The Fiscal Foundations of Safe Assets*

Adrien d'Avernas† Egemen Eren‡ Quentin Vandeweyer§

September 21, 2026

Preliminary draft. Please do not distribute.

Abstract

Government bonds hedge equity risk in some countries and periods and not in others. This paper shows that the difference depends systematically on the government's fiscal position. In a standard consumption-based model with sovereign default risk, a heavier fiscal burden raises the stock-bond correlation, because bad news then raises default risk more than it lowers the risk-free rate. In a panel of 44 advanced and emerging economies since 1960, the stock-bond correlation rises with a country's forward-looking interest expense, our measure of its fiscal burden. It accounts for a quarter of the cross-country variation in average correlations and predicts the correlation within countries. A quasi-experiment in the euro area, where countries share one monetary policy, corroborates this result: after the Greek deficit revision of October 2009, correlations in the periphery rose by 0.44 relative to the core. Two text-based measures of fiscal news, rating actions taken for fiscal reasons and the fiscal salience of newspaper coverage, predict the correlation in the same direction. Within countries, the relationship runs through the credit component of the yield, which we isolate with sovereign credit default swaps.

JEL Classification: E43, E62, G12, H63

Keywords: stock-bond correlation, interest expense, safe assets, fiscal policy, sovereign risk

*We thank seminar participants at the BIS and ESM. The views expressed are those of the authors and do not necessarily reflect those of the Bank for International Settlements. We used AI tools for research assistance, text classification and editorial support. All errors are our own.

†Stockholm School of Economics. Email: adrien.davernas@hhs.se

‡Bank for International Settlements. Email: egemen.eren@bis.org

§University of Chicago Booth School of Business. Email: quentin.vandeweyer@chicagobooth.edu

1 Introduction

A government bond is a safe asset when it appreciates in bad states of the world, hedging equity risk for its holders (Brunnermeier et al., 2024; Gómez-Cram et al., 2024). Whether a bond has this property is a first-order determinant of the risk premium investors require to hold it, and hence of the government's cost of borrowing (Leombroni et al., 2026). Few bonds had it through the 1990s, when stocks and bonds moved together across advanced and emerging economies alike. After 2000, bonds in the advanced economies became hedges against equity risk (Campbell et al., 2020) and bonds in the emerging market economies did not (Figure 1, panel a). After 2022 the hedge weakened across the advanced economies. The same divergence appears within a single currency union. The countries of the euro-area core and periphery share a central bank, a currency and an inflation target, and during the sovereign debt crisis the hedging property of their bonds diverged (Figure 1, panel b).

In this paper we relate the variation in the stock-bond correlation across countries and over time to the government's fiscal burden. We first proxy the burden as the forward-looking interest expense ratio: the debt-to-GDP ratio times the ten-year yield, the interest bill the government would owe if its whole debt stock were refinanced at today's rate.[1] The ratio is computed from government accounts and market yields alone, and it is the observable counterpart of the statistic through which the fiscal channels of our model operate. Where the fiscal burden is low, bonds hedge equity risk. Where it is high, stocks and bonds move together. We establish the pattern in four steps: a panel of 44 advanced and emerging economies that reaches back to 1960, a quasi-experiment in the euro area, two measures of the fiscal position built from text, and a decomposition of the yield inside the ratio.

That statistic comes from a standard consumption-based asset pricing model with sovereign default risk and a convenience yield on government debt. In the model, a single productivity shock moves both bond prices and the government's fiscal position. A bad shock to current output lowers the real risk-free rate, raising bond prices and pulling the stock-bond correlation toward negative territory: the safe-haven channel. The same shock raises the government's fiscal burden, which lowers the bond price and pushes

[1] Furman and Summers (2020), p. 9: "[F]iscal sustainability cannot be assessed by traditional debt-to-GDP ratios but should instead be understood with measures like nominal or real interest as a share of GDP."

Figure 1: Stock-bond correlations across advanced and emerging economies and across the euro-area core and periphery

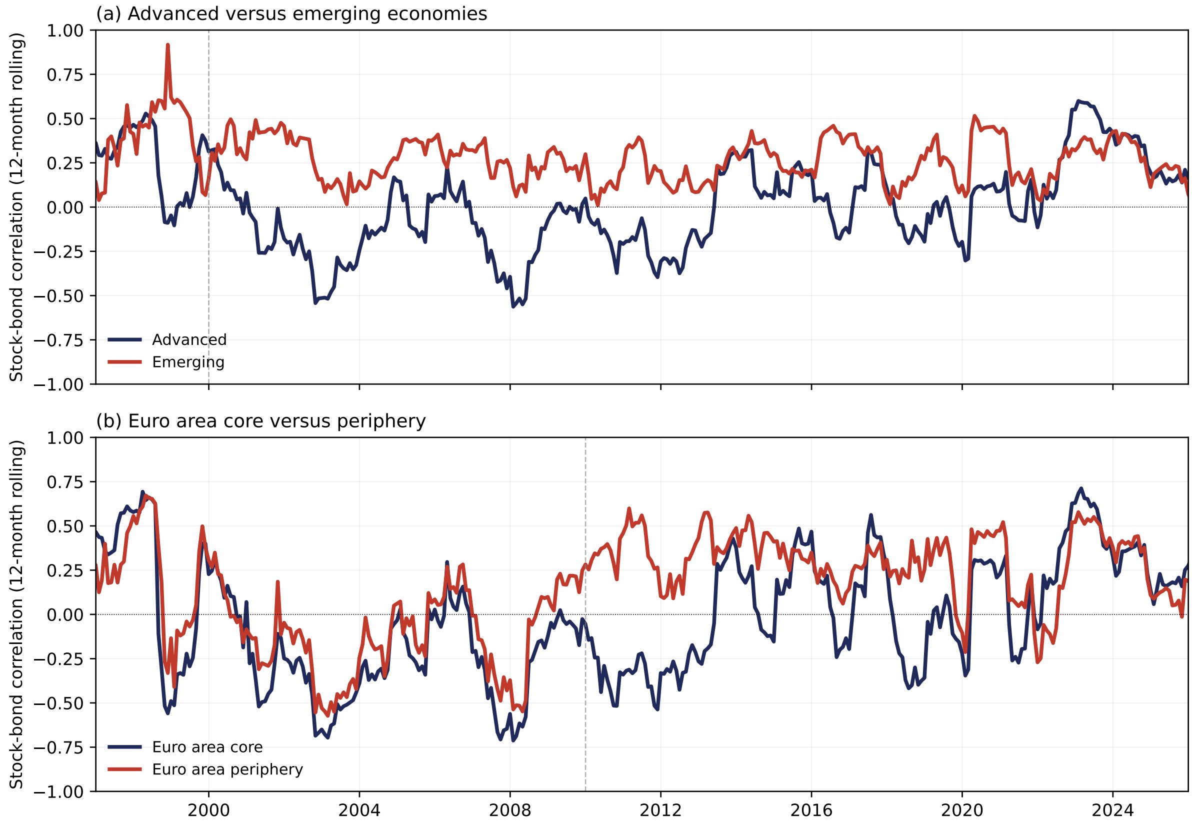

Note: Twelve-month rolling correlation of monthly equity and ten-year zero-coupon government bond returns, averaged within each group. Panel (a) plots the 29 advanced economies against the 15 emerging markets; the emerging average rests on fewer than five countries before 2000. Panel (b) plots the euro-area core (AT, BE, DE, FI, FR, NL, SI) against the periphery (ES, GR, IE, IT, PT). The dashed lines mark the start of 2000 and of 2010. Section 5 excludes Slovenia from the core and restores it in a robustness row.

the correlation the other way through two fiscal channels. The fiscal-capacity channel operates through a higher probability of sovereign default under monetary dominance, when the central bank holds its inflation target, or through higher expected inflation under fiscal dominance, when the budget forces the price level to adjust. The convenience channel operates through a fall in the convenience value of government debt as the debt stock grows. Which dominates depends on the level of the burden, within the range the propositions cover. When it is low, the safe-haven channel prevails and bonds hedge equity risk. When it is high, the two fiscal channels dominate and bonds amplify equity risk. The model summarizes the burden in a single statistic, the forward debt obligation, the debt stock times the gross rate at which it reprices. It motivates a regression in which debt and yields enter as a product. The bond's exposure to news rises with the default hazard, so the association should be stronger where the burden is already high. And under monetary dominance fiscal stress is priced as default risk, so the component of the yield that carries the relationship within a country tells which regime the data are consistent with.

In the panel, the interest expense ratio is positively associated with the stock-bond correlation both across and within countries. Across countries it accounts for a quarter of the variation in average correlations since 1960, and a country whose average interest expense is 3.4 percentage points of GDP higher, the distance between the tenth and ninetieth percentiles, has an average correlation 0.15 higher. Within countries the same sign holds in a monthly panel of correlations of daily returns, under country and month fixed effects, macroeconomic controls and the policy rate, so it does not reduce to the level of the policy rate, and the estimate is positive in 34 of 44 countries. In a horse race on the annual panel the product of debt and yield keeps its coefficient beside its two components, at the ten percent level, while neither debt nor the yield is distinguishable from zero beside it.

In the panel regressions, the regressor is not randomly assigned, and both sides of the regression carry a market price, a yield inside the ratio on the right and a correlation of returns on the left, which limits what the estimates can establish. To move closer to identification, we turn to the euro area, where countries share a central bank, a currency and an inflation target while their fiscal positions diverge. The Greek deficit revision of October 2009 provides a quasi-experiment. The pre-event dynamics are broadly consistent with parallel trends, with the qualifications Section 5 reports. The gap between the periphery and the core opens within two quarters of the announcement. Relative to the period before the announcement, it averages 0.44 over the following eighteen months and peaks at 0.72. The estimates corroborate the pattern visible in the raw series of Figure 1, panel (b): the periphery's correlation rises across the announcement while the core's barely moves. The periphery is Greece, Ireland, Italy, Portugal and Spain, the five countries that sovereign credit markets already priced as the riskier group before the failure of Lehman Brothers, so the grouping is not chosen after the fact. The design does not separate the fiscal news from the policy response it set in motion.

The euro-area episode is a single event. To test the relationship more broadly without relying on a market price to measure the fiscal position, we build two measures of it from text, which we call fiscal shifters: the first from newspaper coverage of a country's public finances, the second from the reasons rating agencies give for their actions. The first is the share of a country's Wall Street Journal coverage that a language model reads as discussing the sustainability of its public finances, its fiscal salience. Within country and within month, a higher share predicts a higher correlation in the following month, and the association is strongest where the fiscal burden is already high, consistent with the state dependence of the model's default hazard. The second counts sovereign rating changes that the agency attributes to the public finances, entering a downgrade as a positive value and an upgrade as a negative one. A fiscal downgrade is associated with a higher correlation in the month it is announced. Neither measure contains a market price, and both carry the positive association with the correlation that the interest expense ratio does. Beside the interest expense ratio the newspaper index keeps a positive association and the rating estimate falls and is indistinguishable from zero. A placebo on the comovement of the country's equity market with the United States market gives no evidence of a general attention channel behind the index, and the newspaper estimate is smaller after 2010.

The last step asks which part of the yield carries the relationship, the question the model's two regimes pose. Where a sovereign credit default swap is quoted, from 2001, we split the yield inside the ratio into a credit component and the rest. Across countries, every component predicts the level of the correlation, as the model's regression implies, since its slope pools the channels without distinguishing the components of the yield. Within countries, only the credit component carries the association. The rest of the yield does not, and where an inflation-linked curve allows a further split, neither the real yield nor breakeven inflation does. We read this pattern as support for the model's monetary-dominance regime, in which interest expense tracks the correlation through the market's assessment of fiscal sustainability.

Related literature. Two strands of work are most closely related. The first studies the determinants of stock-bond correlations. Early work documents the time-varying nature of the relationship (Ilmanen, 2003; Connolly et al., 2005), standard macroeconomic fundamentals leave most of the variation in comovement unexplained (Baele et al., 2010), and macroeconomic news explains a substantial part of its changes but through effects that existing theories do not predict (Duffee, 2023). Baele et al. (2020) document flight-to-safety episodes in stock and bond returns across 23 countries. A prominent set of explanations attributes the sign change at the turn of the century to shifts in inflation dynamics and monetary policy (Campbell et al., 2017; Song, 2017; Campbell et al., 2020;

David and Veronesi, 2013): in Campbell et al. (2020), the comovement of inflation with the output gap turns positive around 2001, so that nominal bonds pay off in downturns, and risk premia amplify the shift, a reading that extends to the renewed inflation risk after 2022 (Pflueger, 2025). These papers focus on the United States time series, where an inflation account and a fiscal account are difficult to separate: the 1980s and the years since 2022 combine adverse inflation dynamics with high fiscal burdens. Their mechanisms run through inflation dynamics and risk premia that a single time series shares, so they do not speak to differences across countries that share a monetary regime, and they say nothing about emerging markets. Antolin-Diaz (2025) studies twelve advanced economies over the postwar period and attributes the shift to negative stock-bond correlations to financial shocks and flight-to-safety dynamics, with the rise of levered intermediaries as the underlying driver and no fiscal variable. Li et al. (2022) are the closest antecedent. In their regime-switching model of United States data, the shift from a monetary to a fiscal regime after 2001 accounts for the change in the correlation, and in the fiscal regime the correlation turns negative, driven by investment shocks. Their object is a policy regime; ours is a continuous measure of the fiscal burden, and in our data a heavier burden goes with a more positive correlation. Dufour et al. (2017) show, on euro-area sovereign bonds over 2004 to 2013, that the sign of a bond's equity beta depends on country risk: low-risk sovereigns hedge equity risk and high-risk ones lost that property in the sovereign debt crisis. Molenaar et al. (2024) document on an international sample that government creditworthiness, along with inflation and real rates, is among the drivers of the correlation. Relative to these papers, we measure fiscal conditions with a single observable statistic derived from a model and test the form in which it enters, on 44 countries including 15 emerging markets and back to 1960, within a monetary union, and in the credit component of the yield.

A second strand studies why some government bonds carry a convenience premium and how that status depends on the issuer's fundamentals relative to competing issuers (Krishnamurthy and Vissing-Jorgensen, 2012; He et al., 2019; Brunnermeier et al., 2024; Caballero and Farhi, 2018; Caballero et al., 2017; Maggiori, 2017; Du et al., 2018). Jiang et al. (2026) argue that the United States government's fiscal position determines whether its debt is truly risk-free, and Jiang et al. (2025) show that convenience yields explain a large share of the variation in euro-area sovereign yields and rise with fiscal surpluses;

Acharya and Laarits (2025) and Cieslak et al. (2024) connect the convenience yield to the stock-bond covariance. We provide cross-country evidence that the hedging property, the foundation of safe-asset status, varies systematically with the issuer's fiscal burden. A third strand studies sovereign default and its pricing: Arellano (2008) models default as countercyclical, Bocola (2016) traces sovereign risk through bank balance sheets, Longstaff et al. (2011) find that global factors dominate sovereign credit spreads, and Cole et al. (2025) model the eurozone crisis as information spillovers that push the yields of risky and safe sovereigns apart. Blanchard (2019), Cochrane (2023) and Hazell and Hobler (2025) connect fiscal conditions to interest rates and inflation; Bohn (1998), Ghosh et al. (2013) and Mehrotra and Sergeyev (2021) develop measures and tests of fiscal sustainability, and Hall and Sargent (2011) decompose the dynamics of United States debt. Gómez-Cram et al. (2024) show that unfunded fiscal expansions can generate a positive bond risk premium and connect fiscal shocks to the stock-bond correlation in the United States, and Gómez-Cram et al. (2025) trace legislative cost news to Treasury yields; we show how the fiscal burden organizes the correlation across 44 countries. Finally, the regime distinction in our model maps onto the monetary and fiscal dominance of the fiscal theory of the price level (Sims, 1994; Woodford, 1994; Leeper, 1991; Bianchi and Melosi, 2019).

Section 2 presents the model and derives its regression form, Section 3 describes the data, Sections 4 to 7 take the four steps in turn, and Section 8 concludes.

2 Model

We construct a parsimonious endowment economy with Epstein–Zin preferences, sovereign default risk, and a convenience flow on government bonds, and isolate the channels through which a single fundamental shock generates stock-bond comovement. A *safe-haven channel*: bad news lowers expected growth and with it the real risk-free rate, raises the bond price, and lowers equity through the dependence on output. A *fiscal-capacity channel*: bad news tightens the government's balance sheet and lowers the bond price, working against the safe-haven channel; under monetary dominance it operates through a higher probability of sovereign default, and under fiscal dominance through a rise in the price level in place of default. A *convenience channel*: bad news raises real debt outstanding relative to output, lowers the marginal value of safe-asset services, and lowers the bond price. Both the fiscal-capacity and convenience channels are driven by the government's fiscal burden; only the safe-haven channel is not. The same balance-sheet condition triggers both regimes, and the three channels are summarized by a single empirical statistic, the forward debt obligation.

2.1 Environment

Output and consumption. A representative household consumes the perishable endowment $C_t = Y_t$, with

$$Y_{t+1} = Y_t \exp(g + \rho\sigma\varepsilon_t + \sigma\varepsilon_{t+1}), \qquad \varepsilon_t \sim \mathcal{N}(0,1) \text{ i.i.d.} \tag{1}$$

The innovation ε_t is a productivity shock with persistence $\rho > 0$: a positive realization raises current output and expected next-period growth.

Preferences. The household has Epstein–Zin recursive utility with unit elasticity of intertemporal substitution:

$$\log V_t = (1-\beta)\log C_t + \frac{\beta}{1-\gamma}\log \mathbb{E}_t V_{t+1}^{1-\gamma}, \tag{2}$$

where $\gamma > 0$ is risk aversion and $\beta \in (0,1)$ the discount factor. Government bonds, B_t (B_t^n) in real (nominal) terms, provide a non-pecuniary service flow $Y_t\,\xi(B_t/Y_t)$ in consumption-equivalent units to their holder (Krishnamurthy and Vissing-Jorgensen, 2012), with $B_t \equiv B_t^n/P_t$ and

$$\xi(B/Y) = \begin{cases} \alpha\, B/Y - \frac{1}{2}\eta\,(B/Y)^2 & \text{if } B/Y \le \alpha/\eta \\ \alpha^2/(2\eta) & \text{if } B/Y > \alpha/\eta \end{cases}, \tag{3}$$

$\alpha \in (0,1)$ and $\eta > 0$. The utility-equivalent flow is $\zeta_t = \lambda_t\, Y_t\, \xi(B_t/Y_t)$, where $\lambda_t \equiv \partial V_t/\partial C_t$ is the marginal utility of consumption, and lifetime utility is

$$U_0 = V_0 + \mathbb{E}_0 \sum_{t\ge 0} \beta^t \zeta_t. \tag{4}$$

We restrict attention to the active range $B_t/Y_t < \alpha/\eta$.

Government. The government runs no spending and collects taxes,

$$s_t \;=\; \tau\, Y_t, \qquad \tau \in (0,1). \tag{5}$$

The government rolls over its debt subject to the nominal flow budget constraint,

$$(1+i_{t-1})\, B^n_{t-1}\,(1-\delta_t) \;=\; P_t\, \tau\, Y_t \;+\; B^n_t, \tag{6}$$

where B^n_t is new issuance of one-period nominal debt. The central bank targets gross inflation $1+\bar{\pi}$, so $P_t = (1+\bar{\pi})P_{t-1}$ when the peg holds. In real terms,

$$B_t \;=\; \frac{(1+i_{t-1})\,(1-\delta_t)}{1+\bar{\pi}}\, B_{t-1} \;-\; \tau\, Y_t. \tag{7}$$

We assume $\alpha + \eta\tau < 1$, which keeps the marginal convenience wedge below one along the equilibrium path.

