sum_j W_{j,t} \left[\frac{1}{\sigma_j} \Sigma_t^{-1} \boldsymbol{\mu}_t + (1 - \theta_j) \mathbf{w}_t \right]. \quad (\text{A8})$$

To simplify, consider only two bonds, i and z , or assume we have already solved for the optimal allocation across the $N - 1$ bonds in \mathcal{I} excluding i . In the latter case, let $\sigma_{z,j,t}$ and $\mu_{z,j,t}$ denote the volatility and excess return of that portfolio. These are indexed by j because that portfolio of $N - 1$ assets could differ across investors. Similarly, denote $\rho_{iz,j,t}$ as the correlation of excess returns between bond i and investor j 's portfolio z .

From Equation (A7), the optimal share for bond i is:

$$x_{j,t}^i = \frac{1}{\sigma_j} \frac{\mu_{i,t}}{\sigma_{i,t}^2} \xi_{j,t}^i + (1 - \theta_j) w_t^i, \quad (\text{A9})$$

where

$$\xi_{j,t}^i \equiv \frac{1}{1 - \rho_{iz,j,t}^2} \left[1 - \frac{\mu_{z,j,t}}{\mu_{i,t}} \frac{\sigma_{i,t}}{\sigma_{z,j,t}} \rho_{iz,j,t} \right]. \quad (\text{A10})$$

The previous $\xi_{j,t}^i$ 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 in Equation (A8), the market-clearing condition becomes:

$$\begin{aligned} B_t^i &= \frac{1}{q_t^i} \sum_j \left(\frac{1}{\sigma_j} \left[\frac{\mu_{i,t}}{\sigma_{i,t}^2} \xi_{j,t}^i \right] + (1 - \theta_j) w_t^i \right) W_{j,t} \\ &\equiv \frac{1}{q_t^i} \sum_j \frac{1}{\sigma_j} \left[\frac{\mathbb{E}_t r_{t+1}^i}{\mathbb{V}_t r_{t+1}^i} \xi_{j,t}^i \right] W_{j,t} + \mathcal{T}_t^i, \end{aligned} \quad (\text{A11})$$

where in the last step we used $\mathcal{T}_t^i = \tau_t B_t^i$, with $\tau_t \equiv \frac{\mathcal{W}_t}{MV_t} \alpha_t^i$, $\mathcal{W}_t \equiv \sum_j W_{j,t} (1 - \theta_j)$, and MV_t the market value of the index. Using $r_{t+1}^i = \frac{\mathcal{R}_{t+1}^i}{q_t^i} - r^f$ and rearranging, we get:

$$q_t^i = \frac{1}{r^f} \mathbb{E}_t (\mathcal{R}_{t+1}^i) - \frac{\tilde{\kappa}_t^i}{r^f} \mathbb{V}_t (\mathcal{R}_{t+1}^i) (B_t^i - \mathcal{T}_t^i), \quad (\text{A12})$$

where $\tilde{\kappa}_t^i \equiv 1 / \left(\sum_j \frac{1}{\sigma_j} W_{j,t} \xi_{j,t}^i \right)$.

The bond price in Equation (A12) is analogous to our baseline pricing equation. The main difference is that with risk-averse investors, the price elasticity is governed by the degree of risk aversion. In our main analysis, instead, we do not specify the micro-foundation behind

this elasticity.²⁸

Value-at-risk Constraint

An identical expression can also be derived for investors who are risk neutral but subject to a value-at-risk (VaR) constraint. These constraints are widely used both in the literature and in regulatory practice (e.g., [Gabaix and Maggiori, 2015](#); [Miranda-Agrippino and Rey, 2020](#)).²⁹

Consider a setting in which investors are heterogeneous and care about both their absolute and relative returns with respect to the index \mathcal{I} . They are risk neutral but subject to a VaR constraint that imposes an upper limit on the amount of risk they can take. Specifically, investor j solves:

$$\begin{aligned} & \max_{\{x_{j,t+1}^1, \dots, x_{j,t+1}^N\}} \mathbb{E}_t \left([\mathbf{x}_{j,t+1} - (1 - \alpha_j) \mathbf{s}_{t+1}]' \cdot \mathbf{r}_{t+1} \right) \\ & \text{subject to } \Phi^2 \mathbb{V}_t \left([\mathbf{x}_{j,t+1} - (1 - \alpha_j) \mathbf{s}_{t+1}]' \cdot \mathbf{r}_{t+1} \right) - 1 \leq 0, \end{aligned}$$

where the parameter Φ^2 captures the intensity of the risk constraint and can be interpreted as a regulatory limit on the amount of risk the investor is allowed to take. Let ϱ_j denote the Lagrange multiplier associated with the VaR constraint. The optimal portfolio is then:

$$\mathbf{x}_{j,t} = \frac{1}{\varrho_j \Phi^2} \boldsymbol{\Sigma}_t^{-1} \boldsymbol{\mu}_t + (1 - \theta_j) \mathbf{w}_t. \quad (\text{A13})$$

This optimal portfolio is identical to Equation (A7), with the only difference being that the risk-aversion parameter σ_j is replaced by the product of the Lagrange multiplier ϱ_j and the regulatory parameter Φ^2 . Following the same steps as before, we can then derive an analogous pricing kernel.

Non-pecuniary Benefits

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

$$\max_{b_t^i} \frac{\mathbb{E}_t(R_{t+1}^i)}{r_f} b_t^i - q_t^i b_t^i + \left(\alpha^i b_t^i - \frac{\varphi_t^i}{2} (b_t^i)^2 \right). \quad (\text{A14})$$

²⁸Unlike the baseline specification, this expression is not fully closed-form because $\xi_{j,t}^i$ depends on bond prices. Nonetheless, it illustrates that this risk-averse setup is conceptually analogous to the framework used in the main text.

²⁹[Adrian and Shin \(2014\)](#) provide a microfoundation for VaR constraints.

This particular specification follows [Greenwood, Hanson, and Vayanos \(2023\)](#). The quadratic term $\alpha^i b_t^i - \frac{\varphi_t^i}{2} (b_t^i)^2$ captures a “convenience benefit” from holding b_t^i units of bond i , such as liquidity or collateral services.

Taking the first-order condition with respect to b_t^i and substituting in the bond supply B_t^i yields the following bond price:

$$q_t^i = \frac{1}{r^f} \mathbb{E}_t \mathcal{R}_{t+1}^i - \varphi_t^i (B_t^i - \xi_t^i), \quad (\text{A15})$$

where $\xi_t^i \equiv \frac{\alpha_t^i}{\varphi_t^i}$. The pricing kernel in the main text is a special case of Equation (A15), in which φ_t^i is determined by the volatility of repayment.

A.3 Definition of Equilibrium

A Recursive Markov Equilibrium is a collection of value functions $\{V(\cdot), V^r(\cdot), V^d(\cdot)\}$; policy functions $\{d(\cdot), B'(\cdot)\}$; 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 (6), (7), and (8), 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 (10).
3. Taking the repayment function as given, bond prices $q(\cdot)$ are consistent with Equation (9).

A.4 Derivation of the Optimality Condition for Debt Issuance

In this appendix, we derive the first-order condition for B' in Equation (15). We start with the government’s problem in case of repayment—Equation (7)—, which we rewrite here for convenience:

$$V^r(y, \tau, b) = \max_{B'} u(c) + \beta \mathbb{E}_{s'|s} \max_{d=\{0,1\}} \left\{ V^r(y', \tau', B'), V^d(y') \right\},$$

subject to: $c = y + q(y, \tau, B') (B' - (1 - \lambda)B) - (\lambda + (1 - \lambda)\nu) B$.