2.2 Equilibrium asset prices

Stochastic discount factor. The Epstein–Zin pricing kernel is

$$M_{t+1} \;=\; \beta\, \frac{C_t}{C_{t+1}} \left(\frac{V_{t+1}}{\left(\mathbb{E}_t V_{t+1}^{1-\gamma}\right)^{1/(1-\gamma)}} \right)^{1-\gamma}, \tag{8}$$

and substituting the equilibrium value function (see Appendix A) yields

$$\log M_{t+1} \;=\log\beta \;-\; (g+\rho\sigma\varepsilon_t) \;-\; \hat{\sigma}\, \varepsilon_{t+1} \tag{9}$$

$$-\; \tfrac{1}{2}(\gamma-1)^2\sigma^2(1+\beta\rho)^2, \tag{10}$$

where $\hat{\sigma} \equiv \sigma[(\gamma-1)(1+\beta\rho)+1]$ is the loading of $\log M_{t+1}$ on the innovation ε_{t+1}.

Real risk-free rate. From $1 = (1+r_t)\, \mathbb{E}_t M_{t+1}$,

$$\log(1+r_t) \;=\; -\log\beta \;+\; g + \rho\sigma\varepsilon_t \;-\; \left(\gamma - \tfrac{1}{2}\right)\sigma^2 \;+\; (1-\gamma)\beta\rho\sigma^2. \tag{11}$$

Good news raises the real rate one-for-one with the predictable component of growth $\rho\sigma$.

Equity. An equity claim has real payoff Y^{λ}_{t+1} at $t+1$ with $\lambda \geq 1$. The nominal price of equity is given by

$$p_t \;=\; P_t\,\mathbb{E}_t\left[M_{t+1}\,Y_{t+1}^{\lambda}\right], \tag{12}$$

and log-normality yields

$$\begin{aligned}\log p_t \;=&\; \log P_t \;+\; \log\beta \;+\; \lambda\log Y_t \;+\; (\lambda-1)(g+\rho\sigma\varepsilon_t)\\ &-\; \tfrac{1}{2}(1-\gamma)^2(1+\beta\rho)^2\sigma^2 \;+\; \tfrac{1}{2}\Big[(1-\gamma)(1+\beta\rho)-1+\lambda\Big]^2\sigma^2,\end{aligned} \tag{13}$$

with slope

$$\frac{\partial\log p_t}{\partial\varepsilon_t} \;=\; \sigma\left[\lambda(1+\rho)-\rho\right] \;>\; 0. \tag{14}$$

The nominal equity price moves in the same direction as the output news.

2.3 Default risk

The government's real liabilities entering t and its real assets at t are

$$L_t \;\equiv\; \frac{(1+i_{t-1})\,B_{t-1}}{1+\bar{\pi}}, \qquad A_t \;\equiv\; \mathbb{E}_t\sum_{j=0}^{\infty} M_{t,t+j}\,\tau\,Y_{t+j}, \qquad M_{t,t+j} \;\equiv\; \prod_{k=1}^{j} M_{t+k}. \tag{15}$$

The real value of liabilities L_t is predetermined at $t-1$; the real value of assets A_t is the present value of the entire stream of future taxes. Default occurs whenever the government is not solvent:

$$\delta_t \;=\; \mathbf{1}\{A_t \;<\; L_t\}. \tag{16}$$

The default event at $t+1$ reduces to $\varepsilon_{t+1} < z_t$ with threshold

$$z_t \;\equiv\; \frac{1}{\sigma}\left[\log\left(\frac{1+i_t}{1+\bar{\pi}}\frac{B_t\,(1-\beta)}{\tau\,Y_t}\right) \;-\; g \;-\; \rho\sigma\varepsilon_t\right]. \tag{17}$$

The threshold rises with the forward debt obligation $(1+i_t)B_t/(1+\bar{\pi})$ and falls with fiscal capacity $\tau Y_t/(1-\beta)$.

2.4 Bond pricing under monetary dominance

The nominal price of a one-period bond is

$$q_t = \frac{1}{1+i_t} = \frac{1}{1-\xi_t'}\,\mathbb{E}_t\left[\frac{M_{t+1}\,(1-\delta_{t+1})}{1+\pi_{t+1}}\right]. \tag{18}$$

The first factor is the convenience premium. The second factor is the discounted expected nominal payoff conditional on no default. Joint log-normality of M_{t+1} and ε_{t+1} reduces the bond price to

$$\begin{aligned} \log q_t &= \log\Big(1-\Phi(z_t+\hat{\sigma})\Big) - \log\Big((1+r_t)(1+\bar{\pi})\Big) \\ &\quad - \log\Big(1-\alpha+\eta\,B_t/Y_t\Big), \end{aligned} \tag{19}$$

where Φ is the standard normal cumulative distribution function and $\hat{\sigma}$ is the change-of-measure shift from the physical to the risk-neutral measure.

Given (17), the bond log-slope is

$$\frac{\partial \log q_t}{\partial \varepsilon_t} = \left(1+\frac{\tau Y_t}{B_t}\right)\frac{\sigma}{1-\Theta_t}\left(\Theta_t + \frac{\alpha-\xi_t'}{1-\xi_t'}\right) - \rho\sigma \tag{20}$$

where

$$\Theta_t = \frac{1}{\sigma}\frac{\phi(z_t+\hat{\sigma})}{1-\Phi(z_t+\hat{\sigma})} \tag{21}$$

is the hazard rate of government default.[2]

We compute the date-t conditional correlation between $\log p_{t+1}$ and $\log q_{t+1}$ on the event of no default at $t+1$, $\varepsilon_{t+1} \geq z_t$, so all conditional moments below are taken under this event. The bond at $t+1$ is then priced as a one-period claim with potential default at $t+2$, the only remaining default risk.

Proposition 1 (Components under monetary dominance)**.** *Under monetary dominance,*

[2] The factor $1/(1-\Theta_{t+1})$ in the bond log-slope is the geometric sum of a yield-default feedback with loop gain Θ_{t+1}: a bad-news shock raises the conditional default probability, the equilibrium yield rises to compensate, the higher yield raises debt service and the default probability further, and so on. For $\Theta_{t+1} < 1$ the series converges and the credit channel is procyclical; for $\Theta_{t+1} > 1$ the slope flips sign and the economy enters the self-fulfilling-default region of Calvo (1988) and Cole and Kehoe (2000). We therefore assume that $\Theta_{t+1} < 1$ always.

conditional on $\varepsilon_{t+1} \geq z_t$,

$$\rho_t^{sb}\,\sigma_t^q/\sigma \;=\; \underbrace{\mathbb{E}_t\left[\frac{\theta_{t+1}}{1-\Theta_{t+1}}\,\Theta_{t+1}\right]}_{\textit{credit}} \;+\; \underbrace{\mathbb{E}_t\left[\frac{\theta_{t+1}}{1-\Theta_{t+1}}\,\frac{\alpha-\xi'_{t+1}}{1-\xi'_{t+1}}\right]}_{\textit{convenience}} \;-\; \underbrace{\rho}_{\textit{safe-haven}},$$

where $\theta_{t+1} \equiv 1 + \tau Y_{t+1}/B_{t+1}$ *is a proxy for fiscal capacity and* $\sigma_t^q \equiv \sqrt{\mathrm{Var}_t(\log q_{t+1})}$.

Credit and convenience are positive contributions while safe-haven is negative, so the sign of ρ_t^{sb} is the sign of the net of the three. The term labeled credit is the fiscal-capacity channel of the introduction: under monetary dominance it operates through the default hazard Θ, and under fiscal dominance the corresponding term in Proposition 2 operates through inflation dilution, not default.

2.5 Bond pricing under fiscal dominance

The central bank holds the inflation peg $\pi_t = \bar{\pi}$ whenever it is consistent with the government's budget. At the peg, real liabilities are $L_t = (1+i_{t-1})B_{t-1}/(1+\bar{\pi})$ and the present value of the maximum surplus stream is $A_t = \tau Y_t/(1-\beta)$. The peg is sustainable when $A_t \geq L_t$. On the complementary event $A_t < L_t$ the budget cannot close at the peg, and one variable must adjust to restore $L_t = A_t$.

Lemma 1 (Common trigger)**.** *The event* $A_t < L_t$ *activates a regime-specific adjustment:*

- *under monetary dominance, the government defaults and inflation is pegged;*
- *under fiscal dominance, inflation adjusts to satisfy the government budget.*

The trigger threshold is identical across regimes; only the adjustment variable differs.

Under fiscal dominance the bond pays one nominal unit at $t+1$ in every state, no default. On the trigger event, the price level rises so that $L_t = A_t$ ex post. The fiscal-dominance price level satisfies

$$1+\pi_t \;=\; \frac{(1+i_{t-1})\,B_{t-1}\,(1-\beta)}{\tau\,Y_t}, \tag{22}$$

whenever $\varepsilon_t \leq z_{t-1}$ ($\pi_t = \bar{\pi}$ otherwise) and the price-level slope is $\partial \log P_t/\partial \varepsilon_t = -\sigma$.

We compute the date-t conditional correlation between $\log p_{t+1}$ and $\log q_{t+1}$ on the event of no trigger at $t+1$, $\varepsilon_{t+1} \geq z_t$. All conditional moments below are taken under this event. The bond at $t+1$ is then priced under fiscal-dominance resolution with potential trigger at $t+2$.

Proposition 2 (Components under fiscal dominance)**.** *Under fiscal-dominance resolution, conditional on* $\varepsilon_{t+1} \geq z_t$*,*

$$\rho_t^{sb}\,\sigma_t^q/\sigma \;=\; \underbrace{\mathbb{E}_t\!\left[\frac{\theta_{t+1}}{1-\Theta_{t+1}}\,\Theta_{t+1}\right]}_{\textit{credit}} \;+\; \underbrace{\mathbb{E}_t\!\left[\frac{\theta_{t+1}}{1-\Theta_{t+1}}\,\frac{\alpha-\xi'_{t+1}}{1-\xi'_{t+1}}\right]}_{\textit{convenience}} \;-\; \underbrace{\rho}_{\textit{safe-haven}}$$
$$-\;\underbrace{\mathbb{E}_t\!\left[\frac{\theta_{t+1}}{1-\Theta_{t+1}}\,\frac{\Psi_{t+1}\Lambda_{t+1}}{1\,+\,\Psi_{t+1}\,\Lambda_{t+1}}\left(1+\frac{\alpha-\xi'_{t+1}}{1-\xi'_{t+1}}\right)\right]}_{\textit{fiscal-capacity correction}}$$

where $\theta_{t+1} \equiv 1+\tau Y_{t+1}/B_{t+1}$ *is a proxy for fiscal capacity,* $\sigma_t^q \equiv \sqrt{\mathrm{Var}_t(\log q_{t+1})}$*,*

$$\Psi_{t+1} \equiv \frac{q_{t+1}-q_{t+1}^{MD}}{q_{t+1}} > 0, \qquad \Lambda_{t+1} \equiv \frac{1}{1-\Theta_{t+1}}\,\frac{\phi(z_{t+1}+\hat{\sigma}-\sigma)}{\sigma\,\Phi(z_{t+1}+\hat{\sigma}-\sigma)} \;-\; 1 > 0,$$

and q_t^{MD} *is the nominal price of a bond under monetary dominance* (18)*.*

The first three terms are identical to Proposition 1. The fourth term, the fiscal-capacity correction, is the contribution specific to fiscal dominance. Under monetary dominance the bond pays zero on the trigger event; under fiscal dominance it pays one nominal unit, diluted by the inflation spike. The correction is negative because the bond is effectively safer. Ψ_{t+1} is the share of q_{t+1} accounted for by this recovery, Λ_{t+1} is the elasticity of that share to news at $t+1$, and the factor $1+(\alpha-\xi'_{t+1})/(1-\xi'_{t+1})$ collects the principal recovery together with the convenience flow.

2.6 Large fiscal capacity and regression

Proposition 3 (Convenience with zero credit risk)**.** *In the limit* $\Phi(z+\hat{\sigma}) \to 0$*, the credit channel (shared by Propositions 1 and 2) and the fiscal-capacity correction of Proposition 2 both vanish, and*

$$\rho_t^{sb}\,\sigma_t^q/\sigma \;=\; \mathbb{E}_t\!\left[\theta_{t+1}\,\frac{\alpha-\xi'_{t+1}}{1-\xi'_{t+1}}\right] \;-\; \rho.$$

The stock-bond correlation is positive if and only if the convenience term exceeds ρ*.*

Large fiscal capacity shuts down default and fiscal-driven inflation, leaving convenience as the only state-dependent driver of ρ_t^{sb} and a function of forward debt obligation.

Proposition 4 (Regression under monetary dominance). *Under* $\Phi(z_{t+1} + \hat{\sigma}) < \Phi(-\sigma)$*, to first order,*

$$\rho_t^{sb}\, \sigma_t^q / \sigma \;\approx\; \beta_0 \,+\, \beta_1 \frac{(1 + i_t)\, B_t}{Y_t}, \tag{23}$$

with $\beta_1 > 0$ *pooling the credit and convenience channels of Proposition 1.*

Proposition 5 (Regression form under fiscal dominance). *For* $\tau\beta/(1 - \beta)$ *sufficiently small, there exists* $c \in (0, \Phi(-\sigma)]$ *such that, under* $\Phi(z_{t+1} + \hat{\sigma}) < c$*, (23) holds with* $\beta_1 > 0$ *pooling the credit, convenience, and fiscal-capacity correction channels of Proposition 2.*

The intercept β_0 combines the safe-haven term with the level of the fiscal channels at the expansion point; the slope β_1 pools the non-safe-haven channels. Together, Propositions 4 and 5 motivate a local regression of the volatility-scaled correlation on the forward debt obligation $(1 + i_t)B_t/Y_t$ under their stated conditions, the observable through which all non-safe-haven channels move the bond-equity correlation.

The bond's exposure to output news depends on how strongly the default probability responds to a shock, and this response is amplified by the primary surplus per unit of debt. The propositions give sufficient conditions for a positive slope of the volatility-scaled correlation and do not establish a reversal beyond the threshold. Under fiscal dominance, the price-level adjustment of (22) provides a partial real recovery on fiscal-stress events. This recovery attenuates the credit and convenience contributions to the bond log-slope, requiring the tighter threshold $c \leq \Phi(-\sigma)$ for $\beta_1 > 0$.

3 Data

We construct a monthly and an annual panel of stock-bond correlations, fiscal variables, and macroeconomic controls for 44 countries, with annual observations over 1960–2025 and the monthly panel extending into 2026. The sample covers 29 advanced economies and 15 emerging market economies.[3] Both panels are unbalanced due to the staggered

[3]Advanced economies include Australia, Austria, Belgium, Canada, Czechia, Denmark, Finland, France, Germany, Greece, Hong Kong, Ireland, Israel, Italy, Japan, Korea, Lithuania, the Netherlands,

entry of countries, but once a country appears its history is nearly complete. The data run through the vendors' latest observations. The monthly panel measures the correlation from daily returns inside each calendar month, which gives 14,282 estimation observations on the common sample of Table 2, which requires the policy rate, against 1,757 in the annual panel constructed from monthly returns; Table 4 uses the 14,414 country-months available without that requirement. The monthly panel holds seven countries in 1990, 22 in 2000, and all 44 from 2010. The annual panel reaches back to 1960 for some countries allowing us to also cover the high-inflation decades.

Stock-bond correlations. Our dependent variable is the realized correlation between stock returns and government bond returns. For each country and month we pair the daily log return on a broad national equity index with the daily return on a synthetic 10-year zero-coupon government bond, and we keep the month if both legs are observed on at least 15 trading days. The annual measure pairs monthly returns on the same two instruments and requires at least ten monthly observations.

Equity and bond returns. We combine data from multiple sources to construct the stock-bond correlation. At daily frequency each leg comes from a single source per country. The bond leg is a 10-year zero-coupon yield: the Gürkaynak et al. (2007) curve for the United States, the Bank of England curve for the United Kingdom, and the published BIS or Bloomberg curve elsewhere. For fourteen countries we fit a Nelson-Siegel curve (Nelson and Siegel, 1987) to Datastream benchmark par yields, with the decay parameter fixed at the value of Diebold and Li (2006), and where a country's par grid is too sparse to support a fit, we use the vendor's converted benchmark quote instead.[4] The equity leg is the CRSP value-weighted return for the United States and the local Datastream or MSCI index elsewhere. We compute daily bond returns with the duration approximation, $r^b_{i,t} = -n\,\Delta y_{i,t}/100$ with $n = 10$: the carry a 10-year bond earns over one day is negligible against daily yield volatility, so we omit it. At the monthly horizon, carry is no longer

New Zealand, Norway, Portugal, Singapore, Slovenia, Spain, Sweden, Switzerland, Taiwan, the United Kingdom and the United States. Emerging market economies include Brazil, Chile, China, Colombia, Hungary, India, Indonesia, Malaysia, Mexico, Peru, the Philippines, Poland, South Africa, Thailand and Turkey.

[4] Lithuania is the only country never fitted, and Appendix B counts the dated substitutions elsewhere; the leave-one-out estimates of Figure C.2, which drop Lithuania with every other country in turn, do not change the results.

negligible, so monthly returns follow the log-linear approximation of Campbell et al. (1997), $r^b_{i,t} = y_{i,t-1}/12 - n\,\Delta y_{i,t}$, which retains it. The monthly histories reach further back by splicing: 29 of the 44 bond series and 35 of the 44 equity series draw on more than one source, always joined on returns and never on yield or price levels.

Forward-looking interest expense. Our right-hand-side variable of interest is the forward-looking interest expense,

$$\mathrm{ie}^{\mathrm{fw}}_{i,t} = \frac{D_{i,t}}{Y_{i,t}} \cdot \frac{y_{10,i,t}}{100}.$$

This is the interest bill the government would owe if its whole debt stock were refinanced at today's yield. It is the empirical counterpart of the forward debt obligation in the model of Section 2, the single statistic that carries its fiscal channels. Realized interest payments are not the object the model points to: the forward debt obligation reprices the whole stock at today's rate, whereas realized payments carry rates fixed at issuance and respond to news only as the debt rolls over. Appendix C.1 moves the regressor from realized payments to the forward-looking measure, and along the way the association rises and is estimated more precisely. We take the debt ratio from the IMF World Economic Outlook, which reports general government gross debt, and extend it back with the IMF Historical Public Debt Database, where most of the pre-1990 observations come from. In the baseline specification, the yield is the same 10-year zero-coupon series that prices the bond leg, and Appendix C.1 rebuilds the measure using different maturities. In the monthly panel the debt ratio is the last annual figure the source had published, so every monthly observation is predetermined.

Macroeconomic controls. The controls are consumer price inflation, real activity growth, and the central bank policy rate, each assembled from a cascade of sources that is identical for every country. Real activity growth is industrial production growth for most of the panel and real GDP growth for the rest. A fourth control, the VIX, is a global time series, so it matters only in specifications without period effects.

Fiscal shifters. Section 6 requires news about the fiscal position that is not itself a market price, and we measure such news in two ways. The first is sovereign rating actions.

We assemble the rating histories of Moody's, S&P, and Fitch into 1,415 actions covering 43 countries from 1974 and map letters to a common notch scale. The shifter is the sum of signed notches across the agencies covering the country in that year, divided by their number, so that a country followed by all three does not register three times the move of a country followed by one. A large language model then reads the newswire coverage of each action, which republishes or paraphrases the agency's statement, and records the origin the agency gave for it and whether it also named a non-fiscal cause. We use these two readings to keep only the actions taken for a stated fiscal reason. The second measure is monthly and comes from newspaper coverage. For each country-month, a large language model reads the Wall Street Journal articles that a separate attribution pass assigns to the country and returns the share that discuss the sustainability of its public finances, the country's fiscal salience. The corpus is 200,391 articles from January 2000. Appendices B.8 and B.9 describe the classifiers, the filters, and the placebos we run against both measures.

Credit and inflation compensation. Section 7 splits the interest bill into a credit share, priced from the sovereign credit default swap (CDS) spread, and a non-default share. The CDS series is the five-year spread, 12,142 country-months for all 44 countries from January 2001. We take the spreads from Markit. Every quote is for a senior unsecured contract on the sovereign. For the fourteen countries with a sovereign inflation-linked curve we also take a five-year real yield, and breakeven inflation is the five-year nominal yield less the real yield. Appendices B.6 and B.7 describe both series.