We define $V^e(y, \tau, B') \equiv \mathbb{E}_{s'|s} \max_{d=\{0,1\}} \left\{ V^r(y', \tau', B'), V^d(y') \right\}$ as the expected continuation value. Let $u_c(c) \equiv \frac{\partial u(c)}{\partial c}$, $q_{B'}(y, \tau, B') \equiv \frac{\partial q(y, \tau, B')}{\partial B'}$, and $V_{B'}^e(y, \tau, B') \equiv \frac{\partial V^e(y, \tau, B')}{\partial B'}$ denote the derivatives with respect to B' . With this notation, the first-order condition with respect to B' is given by

$$u_c(c) [q(y, \tau, B') + q_{B'}(y, \tau, B') (B' - (1 - \lambda)B)] + \beta V_{B'}^e(y', \tau', B') = 0. \quad (\text{A16})$$

For a given optimal default policy, $d(y, \tau, B)$, we can write the expected continuation value as follows:

$$\begin{aligned} V^e(y', \tau', B') &\equiv \mathbb{E}_{s'|s} \max_{d \in \{0,1\}} \left\{ V^r(y', \tau', B'), V^d(y') \right\} \\ &= \mathbb{E}_{s'|s} \left[V^r(y', \tau', B') \times (1 - d(y', \tau', B')) \right] + \mathbb{E}_{s'|s} \left[V^d(y') \times d(y', \tau', B') \right]. \end{aligned} \quad (\text{A17})$$

Based on this last expression, and taking derivatives with respect to B' , we can write the last term in Equation (A16) as

$$\begin{aligned} V_{B'}^e(y', \tau', B') &= \mathbb{E}_{s'|s} \left[V_{B'}^r(y', \tau', B') (1 - d(y', \tau', B')) \right] + \\ &\quad + \mathbb{E}_{s'|s} \left[\left(V^d(y') - V^r(y', \tau', B') \right) d_{B'}(y', \tau', B') \right]. \end{aligned}$$

In what follows, and to simplify the analysis, we omit the last term of the previous equation, which captures how changes in B' affect the outside option value of defaulting.³⁰ Substituting this into the first-order condition, we obtain:

$$u_c(c) \left[q(y, \tau, B') + q_{B'}(y, \tau, B') (B' - (1 - \lambda)B) \right] = -\beta \mathbb{E}_{s'|s} \left[V_{B'}^r(y', \tau', B') (1 - d(y', \tau', B')) \right]. \quad (\text{A18})$$

From the envelope condition, we have that $V_B^r(y, b, \tau) \equiv \frac{\partial V^r(y, \tau, B)}{\partial B}$ is given by:

$$V_B^r(y, \tau, B) = -u_c(c) \left((1 - \lambda) (\nu + q(y, \tau, B)) + \lambda \right). \quad (\text{A19})$$

Evaluating at (y', τ', B') , multiplying by $1 - d' \equiv 1 - d(y', \tau', B')$ and taking expectations, the previous expression can be written as:

$$\begin{aligned} \mathbb{E}_{s'|s} \left[V_{B'}^r(y', \tau', B') (1 - d') \right] &= -\mathbb{E}_{s'|s} \left[u_c(c') \left((1 - \lambda) (\nu + q') + \lambda \right) (1 - d') \right] \\ &= -\mathbb{E}_{s'|s} \left[u_c(c') \mathcal{R}(y', \tau', B') \right] \end{aligned} \quad (\text{A20})$$

where $B'' \equiv B'(y', \tau', B')$ and $q' \equiv q(y', \tau', B'')$ denote next period's bond supply and price, respectively. In the last step of the previous expression, we substituted the repayment function—as defined in Equation (10). Substituting this expression into Equation (A18), we obtain:

$$u_c(c) \left[q(y, \tau, B') + q_{B'}(y, \tau, B') (B' - (1 - \lambda)B) \right] = \beta \mathbb{E}_{s'|s} \left[u_c(c') \mathcal{R}(y', \tau', B') \right]. \quad (\text{A21})$$

In our model the derivative of the pricing function with respect to an additional unit of

³⁰This simplification is made solely to illustrate the mechanism. The full solution of the model does not rely on this assumption.

debt is

$$q_{B'}(y, \tau, B') \equiv \frac{\partial q(y, \tau, B')}{\partial B'} = \frac{1}{r^f} \frac{\partial \mathbb{E}_{s'|s} \mathcal{R}(y, \tau, B')}{\partial B'} + \mathcal{S}(y, \tau, B'; \kappa). \quad (\text{A22})$$

Replacing with this expression in Equation (A21) and rearranging terms, we get Equation (15) in the main text.

A.5 Structural versus Reduced-form Elasticity

From the analysis in the main text, the reduced-form elasticity for active investors, based on an exogenous change in the passive demand, can be written as:

$$\hat{\eta}(y, \tau, B'; \kappa) = \frac{\Delta \log q(\cdot)}{\Delta \log \mathcal{A}'} = (-) \frac{\Delta \log q(\cdot)}{\Delta \mathcal{T}'} (B' - \mathcal{T}(\tau, B')),$$

with $\Delta \mathcal{T}' \equiv \mathcal{T}(\tau', B') - \mathcal{T}(\tau, B')$. For the remainder of this analysis, it is useful to define $\tilde{\Psi}(\cdot) \equiv (-) \frac{1}{r^f} \kappa \mathbb{V}_{s'|s} \mathcal{R}'(\cdot) (B' - \mathcal{T}' - \bar{\mathcal{A}})$, so that $q(\cdot) = \frac{1}{r^f} \mathbb{E}_{s'|s} \mathcal{R}'(\cdot) + \tilde{\Psi}(\cdot)$. After some algebra, one can decompose the previous expression in two components, the structural elasticity and changes in bond prices driven by changes in expected future payoffs. That is,

$$\begin{aligned} \hat{\eta}(\cdot) &= (-) \frac{\Delta q(\cdot)}{\Delta \mathcal{T}'} \frac{B' - \mathcal{T}'}{q(\cdot)} \\ &= (-) \frac{1}{r^f} \frac{\Delta \mathbb{E}_{s'|s} \mathcal{R}'(\cdot)}{\Delta \mathcal{T}'} \frac{B' - \mathcal{T}'}{q(\cdot)} + \frac{\Delta \tilde{\Psi}(\cdot)}{\Delta \mathcal{T}'} \frac{B' - \mathcal{T}'}{q(\cdot)} \\ &= (-) \frac{1}{r^f} \frac{\Delta \mathbb{E}_{s'|s} \mathcal{R}'(\cdot)}{\Delta \mathcal{T}'} \frac{B' - \mathcal{T}'}{q(\cdot)} + \frac{\kappa}{r^f} \left\{ \frac{\Delta \mathbb{V}_{s'|s} \mathcal{R}'(\cdot)}{\Delta \mathcal{T}'} (B' - \mathcal{T}' - \bar{\mathcal{A}}) - \mathbb{V}_{s'|s} \mathcal{R}'(\cdot) \right\} \frac{B' - \mathcal{T}'}{q(\cdot)} \\ &= (-) \underbrace{\frac{\kappa}{r^f} \mathbb{V}_{s'|s} \mathcal{R}'(\cdot) \frac{B' - \mathcal{T}'}{q(\cdot)}}_{\equiv \eta(y, \tau, B'; \kappa)} - \frac{\Delta \mathbb{E}_{s'|s} \mathcal{R}'(\cdot) - \kappa \Delta \mathbb{V}_{s'|s} \mathcal{R}'(\cdot) \times (B' - \mathcal{T}(\tau, B') - \bar{\mathcal{A}})}{\Delta \mathcal{T}' / (B' - \mathcal{T}')} \frac{1}{r^f q(\cdot)}, \end{aligned}$$

where $\Delta \mathbb{E}_{s'|s} \mathcal{R}'(\cdot) \equiv \mathbb{E}_{s'|s} \mathcal{R}(y, \tau', B') - \mathbb{E}_{s'|s} \mathcal{R}(y, \tau, B')$, and $\Delta \mathbb{V}_{s'|s} \mathcal{R}'(\cdot) \equiv \mathbb{V}_{s'|s} \mathcal{R}(y, \tau', B') - \mathbb{V}_{s'|s} \mathcal{R}(y, \tau, B')$ denote changes in expected repayment and variance, given the change in \mathcal{T}' . More succinctly,

$$\hat{\eta}(y, \tau, B'; \kappa) = \eta(y, \tau, B'; \kappa) + \alpha(y, \tau, B'; \kappa), \quad (\text{A23})$$

where the function $\alpha(\cdot)$ is given by

$$\alpha(y, \tau, B'; \kappa) = (-) \frac{\Delta \mathbb{E}_{s'|s} \mathcal{R}'(\cdot) - \kappa \Delta \mathbb{V}_{s'|s} \mathcal{R}'(\cdot) \times (B' - \mathcal{T}(\tau, B') - \bar{\mathcal{A}})}{\Delta \mathcal{T}' / (B' - \mathcal{T}(\tau, B'))} \frac{1}{r^f q(\cdot)}.$$

From this decomposition, it is clear that the information contained in bond price changes alone is not enough to identify the slope parameter κ (and thus the structural elasticity) as they are confounded by changes in expected future payoffs.

In the main text, for ease of exposition, we consider a case in which $\alpha(y, \tau, B'; \kappa) \approx (-) \frac{\Delta \mathbb{E}_{s'|s} \mathcal{R}'(\cdot)}{\Delta \mathcal{T}' / (B' - \mathcal{T}(\tau, B'))} \frac{1}{r^f q(\cdot)}$. This approximation holds in our quantitative analysis, since our

parameterization for $\bar{\mathcal{A}}$ implies that $B' - \mathcal{T}(\tau, B') - \bar{\mathcal{A}}$ is close to zero.

B Empirical Appendix

B.1 Data and Details behind the EMBIGD

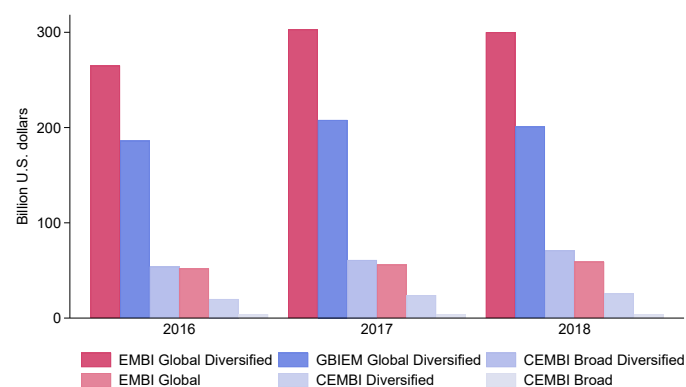
As described in the main text, we collect data from three main sources: Datastream, J.P. Morgan, and Morningstar. Datastream provides daily bond prices, spreads, and daily returns for all the bonds included in the EMBIGD. From J.P. Morgan Markets, we obtain index weights, face amounts of bonds in the index, and various bond characteristics, such as maturity and duration.

The EMBIGD includes bonds with an original maturity of at least 2.5 years and a minimum face amount outstanding of US\$500 million. To be classified as an emerging economy, a country’s gross national income (GNI) per capita must remain below an Index Income Ceiling (IIC) for three consecutive years. The IIC is defined by J.P. Morgan and updated annually based on the growth rate of the World GNI per capita (Atlas method, current US\$), as published by the World Bank. Bonds must settle internationally and have accessible and verifiable bid and ask prices. Once included, they may remain in the index until 12 months before maturity. Local law instruments are not eligible.

Figure B1 shows the assets under management (AUM) tracking the EMBIGD and other major benchmark indexes. Among emerging-market bond indexes, the EMBIGD is the most widely used and, as of 2018, was followed by funds managing approximately US\$300 billion. To put this in perspective, Figure B2 reports 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. For most countries, this share exceeds 60%—highlighting the index’s central role in external sovereign debt markets.

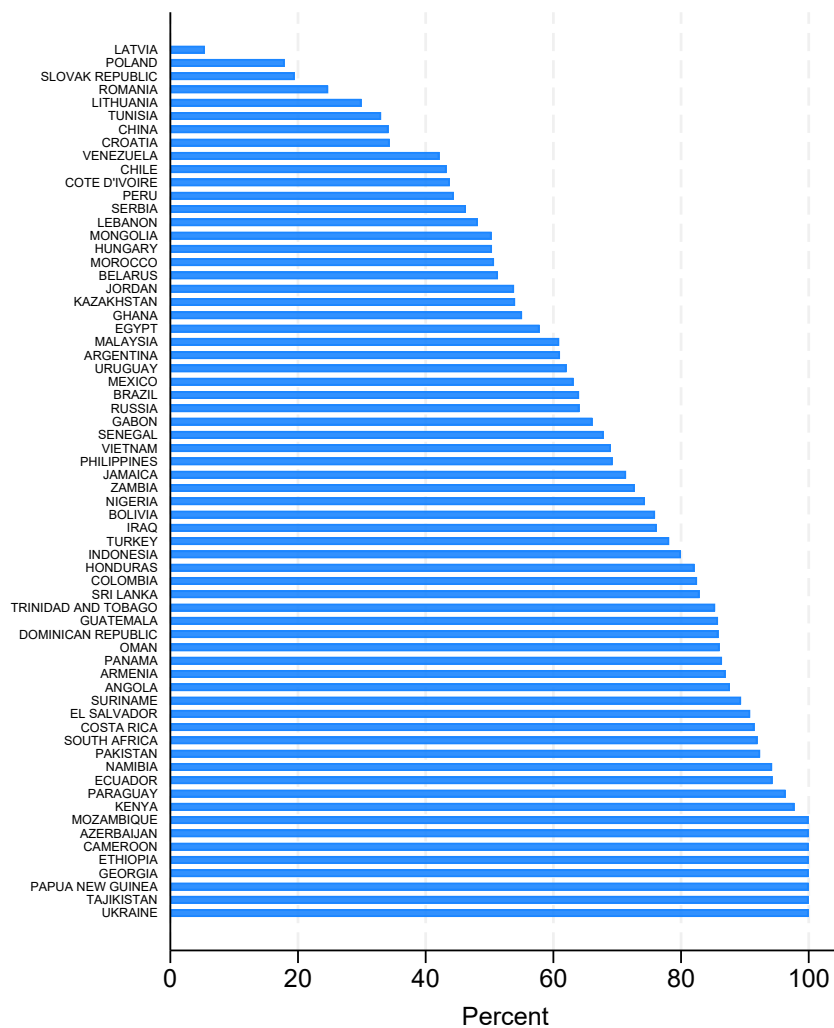
For our main analysis, we apply the following data cleaning procedures. First, we exclude extreme values of daily returns, stripped spreads, and $Z_{c,t}$. Specifically, we drop stripped spreads below 0 or above 5,000 basis points, as well as observations below the 5th or above the 95th percentile of 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 B1: AUM benchmarked to emerging economies bond indexes



Note: The figure shows assets under management (in billions of U.S. dollars) benchmarked to emerging market bond indexes in 2018.

Appendix Figure B2: Share of U.S. dollar-denominated emerging economies sovereign debt



Note: Bars show the share of U.S. dollar-denominated sovereign debt in the EMBI Global index relative to each country's international government debt securities. Averages are computed by first calculating the share by country and year, then averaging across 2016–2022.