Table 1 summarizes the main variables of the paper. Panel A describes the monthly estimation panel, which carries the baseline regressions, and Panel B the annual panel, which reaches back to 1960 and carries the horse race between interest expense and its components. Panel C describes the fiscal shifters of Section 6 and the yield components of Section 7, each on the sample its section estimates on. For every variable the table reports the number of observations, the mean, the overall standard deviation with its between- and within-country parts, and the tenth and ninetieth percentiles. Both regimes are common: in each panel the correlation's tenth percentile is negative and its ninetieth positive. The between- and within-country standard deviations of $\mathrm{ie}^{\mathrm{fw}}$ are of similar size, so the panel carries cross-country level differences and within-country movements in equal measure.

Table 1: Summary statistics

Variable	N	Mean	SD	Between	Within	p10	p90
A. Monthly panel							
Stock-bond correlation	14,632	-0.01	0.34	0.13	0.32	-0.47	0.45
Interest expense, forward-looking (% of GDP)	14,490	2.79	2.57	1.86	1.89	0.44	5.77
Debt / GDP (%)	14,490	63.70	38.48	32.96	19.31	28.08	108.22
10-year zero-coupon yield (% p.a.)	14,632	4.75	3.29	2.58	2.36	0.99	8.96
Policy rate (% p.a.)	14,442	3.67	3.83	2.78	2.98	0.05	8.00
CPI inflation (% p.a.)	14,620	3.02	4.19	3.19	3.33	0.20	6.25
Real activity growth (% p.a.)	14,514	2.48	10.86	2.29	10.63	-4.12	8.94
B. Annual panel							
Stock-bond correlation	1,788	0.13	0.39	0.15	0.37	-0.42	0.63
Interest expense, forward-looking (% of GDP)	1,734	3.08	2.89	1.95	2.27	0.63	6.43
Debt / GDP (%)	1,737	57.12	36.61	26.01	26.66	17.95	103.92
10-year zero-coupon yield (% p.a.)	1,785	5.92	3.77	2.36	3.26	1.56	11.14
Policy rate (% p.a.)	1,746	4.88	4.33	2.64	3.81	0.18	10.17
CPI inflation (% p.a.)	1,788	3.98	4.41	3.07	3.84	0.57	8.72
Real activity growth (% p.a.)	1,780	3.05	6.61	2.03	6.36	-2.34	8.68
C. Fiscal shifters and yield components							
Rating shifter, all actions (notches)	1,391	-0.02	0.61	0.10	0.61	-0.50	0.33
Rating shifter, fiscal origin (notches)	1,391	0.00	0.19	0.03	0.18	0.00	0.00
Fiscal salience (share of articles)	11,962	0.03	0.09	0.04	0.08	0.00	0.09
5-year CDS spread (% p.a.)	11,013	0.83	1.07	0.63	0.86	0.08	1.91
Credit share of IE^{fw} (% of GDP)	10,957	0.55	1.10	0.59	0.94	0.03	1.23
Non-default share of IE^{fw} (% of GDP)	10,957	1.48	1.66	1.29	1.13	-0.23	3.74
5-year real yield (% p.a.)	2,687	1.06	2.19	1.72	1.47	-1.38	4.02
5-year breakeven inflation (% p.a.)	2,687	2.58	1.71	1.60	0.92	0.94	5.08

Note: This table reports summary statistics for the main variables of the paper. Panels A and B are the two estimation panels, each cut to the country-periods that carry a dependent variable: in the monthly panel the stock-bond correlation is the non-overlapping monthly correlation of daily equity and daily 10-year zero-coupon bond returns, and in the annual panel the calendar-year correlation of monthly returns. Panel C reports the fiscal shifters of Section 6 and the yield components of Section 7, each on the sample its section estimates on. Between is the standard deviation of country means and within the standard deviation of deviations from them. Real activity is industrial production growth for most of the panel and real GDP growth for the rest.

Appendix B documents the construction in full. For each country it states which source supplies each series and why, how we convert benchmark par yields to a zero-coupon curve where no published curve is available, and how we splice the monthly bond history across sources back to 1960. It tests the spliced history segment against the series that supersedes it on the annual correlation itself. It summarizes the screens and dated exclusions we apply, which the replication package documents one by one with their reasons. Appendix C.5 re-estimates the main tables with every screen and correction switched off.

4 Interest expense and the stock-bond correlation

The model ties the stock-bond correlation to the product of the debt stock and the rate at which the debt reprices. Its observable counterpart is the forward-looking interest expense ratio $\mathrm{ie}^{\mathrm{fw}}$ defined in Section 3. This section estimates the relationship between $\mathrm{ie}^{\mathrm{fw}}$ and the correlation on both panels: the monthly panel, where the correlation is measured from daily returns and which carries the baseline, and the annual panel, which reaches back to 1960 and is where the product can be raced against its components (i.e. debt-to-GDP ratio and the yield). All right-hand-side variables are lagged one period, and standard errors are two-way clustered by country and by period.

Figure 2 plots the annual panel across and within countries. Countries with higher average $\mathrm{ie}^{\mathrm{fw}}$ have more positive average correlations; the cross-country slope is 0.039 with an R^2 of 0.25. The yield leg is what separates countries with similar debt stocks. Japan carries the second-highest average debt ratio in the sample, yet low yields hold its average $\mathrm{ie}^{\mathrm{fw}}$ near two percent of GDP and its correlation is negative on average. Brazil combines moderate debt with high yields and sits at the top of the distribution on both axes. Within countries the fitted lines slope upward as well, with a pooled within-country slope of 0.048.[5] These slopes carry no fixed effects, no controls and no lag; the rest of the section shows that the pattern survives even with two-way fixed effects, macroeconomic controls including monetary policy rates.

We estimate

$$\rho^{sb}_{i,t} \;=\; \alpha_i \;+\; \delta_t \;+\; \beta\, \mathrm{ie}^{\mathrm{fw}}_{i,t-1} \;+\; X'_{i,t-1}\gamma \;+\; \varepsilon_{i,t}, \tag{24}$$

where α_i and δ_t are country and period fixed effects, and $X_{i,t-1}$ contains consumer price inflation and industrial production growth. Table 2 builds the specification one change at a time, with columns (1) through (6) on a common sample. The pooled coefficient is 0.040. Either set of fixed effects alone leaves it near that level; both together take it to 0.013, so a third of the pooled loading survives within country and within month. The macro controls move it to 0.0137, our baseline (column 5). The central bank policy rate leaves the estimate unchanged (column 6), so the relationship does not reduce to the level of interest rates. The five-year correlation and the five-year regressor give 0.0133 (column 7). The

[5] The six countries in panel (b) of Figure 2 are illustrations. Their median slope, +0.076, is close to the median of +0.079 across all 44 countries, and the descriptive within slope is positive in 35 of the 44. These slopes use the contemporaneous ratio on the plotted sample; Table C.4 reports the lagged regression for every country with its standard error. Three of the six rest on thirty or fewer country-years.

Figure 2: Interest expense and the stock-bond correlation, across and within countries

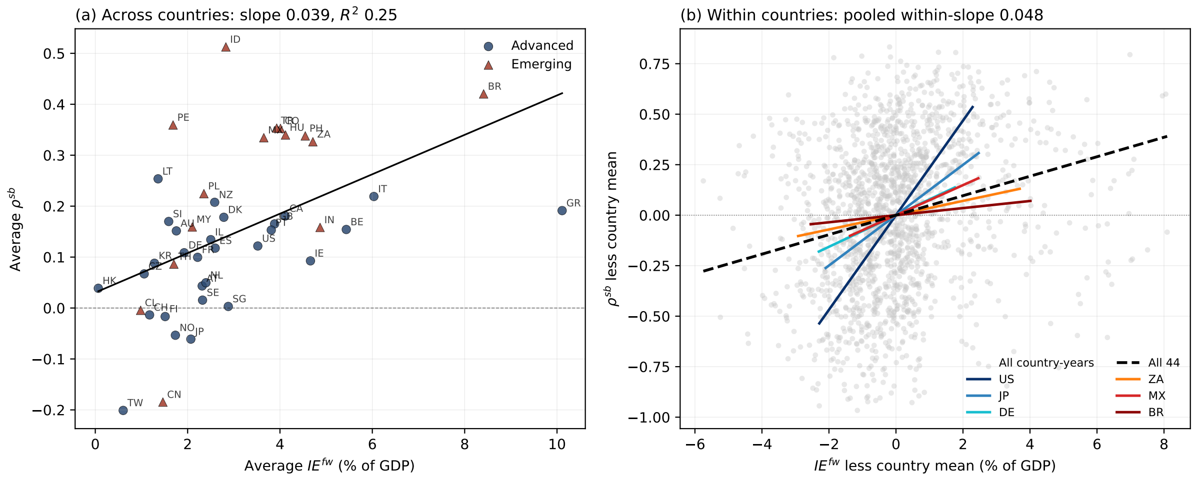

Note: Panel (a) plots each country's time-series average of ie^{fw} against its time-series average of ρ^{sb} on the annual panel, with markers distinguishing advanced economies from emerging markets and a fitted line through all 44 countries. Panel (b) plots the same country-year pairs in deviations from country means, with fitted lines for six countries chosen for long coverage and a spread of advanced and emerging markets, and a dashed line for the pooled within-country slope over all 44 countries, 0.048. Panel (b) omits 18 country-years above the 99th percentile of ie^{fw} for legibility; every regression in the paper uses the full sample.

economic magnitude in the baseline is such that a one-within-country-standard-deviation increase in ie^{fw}, 1.9 percentage points of GDP (Table 1), moves the predicted correlation by about 0.03, about eight percent of the correlation's within-country standard deviation.

The pooled association draws on differences across countries and on movements common to all countries over time, as well as on the variation that both sets of fixed effects leave. Country effects alone still use the common movements, and month effects alone still use the differences across countries, so with either set the coefficient stays near its pooled value. The association is stronger in the variation that the second set removes than in the variation left by both, whichever set enters first.[6] With both sets, interest expense and the correlation are each adjusted for country and month effects together, and column (4) measures the association between what remains. The baseline includes both sets to control for fixed country characteristics and for shocks common to all countries.

In the model the fiscal burden is multiplicative: debt matters through the rate at which it reprices, and the rate matters in proportion to the stock it applies to. The debt ratio is published once a year, so the horse race between the product and its components runs on the annual panel, the frequency at which both factors move. Table 3 reports it.

[6]With country effects alone, the coefficient of 0.037 is a weighted average of 0.092 on the variation that month effects go on to remove and 0.013 on the rest, with weights of 0.30 and 0.70. With month effects alone, the coefficient of 0.033 averages 0.050 on the variation that country effects go on to remove and 0.013 on the rest, with weights of 0.55 and 0.45. The weights are shares of the variance of lagged interest expense left by the included effects, and both identities hold exactly before rounding.

Table 2: Interest expense and the stock-bond correlation, monthly panel

VARIABLES	(1) $\rho^{sb,10Y}_{i,t}$	(2) $\rho^{sb,10Y}_{i,t}$	(3) $\rho^{sb,10Y}_{i,t}$	(4) $\rho^{sb,10Y}_{i,t}$	(5) $\rho^{sb,10Y}_{i,t}$	(6) $\rho^{sb,10Y}_{i,t}$	(7) $\rho^{sb,5Y}_{i,t}$
$IE^{fw}_{i,t-1}$	0.0398***	0.0369***	0.0334***	0.0130***	0.0137***	0.0137***	0.0133***
	(0.0097)	(0.0120)	(0.0076)	(0.0031)	(0.0029)	(0.0029)	(0.0040)
Macro controls	No	No	No	No	Yes	Yes	Yes
Policy rate control	No	No	No	No	No	Yes	No
Country FE	No	Yes	No	Yes	Yes	Yes	Yes
Month FE	No	No	Yes	Yes	Yes	Yes	Yes
Observations	14,282	14,282	14,282	14,282	14,282	14,282	13,903
R^2	0.089	0.157	0.343	0.439	0.442	0.442	0.432

Note: This table reports panel regressions of the monthly stock-bond correlation, the non-overlapping correlation of daily equity and daily 10-year zero-coupon bond returns within a calendar month, on forward-looking interest expense lagged one month. Columns (1) to (6) are estimated on the common sample of column (6), so a change across columns is the specification and never the sample; column (7) replaces the dependent variable with its five-year twin and is estimated on the rows that variable reaches. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1.

Alone, $\mathrm{ie}^{\mathrm{fw}}$ enters at 0.0145, significant at the one percent level; the debt ratio alone is indistinguishable from zero, and the yield alone gives 0.0133. When the product and both of its components enter together in column (4), $\mathrm{ie}^{\mathrm{fw}}$ keeps a coefficient of 0.0131, now significant at the ten percent level, while neither component is distinguishable from zero.

Table 4 asks whether the relationship is stable over time. The columns are nested, each keeping every country-month from the stated year onward, and the two panels estimate the same coefficient under pooled OLS and under the baseline model on identical rows. The within estimate is positive and significant in every column. The pooled estimate is larger in every column. The within estimate is larger on the 2015 sample, 0.032 against 0.014 on the full sample, and the pooled estimate rises by less, so the ratio of the two falls from 2.9 to 1.6. The slope across country averages is stable, between 0.044 and 0.051 on the five samples, so what has changed in the recent data is the within-country association, which has moved toward the cross-country one.

The aggregate series in Figure 1 show the same movement over time, and the fiscal variable moves with it. Over the 1990s yields fell and, from the middle of the decade, debt ratios fell with the fiscal consolidations of the period, so the average forward-looking interest burden of the advanced economies fell from 5.7 percent of GDP in 1990 to 3.1 in 1998. Their average correlation, positive through the middle of the decade, was negative by 1998, and across all 29 advanced economies in the panel it stayed, apart from 1999, at or below zero through 2021, years in which low yields held the burden down, to 0.3 percent of GDP in 2020. After 2020 the sequence ran in reverse: debt ratios were higher

Table 3: Interest expense against its components, annual panel

VARIABLES	(1) $\rho_{i,y}^{sb,10Y}$	(2) $\rho_{i,y}^{sb,10Y}$	(3) $\rho_{i,y}^{sb,10Y}$	(4) $\rho_{i,y}^{sb,10Y}$
$IE_{i,t-1}^{fw}$	0.0145***	—	—	0.0131*
	(0.0050)			(0.0073)
$(D/Y)_{i,t-1}$	—	0.0003	—	-0.0002
		(0.0006)		(0.0006)
$y_{i,t-1}^{10}$	—	—	0.0133**	0.0038
			(0.0061)	(0.0068)
Macro controls	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	1,757	1,757	1,757	1,757
R^2	0.375	0.371	0.374	0.375

Note: This table reports panel regressions of the annual stock-bond correlation, the calendar-year correlation of monthly equity and monthly 10-year zero-coupon bond returns, on forward-looking interest expense and its components, each lagged one year. IE^{fw} is the product of debt to GDP and the 10-year yield. All four columns are estimated on the common sample the saturated column defines and carry country and year effects and the two macroeconomic controls. A dash marks a regressor the column does not enter. Standard errors, in parentheses, are two-way clustered by country and by year. *** p<0.01, ** p<0.05, * p<0.1.

and yields rose, the average burden reached 2.5 percent of GDP in 2023, and the average correlation turned positive in 2022 and remained positive through 2025. The United States alone traces the same path, from a correlation near 0.4 at a burden above five percent of GDP in 1990, through two decades in which the correlation was negative in all but two years and the burden drifted down from near four percent of GDP to two, to a positive correlation in 2022 to 2024 as the burden returned to five percent.[7] Appendix C.1 gives the slower-moving evidence behind this paragraph: Figure C.3 re-sorts the countries every month on their lagged interest expense, and the top third sits above the bottom third in all but one month since 1997; Figure C.4 traces the United States since 1962.

The within estimate is moderate because the specification asks a demanding question of the data. Month effects absorb the component of the correlation that all countries share in a month, which is where its largest movements occur, and together with the ratio they explain a third of its variance (Table 2, column (3)). The debt ratio changes once a year, so what identifies the coefficient is the yield moving against a slowly moving debt stock. Inflation is the leading account of the stock-bond correlation in United States

[7]The averages in this paragraph are of the monthly panel by calendar year; the 1990s figures are for the ten advanced economies observed by 1991, seven of which are present in 1990, so that entry into the panel does not move them. In the United States alone, the time-series regression of Table C.3 gives a slope of 0.21 (0.02) with an R^2 of 0.43, the highest in the table, and across the 29 advanced economies the slope is positive in 24 and significant at the five percent level in 15.

Table 4: The monthly baseline on later samples

Sample	(1) Full	(2) 1990+	(3) 2000+	(4) 2010+	(5) 2015+
A. Pooled OLS					
$IE^{fw}_{i,t-1}$	0.0399***	0.0387***	0.0330***	0.0354***	0.0504***
	(0.0097)	(0.0093)	(0.0082)	(0.0108)	(0.0061)
B. Two-way fixed effects and controls					
$IE^{fw}_{i,t-1}$	0.0138***	0.0145***	0.0171***	0.0156**	0.0317***
	(0.0029)	(0.0029)	(0.0038)	(0.0064)	(0.0064)
Observations	14,414	14,036	12,463	8,279	5,666
Months	653	433	313	193	133

Note: This table reports the monthly baseline on cumulative, nested samples, so it answers whether the result survives dropping early data. The cut is on the dependent variable's date and the boundary is inclusive, so the 1990 column begins in January 1990 and uses the December 1989 regressor on its first row. Both specifications within a column are estimated on the same rows, so the two panels differ by model and never by sample; every column carries all 44 countries. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1.

data (Campbell et al., 2017, 2020), and the panel agrees with it across countries: without country effects, higher inflation goes with a more positive correlation. Within a country and a month, on the country-months with inflation at most 10 percent, its coefficient is zero alone or beside the fiscal variable, whose coefficient is unchanged (Table C.5). On the within variation the fiscal variable carries the larger standardized effect, eight percent of a within-country standard deviation against six for inflation and under one for industrial production. A one-standard-deviation move is also a modest yardstick for a regressor with large realized swings. The median country's own range of ie^{fw}, from its tenth to its ninetieth percentile, is three percentage points of GDP and moves the predicted correlation by 0.04. Portugal's deterioration between 2008 and 2012, 8.5 points of GDP, implies 0.12, more than a third of a within-country standard deviation, over years in which its correlation went from −0.34 to +0.34. Across countries, on the monthly panel's country means, the association is three times as strong. Between the tenth and ninetieth percentiles of average interest expense, 3.4 points of GDP apart, the between-country slope predicts a gap in average correlation of 0.15 against 0.05 at the within estimate, and the observed gap is 0.28.[8]

Appendix C.1 collects the robustness checks for this section. Figure C.1 varies the construction of ie^{fw} in both of its choices. Along a blend running from realized interest payments, the current interest expense over GDP, to the forward-looking ratio, the estimate

[8]Section 5 reports the euro-area quasi-experiment. After the Greek deficit revision of October 2009, the change in the gap between the periphery and the core, relative to the period before the revision, averages 0.44 over the following eighteen months and peaks at 0.72.

is indistinguishable from zero at the backward-looking end and rises monotonically toward ie^{fw}, where it is also estimated most precisely, the end the model's forward debt obligation points to. Across the maturity of the yield inside the product the estimate is similar at every maturity and peaks at the ten years we use. Figure C.2 re-estimates the baseline dropping each country in turn; no exclusion changes the sign on either panel, and the largest move is 14 percent of the estimate. Table C.6 shows the association runs one way on the annual panel: lagged interest expense predicts the correlation beside the correlation's own lag, and the lagged correlation does not predict interest expense. Table C.1 replaces the dependent variable with its Fisher transformation and with the volatility-scaled correlation from the model's regression form, and the estimate stays positive and significant at the one percent level in every column. Table C.2 swaps the monthly debt leg for a linear interpolation, and the two constructions agree to within 0.0002 at ten years. Tables C.3 and C.4 report the regression country by country at both frequencies: the estimate is positive in 34 of 44 countries at each frequency, and the country estimates correlate at 0.80 across the two. The positive coefficients of Tables 2 and 3 survive removing the discretionary screens. Appendix B summarizes the discretionary screens, masks and corrections we apply to the raw data, which the replication package documents one by one, and Appendix C.5 re-estimates Tables 2 through 4 on a dataset rebuilt from the same raw inputs with all but the two estimation floors switched off. The cleaning mostly lowers the ten-year estimates: without the screens the coefficient is larger in every ten-year column of Table 2 and in both columns of Table 3 in which it enters, by up to 1.2 standard errors of the baseline, and smaller only in the five-year column; it keeps its sign and remains significant in every column of both builds.

The interest expense ratio could be a yield with a label. Its monthly variation comes mostly from the yield moving against a slowly changing debt stock, and a yield moves with monetary policy and the term premium as well as with fiscal risk. Three pieces of evidence speak to this concern. The policy rate leaves the association unchanged (Table 2). Section 5 compares countries that share monetary policy, a currency and an inflation target, and Section 6 measures the fiscal position from text with no market price in it; both find the association. Section 7 splits the yield and finds that within countries the association sits in the compensation for sovereign default, while the rest of the yield carries a precisely estimated zero. What these results leave open is how the yield's association scales with the debt stock, the question the annual horse race answers only at the ten percent level.

Section 5 takes a step further toward identifying the effect within the context of the euro area, where countries share a central bank, but differ in fiscal positions. Section 6 builds two measures of the fiscal position from text, newspaper fiscal salience and signed rating actions, and finds the same positive association. Section 7 decomposes the yield inside $\mathrm{ie}^{\mathrm{fw}}$ into its components and asks which carries the association.

5 A quasi-experiment: the euro area after October 2009

The panel regressions of Section 4 associate interest expense with the stock-bond correlation, but the regressor is not randomly assigned. The ideal comparison would contrast countries that experienced an unanticipated, sustained shift in fiscal sustainability with otherwise similar countries, with outcomes moving in parallel beforehand and a transmission channel running through sovereign credit, not through monetary or inflation factors. The euro area in 2009–2011 matches more of these features than any other episode in our sample: fiscal positions diverged within a group of countries that share a central bank, a currency, and an inflation target, so that within-area differences in sovereign yields reflect country-specific components, predominantly sovereign credit. An account based on the covariance regime of inflation operates at the level of the monetary regime, which these countries share; it can move the area-wide correlation but cannot by itself generate a divergence between core and periphery within it.[9] On 20 October 2009 the incoming Greek government announced that the 2009 deficit would be 12.7 percent of GDP, roughly twice the figure previously reported. The revision triggered a sequence of sovereign rating downgrades and repriced periphery sovereign credit. The euro-area crisis literature converges on this revision as the trigger of the sovereign debt crisis (Lane, 2012; Krishnamurthy et al., 2018), and we adopt it as the event date.

The sample is the eleven euro-area members with continuous daily data through the

[9] The country-specific inflation risk that emerged within the area, redenomination risk, is not an independent inflation account: pressure to leave the currency union arises from fiscal stress, so redenomination premia are a form the fiscal channel takes. Krishnamurthy et al. (2018) document redenomination premia in Spanish and Portuguese sovereign yields around the ECB's interventions.

event window: Greece, Ireland, Italy, Portugal and Spain as the periphery, and Austria, Belgium, Germany, Finland, France and the Netherlands as the core.[10] The dependent variable is the correlation of daily equity log returns with daily ten-year zero-coupon bond returns, computed inside event bins of three months anchored on the announcement, requiring at least 40 trading days per bin. Bin +1 covers 20 October 2009 to 20 January 2010; the announcement falls on a bin boundary, so no period is discarded as contaminated. With six bins on each side the panel is balanced at 132 country-bin observations, and the post-event window closes in April 2011, before the Securities Markets Programme reached Italian and Spanish debt and before Outright Monetary Transactions, so the estimates do not depend on the later ECB regime.[11] We estimate

$$\rho^{sb}_{i,k} \;=\; \alpha_i \;+\; \delta_k \;+\; \sum_{k\neq -1} \beta_k \cdot \Big(\text{Periphery}_i \times \mathbf{1}[\text{bin} = k]\Big) \;+\; \varepsilon_{i,k}, \tag{25}$$

where α_i and δ_k are country and event-bin fixed effects and $k = -1$ is the reference bin. Standard errors are country-clustered with a small-sample correction for the eleven clusters (Cameron et al., 2008), and we corroborate every reading with randomization inference over all $\binom{11}{5} = 462$ ways of assigning five of the eleven countries to the periphery, which is exact at eleven clusters and does not rely on asymptotics in their number.

Figure 3 reports the path. The coefficient in the first bin after the announcement is small, 0.12 with a randomization p-value of 0.10, consistent with a trigger whose pricing implications took a quarter to reach the realized correlation. The gap appears in the second bin and persists through the end of the window. The largest coefficient, 0.72, falls in the bin from 20 April 2010, the window containing the first Greek assistance package. Averaged over the post-event window, the change in the periphery-core gap relative to the reference bin is 0.44, and no assignment of five countries produces a larger estimate than the actual periphery, a randomization p-value of 0.002.[12]

The periphery label was produced by the crisis itself, so we verify that it is not chosen after the fact. Sorting the eleven countries on their average five-year sovereign CDS spread over the first half of 2008, before the failure of Lehman Brothers and fifteen months

[10] Slovenia, the twelfth long-standing member, has no bond quotations in the reference window and is excluded from the reported specification; Appendix C.2 restores it.

[11] The window also closes before the second Greek assistance package (March 2012) and the "whatever it takes" speech (July 2012).

[12] The standard error of the post-event average is 0.09.

Figure 3: The euro area around the October 2009 Greek deficit revision

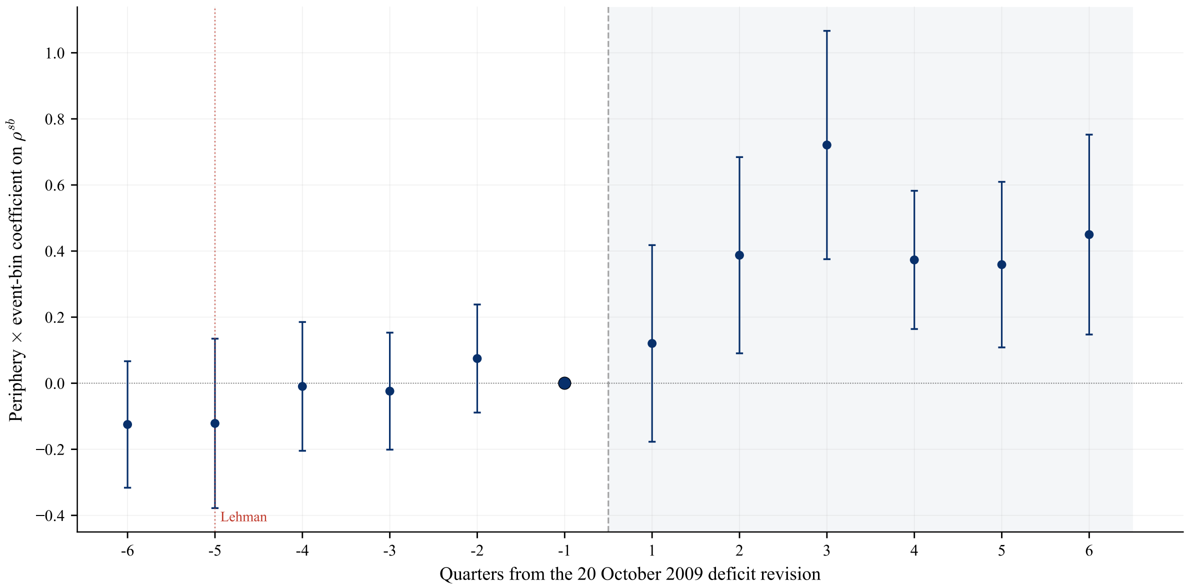

Note: Each point is the coefficient on the interaction of a periphery indicator with an event-bin indicator, from a regression with country and bin fixed effects; bin -1 is the reference. The dependent variable is the correlation of daily equity log returns with daily 10-year zero-coupon government bond returns inside each bin, requiring at least 40 trading days. Bins are three-month windows anchored on the announcement of 20 October 2009, six on either side. The sample is the 11 euro-area members with continuous data through the window, 132 country-bin observations; the periphery is Greece, Ireland, Italy, Portugal and Spain. Bars are 95 percent confidence intervals from country-clustered standard errors with a small-sample correction, a $t(10)$ critical value and a $\sqrt{G/(G-1)}$ inflation. The dotted vertical line marks the failure of Lehman Brothers, the dashed line the announcement, and the shaded area the post-event window.

before the revision, puts the same five countries above the median: the grouping the crisis literature settled on is the one predetermined credit pricing selects. The classification is also not knife-edge. Moving Belgium, the country closest to the median, into the periphery changes the difference-in-differences estimate from 0.436 to 0.444 and leaves the path in Figure 3 materially unchanged. And the indicator is not doing the work: replacing it with the predetermined spread itself gives 0.243 per standard deviation, and with the periphery 1.71 standard deviations above the core the implied group difference is 0.42 against the reported 0.44 (Figure C.5).

The pre-event dynamics are broadly consistent with parallel trends, with two qualifications. The five pre-event estimates run from -0.12 to 0.07, none distinguishable from zero under country-clustered inference, and four of the five have randomization p-values between 0.23 and 0.92. The earliest bin rejects at the five percent level under randomization inference (p 0.035), but its estimate of -0.12 is opposite in sign to the post-event divergence and sits six bins before the event, so we do not read it as anticipation. A linear trend fitted to the pre-event coefficients gives 0.034 per bin; the randomization p-value of a zero-trend test is 0.069, the fitted drift is smaller than the single-bin jump at $k = 2$, and a joint test on all five pre-event bins gives p 0.22. These tests cover a demanding period: the pre-announcement window opens in April 2008 and spans the failure of Lehman Brothers, with equity volatility across the eleven countries averaging 37 percent annualized against 24 percent over 1999–2019, so the period in which the periphery-core gap fails to move is a turbulent one.

Two features of the episode limit what the design can claim. First, the largest coefficient falls in the window containing the first Greek assistance package, so the estimate is the response of correlations to the fiscal news and to the policy reaction it set in motion; the design does not separate the two. Second, flight to quality can inflate the gap: if investors responded to periphery stress by moving into core debt, the core correlation falls with the periphery's own shock, and part of the estimated gap would reflect safe-haven flows into the control group. The non-euro-area advanced economies also moved toward more negative correlations over 2010 and 2011, so a global risk-off channel operated alongside the fiscal one. In the raw series, however, the gap is driven by the periphery, whose average correlation moves from -0.17 to $+0.23$ across the announcement while the core's is nearly unchanged, from -0.37 to -0.35, so the estimate is not a safe-haven repricing of core debt. We do not claim the design eliminates these channels, only that the sign and timing of the gap survive after acknowledging them.

Appendix C.2 re-estimates the design under sixteen variations, collected in Figure C.5. Across the thirteen variations of treatment, window, sample and dependent variable the estimate lies between 0.37 and 0.49, every interval excludes zero, and dropping Greece, whose deficit is the subject of the announcement, gives 0.40. Three placebo rows anchor a false boundary two, three and four bins before the announcement; none is distinguishable from zero. The within-area evidence is consistent with fiscal positions transmitted through sovereign credit risk, the account Section 7 tests directly by decomposing the yield inside $\mathrm{ie}^{\mathrm{fw}}$ into its credit and non-default components.

6 Fiscal shifters

Section 4 measures the fiscal state from prices, a debt stock multiplied by a yield, and the correlation it predicts is built from returns. Both sides of that regression are market prices. Suppose investors demand a higher return on a country's assets for reasons other than its public finances. The yield rises, and the regressor, a debt stock multiplied by that yield, rises with it; if the same repricing moves the country's stocks and bonds in one direction together, the correlation rises as well. The association would then come from a common repricing and not necessarily from the fiscal channel of the model. We measure the flow of fiscal news from text: how much of a country's Wall Street Journal coverage discusses the sustainability of its public finances, which we call its fiscal salience, and whether a rating agency acting on the country cites the public finances as its reason. The first measure is dense, one share per country-month; the second is sparse but carries the agency's stated reason as the newswire reports it. We ask three things of them: whether a higher value of each is associated with a higher stock-bond correlation within country and month, as $\mathrm{ie}^{\mathrm{fw}}_{i,t-1}$ is in Section 4; whether that association is stronger where the fiscal burden is already high, as the model implies for the response to news at a high default hazard, a question a regressor that measures the state itself cannot pose; and whether the fiscal content of the text adds to what general bad news predicts and to what $\mathrm{ie}^{\mathrm{fw}}_{i,t-1}$ predicts. Section 3 and Appendices B.8 and B.9 describe both measures. Within country and within month, higher salience predicts a higher correlation the next month, and a signed fiscal rating action, downgrades positive, is associated with a higher correlation in its own month. The association of the index is strongest where the fiscal burden is already high. Beside $\mathrm{ie}^{\mathrm{fw}}_{i,t-1}$ the index keeps a positive association and the rating estimate falls and is indistinguishable from zero. Standard errors in the panel regressions are two-way clustered by country and by month.

The fiscal salience index for a country-month, $\mathrm{FSI}_{i,t}$, is the share of the articles attributed to the country in which the classifier finds a substantive discussion of the sustainability of the public finances. An attribution pass assigns each article to the countries it discusses, and the label records whether the topic is present, whatever tone the article takes. Because the index is a share of the month's coverage, the logarithm of one plus the article count enters every specification, so that the coefficient compares shares at a given volume of coverage. Both the index and the count are lagged one month, so the index is measured before the month whose correlation it predicts. The panel covers 44 countries from February 2000 to January 2026, 11,802 country-months. We estimate

$$\rho^{sb}_{i,t} \;=\; \alpha_i \;+\; \delta_t \;+\; \beta\,\mathrm{FSI}_{i,t-1} \;+\; \theta\log(1+n_{i,t-1}) \;+\; X'_{i,t}\gamma \;+\; \varepsilon_{i,t}, \tag{26}$$

where α_i and δ_t are country and month fixed effects, $n_{i,t-1}$ is the article count, and $X_{i,t}$ contains consumer price inflation and industrial production growth, which enter contemporaneously on this panel.[13]

Table 5 builds the specification as Table 2 does, in three steps, with columns (1) through (3) on a common sample. The pooled coefficient is 0.121, significant only at the ten percent level. Country and month effects together take it to 0.105, and the macro controls move it to 0.101, our baseline (column 3), significant at the one percent level. A one-standard-deviation move in the index within country and month, 0.16, shifts the predicted correlation by about 0.016, six percent of the correlation's own standard deviation within country and month, 0.26. Dropping 2008 and 2009 tests whether the association is confined to the financial crisis and leaves the estimate at 0.094 (column 4), close to the baseline. The last column asks whether the association has weakened since the crisis. From 2010 onward it is 0.056 (column 5), significant at the ten percent level, about half the baseline.

Table 5: Newspaper fiscal salience and the stock-bond correlation

VARIABLES Sample	(1) $\rho^{sb,10Y}_{i,t}$ Full	(2) $\rho^{sb,10Y}_{i,t}$ Full	(3) $\rho^{sb,10Y}_{i,t}$ Full	(4) $\rho^{sb,10Y}_{i,t}$ Excl. 2008–2009	(5) $\rho^{sb,10Y}_{i,t}$ 2010+
$FSI_{i,t-1}$	0.1210*	0.1053***	0.1014***	0.0945***	0.0560*
	(0.0652)	(0.0346)	(0.0342)	(0.0332)	(0.0315)
Attention control	Yes	Yes	Yes	Yes	Yes
Country FE	No	Yes	Yes	Yes	Yes
Month FE	No	Yes	Yes	Yes	Yes
Macro controls	No	No	Yes	Yes	Yes
Observations	11,802	11,802	11,802	10,815	7,662
Countries	44	44	44	44	44
R^2	0.040	0.396	0.398	0.399	0.410

Note: This table reports panel regressions of the monthly stock-bond correlation, the non-overlapping correlation of daily equity and daily 10-year zero-coupon bond returns within a calendar month, on the newspaper fiscal-salience index lagged one month. The panel covers 44 countries over 2000-02–2026-01. A country-month enters when it carries at least 15 trading days. The index is the share of that country-month's Wall Street Journal articles in which a language model finds a substantive discussion of the sustainability of the public finances, the topic and not its tone: the denominator counts the articles attributed to the country in the month and the numerator counts those of them the model flags. Articles are assigned to countries by a separate attribution pass of a language model, on both sides of the share, under attribution rule attr_v1.0. The attention control, the logarithm of one plus the article count, enters every column, including the pooled one, because the index is a share whose denominator is that count. The macro controls on this panel are contemporaneous, while Section 4's monthly controls are lagged; a footnote in Section 6 re-estimates the reported column with lagged controls. The global series (`vix`, `ebp`) take one value per month, so the month effects absorb them in every fixed-effect column and they are carried there for parity with the pooled column alone. Columns (1), (2) and (3) are estimated on one sample fixed before the first column, over every regressor in the table and over the lagged interest-expense measure that Appendix C.3 adds to the reported column, so a change across them is the specification; columns (4) and (5) restrict the sample as the headings state. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1. P-values behind the stars use the standard normal reference for the clustered t-statistic in the Python build and a t-distribution with the smaller cluster count less one degrees of freedom in the Stata build.

[13] In Section 4 and in equation (27) the same two controls are lagged one month. Lagging them one month here tests sensitivity to control timing and leaves the reported column at 0.101 with a standard error of 0.034, on the 11,800 country-months that carry the lags.

Table 6 asks whether the association depends on the level of fiscal burden. In the model the bond's exposure to news rises with the default hazard, which rises with the forward debt obligation (equation 17), so the association should be larger where the burden is already high. Panels A and B re-estimate the baseline inside four bins of the lagged fiscal burden, ie^{fw} in Panel A and the debt ratio in Panel B. The bins are cut on an expanding window, with cut points computed on the observations dated before each country-month. In each panel column (5) estimates all bins jointly and tests whether the slopes are equal. The association is strongest where the burden is highest. In the bins sorted on ie^{fw} the estimate rises from -0.029 in the lowest bin to 0.189 in the top bin. There it is significant at the one percent level under two-way clustering and under a wild cluster bootstrap on the 26 countries in the bin (Cameron et al., 2008), and equal slopes are rejected at the one percent level. The yield in ie^{fw} could sort nonfiscal repricing into the top bin. In the bins sorted on the debt ratio, the top bin gives 0.200 and equal slopes are again rejected at the one percent level. On either expanding-window sort the three lower bins are individually indistinguishable from zero. Where the rise begins depends on how the bins are cut. With the bins cut once on the whole sample, the estimate increases across all four bins on both sorting variables, and the third bin is also positive at the five percent level under two-way clustering with a normal reference distribution, so the rise begins below the top bin (Table C.8). On the debt sort that third bin has a wild cluster bootstrap p-value of 0.066. The pattern is consistent with the state dependence of the model's default hazard.

The second measure uses rating actions taken for fiscal reasons. Figure 4 shows the cross-country pattern behind it: countries with lower average sovereign ratings have higher average stock-bond correlations. The 14 countries whose average rating sits within one notch of AAA have an average stock-bond correlation of 0.07, and the lowest-rated quarter of the sample, 11 countries, averages 0.33. The regressions below ask whether fiscal rating actions are associated with the correlation within country and month, once the country effects absorb these cross-country levels. The rating shifter is the signed notch count of Section 3, downgrades positive, divided by the number of agencies covering the country that year. The newswire coverage republishes or paraphrases the agency's statement. An action enters under the filter F_C, which keeps it only where that coverage attributes to the agency a fiscal rationale and no cyclical, banking, external or political cause beside it.