B.2 Diversification Methodology and Timeline

Relative to a market capitalization-weighted index, the EMBIGD employs a diversification methodology designed to achieve a more balanced distribution of country weights. This approach prevents countries with large market capitalizations from dominating the index. To accomplish this, the methodology 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 face value of bonds across countries included in the index, referred to as the Index Country Average (ICA). It is defined as:

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

where $B_{c,t}$ denotes the total face value of bonds from country c included in the index at time t , and C 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:

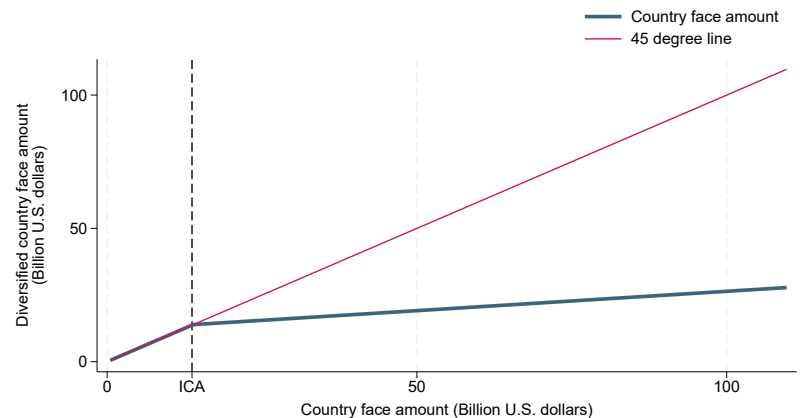
$$f_{c,t}B_{c,t} = \begin{cases} ICA_t \times 2 & \text{if } B_{c,t} = \max\{B_{c,t}\}_{\forall c} \\ ICA_t + \frac{ICA_t}{\max\{B_{c,t}\}_{\forall c} - ICA_t} (B_{c,t} - ICA_t) & \text{if } B_{c,t} > ICA_t \\ B_{c,t} & \text{if } B_{c,t} \leq ICA_t. \end{cases} \quad (B1)$$

This formula implies that $f_{c,t} = 1$ for countries whose face amount is below the index average. For the other countries, $f_{c,t} < 1$, which implies that a proportional reduction is applied uniformly across all of their bonds in the index.³¹ The index's country weights are then computed as $\frac{f_{c,t}B_{c,t}}{\sum_{n \in \mathcal{I}} f_{n,t}B_{n,t}}$. To illustrate this methodology, Figure B3 compares the diversified and non-diversified face amounts across different values of $B_{c,t}$, using data from December 2018 to compute the ICA.

Figure B4 lists the countries included in the EMBIGD and their associated index weights (red bars). For comparison, it also shows the weights for the EMBI Global (EMBIG), which follows the same bond inclusion criteria. The key difference is weighting: the EMBIG uses market capitalization, while the EMBIGD applies a cap rule to limit the weights of countries with above-average debt outstanding.

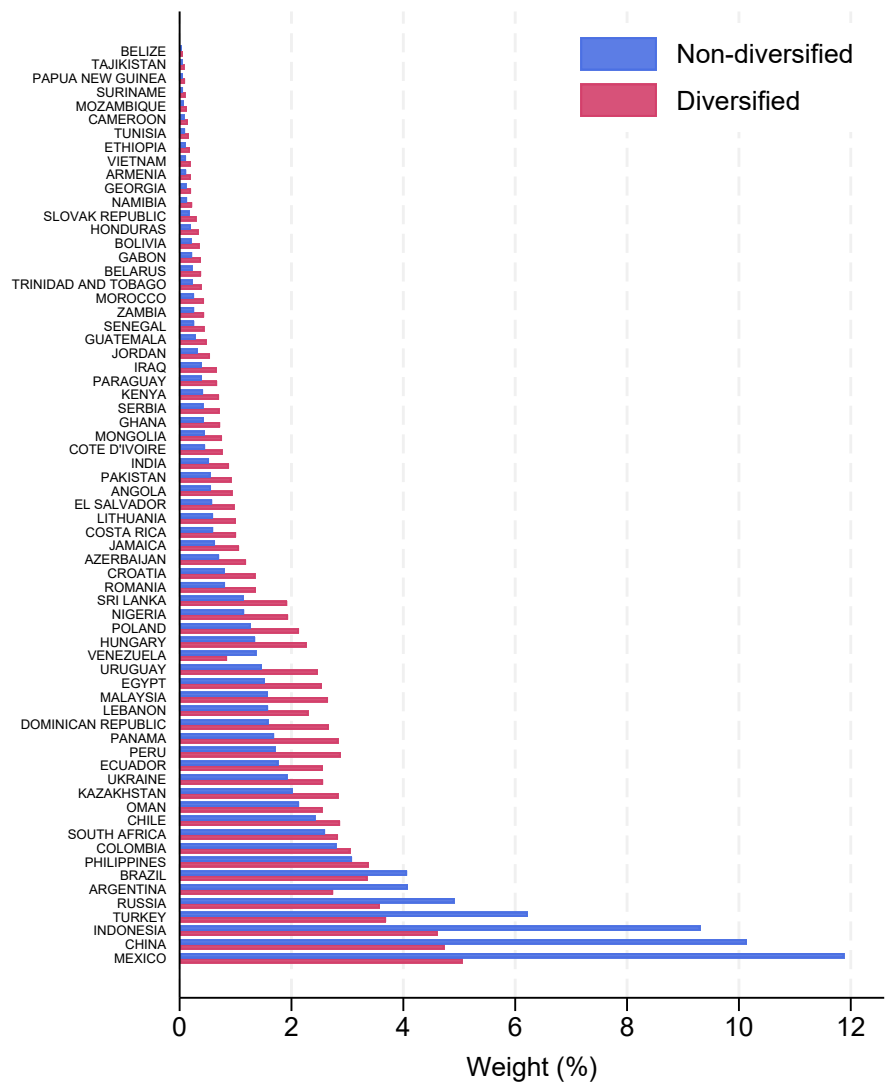
³¹In addition to this rule, country weights are capped at 10%. Any excess weight above this cap will be redistributed pro rata to smaller countries below the cap, across all bonds from countries not capped at 10%.

Appendix Figure B3: Effect of the diversification methodology on the country face amount



Note: The figure compares the diversified and non-diversified face amounts across different values of $B_{c,t}$, using data from December 2018 to compute the ICA.

Appendix Figure B4: EMBIGD: Country-level weights

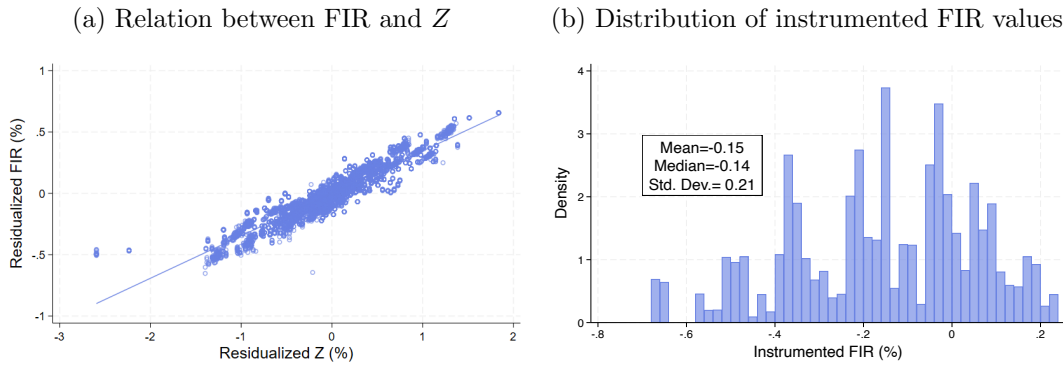


Note: The figure shows the EMBI Global Diversified and Non-diversified weights for December 2018.

B.3 Timeline of Events and the Instrumented FIR

In Panel (a) of Figure B5, we present 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 relationship, with an R^2 of 86%. Panel (b) displays the distribution of the instrumented FIR measure, which ranges from -0.7% to 0.25% and is skewed toward negative values. Since our analysis focuses on countries with a constant face amount, the predominance of bond inclusions over exclusions from other countries generally leads to a reduction in the index weight assigned to our sample countries.