Table 6: Newspaper fiscal salience by bin of the lagged fiscal burden

VARIABLES	(1) $\rho_{i,t}^{sb,10Y}$	(2) $\rho_{i,t}^{sb,10Y}$	(3) $\rho_{i,t}^{sb,10Y}$	(4) $\rho_{i,t}^{sb,10Y}$	(5) $\rho_{i,t}^{sb,10Y}$
Panel A. $IE_{i,t-1}^{fw}$*, expanding window*					
Observed range	[-0.63,2.47]	[1.01,3.37]	[1.90,5.44]	[3.07,54.59]	
Mean $IE_{i,t-1}^{fw}$	0.78	1.95	2.81	5.90	2.53
$FSI_{i,t-1}$	-0.0286	0.0513	0.0780	0.1886***	-0.0628
	(0.0319)	(0.0477)	(0.0491)	(0.0627)	(0.0382)
Wild cluster p				[0.008]	
× Bin 2					0.0875 (0.0599)
× Bin 3					0.2064*** (0.0552)
× Bin 4					0.3088*** (0.0740)
Equal slopes, p					<0.0001
Observations	4,433	2,536	2,250	2,541	11,760
Countries	36	39	36	26	44
Panel B. $Debt/GDP_{i,t-1}$*, expanding window*					
Observed range	[0.05,45.93]	[34.94,60.35]	[49.70,81.29]	[66.41,258.37]	
Mean Debt/GDP$_{i,t-1}$	25.19	43.43	61.99	110.67	64.00
$FSI_{i,t-1}$	0.0078	-0.0376	0.0675	0.1999***	0.0140
	(0.0386)	(0.0342)	(0.0450)	(0.0480)	(0.0281)
Wild cluster p	[0.836]		[0.101]	[0.001]	
× Bin 2					-0.0328 (0.0386)
× Bin 3					0.0702 (0.0811)
× Bin 4					0.2406*** (0.0618)
Equal slopes, p					0.0001
Observations	2,653	2,988	2,489	3,630	11,760
Countries	27	33	28	22	44
Attention control	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes
Macro controls	Yes	Yes	Yes	Yes	Yes

Note: Each panel re-estimates the headline column of Table 5 inside bins of the sorting variable named in the panel heading, and column (5) estimates all bins jointly with the index's slope free to differ by bin. The interaction rows report the difference from Bin 1, with the standard error in parentheses beside the coefficient, and Equal slopes reports the joint test that those differences are zero. Bins are cut across the pooled estimation sample and do not partition countries: a country moves between bins as its sorting variable moves. Intervals are half open and closed on the right, so a value sitting exactly on a cut enters the lower bin alone. Panels A and B assign each country-month to the bin its state falls in under cut points computed on the observations dated strictly before that month; because the cut points move, a bin's observed range there overlaps its neighbors; Panel A and Panel B leave 42 early rows unbinned each, because their history is too short to cut. Table C.8 repeats both panels with the cut points computed once on the whole sample. The debt series carries the preceding December's annual value through the following calendar year, so 0 of 1,188 country-years take more than one monthly value and ties on a cut point are the ordinary case in Panel B. A bin whose country count falls below 30 carries a wild cluster restricted bootstrap-t p in brackets, 9,999 draws with Rademacher weights clustered on country. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1. P-values behind the stars use the standard normal reference for the clustered t-statistic in the Python build and a t-distribution with the smaller cluster count less one degrees of freedom in the Stata build.

The shifter is dated to the month of the action, and the macro controls are lagged one month as in Section 4. The panel covers 44 countries from February 1974 to January 2026, 14,385 country-months; one of the 44 countries, Switzerland, records no rating action in the panel. A filter-passing action falls in 205 of them and the shifter moves in 195; in the other ten the only passing actions are moves into or out of review status, which the notch scale scores at zero. The newswire archive used for classification begins in January 1996, so no action enters before then, and the earlier months stay in the sample with the shifter at zero.[14] Downgrades and upgrades need not have equal and opposite

[14]Estimating only on the months the archive covers gives 0.162 with a standard error of 0.063 on 13,271

Figure 4: Sovereign rating and the stock-bond correlation across countries

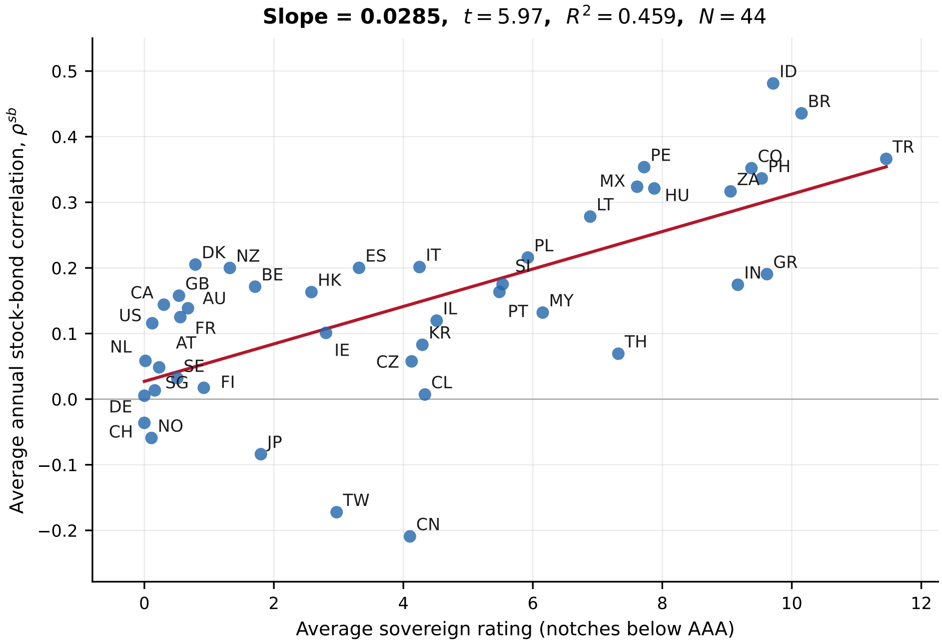

Note: Each point is one country, labeled by its ISO code. The horizontal axis is the time-series average of the year-end sovereign rating, in notches below AAA. At each year end we average the recorded ratings of Moody's, S&P and Fitch with equal weights; an agency enters from its first dated rating in our data. The vertical axis is the time-series average of the calendar-year correlation ρ^{sb} of Section 4 over the same country-years, in the annual panel of this section's rating actions. The line is an ordinary least squares fit with equal country weights. The figure reports the slope, its t-statistic on the conventional OLS standard error, and R^2.

associations. Both rating tables report the net action and, in a second panel, its two arms, Action = Down − Up, with the p-value of the symmetry restriction $\beta_{\text{Down}} + \beta_{\text{Up}} = 0$ that the single coefficient imposes. We estimate

$$\rho^{sb}_{i,t} \;=\; \alpha_i \;+\; \delta_t \;+\; \beta\,\text{Action}^{F_C}_{i,t} \;+\; X'_{i,t-1}\gamma \;+\; \varepsilon_{i,t}, \tag{27}$$

where α_i and δ_t are country and month fixed effects and $X_{i,t-1}$ contains consumer price inflation and industrial production growth.

Table 7 associates an agency's fiscal reassessment with comovement in the month it acts, consistent with the model's pricing of fiscal sustainability as default risk. The pooled coefficient is 0.126, significant at the ten percent level. Month effects absorb what every country shares in a month and leave the slope imprecise; country effects remove persistent national differences and sharpen it. With both sets of effects it is 0.158, and the macro controls leave it there, our baseline (column 5), significant at the five percent level. In Panel B the downgrade arm is the larger in absolute value and the one significant at the

country-months, so the earlier months, in which the shifter is zero, do not carry the estimate.

one percent level, which fits a bond more sensitive to news where the default hazard is high, if a downgrade marks the more stressed state. Symmetry is not rejected, with a p-value of 0.505, on an upgrade arm too imprecise to reject it. One notch of downgrade by every agency covering a country is associated with a shift in the predicted correlation of 0.19, about seven tenths of the correlation's standard deviation within country and month, 0.26. The shifter moves in only 195 of 14,385 country-months and has a standard deviation within country and month of 0.046, so a one-standard-deviation move in the net action shifts the predicted correlation by 0.007, about three percent of the correlation's.

Table 7: Sovereign rating actions and the stock-bond correlation

VARIABLES	(1) $\rho^{sb,10Y}_{i,t}$	(2) $\rho^{sb,10Y}_{i,t}$	(3) $\rho^{sb,10Y}_{i,t}$	(4) $\rho^{sb,10Y}_{i,t}$	(5) $\rho^{sb,10Y}_{i,t}$
Panel A. Net action					
$Action^{F_C}_{i,t}$	0.1260*	0.1388	0.1343**	0.1584**	0.1580**
	(0.0749)	(0.0933)	(0.0571)	(0.0630)	(0.0635)
Panel B. Two arms					
$Down^{F_C}_{i,t}$	0.2338***	0.3044***	0.1294**	0.1861***	0.1878***
	(0.0525)	(0.0581)	(0.0511)	(0.0586)	(0.0592)
$Up^{F_C}_{i,t}$	0.0384	0.1154	-0.1419	-0.1162	-0.1127
	(0.1481)	(0.1286)	(0.1200)	(0.1038)	(0.1031)
Symmetry p	0.084	0.003	0.923	0.536	0.505
Country FE	No	No	Yes	Yes	Yes
Month FE	No	Yes	No	Yes	Yes
Macro controls	No	No	No	No	Yes
Observations	14,385	14,385	14,385	14,385	14,385
R^2	0.000	0.291	0.114	0.436	0.438

Note: This table reports panel regressions of the monthly stock-bond correlation, the non-overlapping correlation of daily equity and daily 10-year zero-coupon bond returns within a calendar month, on a monthly measure of sovereign rating actions. The measure counts the rating notches the agencies moved a country in a month, signed so that downgrades enter *Down* and upgrades enter *Up*, divided by the number of agencies covering that country in that year, so a coefficient is the association with one notch of movement by the country's full set of raters. Panel A uses the net action $Action = Down - Up$, which is positive in the downgrade direction. Panel B frees the two arms and reports the p-value of the restriction Panel A imposes, that $b_{Down} + b_{Up} = 0$. An action enters under the F_C filter, which keeps it only where the agency's own stated rationale is fiscal and the agency named no cyclical, banking, external or political cause beside it; the filter and the classification behind it are described in Appendix B. The sample is 44 countries over 1974-02–2026-01, 14,385 country-months; a filter-passing action falls in 205 of them and moves the shifter by at least a fraction of a notch in 195. The estimation sample is fixed before column (1), over every regressor in the table and over the lagged interest-expense measure that Appendix C.3 adds to column (5), so a change across columns is the specification and never the sample. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1. P-values behind the stars use the standard normal reference for the clustered t-statistic in the Python build and a t-distribution with the smaller cluster count less one degrees of freedom in the Stata build.

Appendix C.3 collects the robustness checks for both measures. We choose which rating actions count as fiscal, and that choice could select the actions that happen to line up with the correlation. We therefore compare three nested rationale filters. Admitting every action, whatever reason the agency gave, gives 0.120. Requiring a fiscal rationale with no other cause named beside it gives 0.158 (Table C.9). The association survives without the filter and strengthens with it. Within the filtered set it is carried by the actions citing the structural fiscal position, the deficit, the debt level or the interest burden, while the actions citing fiscal governance give an imprecise estimate (Table C.9). The filter uses agencies' own stated fiscal rationales to select the news the model is about: reassessments of the government's ability to pay.

The first check on the index asks whether it captures country attention that raises comovement across assets. With country and month effects and no macro controls, the index gives 0.003, indistinguishable from zero, for the correlation of the country's equity returns with United States equity returns. The same specification gives 0.105 on the stock-bond correlation (Table C.14). A general attention channel would move both correlations, and the placebo gives no evidence of one. The index could be counting bad news about the country, and bad news of any kind can move stocks and bonds together. A Loughran-McDonald dictionary negativity score from the same articles predicts the correlation on its own. Entered beside the index, it does not absorb the index, which keeps 0.093 (Table C.11). Policy uncertainty is a third candidate for what the index measures. On the 19 countries with an economic policy uncertainty index, controlling for it moves the newspaper index from 0.260 to 0.256 on the same observations (Table C.12).

To test whether fiscal flags depend on country identity, a second pass replaces country names, currencies, officials, institutions and dates with placeholders and reclassifies the articles. Article attribution to countries is held fixed across arms. On rewritten text the index gives 0.049 against 0.075 on unmodified text from the same country-months. The point estimate falls by a third and stays positive at the ten percent level (Table C.13).[15] The decline suggests that part of the association depends on the classifier knowing the country. Assigning articles to countries they do not discuss adds noise to both numerator and denominator. Requiring a country alias on both sides of the share keeps the estimate positive and significant at the one percent level. Correctly attributing an article does not ensure that its fiscal discussion concerns the same country's public finances. Tying the numerator's fiscal discussion to the country gives 0.351 on that narrower share, significant at the one percent level. The estimate is larger where the share rests on more articles, consistent with attenuation from measurement error in shares with thin denominators (Table C.16).

The association appears in both the advanced and the emerging group (Table C.15). Leaving each country out in turn moves the baseline index estimate between 0.084 and

[15] Some identifiers survive the pass, and Appendix B.9 counts them. Dropping the article-country pairs in which one survives gives 0.043 (0.025) on 4,424 country-months.

0.108 and the rating estimate between 0.090, without Greece, and 0.179, without Indonesia (Table C.17). No drop changes either sign. Removing Greece moves the rating estimate the most and leaves it indistinguishable from zero. The rating measure can be sorted the same way, to ask whether its association is also stronger where the fiscal burden is already high. Rating actions fall in only 205 country-months, so the bins of Table 6 carry few of them each. Using the same sorting variables and expanding-window cuts on the rating sample, only the top bin is significant at the five percent level under two-way clustering on either sort. Under the wild cluster bootstrap it is significant only on the debt sort. The test of equal slopes does not reject at the five percent level in any panel and rejects at ten percent in the interest-expense sort alone (Table C.10).

The two measures agree in sign, but the yield in $\mathrm{ie}^{\mathrm{fw}}_{i,t-1}$ may already price the fiscal information they contain. Added to the reported columns of Tables 5 and 7, on identical rows, $\mathrm{ie}^{\mathrm{fw}}_{i,t-1}$ leaves the index positively associated with the correlation and the rating coefficient smaller and indistinguishable from zero.[16] The comparison neither isolates the yield's contribution nor shows that the rating measure adds no information.

7 Decomposing interest expense: credit and non-default components

The fiscal variable of Section 4 multiplies a debt stock by a yield, and a yield bundles compensation for default with the real rate, expected inflation and the term premium. The model divides its channels along that line: under monetary dominance fiscal stress is priced as default risk, under fiscal dominance as inflation, and the convenience and safe-haven channels, which the decomposition does not isolate, sit in the real component. Where a sovereign credit default swap is quoted we split the interest bill into a credit share, $(D/Y)_{i,t-1} \cdot CDS_{i,t-1}/100$, and a non-default share, $(D/Y)_{i,t-1} \cdot (y^5_{i,t-1} - CDS_{i,t-1})/100$, which sum to the five-year $\mathrm{ie}^{\mathrm{fw}}_{i,t-1}$.[17] For the fourteen countries with an inflation-linked

[16]With the controls and fixed effects of the reported columns and $\mathrm{ie}^{\mathrm{fw}}_{i,t-1}$ added, the index gives 0.081 (0.032) against 0.101 (0.034) alone, on 11,802 country-months, and the rating shifter gives 0.064 (0.058) against 0.158 (0.064) alone, on 14,385; $\mathrm{ie}^{\mathrm{fw}}_{i,t-1}$ enters at 0.0139 (0.0039) beside the index and at 0.0134 (0.0031) beside the shifter. Table C.7 in Appendix C.3 reports both pairs.

[17]Every spread is a Markit composite on a senior unsecured five-year contract, quoted in dollars, and in euros for the United States (Appendix B.6). A dollar spread differs from local-currency default risk by a quanto component that prices expected depreciation on default (Du and Schreger, 2016; Kremens and Martin, 2019), and the spread on a euro-area sovereign also prices redenomination (Krishnamurthy et al.,

curve the non-default share splits further into a real yield and breakeven inflation. Without country effects every component of the yield enters positively; within countries only the credit share does.[18]

Table 8 reports the decomposition for all 44 countries from 2001.[19] Without country effects, in columns (1) through (4), both shares enter positively: across countries the level of the interest bill predicts the level of the correlation whichever component of the yield carries it, the level relationship of Table 4.

Country effects separate the shares. In the baseline column (6) the credit share enters at 0.055, significant at the one percent level under two-way clustered inference and at five percent under the wild cluster bootstrap, while the non-default share is indistinguishable from zero; the difference between the two, estimated as one coefficient by entering the credit share beside ie^{fw}, has a bootstrap p-value of 0.007. The zero is precise: the two shares have similar within-country dispersion, and the confidence interval on the non-default share excludes both its own pooled loading in column (4) and the credit loading. On this sample the five-year ie^{fw} alone carries twice its full-sample coefficient of Table 2 within country and month, and the decomposition assigns it to the credit share. Column (7) re-estimates the baseline at annual frequency, since only nine percent of the credit share's within-country variance occurs month to month inside a year, and the coefficient is essentially unchanged on a twelfth of the observations.

The two shares differ in what the fixed effects leave. Half of the non-default share's within-country variance is common to all countries in a month, and the rest follows the country's own policy rate and term premium; the policy-rate control of Table 2 leaves the within estimate unchanged. A tenth of the credit share's within-country variance is common, and 84 percent of it is the euro periphery's. After 2021 the within-country dispersion of the credit share falls to a quarter of its earlier size, and the credit association is concentrated in the years before.

2018), which Section 5 treats as a form the fiscal channel takes. The quanto component is partly absorbed by the breakeven term of Table 9, since expected depreciation on default also enters local-currency inflation compensation.

[18] The swap is a five-year contract, so the decomposition splits the five-year yield and is estimated on the five-year correlation; column (7) of Table 2 shows the relationship essentially unchanged at that maturity.

[19] The sample excludes country-months whose lagged spread exceeds 1,000 basis points, where a sovereign contract prices a high probability of restructuring and the linear reading no longer applies; the exclusion removes 34 estimation rows, 26 of them Greek. Appendix C.4 tightens it to 500 and to 200 basis points, and the credit coefficient rises to 0.10 and then 0.14, so the baseline sits at the conservative end. No spread in the inflation-linked panel of Table 9 reaches the threshold.

Table 8: Credit and non-default interest expense and the five-year stock-bond correlation

VARIABLES	(1) $\rho_{i,t}^{sb,5Y}$	(2) $\rho_{i,t}^{sb,5Y}$	(3) $\rho_{i,t}^{sb,5Y}$	(4) $\rho_{i,t}^{sb,5Y}$	(5) $\rho_{i,t}^{sb,5Y}$	(6) $\rho_{i,t}^{sb,5Y}$	(7) $\rho_{i,y}^{sb,5Y}$
Credit	0.0687***	0.0683***	0.0700***	0.0668***	0.0511***	0.0549***	0.0525***
	(0.0175)	(0.0175)	(0.0190)	(0.0184)	(0.0111)	(0.0152)	(0.0090)
						[0.0365]	
Non-default	0.0379***	0.0325***	0.0312***	0.0434***	-0.0013	-0.0100	-0.0115
	(0.0088)	(0.0102)	(0.0104)	(0.0091)	(0.0119)	(0.0110)	(0.0129)
						[0.3713]	
Macro controls	No	Yes	Yes	Yes	Yes	Yes	Yes
VIX control	No	No	Yes	No	No	No	No
Country FE	No	No	No	No	Yes	Yes	Yes
Time FE	No	No	No	Yes	No	Yes	Yes
Observations	10,848	10,848	10,848	10,848	10,848	10,848	911
R^2	0.092	0.098	0.112	0.286	0.209	0.397	0.587

Note: This table reports panel regressions of the monthly five-year stock-bond correlation on the credit and non-default shares of forward-looking interest expense, $(D/Y)_{i,t-1} \cdot CDS_{i,t-1}/100$ and $(D/Y)_{i,t-1} \cdot (y^5_{i,t-1} - CDS_{i,t-1})/100$, both in percent of GDP and summing to $IE^{fw}_{i,t-1}$. The sample is 44 countries over 2001–2026, excluding country-months whose lagged spread exceeds 1,000 basis points, and is fixed before column (1). Column (7) re-estimates column (6) at annual frequency, with annual averages of the regressors and controls lagged one year and the annual-from-daily five-year correlation as the dependent variable; the Time FE row is calendar months in columns (4) and (6) and years in column (7). Standard errors, in parentheses, are two-way clustered by country and by month. Brackets in column (6) report the p-value of a wild cluster restricted bootstrap-t clustered by country, over 9,999 draws with seed 20260824. *** p<0.01, ** p<0.05, * p<0.1.