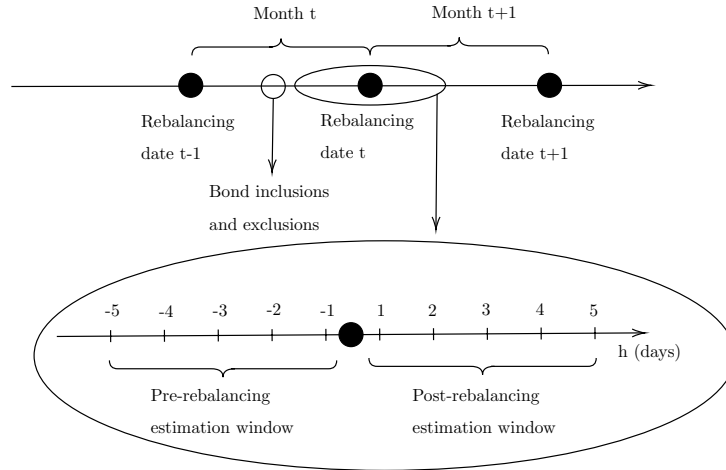
Appendix Figure B5: Flows implied by rebalancing (FIR)



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. The FIR is computed as described in Equation (18) and Z is calculated according to Equation (19). Panel (b) shows a histogram of the instrumented FIR. The sample period is 2016–2018.

Figure B6 illustrates the timeline of events within each rebalancing episode. Changes in EMBIGD weights take place on the last U.S. business day of each month. In our regressions, we define the event window as the 5-day period surrounding each rebalancing date.

Appendix Figure B6: Timeline of events



Note: The figure describes the timeline of events within each rebalancing episode.

B.4 Additional Analysis and Results

We present a set of robustness checks that complement our main analysis. First, Table B1 reports estimates from an OLS regression in which the FIR is not instrumented. Table B2 shows results from an alternative specification where bond prices are held constant when constructing the FIR, using prices from the previous rebalancing period. In this case, as in the baseline, variation comes only from quantities—not from prices. In both exercises, the estimated effects are slightly smaller than those reported in the main text.

Our findings are also robust to alternative event windows around rebalancing dates (Table B3). Consistent with the results in Figure 3, estimates decrease slightly in magnitude for shorter windows and increase as the window is extended. In addition, the results remain unchanged when quasi-sovereign bonds are excluded from the analysis (Table B4).

Appendix Table B1: OLS regression

Dependent Variable: Log Price						
	[-5:+5]			No h=-1		
FIR X Post	0.197*** (0.076)	0.198** (0.077)	0.197*** (0.076)	0.134 (0.086)	0.236*** (0.076)	0.167* (0.092)
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 the FIR measure around rebalancing events. Standard errors are clustered at the country-month level, and the sample covers the 2016–2018 period. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. See the footnote in Table 3 for additional details.

Appendix Table B2: OLS regression - Fixed bond prices

Dependent Variable: Log Price						
	[-5:+5]			No h=-1		
FIR X Post	0.230*** (0.075)	0.232*** (0.076)	0.231*** (0.075)	0.142* (0.086)	0.266*** (0.075)	0.177* (0.092)
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 an alternative FIR measure that holds bond prices constant. Standard errors are clustered at the country-month level, and the sample period is 2016–2018. *, **, and *** denote statistically significant at the 10%, 5%, and 1% level, respectively. See footnote in Table 3 for additional details.

Appendix Table B3: Log price and FIR: different event windows

Panel A-Dependent Variable: Log Price					Panel B-Dependent Variable: Log Price (Excl. h=-1)				
	[-2:+2]	[-3:+3]	[-4:+4]	[-5:+5]		[-2:+1]	[-3:+2]	[-4:+3]	[-5:+4]
FIR X Post	0.146*** (0.053)	0.197*** (0.071)	0.221** (0.086)	0.231** (0.099)	FIR X Post	0.220*** (0.056)	0.257*** (0.074)	0.271*** (0.087)	0.263*** (0.098)
Bond-Month FE	Yes	Yes	Yes	Yes	Bond-Month FE	Yes	Yes	Yes	Yes
Observations	42,217	63,327	84,435	105,548	Observations	21,106	42,216	63,325	84,433
N. of Bonds	738	738	738	738	N. of Bonds	738	738	738	738
N. of Countries	68	68	68	68	N. of Countries	68	68	68	68
N. of Clusters	1,576	1,576	1,576	1,576	N. of Clusters	1,576	1,576	1,576	1,576
F (FS)	1,660	1,662	1,664	1,666	F (FS)	1,667	1,667	1,669	1,670

Note: The table presents 2SLS estimates of log bond prices on the instrumented FIR measure, around rebalancing events. Each column reports the estimates for different h -day symmetric windows. Standard errors are clustered at the country-month level, and the sample period is 2016–2018. *, **, and *** denote statistically significant at the 10%, 5%, and 1% level, respectively. See footnote in Table 3 for additional details.

Appendix Table B4: Log price and FIR: dropping quasi-sovereign bonds

Dependent Variable: Log Price				
FIR X Post	0.249** (0.107)	0.249** (0.108)	0.249** (0.107)	0.175* (0.103)
Bond FE	Yes	Yes	No	No
Month FE	Yes	No	No	No
Bond Characteristics-Month FE	No	Yes	No	No
Country-Month FE	No	Yes	No	No
Bond-Month FE	No	No	Yes	Yes
Month-Post FE	No	No	No	Yes
Bond Controls	No	Yes	No	No
Observations	73,140	73,100	73,140	73,140
N. of Bonds	430	430	430	430
N. of Countries	65	65	65	65
N. of Clusters	1,513	1,512	1,513	1,513
F (FS)	0	3,151	3,231	1,099

Note: The table presents 2SLS estimates of log bond prices on the instrumented FIR measure, around rebalancing events. We exclude from the analysis quasi-sovereign bonds. Results correspond to a 5-day symmetric window. Standard errors are clustered at the country-month level, and the sample period is 2016–2018. *, **, and *** denote statistically significant at the 10%, 5%, and 1% level, respectively. See footnote in Table 3 for additional details.

B.5 Analysis across Bond Maturities

In Table B5, we examine heterogeneity in bond price reactions across maturity buckets. Bonds are classified into three groups—short, medium, and long maturity—based on the 33rd and 66th percentiles of the maturity distribution. Short-term bonds have less than 4.8 years remaining, long-term bonds have more than 8.8 years, and the remainder are classified as medium maturity. We then re-estimate our baseline specification separately for each group, using the same event-study setup and instrumented FIR measure.

We find that price responses to the FIR increase sharply with bond maturity. For short-term bonds, a 1 p.p. increase in the FIR is associated with a modest and statistically insignificant price effect. In contrast, long-term bonds exhibit price reactions that are about four times larger—around 0.41%. These results are consistent with [Broner, Lorenzoni, and Schmukler \(2013\)](#), who show that short-term bonds are more tightly anchored to policy rates and may be less sensitive to broader shifts in investor flows. As argued in the main text, the stronger response for long-term bonds may also reflect endogenous changes in expected repayment induced by the supply-shifting shock.

Appendix Table B5: Demand elasticities across maturities

Dependent Variable: Log Price						
	Long Maturity		Medium Maturity		Short Maturity	
FIR X Post	0.411**	0.412**	0.303***	0.303***	0.095	0.093
	(0.198)	(0.196)	(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
F (FS)	502	989	678	2,569	729	1,274

Note: The table presents 2SLS estimates of log bond prices on the instrumented FIR measure 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. The coefficients for *Post* and FIR are included in the estimation but not reported for brevity. Standard errors are clustered at the country-month level. *, **, and *** denote statistically significant at the 10%, 5%, and 1% level, respectively.