Table 9 repeats the decomposition on the fourteen countries with an inflation-linked government curve, where the non-default share splits into a real yield and breakeven inflation net of default compensation. Pooled, all three components enter positively. Within country and month the credit share enters at 0.139 and survives the wild cluster bootstrap, enumerated exactly over the fourteen clusters, while the real yield and breakeven are indistinguishable from zero; with country effects alone both are positive and imprecise, and the month effects, which absorb two fifths of each one's within-country variance, remove them. The non-default zero of Table 8 is therefore not two offsetting effects: neither the level of real rates nor priced inflation compensation is associated with the correlation within a country.

In the model the forward debt obligation is the observable through which every channel but the safe-haven one moves the correlation, and its slope pools the credit and convenience channels without distinguishing the components of the coupon (Propositions 4 and 5); the pooled estimates have that form. State dependence enters through the default hazard under monetary dominance, which the swap prices, through the price level under fiscal dominance, which would surface in breakeven inflation, and through the debt ratio in

Table 9: Credit, the real yield and breakeven inflation, countries with an inflation-linked curve

VARIABLES	(1) $\rho_{i,t}^{sb,5Y}$	(2) $\rho_{i,t}^{sb,5Y}$	(3) $\rho_{i,t}^{sb,5Y}$	(4) $\rho_{i,t}^{sb,5Y}$	(5) $\rho_{i,t}^{sb,5Y}$	(6) $\rho_{i,t}^{sb,5Y}$
Credit	0.1456***	0.1401***	0.1414***	0.1523***	0.1046*	0.1391***
	(0.0373)	(0.0349)	(0.0315)	(0.0272)	(0.0538)	(0.0288)
						[0.0305]
Real yield	0.0486***	0.0419***	0.0407***	0.0512***	0.0285	-0.0150
	(0.0118)	(0.0100)	(0.0104)	(0.0127)	(0.0224)	(0.0133)
						[0.2731]
Breakeven	0.0561***	0.0368**	0.0299**	0.0115	0.0480	-0.0308
	(0.0115)	(0.0135)	(0.0130)	(0.0162)	(0.0382)	(0.0200)
						[0.2612]
Macro controls	No	Yes	Yes	Yes	Yes	Yes
VIX control	No	No	Yes	No	No	No
Country FE	No	No	No	No	Yes	Yes
Month FE	No	No	No	Yes	No	Yes
Observations	2,291	2,291	2,291	2,291	2,291	2,291
R^2	0.217	0.230	0.250	0.558	0.281	0.606

Note: This table reports panel regressions of the monthly five-year stock-bond correlation on the three components of forward-looking interest expense for the 14 countries with an inflation-linked government curve: $(D/Y)_{i,t-1} \cdot r^5_{i,t-1}/100$, a real yield; $(D/Y)_{i,t-1} \cdot (y^5_{i,t-1} - r^5_{i,t-1} - CDS_{i,t-1})/100$, breakeven inflation net of default compensation; and $(D/Y)_{i,t-1} \cdot CDS_{i,t-1}/100$, credit. All are in percent of GDP and sum to $IE^{fw}_{i,t-1}$. The sample is 2,291 country-months over 2004–2026, fixed before column (1). Standard errors, in parentheses, are two-way clustered by country and by month. Brackets in column (6) report the p-value of a wild cluster restricted bootstrap-t clustered by country, enumerated over all 16,384 sign vectors. *** p<0.01, ** p<0.05, * p<0.1.

the convenience channel (Proposition 3), which moves once a year and multiplies every share alike, so monthly variation cannot separate it. The within-country estimates load on the credit share and on neither the real yield nor the breakeven, consistent with the monetary-dominance regime, in which fiscal stress is priced as default risk, and they support the interpretation that interest expense tracks the correlation through the market's assessment of fiscal sustainability.

8 Conclusion

The hedging property of government bonds is recent, uneven and fragile. It emerged in the advanced economies around 2000, did not appear in the emerging-market average, and weakened again after 2022. The fiscal cost of carrying the public debt is associated with these differences across countries and over time. Where the forward-looking interest expense ratio is low, sovereign bonds appreciate in equity downturns and hedge; where it is high, bonds and equities move together. A model with a safe-haven channel, a fiscal-capacity channel and a convenience channel generates this pattern under its stated conditions, and it motivates the multiplicative form, debt times the rate at which it reprices, which the data do not reject against either factor alone.

The pattern is not an artifact of any one design. It appears across countries and within them, keeps its coefficient, at the ten percent level, in a horse race in which neither the debt stock nor the yield loads beside it, reappears in two fiscal shifters read from newspaper coverage and rating actions taken for fiscal reasons, and emerges from a quasi-experiment among countries that share monetary policy, a currency and an inflation target. Across countries every component of the yield predicts the level of the correlation; within the countries with sovereign credit default swaps, the association loads on the compensation for sovereign default risk, consistent with the model's monetary-dominance regime. The same ratio tracks the aggregate movements of the last three decades, the fiscal consolidations and falling yields of the 1990s, when the hedge appeared, and the higher debt and rising yields after 2020, when it weakened. What sets the burden at which a sovereign's bonds stop hedging, and why the United States and Japan keep the property at debt levels that place other countries in the comovement regime, our estimates do not settle.

References

Acharya, Viral V. and Toomas Laarits, "When Do Treasuries Earn the Convenience Yield? A Hedging Perspective," *Journal of Finance*, 2025. Forthcoming.

Anderson, Nicola and John Sleath, "New Estimates of the Real and Nominal Yield Curves," Working Paper 126, Bank of England 2001.

Antolin-Diaz, Juan, "How Did Government Bonds Become Safe?," Working Paper 2025.

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

Baele, Lieven, Geert Bekaert, and Koen Inghelbrecht, "The Determinants of Stock and Bond Return Comovements," *Review of Financial Studies*, 2010, *23* (6), 2374–2428.

_ , _ , _ , and Min Wei, "Flights to Safety," *Review of Financial Studies*, 2020, *33* (2), 689–746.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis, "Measuring Economic Policy Uncertainty," *Quarterly Journal of Economics*, 2016, *131* (4), 1593–1636.

Bianchi, Francesco and Leonardo Melosi, "The Dire Effects of the Lack of Monetary and Fiscal Coordination," *Journal of Monetary Economics*, 2019, *104*, 1–22.

Blanchard, Olivier, "Public Debt and Low Interest Rates," *American Economic Review*, 2019, *109* (4), 1197–1229.

Bocola, Luigi, "The Pass-Through of Sovereign Risk," *Journal of Political Economy*, 2016, *124* (4), 879–926.

Bohn, Henning, "The Behavior of U.S. Public Debt and Deficits," *Quarterly Journal of Economics*, 1998, *113* (3), 949–963.

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

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

_ , _ , and Pierre-Olivier Gourinchas, "The Safe Assets Shortage Conundrum," *Journal of Economic Perspectives*, 2017, *31* (3), 29–46.

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

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Campbell, John Y., Adi Sunderam, and Luis M. Viceira, "Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds," *Critical Finance Review*, 2017, *6* (2), 263–301.

_ , Andrew W. Lo, and A. Craig MacKinlay, *The Econometrics of Financial Markets*, Princeton University Press, 1997.

_ , Carolin Pflueger, and Luis M. Viceira, "Macroeconomic Drivers of Bond and Equity Risks," *Journal of Political Economy*, 2020, *128* (8), 3148–3185.

Cieslak, Anna, Wenhao Li, and Carolin Pflueger, "Inflation and Treasury Convenience," Working Paper 2024.

Cochrane, John H., *The Fiscal Theory of the Price Level*, Princeton University Press, 2023.

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

_ , Daniel Neuhann, and Guillermo Ordoñez, "Information Spillovers and Sovereign Debt: Theory Meets the Eurozone Crisis," *Review of Economic Studies*, 2025, *92* (1), 197–237.

Connolly, Robert, Chris Stivers, and Licheng Sun, "Stock Market Uncertainty and the Stock-Bond Return Relation," *Journal of Financial and Quantitative Analysis*, 2005, *40* (1), 161–194.

David, Alexander and Pietro Veronesi, "What Ties Return Volatilities to Price Valuations and Fundamentals?," *Journal of Political Economy*, 2013, *121* (4), 682–746.

Diebold, Francis X. and Canlin Li, "Forecasting the Term Structure of Government Bond Yields," *Journal of Econometrics*, 2006, *130* (2), 337–364.

Du, Wenxin and Jesse Schreger, "Local Currency Sovereign Risk," *Journal of Finance*, 2016, *71* (3), 1027–1070.

_ , Joanne Im, and Jesse Schreger, "The U.S. Treasury Premium," *Journal of International Economics*, 2018, *112*, 167–181.

Duffee, Gregory R., "Macroeconomic News and Stock–Bond Comovement," *Review of Finance*, 2023, *27* (5), 1859–1882.

Dufour, Alfonso, Andrei Stancu, and Simone Varotto, "The Equity-Like Behaviour of Sovereign Bonds," *Journal of International Financial Markets, Institutions and Money*, 2017, *48*, 25–46.

Furman, Jason and Lawrence H. Summers, "A Reconsideration of Fiscal Policy in the Era of Low Interest Rates," Discussion Draft, Brookings Institution 2020.

Ghosh, Atish R., Jun I. Kim, Enrique G. Mendoza, Jonathan D. Ostry, and Mahvash S. Qureshi, "Fiscal Fatigue, Fiscal Space and Debt Sustainability in Advanced Economies," *Economic Journal*, 2013, *123* (566), F4–F30.

Gilchrist, Simon and Egon Zakrajšek, "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 2012, *102* (4), 1692–1720.

Gómez-Cram, Roberto, Howard Kung, and Hanno Lustig, "Government Debt in Mature Economies. Safe or Risky?," Working Paper 2024.

_ , _ , and _ , "Can Treasury Markets Add and Subtract?," NBER Working Paper 33604, National Bureau of Economic Research 2025.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, "The U.S. Treasury Yield Curve: 1961 to the Present," *Journal of Monetary Economics*, 2007, *54* (8), 2291–2304.

Hall, George J. and Thomas J. Sargent, "Interest Rate Risk and Other Determinants of Post-WWII U.S. Government Debt/GDP Dynamics," *American Economic Journal: Macroeconomics*, 2011, *3* (3), 192–214.

Hazell, Jonathon and Stephan Hobler, "Do Deficits Cause Inflation? A High Frequency Narrative Approach," Working Paper 2025.

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

Ilmanen, Antti, "Stock-Bond Correlations," *Journal of Fixed Income*, 2003, *13* (2), 55–66.

Jiang, Zhengyang, Hanno Lustig, Stijn Van Nieuwerburgh, and Mindy Z. Xiaolan, "Bond Convenience Yields in the Eurozone Currency Union," NBER Working Paper 34307, National Bureau of Economic Research 2025.

_ , _ , _ , and _ , "Manufacturing Risk-Free Government Debt," *Journal of Financial Economics*, 2026, *176.*

Kremens, Lukas and Ian Martin, "The Quanto Theory of Exchange Rates," *American Economic Review*, 2019, *109* (3), 810–843.

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

_ , Stefan Nagel, and Annette Vissing-Jorgensen, "ECB Policies Involving Government Bond Purchases: Impact and Channels," *Review of Finance*, 2018, *22* (1), 1–44.

Lane, Philip R., "The European Sovereign Debt Crisis," *Journal of Economic Perspectives*, 2012, *26* (3), 49–68.

Leeper, Eric M., "Equilibria Under 'Active' and 'Passive' Monetary and Fiscal Policies," *Journal of Monetary Economics*, 1991, *27* (1), 129–147.

Leombroni, Matteo, Carolin Pflueger, and Adi Sunderam, "Risk and Return in Government Bonds," Working Paper 2026.

Li, Erica X.N., Tao Zha, Ji Zhang, and Hao Zhou, "Does Fiscal Policy Matter for Stock-Bond Return Correlation?," *Journal of Monetary Economics*, 2022, *128*, 20–34.

Longstaff, Francis A., Jun Pan, Lasse Heje Pedersen, and Kenneth J. Singleton, "How Sovereign Is Sovereign Credit Risk?," *American Economic Journal: Macroeconomics*, 2011, *3* (2), 75–103.

Loughran, Tim and Bill McDonald, "When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks," *Journal of Finance*, 2011, *66* (1), 35–65.

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

Mehrotra, Neil R. and Dmitriy Sergeyev, "Debt Sustainability in a Low Interest Rate World," *Journal of Monetary Economics*, 2021, *124*, S1–S18.

Molenaar, Roderick, Edouard Sénéchal, Laurens Swinkels, and Zhenping Wang, "Empirical Evidence on the Stock-Bond Correlation," *Financial Analysts Journal*, 2024, *80* (3), 17–36.

Nelson, Charles R. and Andrew F. Siegel, "Parsimonious Modeling of Yield Curves," *Journal of Business*, 1987, *60* (4), 473–489.

Pflueger, Carolin, "Back to the 1980s or Not? The Drivers of Inflation and Real Risks in Treasury Bonds," 2025.

Sims, Christopher A., "A Simple Model for Study of the Determination of the Price Level and the Interaction of Monetary and Fiscal Policy," *Economic Theory*, 1994, *4* (3), 381–399.

Song, Dongho, "Bond Market Exposures to Macroeconomic and Monetary Policy Risks," *Review of Financial Studies*, 2017, *30* (8), 2761–2817.

Woodford, Michael, "Monetary Policy and Price Level Determinacy in a Cash-in-Advance Economy," *Economic Theory*, 1994, *4* (3), 345–380.

A Proofs and derivations

A.1 Stochastic discount factor

The representative household has Epstein–Zin preferences with unit elasticity of intertemporal substitution,

$$V_t \;=\; C_t^{1-\beta}\left(\mathbb{E}_t V_{t+1}^{1-\gamma}\right)^{\beta/(1-\gamma)}, \qquad \beta\in(0,1),\;\gamma>0, \tag{28}$$

and faces the budget $C_t + q_t\,B_{t+1} = W_t + B_t\,R_t$, where B_{t+1} is the position in a real asset with gross return R_{t+1} and price q_t. In equilibrium $C_t = Y_t$ with $\log(Y_{t+1}/Y_t) = g + \rho\sigma\varepsilon_t + \sigma\varepsilon_{t+1}$ and $\varepsilon_{t+1}\sim\mathcal{N}(0,1)$ i.i.d.

Differentiating (28) holding V_{t+1} fixed,

$$\lambda_t \;\equiv\; \partial V_t/\partial C_t \;=\; (1-\beta)\,V_t/C_t. \tag{29}$$

Substituting the budget into the household problem, the first-order condition for B_{t+1} is $q_t\,\lambda_t = \partial V_t/\partial B_{t+1}$, and the envelope theorem applied one period ahead gives $\partial V_{t+1}/\partial B_{t+1} = R_{t+1}\,\lambda_{t+1}$.

Differentiating (28) with respect to B_{t+1} through the continuation value yields

$$\begin{aligned}\frac{\partial V_t}{\partial B_{t+1}} \;&=\; \frac{\beta}{1-\gamma}\,C_t^{1-\beta}\left(\mathbb{E}_t V_{t+1}^{1-\gamma}\right)^{\beta/(1-\gamma)-1}\,\mathbb{E}_t\!\left[(1-\gamma)\,V_{t+1}^{-\gamma}\,\frac{\partial V_{t+1}}{\partial B_{t+1}}\right]\\ &=\; \beta\,V_t\left(\mathbb{E}_t V_{t+1}^{1-\gamma}\right)^{-1}\mathbb{E}_t\!\left[V_{t+1}^{-\gamma}\,R_{t+1}\,\lambda_{t+1}\right].\end{aligned}$$

Replacing λ_{t+1} via (29) gives $q_t = \mathbb{E}_t[M_{t+1}R_{t+1}]$ with

$$M_{t+1} \;=\; \beta\,\frac{C_t}{C_{t+1}}\left(\frac{V_{t+1}}{(\mathbb{E}_t V_{t+1}^{1-\gamma})^{1/(1-\gamma)}}\right)^{1-\gamma}. \tag{30}$$

Conjecture $\log V_t = \log Y_t + a + b\,\varepsilon_t$. Then,

$$\log V_{t+1} = \log Y_t + g + \rho\sigma\varepsilon_t + a + (\sigma+b)\,\varepsilon_{t+1}, \tag{31}$$

and the log-normal moment-generating function gives

$$\frac{1}{1-\gamma}\log \mathbb{E}_t V_{t+1}^{1-\gamma} \;=\; \log Y_t + g + \rho\sigma\varepsilon_t + a + \tfrac{1}{2}(1-\gamma)(\sigma+b)^2. \tag{32}$$

Taking logs of (28) and matching coefficients on ε_t and on the constant,

$$b \;=\; \beta\rho\sigma, \qquad a \;=\; \frac{\beta}{1-\beta}\left(g \;+\; \frac{1}{2}(1-\gamma)\,\sigma^2(1+\beta\rho)^2\right). \tag{33}$$

Thus,

$$\log\!\left(V_{t+1}/\left(\mathbb{E}_t V_{t+1}^{1-\gamma}\right)^{1/(1-\gamma)}\right) \;=\; \sigma(1+\beta\rho)\,\varepsilon_{t+1} \;-\; \tfrac{1}{2}(1-\gamma)\,\sigma^2(1+\beta\rho)^2. \tag{34}$$

Given

$$\log(C_{t+1}/C_t) = g + \rho\sigma\varepsilon_t + \sigma\varepsilon_{t+1}, \tag{35}$$

we get

$$\begin{aligned}\log M_{t+1} \;&=\log\beta \;-\; (g+\rho\sigma\varepsilon_t) \;-\; \hat{\sigma}\,\varepsilon_{t+1}\\ &\quad-\; \tfrac{1}{2}(\gamma-1)^2\,(1+\beta\rho)^2\sigma^2,\end{aligned} \tag{36}$$

where $\hat{\sigma} \equiv \sigma[(\gamma-1)(1+\beta\rho)+1]$ is the loading of $\log M_{t+1}$ on the innovation ε_{t+1}.

A.2 Real risk-free rate

From (36), $\log M_{t+1}$ is conditionally Gaussian, so

$$\log \mathbb{E}_t M_{t+1} \;=\; \log\beta \;-\; (g+\rho\sigma\varepsilon_t) \;+\; \left(\gamma-\tfrac{1}{2}\right)\sigma^2 \;-\; (1-\gamma)\beta\rho\sigma^2. \tag{37}$$

From $1 = (1+r_t)\,\mathbb{E}_t M_{t+1}$,

$$\log(1+r_t) \;=\; -\log\beta \;+\; (g+\rho\sigma\varepsilon_t) \;-\left(\gamma-\tfrac{1}{2}\right)\sigma^2 \;+\; (1-\gamma)\beta\rho\sigma^2. \tag{38}$$

A.3 Equity price

The equity claim has real payoff Y_{t+1}^{λ}, real price $\mathbb{E}_t[M_{t+1}\, Y_{t+1}^{\lambda}]$, and nominal price $p_t = P_t\, \mathbb{E}_t[M_{t+1}\, Y_{t+1}^{\lambda}]$. With $\log(M_{t+1} Y_{t+1}^{\lambda}) = \log M_{t+1} + \lambda \log Y_{t+1}$, substituting (36) and $\log Y_{t+1} = \log Y_t + g + \rho\sigma\varepsilon_t + \sigma\varepsilon_{t+1}$,

$$\begin{aligned} \log\!\Big(M_{t+1} Y_{t+1}^{\lambda}\Big) \;&=\; \log\beta + \lambda \log Y_t + (\lambda-1)(g+\rho\sigma\varepsilon_t) \;-\; \tfrac{1}{2}(\gamma-1)^2(1+\beta\rho)^2\sigma^2 \\ &\quad+\; \Big[\lambda-\gamma-(\gamma-1)\beta\rho\Big]\sigma\varepsilon_{t+1}. \end{aligned} \tag{39}$$

Hence,

$$\log p_t \;=\; \log P_t \;+\; \log\beta + \lambda\log Y_t + (\lambda-1)(g+\rho\sigma\varepsilon_t) \tag{40}$$

$$-\; \tfrac{1}{2}(\gamma-1)^2(1+\beta\rho)^2\sigma^2 \;+\; \tfrac{1}{2}\Big[\lambda-\gamma-(\gamma-1)\beta\rho\Big]^2\sigma^2. \tag{41}$$

Under the peg $\log P_t$ is deterministic, so the slope is

$$\frac{\partial \log p_t}{\partial \varepsilon_t} \;=\; \sigma\,\Big[\lambda(1+\rho)-\rho\Big] > 0 \tag{42}$$

since $\lambda \geq 1$ and $\rho \geq 0$.