B.6 FIR and Changes in Credit Default Swaps

In Table B6, we assess whether changes in the FIR affect not only bond prices but also measures of credit risk. Specifically, we regress 5-year CDS spreads on the FIR using the same 2SLS setup and five-day window as in the main text. As in the baseline specification, we include a post-rebalancing dummy interacted with the instrumented FIR and estimate variants with and without fixed effects.

Across specifications, we find that an increase in the FIR is associated with a statistically significant decline in CDS spreads, albeit at the 10% level. The point estimates imply that a 1 p.p. increase in the FIR reduces CDS spreads by about 0.8% in log terms. While somewhat noisy, these effects suggest that the supply shock affects not only market prices but also investor perceptions of sovereign credit risk.

Together with the differential bond price reactions across maturities, the decline in CDS spreads further supports the idea that FIR-induced price increases partly reflect improvements in expected future repayments, reinforcing the need to distinguish between reduced-form and structural elasticities when interpreting price responses.

Appendix Table B6: Log CDS and FIR

Dependent variable: log CDS			
FIR X Post	-0.796 *	-0.796 *	-0.796 *
	(0.448)	(0.449)	(0.448)
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: This table shows 2SLS estimates of five-year log CDS of countries on the FIR measure (Equation (18)), instrumented by Z (Equation (19)), around rebalancing dates. The first- and second-stage equations are described in Equation (20). The estimations use a symmetric five-trading-day window, with $Post$ as an indicator variable (equal to 1 for the five trading days after rebalancing, and 0 otherwise). Month fixed effects are dummy variables equal to 1 for each rebalancing month, and 0 otherwise. Standard errors are clustered at the country-month level, and the sample period is 2016–2018. *, **, and *** denote statistically significant at the 10%, 5%, and 1% level, respectively.

B.7 Yield to Maturity Responses to Rebalancing Shocks

We can use the estimated price responses in Table 3 to compute a back-of-the-envelope semi-elasticity of bond yields with respect to the FIR. Using the approximation $\frac{\Delta Y_t^i}{FIR_{c,t}} \approx \frac{\hat{\eta}^i(1+Y_t^i)}{D_t^i}$, where Y_t^i and D_t^i denote the bond’s yield and duration, respectively, and plugging in the average values from our sample (5% yield and 6.4 duration), we obtain a semi-elasticity of approximately 5 basis points.

In Table B7, we re-estimate the baseline specification in Equation (20) using yield to maturity as the dependent variable, finding results that closely align with those based on bond prices. In Table B8, we extend the analysis by dividing the sample into three groups based on country spread levels and re-estimating the baseline specification with yield to maturity as the dependent variable. The results remain quantitatively similar to those of our baseline across all three groups.

Appendix Table B7: Effects of FIR on yield to maturity

Dependent Variable: Yield to Maturity						
	[-5:+5]				No h=-1	
FIR X Post	-3.511** (1.671)	-3.541** (1.693)	-3.543** (1.668)	-5.485** (2.604)	-4.035** (1.667)	-5.769** (2.588)
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
F (FS)	645	1,640	1,688	470	1,690	470

Note: The table presents 2SLS estimates of yield to maturity on the instrumented FIR measure around rebalancing events. *Post* is an indicator equal to 1 for the five trading days following each rebalancing, and 0 otherwise. *Bond Characteristics* refer to fixed effects formed by interacting maturity bins, credit rating grades, and bond type 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, and bond type indicators distinguish between sovereign and quasi-sovereign bonds. *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. Coefficients for *Post* and FIR are included in the regression but omitted from the table for brevity. Standard errors are clustered at the country-month level. The sample covers the 2016–2018 period. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Appendix Table B8: Effects of FIR on yield to maturity: The role of default risk

Dependent Variable: Yield to Maturity						
	High Spread		Median Spread		Low Spread	
FIR X Post	-10.855** (4.467)	-6.190** (2.797)	-4.065** (2.043)	-3.693* (1.998)	-2.839* (1.592)	-2.356 (1.569)
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	26,620	26,618	26,488	26,488	26,472	26,472
N. of Bonds	379	379	442	442	361	361
N. of Countries	57	57	50	50	42	42
N. of Clusters	976	976	818	818	609	609
F (FS)	480	2,267	390	714	386	871

Note: The table presents 2SLS estimates of yield to maturity on the instrumented FIR measure 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. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

C Quantitative Appendix

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

C.1 Secondary Markets

The empirical elasticity computed in Section 3 exploits exogenous variation in the passive demand in a small window around announcements of changes in the EMBIGD weights. To tightly link our model with the empirical analysis, our baseline model in Section 2 already incorporates a passive demand and exogenous changes in index weights, τ . There is, however, a frequency disconnect in the sense that the model is calibrated at quarterly frequency, making it unsuitable for quantifying high-frequency price reactions to changes in τ .

To address this frequency disconnect, we introduce secondary markets in the model. This extension allows us to capture the high-frequency and short-term nature of our empirical elasticity. In particular, we consider two instances of trading in secondary markets within each period: before and after the realization of the index weights, τ' . The timing assumption 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 market open. Let $q^{SM,0}(y, \tau, B')$ denote the opening price.*
4. *The next-period index weights τ' are realized.*
5. *The secondary market closes. Let $q^{SM,1}(y, \tau', B')$ denote the bond closing price.*

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 in this instance is given by

$$q^{SM,0}(y, \tau, B') = \beta^* \mathbb{E}_{y', \tau' | y, \tau} \mathcal{R}(y', \tau', B') \Psi^{SM,0}(y, \tau, B'). \quad (C1)$$

The term $\mathbb{E}_{y', \tau' | y, \tau} \mathcal{R}(y', \tau', B')$ represents the expected next-period repayment of the bond, conditional on the information available when the secondary market opens. Following the derivation in sections 2.2 and 2.3, the downward-sloping component of the price function is

$$\Psi^{SM,0}(y, \tau, B') = 1 - \kappa_0 \frac{\mathbb{V}_{\{y', \tau'\} | \{y, \tau\}} \mathcal{R}(y', \tau', B')}{\mathbb{E}_{\{y', \tau'\} | \{y, \tau\}} \mathcal{R}(y', \tau', B')} (B' - \mathcal{T}(\tau, B') - \bar{A}) \quad (C2)$$

Notice that $q^{SM,0}(y, \tau, B')$ coincides with the price in the primary market $q(y, \tau, B')$, 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 the new index weights τ' are realized. In this case, the bond price is

$$q^{SM,1}(B', y, \tau') = \beta^* \mathbb{E}_{y'|y} \mathcal{R}(y', \tau', B') \Psi^{SM,1}(y, \tau', B'). \quad (C3)$$

The term $\mathbb{E}_{y'|y} \mathcal{R}(y', \tau', B')$ is the expected next-period repayment of the bond, conditional on the information available when the secondary market closes. This term is analogous to the one in Equation (C1), but incorporates the information provided by the realization of τ' . Similarly, the downward-sloping component of the price function is given by

$$\Psi^{SM,1}(y, \tau', B') \equiv 1 - \kappa_0 \frac{\mathbb{V}_{y'|y} \mathcal{R}(y', \tau', B')}{\mathbb{E}_{y'|y} \mathcal{R}(y', \tau', B')} (B' - \mathcal{T}(\tau', B') - \bar{\mathcal{A}}). \quad (C4)$$

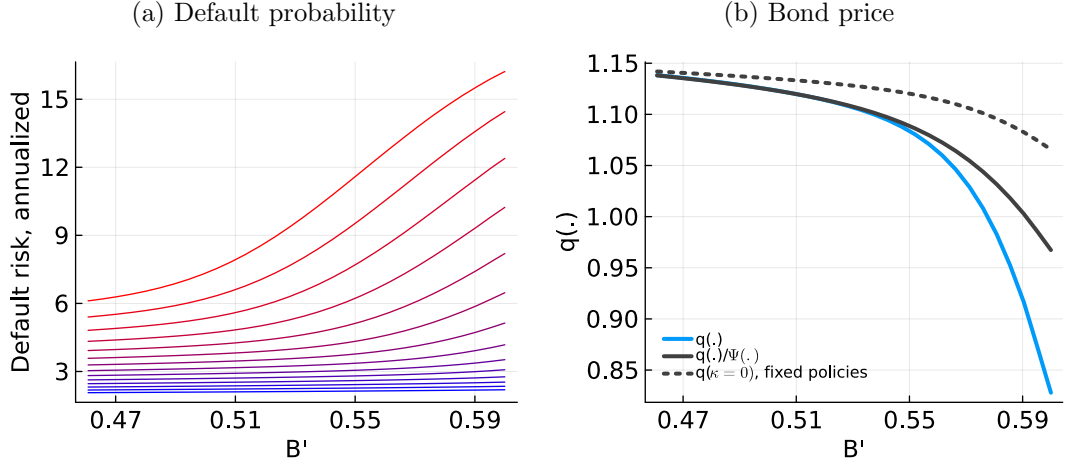
Notice that, the only difference between $q^{SM,1}$ and $q^{SM,0}$ arises from the update of τ' since both the endowment and the stock of debt remain fixed while the secondary market is open. Moreover, in the absence of secondary markets, the timing assumption is exactly the same as in the baseline model.

C.2 Solution Method

We employ a global solution method to solve our quantitative 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 consisting of 250 equally spaced points between $\underline{B} = 0$ and $\bar{B} = 1.2$. We ensure that \bar{B} is sufficiently large so that it never binds in our simulations. The steps of the algorithm are as follows:

1. We start with a guess for the value functions $V^r(y, \tau, B)$ and $V^d(y)$. We also guess the bond price function $q(y, \tau, B')$ as a function of the end-of-period stock of debt, B' .
2. Based on these guesses, we solve for the optimal bond policy $B'(y, \tau, B)$, as described in Equation (7). To this end, we use an optimizing algorithm 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, $V^d(y)$, as given by Equation (8).
4. We solve for the government's optimal default choice, as shown in Equation (6). As standard in the literature, we convexify the default decision to achieve convergence. We assume that in each period, the government's value function $V^d(y)$ is subject to an

Appendix Figure C1: Default risk and bond prices



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' . The solid blue line shows the bond pricing kernel $q(y, \tau, B')$, evaluated at the mean values of y and τ . The solid black line depicts $\frac{1}{r_f} \mathbb{E}_{s'|s}(\mathcal{R}'(\cdot))$ —that is, $q(\cdot)/\Psi(\cdot)$. The dashed black line shows the bond price in a counterfactual scenario in which government policies are those of our baseline but bonds are priced by perfectly elastic investors.

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.

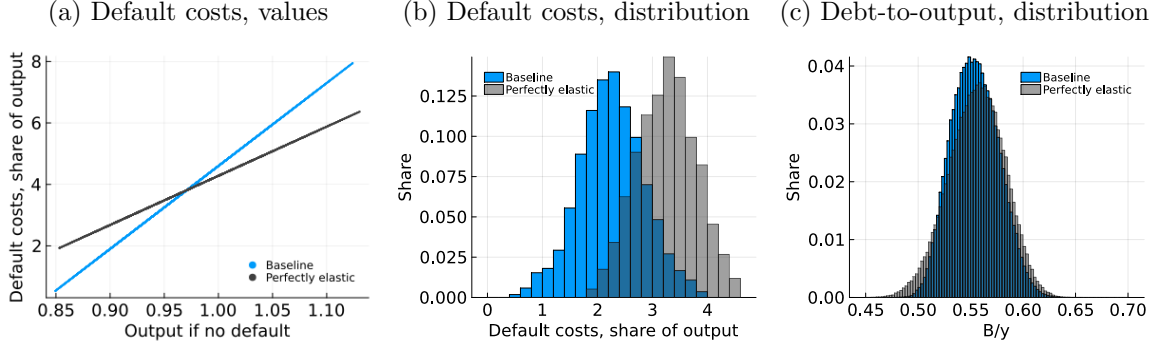
5. Taking the policy functions $B'(y, \tau, B)$ and $d(y, \tau, B)$ as given, we update $q(y, \tau, B')$ according to Equation (9). We use cubic splines to evaluate the right-hand side of the pricing equation at $B'' \equiv B'(y', \tau', B')$.
6. We iterate over the previous steps until convergence of the value functions and the bond price function.

C.3 Additional Figures and Tables

Figure C1 illustrates the default probability and the bond price function $q(y, \tau, B')$ across different values of B' and y . Panel (a) shows that default risk rises with B' and falls with y : the government is more likely to default in states with high debt and low output. Accordingly, Panel (b) shows that the bond price decreases with B' and increases with y (blue line).

The figure also highlights the role of a downward-sloping demand on bond prices. The solid black line in Panel (b) shows the expected repayment per unit of the bond, $\frac{1}{r_f} \mathbb{E}_{s'|s}(\mathcal{R}'(\cdot))$, which is equal to $q(\cdot)/\Psi(\cdot)$. At low B' levels—where default risk and the variance of payments are minimal—bond prices remain largely unaffected by the downward-sloping demand term

Appendix Figure C2: Default costs: Comparison with fully recalibrated perfectly elastic model



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.

$\Psi(\cdot)$. However, as B' increases, the larger repayment volatility lowers the $\Psi(\cdot)$ term, which pushes the bond price $q(\cdot)$ well below $\frac{1}{rT} \mathbb{E}_{s'|s}(\mathcal{R}'(\cdot))$. The dotted line depicts bond prices under a counterfactual scenario in which we maintain the baseline $\{B'(\cdot), d(\cdot)\}$ policies but assume that bonds are priced by perfectly elastic investors—i.e., $\kappa = 0$. This adjustment affects not only the current value of $\Psi(\cdot)$ but also its future path. In this case, bond prices exhibit much lower curvature relative to our baseline. This simple analysis highlights the dynamic effects of a downward-sloping demand on bond prices: not only does the current $\Psi(\cdot)$ matter for pricing, but so does the entire expected future path of $\Psi(\cdot)$.