A.4 Asset value $\bar{A}_t$

The real fiscal asset value is $A_t = \mathbb{E}_t \sum_{j\geq 0} M_{t,t+j}\, \tau Y_{t+j}$, which factors as $A_t = \tau Y_t\, \bar{A}_t$ with

$$\bar{A}_t \;=\; \mathbb{E}_t\sum_{j\geq 0} M_{t,t+j}\, Y_{t+j}/Y_t \;=\; 1 \;+\; \mathbb{E}_t\sum_{j\geq 1} M_{t,t+j}\, Y_{t+j}/Y_t. \tag{43}$$

The Bellman recursion is

$$A_t = \tau Y_t + \mathbb{E}_t[M_{t+1}A_{t+1}], \tag{44}$$

Dividing by τY_t gives

$$\bar{A}_t \;=\; 1 \;+\; \mathbb{E}_t\Big[M_{t+1}(Y_{t+1}/Y_t)\,\bar{A}_{t+1}\Big]. \tag{45}$$

Adding (36) and $\log(Y_{t+1}/Y_t) = g + \rho\sigma\varepsilon_t + \sigma\varepsilon_{t+1}$,

$$\begin{aligned}\log\Big[M_{t+1}(Y_{t+1}/Y_t)\Big] \;&= \log\beta \;-\; (\gamma-1)(1+\beta\rho)\sigma\varepsilon_{t+1} \\ &\qquad -\; \tfrac{1}{2}(\gamma-1)^2(1+\beta\rho)^2\sigma^2. \end{aligned} \tag{46}$$

Taking the conditional expectation,

$$\mathbb{E}_t\Big[M_{t+1}(Y_{t+1}/Y_t)\Big] \;=\; \beta. \tag{47}$$

Conjecture $\bar{A}_t = \bar{A}$, a constant. Verify by substituting into the Bellman:

$$\bar{A} \;=\; 1 \,+\, \mathbb{E}_t\Big[M_{t+1}(Y_{t+1}/Y_t)\,\bar{A}\Big] \;=\; 1 \,+\, \bar{A}\,\mathbb{E}_t\Big[M_{t+1}(Y_{t+1}/Y_t)\Big] \;=\; 1 \,+\, \beta\,\bar{A}, \tag{48}$$

which solves for

$$\bar{A} \;=\; \frac{1}{1-\beta}. \tag{49}$$

A.5 Debt-to-output sensitivity

Given $\delta_t = 0$, the real flow budget gives $B_t = (1+i_{t-1})B_{t-1}/(1+\bar{\pi}) - \tau Y_t$, with i_{t-1}, B_{t-1} predetermined. Differentiating,

$$\frac{\partial \log(B_t/Y_t)}{\partial \varepsilon_t} \;=\; -\sigma\left(1 + \frac{\tau Y_t}{B_t}\right). \tag{50}$$

A.6 Bond Euler equation with convenience

The first-order condition for B_t equates the marginal cost λ_t to the marginal benefit:

$$\lambda_t \;=\; (1+i_t)\,\mathbb{E}_t\Big[\,\lambda_t\,M_{t+1}(1-\delta_{t+1})/(1+\pi_{t+1})\Big] \;+\; \partial\zeta_t/\partial B_t. \tag{51}$$

Differentiating (3),

$$\frac{\partial\zeta_t}{\partial B_t} \;=\; \lambda_t\,Y_t\,\xi'(B_t/Y_t)/Y_t \;=\; \lambda_t\,\Big[\alpha - \eta\,B_t/Y_t\Big], \tag{52}$$

The scaling λ_t cancels out:

$$q_t = \frac{1}{1+i_t} = \frac{1}{1-\alpha+\eta\, B_t/Y_t}\,\mathbb{E}_t\!\left[\frac{M_{t+1}(1-\delta_{t+1})}{1+\pi_{t+1}}\right], \tag{53}$$

with multiplier $1/(1-\alpha+\eta(B_t/Y_t)) > 1$ raising the bond price above the no-convenience benchmark.

Lemma 2. *For $Z \sim \mathcal{N}(0,1)$ and constants a, z,*

$$\frac{\mathbb{E}\Big[\exp(aZ)\mathbf{1}\{Z<z\}\Big]}{\mathbb{E}\Big[\exp(aZ)\Big]} = \Phi(z-a). \tag{54}$$

Proof. $\int_{-\infty}^{z} \exp(au)\phi(u)\,du = \exp(a^2/2)\int_{-\infty}^{z}\phi(u-a)\,du = \exp(a^2/2)\Phi(z-a)$ via the change of variable $v = u-a$ and completing the square $au - u^2/2 = -(u-a)^2/2 + a^2/2$. Dividing by $\int \exp(au)\phi(u)\,du = \exp(a^2/2)$ gives the result. □

A.7 Bond price under monetary dominance

The bond Euler (53) with $\delta_{t+1} = \mathbf{1}\{\varepsilon_{t+1} < z_t\}$ and the peg $\pi_{t+1} = \bar{\pi}$ deterministic gives, by Lemma 2,

$$\frac{\mathbb{E}_t[M_{t+1}\,\mathbf{1}\{\varepsilon_{t+1}<z_t\}]}{\mathbb{E}_t M_{t+1}} = \Phi(z_t+\hat{\sigma}), \tag{55}$$

so

$$\mathbb{E}_t[M_{t+1}(1-\delta_{t+1})] = \frac{1-\Phi(z_t+\hat{\sigma})}{1+r_t}. \tag{56}$$

Thus,

$$q_t = \frac{1}{1+i_t} = \frac{1}{1-\alpha+\eta\, B_t/Y_t}\,\frac{1-\Phi(z_t+\hat{\sigma})}{(1+r_t)}. \tag{57}$$

A.8 Default threshold and equilibrium fixed point

The default condition $A_{t+1} < L_{t+1}$ at $t+1$, with $A_{t+1} = \tau Y_{t+1}/(1-\beta)$ from (49) and $L_{t+1} = (1+i_t)B_t/(1+\bar{\pi})$, reduces to

$$\frac{\tau}{1-\beta} Y_t \exp\left(g + \rho\sigma\varepsilon_t + \sigma\varepsilon_{t+1}\right) \; < \; \frac{1+i_t}{1+\bar{\pi}} B_t, \tag{58}$$

which on taking logs is $\varepsilon_{t+1} < z_t$ with

$$z_t \; = \; \frac{1}{\sigma}\left[\log\left(\frac{1+i_t}{1+\bar{\pi}}\frac{B_t\,(1-\beta)}{\tau\, Y_t}\right) \; - \; g \; - \; \rho\sigma\varepsilon_t\right]. \tag{59}$$

Substituting $1+i_t$ from (18) into (59) yields the implicit equation for z_t:

$$\begin{aligned}
&\sigma z_t \; + \; \log\bigl(1 - \Phi(z_t + \hat{\sigma})\bigr) \\
&\quad = \log(1+r_t) \; + \; \log(B_t/Y_t) \; + \; \log(1-\beta) \; - \; \log\tau \; - \; g \; - \; \rho\sigma\varepsilon_t \\
&\qquad + \; \log\bigl(1 - \alpha + \eta\, B_t/Y_t)\bigr).
\end{aligned} \tag{60}$$

Differentiate (60) with respect to ε_t:

$$(\Theta_t - 1)\,\frac{\partial z_t}{\partial \varepsilon_t} = \left(1 + \frac{\tau Y_t}{B_t}\right)\left(1 + \frac{\eta}{(1-\alpha)\, Y_t/B_t + \eta}\right). \tag{61}$$

where

$$\Theta_t = \frac{1}{\sigma}\frac{\phi(z_t + \hat{\sigma})}{1 - \Phi(z_t + \hat{\sigma})} \tag{62}$$

is the hazard rate.

A.9 Bond log-slope under monetary dominance

From (57), we can derive

$$\frac{\partial \log q_t}{\partial \varepsilon_t} \; = \; -\,\sigma\Theta_t\frac{\partial z_t}{\partial \varepsilon_t} \; - \; \rho\sigma \; + \; \sigma\left(1 + \frac{\tau Y_t}{B_t}\right)\frac{\eta}{(1-\alpha)\, Y_t/B_t + \eta}. \tag{63}$$

Substituting for $\partial z_t/\partial \varepsilon_t$ yields

$$\frac{\partial \log q_t}{\partial \varepsilon_t} \;=\; \left(1+\frac{\tau Y_t}{B_t}\right)\frac{\sigma}{1-\Theta_t}\left(\Theta_t \;+\; \frac{\eta}{(1-\alpha)\,Y_t/B_t+\eta}\right) \;-\; \rho\sigma. \tag{64}$$

A.10 Bond log-slope under fiscal dominance

Under fiscal dominance the bond pays one nominal at $t+1$ in every state, while the price level rises on the trigger event δ_{t+1} so that $L_{t+1} = A_{t+1} = \tau Y_{t+1}/(1-\beta)$. The period-$(t+1)$ deflator is therefore

$$\frac{1}{1+\pi_{t+1}} \;=\; \frac{1-\delta_{t+1}}{1+\bar{\pi}} \;+\; \frac{\tau Y_{t+1}\,\delta_{t+1}}{(1-\beta)(1+i_t)\,B_t}. \tag{65}$$

Substituting (65) into (53) with $\delta_{t+1} = \mathbf{1}\{\varepsilon_{t+1} < z_t\}$, applying Lemma 2 to the first integral and $\mathbb{E}_t[M_{t+1}(Y_{t+1}/Y_t)] = \beta$ from (47) to the second,

$$q_t\,(1-\xi_t') \;=\; \frac{1-\Phi(z_t+\hat{\sigma})}{(1+r_t)(1+\bar{\pi})} \;+\; \frac{\tau\beta\,Y_t\,\Phi(z_t+\hat{\sigma}-\sigma)}{(1-\beta)(1+i_t)\,B_t}. \tag{66}$$

Identifying $q_t^{\mathrm{MD}}(1-\xi_t')$ from (57) in the first term, using $1/(1+i_t) = q_t$, and dividing the second by $q_t(1-\xi_t')$,

$$q_t \;=\; \frac{q_t^{\mathrm{MD}}}{1-\Psi_t}, \qquad \Psi_t \;\equiv\; \frac{q_t - q_t^{\mathrm{MD}}}{q_t} \;=\; \frac{\tau\beta\,\Phi(z_t+\hat{\sigma}-\sigma)}{(1-\beta)\,(B_t/Y_t)\,(1-\xi_t')}. \tag{67}$$

Substituting $\log q_t = \log q_t^{\mathrm{MD}} - \log(1-\Psi_t)$ into (59), (60) becomes

$$\sigma z_t \;+\; \log\!\big(1-\Phi(z_t+\hat{\sigma})\big) \;=\; \big[\text{RHS of (60)}\big] \;+\; \log(1-\Psi_t). \tag{68}$$

Differentiating (68) in ε_t with

$$\frac{\partial \log \Psi_t}{\partial \varepsilon_t} \;=\; \frac{\phi(z_t+\hat{\sigma}-\sigma)}{\Phi(z_t+\hat{\sigma}-\sigma)}\,\frac{\partial z_t}{\partial \varepsilon_t} \;+\; \sigma\,\theta_t\,\frac{1+\alpha-2\xi_t'}{1-\xi_t'}, \qquad \theta_t \equiv 1+\tau Y_t/B_t,$$

yields

$$\frac{\partial z_t}{\partial \varepsilon_t} \;=\; \frac{1}{1+\Psi_t\,\Lambda_t}\,\frac{\partial z_t^{\mathrm{MD}}}{\partial \varepsilon_t}, \qquad \Lambda_t \;\equiv\; \frac{1}{1-\Theta_t}\,\frac{\phi(z_t+\hat{\sigma}-\sigma)}{\sigma\,\Phi(z_t+\hat{\sigma}-\sigma)} \;-\; 1, \tag{69}$$

with $\partial z_t^{\mathrm{MD}}/\partial\varepsilon_t$ as given in Appendix A.8. Taking logs of $q_t = q_t^{\mathrm{MD}}/(1-\Psi_t)$, differentiating, and substituting (69),

$$\frac{\partial \log q_t}{\partial \varepsilon_t} \;=\; -\rho\sigma \;+\; \frac{\sigma\theta_t}{1-\Theta_t}\left(\Theta_t \;+\; \frac{\alpha-\xi_t'}{1-\xi_t'}\right) \;-\; \frac{\sigma\theta_t}{1-\Theta_t}\,\frac{\Psi_t\,\Lambda_t}{1+\Psi_t\,\Lambda_t}\left(1+\frac{\alpha-\xi_t'}{1-\xi_t'}\right). \tag{70}$$

Under $\Theta_t < 1$, the inequality $\phi(x)/\Phi(x) > -x$ for $x<0$ applied to $z_t+\hat{\sigma}-\sigma$ and $-(z_t+\hat{\sigma})$ gives $\Lambda_t > 0$. Multiplying Ψ_t from (67) by Λ_t from (69), the $\Phi(z_t+\hat{\sigma}-\sigma)$ factor of Ψ_t cancels with the $1/\Phi(z_t+\hat{\sigma}-\sigma)$ inside Λ_t, leaving

$$\Psi_t\,\Lambda_t \;=\; \frac{\tau\beta\left[\phi(z_t+\hat{\sigma}-\sigma)/[\sigma(1-\Theta_t)] \;-\; \Phi(z_t+\hat{\sigma}-\sigma)\right]}{(1-\beta)\,(B_t/Y_t)\,(1-\xi_t')}. \tag{71}$$

A.11 Proof of Proposition 1

Lemma 3. *For continuously differentiable h with $h\,\phi$ vanishing at $\pm\infty$,*

$$\mathrm{Cov}_t(\varepsilon_{t+1},\,h(\varepsilon_{t+1})) \;=\; \int h(\varepsilon)\,\varepsilon\,\phi(\varepsilon)\,d\varepsilon \;=\; -\!\int h(\varepsilon)\,\phi'(\varepsilon)\,d\varepsilon \;=\; \mathbb{E}_t[h'(\varepsilon_{t+1})],$$

using $\phi'(\varepsilon) = -\varepsilon\phi(\varepsilon)$ and integration by parts.

From (41), $\log p_{t+1}$ is affine in ε_{t+1} with slope $\sigma[\lambda(1+\rho)-\rho]$, so $\sigma_t^p = \sigma[\lambda(1+\rho)-\rho]$ and $\rho_t^{sb}\,\sigma_t^q = \mathrm{Cov}_t(\varepsilon_{t+1},\log q_{t+1})$.

The bond price at $t+1$ is (57) evaluated one period forward, and its log-slope is (64) evaluated one period forward:

$$\frac{\partial \log q_{t+1}}{\partial \varepsilon_{t+1}} \;=\; -\rho\sigma \;+\; \frac{\sigma\,\theta_{t+1}}{1-\Theta_{t+1}}\left(\Theta_{t+1} \;+\; \frac{\alpha-\xi_{t+1}'}{1-\xi_{t+1}'}\right).$$

Applying Lemma 3 with $h = \log q_{t+1}$,

$$\rho_t^{sb}\,\sigma_t^q/\sigma \;=\; \mathbb{E}_t\!\left[\frac{\theta_{t+1}\,\Theta_{t+1}}{1-\Theta_{t+1}}\right] \;+\; \mathbb{E}_t\!\left[\frac{\theta_{t+1}}{1-\Theta_{t+1}}\,\frac{\alpha-\xi_{t+1}'}{1-\xi_{t+1}'}\right] \;-\; \rho. \tag{72}$$

A.12 Proof of Proposition 2

As in Appendix A.11, $\rho_t^{sb}\,\sigma_t^q = \mathrm{Cov}_t(\varepsilon_{t+1},\log q_{t+1})$. The bond price at $t+1$ is now resolved under fiscal dominance: (66) evaluated one period forward, with log-slope (70) evaluated one period forward,

$$\frac{\partial \log q_{t+1}}{\partial \varepsilon_{t+1}} = -\rho\sigma + \frac{\sigma\,\theta_{t+1}}{1-\Theta_{t+1}}\left(\Theta_{t+1} + \frac{\alpha-\xi'_{t+1}}{1-\xi'_{t+1}}\right) - \frac{\sigma\,\theta_{t+1}}{1-\Theta_{t+1}}\,\frac{\Psi_{t+1}\,\Lambda_{t+1}}{1+\Psi_{t+1}\,\Lambda_{t+1}}\left(1 + \frac{\alpha-\xi'_{t+1}}{1-\xi'_{t+1}}\right).$$

Applying Lemma 3 with $h = \log q_{t+1}$,

$$\rho_t^{sb}\,\sigma_t^q/\sigma = \mathbb{E}_t\left[\frac{\theta_{t+1}\,\Theta_{t+1}}{1-\Theta_{t+1}}\right] + \mathbb{E}_t\left[\frac{\theta_{t+1}}{1-\Theta_{t+1}}\,\frac{\alpha-\xi'_{t+1}}{1-\xi'_{t+1}}\right] - \rho - \mathbb{E}_t\left[\frac{\theta_{t+1}}{1-\Theta_{t+1}}\,\frac{\Psi_{t+1}\,\Lambda_{t+1}}{1+\Psi_{t+1}\,\Lambda_{t+1}}\left(1+\frac{\alpha-\xi'_{t+1}}{1-\xi'_{t+1}}\right)\right]. \tag{73}$$

The first three terms match (72); the final term is the fiscal-capacity correction.

A.13 Proof of Proposition 3

In the limit of the proposition, the four Gaussian quantities $\Phi(z_{t+1}+\hat{\sigma})$, $\phi(z_{t+1}+\hat{\sigma})$, $\Phi(z_{t+1}+\hat{\sigma}-\sigma)$, and $\phi(z_{t+1}+\hat{\sigma}-\sigma)$ are zero: the first by the proposition's assumption, the other three because Φ is increasing and ϕ peaks at the origin. The hazard Θ_{t+1} and the fiscal-capacity object $\Psi_{t+1}\Lambda_{t+1}$ in (71) are linear in these four quantities over bounded denominators, hence zero. The credit and fiscal-capacity correction channels of (73) drop out, leaving the convenience and safe-haven channels:

$$\rho_t^{sb}\,\sigma_t^q/\sigma = \mathbb{E}_t\left[\frac{\theta_{t+1}\,(\alpha-\xi'_{t+1})}{1-\xi'_{t+1}}\right] - \rho.$$

A.14 Proof of Proposition 4

Evaluate forward objects at $\varepsilon_t = 0$, where the output process of Appendix A.1 reduces to $\log(Y_{t+1}/Y_t) = g + \sigma\,\varepsilon_{t+1}$. Define

$$\ell_t \equiv \frac{(1+i_t)\,B_t}{(1+\bar{\pi})\,Y_t}, \qquad u_{t+1} \equiv \ell_t\,\exp(-g-\sigma\,\varepsilon_{t+1}),$$

so $\partial u_{t+1}/\partial \ell_t = u_{t+1}/\ell_t > 0$. The flow budget of Appendix A.5 at $t+1$ with $\delta_{t+1} = 0$ gives $B_{t+1}/Y_{t+1} = u_{t+1} - \tau$, so

$$\theta_{t+1} \;=\; \frac{u_{t+1}}{u_{t+1}-\tau}, \qquad 1-\xi'_{t+1} \;=\; 1-\alpha-\eta\tau+\eta\, u_{t+1}, \qquad \alpha-\xi'_{t+1} \;=\; \eta\,(u_{t+1}-\tau), \tag{74}$$

and

$$\theta_{t+1}\,\frac{\alpha-\xi'_{t+1}}{1-\xi'_{t+1}} \;=\; \frac{u_{t+1}}{u_{t+1}-\tau}\,\frac{\eta\,(u_{t+1}-\tau)}{1-\alpha-\eta\tau+\eta\, u_{t+1}} \;=\; \frac{\eta\, u_{t+1}}{1-\alpha-\eta\tau+\eta\, u_{t+1}}. \tag{75}$$

Evaluating (60) one period forward, substituting $\log(1+r_{t+1})$ from (38) and $B_{t+1}/Y_{t+1} = u_{t+1} - \tau$, the $g + \rho\sigma\,\varepsilon_{t+1}$ entering through $\log(1 + r_{t+1})$ cancels the $-g - \rho\sigma\,\varepsilon_{t+1}$ on the right-hand side and leaves

$$\sigma\, z_{t+1} \;+\; \log\!\Big(1-\Phi(z_{t+1}+\hat\sigma)\Big) \;=\; \log(u_{t+1}-\tau) \;+\; \log\!\Big(1-\alpha+\eta\,(u_{t+1}-\tau)\Big) \;+\; \text{const}, \tag{76}$$

so z_{t+1} depends on ε_{t+1} only through u_{t+1}. Differentiating (76) in u_{t+1} and using (74),

$$\frac{\partial z_{t+1}}{\partial u_{t+1}} \;=\; \frac{1}{\sigma\,(1-\Theta_{t+1})\,(u_{t+1}-\tau)}\left[1 \;+\; \frac{\alpha-\xi'_{t+1}}{1-\xi'_{t+1}}\right] \;>\; 0 \tag{77}$$

under $\Theta_{t+1} < 1$.