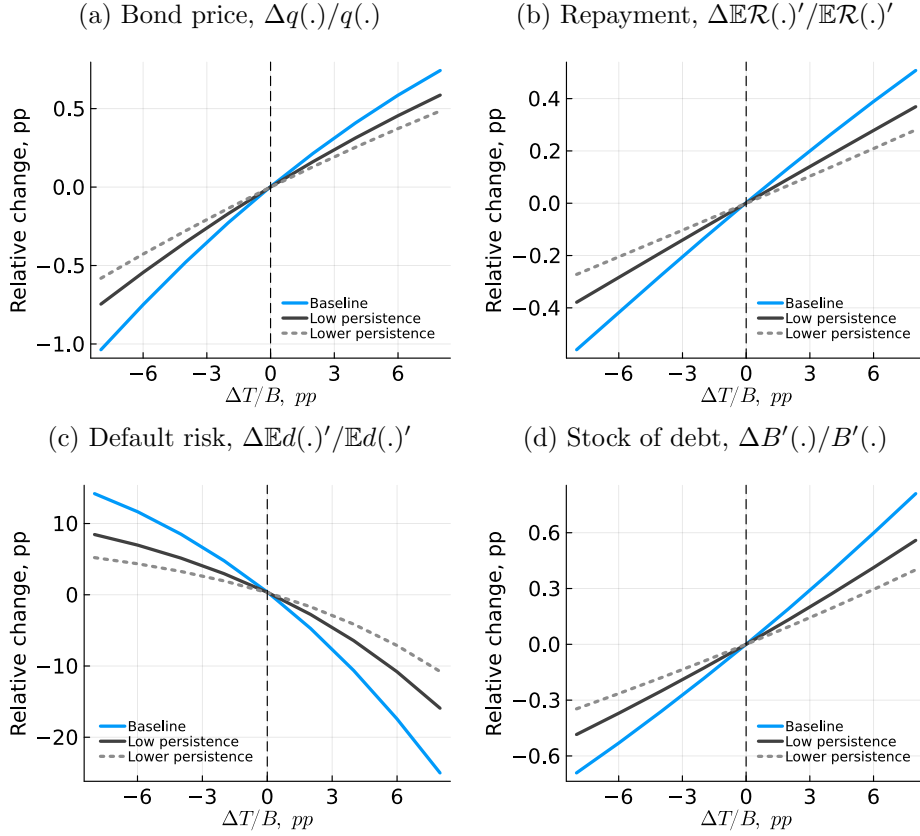
Figure 7 in the main text compares default costs and the implied debt-to-output distribution between our baseline model and a perfectly elastic case in which we recalibrate $\Delta\{d_0, d_1\}$ to match the average and volatility of spreads. Figure C2 presents a similar exercise, but also recalibrates the discount factor β to match the targeted average debt-to-output ratio in the data. The implied default cost parameters (Panels a and b) are similar to those in Figure 7. By slightly increasing β (from 0.948 to 0.95), the perfectly elastic case generates a debt-to-output distribution comparable to that of our baseline model (as shown in Panel c).

C.4 Decomposing the Demand Elasticity

What accounts for the gap between the reduced-form and structural elasticity? To answer this question, we examine the mechanisms driving changes in the expected repayment function following a shift in τ .

Figure C3 illustrates the effects of shifts in passive demand. To construct these panels,

Appendix Figure C3: Effects of Changes in demand on prices and policies



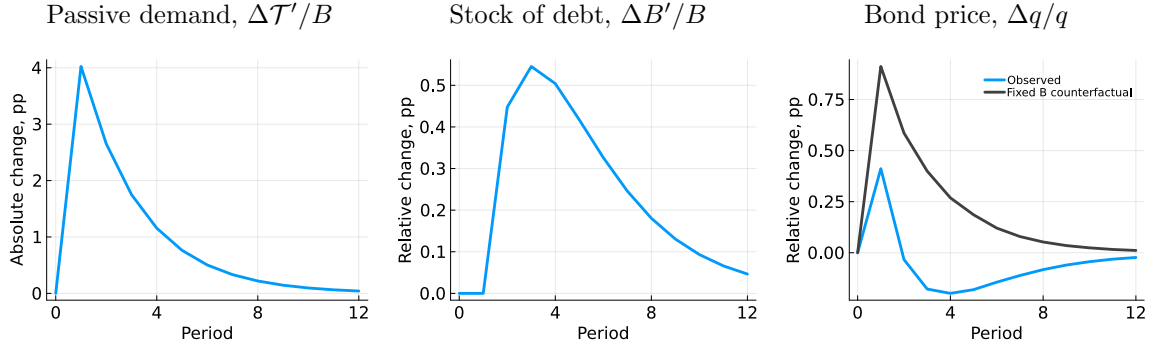
Note: The figure shows how changes in the passive demand (i.e., FIR) affect bond prices, expected repayment, one-period ahead default risk, and the bond supply. The blue lines show results under our baseline calibration. The gray lines show results for parameterizations in which we decrease the persistence of the FIR. For these cases, we set $\rho_\tau = 0.50$ and $\rho_\tau = 0.30$. In all cases, we evaluate these changes at the mean value for endowment and debt.

we evaluate, at the beginning of the period, bond prices $q^{SM,0}(y, \tau, B')$, expected repayment $\mathbb{E}_{y', \tau' | y, \tau} \mathcal{R}'(y, \tau, B')$, default risk $\mathbb{E}_{y', \tau' | y, \tau} d'(y, \tau', B')$, and the bond policy $B'(y, \tau, B)$ at the mean values of y , τ , and B . We then recompute these variables at the end of the period for different realizations of τ' , namely $q^{SM,1}(y, \tau', B')$, $\mathbb{E}_{y' | y} \mathcal{R}'(y, \tau', B')$, $\mathbb{E}_{y' | y} d'(y, \tau', B')$, and $B'(y, \tau', B)$. The y -axes report the relative change for each variable, while the x -axes display changes in passive demand relative to the debt stock, $\Delta \mathcal{T}/B = \tau' - \tau$.

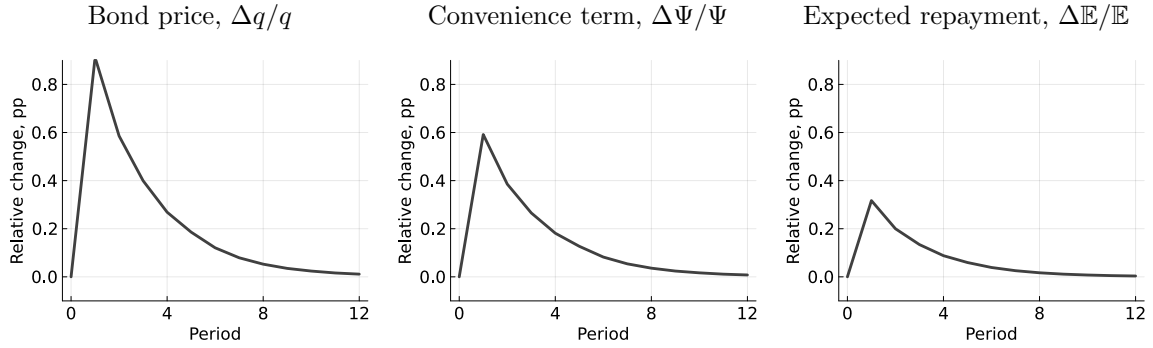
The blue lines correspond to our baseline calibration. The solid and dashed gray lines represent alternative scenarios with lower persistence in the $\{\tau\}$ process. Panels (a) and (b) reveal a clear monotonic relationship between changes in passive demand, bond prices, and expected repayment. The change in expected repayment, which is large relative to the total price response, explains much of the gap between the reduced-form and model-implied elasticities, as described in Equation (17). As the persistence of τ declines, the effects on bond prices—and especially on expected repayment—become more muted.

Appendix Figure C4: Impulse response to an increase in the 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 shows a decomposition for bond price changes across time, in a counterfactual in which the stock of debt remains fixed.

The change in expected repayment reflects changes in both default risk and future debt issuance. Panel (c) shows that default risk falls as passive demand increases, while Panel (d) illustrates the government's debt response to τ' . The debt adjustment is positive but less than proportional: a 5% increase (decrease) in the passive share leads to a debt expansion (contraction) of less than 1%. Investors anticipate this response and adjust their expectations, which affects current bond prices. As the persistence of the $\{\tau\}$ process declines, the magnitude of these forces decreases.

To further explore the dynamics of bond prices and debt, Figure C4 presents the impulse response to a positive τ shock, simulated from the model. The top panel shows that debt rises on impact and remains elevated, while bond prices increase by about 0.40% initially. Over time, as τ reverts, bond prices gradually decline due to the larger bond supply and the resulting increase in default risk.

For comparison, the black lines show a counterfactual in which bond issuance is held fixed. In that case, the initial bond price response is nearly twice as large, at around 0.80%. The bottom panel decomposes this response, highlighting the role of the inconvenience term

component, $\Psi(\cdot)$. Since the passive demand shock is persistent, Ψ rises initially, then declines as the shock dissipates. More importantly, the expected dynamics of $B'(\cdot)$ and $d(\cdot)$ shape the trajectory of expected repayments, as shown in the right-hand panel.