Differentiating $\Theta_{t+1} = (1/\sigma)\phi(z_{t+1} + \hat\sigma)/[1 - \Phi(z_{t+1} + \hat\sigma)]$ in z_{t+1},

$$\frac{\partial \Theta_{t+1}}{\partial z_{t+1}} \;=\; \Theta_{t+1}[\sigma\,\Theta_{t+1} \;-\; (z_{t+1}+\hat\sigma)] \;>\; 0,$$

where positivity uses $\phi(x)/[1 - \Phi(x)] > x$ for all $x \in \mathbb{R}$. Chaining with (77) and $\partial u_{t+1}/\partial \ell_t = u_{t+1}/\ell_t > 0$, both Θ_{t+1} and $\Theta_{t+1}/(1 - \Theta_{t+1})$ are strictly increasing in ℓ_t.

The bond log-slope (70) evaluated one period forward is

$$\begin{aligned}\frac{\partial \log q_{t+1}}{\partial \varepsilon_{t+1}} \;=\; -\,\rho\sigma \;+\; \frac{\sigma\,\theta_{t+1}}{1-\Theta_{t+1}}\left(\Theta_{t+1} \;+\; \frac{\alpha-\xi'_{t+1}}{1-\xi'_{t+1}}\right)\\ -\;\frac{\sigma\,\theta_{t+1}}{1-\Theta_{t+1}}\,\frac{\Psi_{t+1}\,\Lambda_{t+1}}{1+\Psi_{t+1}\,\Lambda_{t+1}}\left(1 \;+\; \frac{\alpha-\xi'_{t+1}}{1-\xi'_{t+1}}\right),\end{aligned} \tag{78}$$

with the $\Psi_{t+1}\Lambda_{t+1}$ term absent under monetary dominance by (64). From (42), $\log p_{t+1}$

is affine in ε_{t+1} with slope $\sigma[\lambda(1+\rho)-\rho]$, so $\rho_t^{sb}\,\sigma_t^q = \mathrm{Cov}_t(\varepsilon_{t+1}, \log q_{t+1})$, and Lemma 3 with $h = \log q_{t+1}$ gives

$$G(\ell_t) \;\equiv\; \rho_t^{sb}\,\sigma_t^q/\sigma \;=\; \frac{1}{\sigma}\,\mathbb{E}_t\!\left[\frac{\partial \log q_{t+1}}{\partial \varepsilon_{t+1}}\right], \qquad S(\ell_t) \;\equiv\; \sigma_t^q \;=\; \sqrt{\mathrm{Var}_t(\log q_{t+1})} \;>\; 0. \tag{79}$$

Each term of (78) is a function of ℓ_t and ε_{t+1} alone via u_{t+1}, z_{t+1}, Θ_{t+1}, ξ'_{t+1} (and $\Psi_{t+1}, \Lambda_{t+1}$ under fiscal dominance), so after integrating ε_{t+1} out, G and S are smooth functions of ℓ_t.

From (75), the convenience term in (78) satisfies

$$\frac{\theta_{t+1}}{1-\Theta_{t+1}}\,\frac{\alpha-\xi'_{t+1}}{1-\xi'_{t+1}} \;=\; \frac{\eta\, u_{t+1}}{(1-\alpha-\eta\tau+\eta\, u_{t+1})\,(1-\Theta_{t+1})},$$

a product of two strictly positive factors, both strictly increasing in ℓ_t under $\alpha+\eta\tau<1$, so its expectation is strictly increasing in ℓ_t. The credit term factors as $\theta_{t+1}\,\Theta_{t+1}/(1-\Theta_{t+1}) = [u_{t+1}/(u_{t+1}-\tau)]\,[\Theta_{t+1}/(1-\Theta_{t+1})]$, combining a factor decreasing in u_{t+1} with one strictly increasing in ℓ_t. Differentiating, its derivative in ℓ_t is positive when

$$\frac{\partial \log \Theta_{t+1}}{\partial \log \ell_t} \;>\; (1-\Theta_{t+1})(\theta_{t+1}-1), \tag{80}$$

i.e., the default-channel elasticity dominates the fiscal-capacity erosion captured by θ_{t+1} approaching 1 at high debt. Chaining (77) with $\partial u_{t+1}/\partial \ell_t = u_{t+1}/\ell_t$ and $\partial \log \Theta_{t+1}/\partial z_{t+1} = \sigma\,\Theta_{t+1} - (z_{t+1}+\hat{\sigma})$,

$$\frac{\partial \log \Theta_{t+1}}{\partial \log \ell_t} \;=\; \frac{u_{t+1}[\sigma\,\Theta_{t+1}-(z_{t+1}+\hat{\sigma})]\big[1+(\alpha-\xi'_{t+1})/(1-\xi'_{t+1})\big]}{\sigma\,(1-\Theta_{t+1})\,(u_{t+1}-\tau)}.$$

Using $1+(\alpha-\xi'_{t+1})/(1-\xi'_{t+1}) \geq 1$, $u_{t+1}>\tau$, and $1-\Theta_{t+1}<1$, the inequality (80) is implied by

$$\Phi(z_{t+1}+\hat{\sigma}) \;<\; \Phi(-\sigma), \tag{81}$$

the assumption of Proposition 4. Under (81), the credit term of (78) is strictly increasing in ℓ_t. Combined with the strict monotonicity of the convenience term, both terms of the monetary-dominance slope are strictly increasing in ℓ_t, and $G'(\bar{\ell})>0$.

The regressor of (23) is $(1+i_t)\,B_t/Y_t = (1+\bar{\pi})\,\ell_t$. Fix a reference $\bar{\ell}$. By a first-order

Taylor expansion of $G(\ell_t)$ in $\ell_t - \bar{\ell}$,

$$\rho_t^{sb}\,\sigma_t^q/\sigma \;=\; G(\ell_t) \;\approx\; G(\bar{\ell}) \;+\; G'(\bar{\ell})\,(\ell_t - \bar{\ell}).$$

Substituting $\ell_t - \bar{\ell} = [(1+i_t)B_t/Y_t - (1+\bar{\pi})\bar{\ell}]/(1+\bar{\pi})$ and identifying coefficients with (23),

$$\beta_1 \;=\; \frac{G'(\bar{\ell})}{1+\bar{\pi}}, \qquad \beta_0 \;=\; G(\bar{\ell}) \;-\; (1+\bar{\pi})\,\bar{\ell}\,\beta_1. \tag{82}$$

Since $G'(\bar{\ell}) > 0$ from the preceding, $\beta_1 > 0$.

A.15 Proof of Proposition 5

The Taylor identification (82) of Appendix A.14 gives $\beta_1 = G'(\bar{\ell})/(1+\bar{\pi})$ at the reference debt level $\bar{\ell}$. The flow budget of Appendix A.5 at $t+1$ with $\delta_{t+1} = 0$ gives $B_{t+1}/Y_{t+1} = u_{t+1} - \tau$, so the no-default support requires $u_{t+1} > \tau$; at $\varepsilon_{t+1} = 0$ this reads $\ell_t > \tau e^g$, and we take $\bar{\ell}$ in this range.

Under fiscal dominance, the bond log-slope (70) is the monetary-dominance log-slope (78) minus

$$F_{t+1} \;\equiv\; \frac{\theta_{t+1}}{1-\Theta_{t+1}}\,\frac{\Psi_{t+1}\,\Lambda_{t+1}}{1+\Psi_{t+1}\,\Lambda_{t+1}}\left[1 + \frac{\alpha - \xi'_{t+1}}{1-\xi'_{t+1}}\right], \tag{83}$$

so $G(\ell_t) = G^{MD}(\ell_t) - \mathbb{E}_t[F_{t+1}]$ with $G^{MD}(\ell_t)$ as in Appendix A.14. The $\Psi\Lambda$ closed form (71) of Appendix A.10 evaluated one period forward and using $B_{t+1}/Y_{t+1} = u_{t+1} - \tau$ from Appendix A.5 gives the form used below.

The fiscal-capacity correction at the lower edge of the admissible range, $\mathbb{E}_t[\partial F_{t+1}/\partial \ell_t]|_{\ell_t = \tau e^g}$, is proportional to $\tau\beta/(1-\beta)$ via the prefactor of (71), while $G^{MD\,\prime}(\tau e^g)$ does not depend on β. The hypothesis of the proposition therefore admits an explicit threshold $\bar{K}(\alpha, \eta, \tau, \sigma)$ such that $\tau\beta/(1-\beta) < \bar{K}$ delivers

$$\mathbb{E}_t\!\left[\frac{\partial F_{t+1}}{\partial \ell_t}\right]\Bigg|_{\ell_t = \tau e^g} \;<\; G^{MD\,\prime}(\tau e^g)\,. \tag{84}$$

The decomposition gives $G'(\tau e^g) = G^{MD\,\prime}(\tau e^g) - \mathbb{E}_t[\partial F_{t+1}/\partial \ell_t]\Big|_{\ell_t = \tau e^g} > 0$, and the smoothness of G in ℓ_t via (83) and (71) extends the inequality to an interval $(\tau e^g, \bar{\ell}^*)$ where

$$\bar{\ell}^* \equiv \sup\{\, \bar{\ell} > \tau e^g \,:\, G'(\ell_t) > 0 \text{ for all } \ell_t \in (\tau e^g, \bar{\ell})\,\}. \tag{85}$$

Writing $z_{t+1}(\ell_t, \varepsilon_{t+1})$ for the trigger map of (76), set

$$c \equiv \Phi(z_{t+1}(\bar{\ell}^*, 0) + \hat{\sigma}). \tag{86}$$

At $\varepsilon_{t+1} = 0$, $u_{t+1} = \ell_t e^{-g}$ gives $\partial u_{t+1}/\partial \ell_t = e^{-g}$, and chaining with $\partial z_{t+1}/\partial u_{t+1} > 0$ from (77) delivers $\partial z_{t+1}(\ell_t, 0)/\partial \ell_t > 0$, so the condition $\Phi(z_{t+1} + \hat{\sigma}) < c$ is equivalent to $\bar{\ell} < \bar{\ell}^*$, on which $G'(\bar{\ell}) > 0$ by construction of $\bar{\ell}^*$. (82) identifies $\beta_1 = G'(\bar{\ell})/(1 + \bar{\pi}) > 0$.

B Data description and construction

This appendix describes how we build each variable, from the raw vendor extract to the regression input. Section 3 gives the compact account; the replication package rebuilds every panel from the sources and records each rule where it is applied.

Three conventions apply throughout. Yields are zero-coupon rates in percent per year. The Federal Reserve and Bank of England curves are continuously compounded by their documentation, and the BIS curves and our own fits test the same against the vendor's par quotes; the Bloomberg curves enter as delivered, and against the same par quotes most of them sit above a continuously compounded curve by the gap between annual and continuous compounding at the prevailing yield, about a tenth of a point at 5 percent. Converting them lowers the fiscal regressor in those countries by 4 percent on average and moves neither baseline by more than about a tenth of a standard error, so we leave the vendor's number. We assign each country a single source for a given daily series and keep it, so a bad date becomes missing and is not filled from a second provider, because a series stitched across two vendors' curve-fitting conventions carries joins that look like price movements. We apply every screen identically to every country.

The screens, masks, exclusions and corrections on the two return legs are declared and individually switchable, 28 rules in all, and Appendix C.5 re-estimates the paper's main tables on a full rebuild of the data with all of them off but the two estimation floors. Section B.10 states what the rules remove in total and what the rebuild shows.

B.1 Country sample and coverage

Our panel is 44 countries. Following the International Monetary Fund's World Economic Outlook classification, these are 29 advanced economies (AT, AU, BE, CA, CH, CZ, DE, DK, ES, FI, FR, GB, GR, HK, IE, IL, IT, JP, KR, LT, NL, NO, NZ, PT, SE, SG, SI, TW, US) and 15 emerging markets (BR, CL, CN, CO, HU, ID, IN, MX, MY, PE, PH, PL, TH, TR, ZA). We hold each country's current classification fixed over the sample, so that a country that changed groups partway through does not enter a panel with country fixed effects as two different countries; the estimation panel is all 44 countries.

The monthly dependent variable is the non-overlapping monthly correlation of daily equity and daily ten-year zero-coupon bond returns, and it carries the baseline regressions. The annual dependent variable is the calendar-year correlation of monthly returns; it reaches decades further back and carries the horse race between interest expense and its components, because the debt ratio varies year to year. Figure B.1 shows the coverage of the two side by side. A few countries have interior gaps, years inside their coverage that carry no cell. There are 35 such years, and in every one of them it is the bond leg that fails, never the equity leg; all 35 fall below our ten-month minimum, made by sparse deep-history quotation or by our own screens.

B.2 Government bond returns

The bond leg is the return on a synthetic 10-year zero-coupon government bond. We build two return series. The daily return, which feeds the monthly panel, is differenced from a daily zero-coupon yield panel in which each country has a single source. The monthly return, which feeds the annual correlation and must reach back to 1960, is spliced from up to three sources per country, and this is the only place Global Financial Data enters the bond leg. Table B.1 reports each country's daily source, its coverage, and how much of its monthly history rests on the deep-history source.

Daily panel. Each country's daily zero-coupon yield comes from one source for its entire daily history. The United States is the Federal Reserve Board's Gürkaynak-Sack-Wright curve (Gürkaynak et al., 2007), daily from 1971, and the United Kingdom the Bank of England's Government Liability Curve (Anderson and Sleath, 2001), daily from 1986. Every other country takes the published zero-coupon curve with the longest active

Figure B.1: Coverage of the two dependent variables, ordered by annual-panel entry

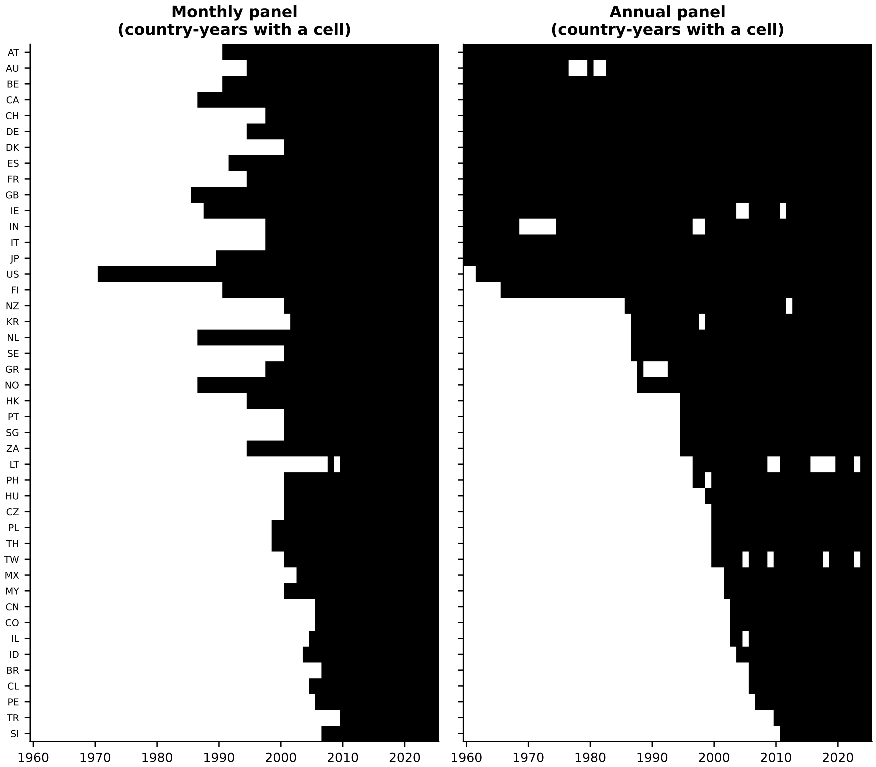

Note: A shaded cell is a country-year carrying at least one stock-bond correlation at that frequency; a blank cell inside a country's span is an interior gap.

history, among the Bank for International Settlements estimated yields and the Bloomberg fair-value curve; this gives BIS for 4 countries and Bloomberg for 23. Where no published curve is usable we fit the Datastream benchmark par curve to zero-coupon ourselves, which is 14 countries. Lithuania is the one remaining case: its par grid ends at the 10-year point, so we convert the vendor's 10-year redemption yield directly. No ten-year zero-coupon yield is observed anywhere, because governments issue coupon bonds, so every source above is itself an estimate. We take a published curve where one exists, so that the estimation is the source's and is documented by the source, and we record, for every observation, which of the two supplied it.

Zero-coupon conversion. For the fitted countries we convert benchmark par yields at the observed maturities to a zero-coupon curve by nonlinear least squares on the Nelson-Siegel (Nelson and Siegel, 1987) form with the fixed decay of Diebold and Li (2006), restated for maturities in years ($\lambda = 0.7308$ per year, twelve times their per-month 0.0609), choosing the factors so that the par yields implied by the fitted curve match the observed ones. We impose par consistency because a fit that treats par yields as zero-coupon yields introduces a level bias correlated with the yield level, and therefore with the fiscal regressor. We evaluate the curve at the 10-year point, and at the 5-year point where sovereign credit default swaps are the comparison instrument.

A Nelson-Siegel curve has three free parameters, so a date whose grid quotes fewer than three usable maturities does not identify one. On those dates we fit nothing and carry the vendor's own par redemption yield, converted from annual to continuous compounding, at the maturities the vendor priced, tenors shorter than a coupon period aside, and nothing at any maturity the vendor did not price. No level is carried from one maturity to another. At the maturities the paper reads, 6.0 percent of the daily ten-year panel and 0.7 percent of the daily five-year panel are a converted vendor quote of this kind, and we ship the census of every such cell. On an upward-sloping curve a par yield sits below the zero-coupon yield, so the substitution lowers the yield and the fiscal regressor built on it, which works against the paper.

Where we discard the fit. Sovereign crisis curves can distort enough that the fit stops tracking the point it is anchored on. We test every fitted date against that country's own 10-year par quote, converted to continuous compounding before the comparison, and where the fitted 10-year departs from it by more than 3 percentage points we take the converted quote instead. The test locates dated windows and does not act on single dates. We declare two windows: Greece from 26 April 2011 to 14 April 2014, the sovereign-debt crisis and the restructuring. And Turkey from 2018, where the test is the wrong instrument: a three-factor curve with a fixed decay cannot represent a par grid as steeply inverted as Turkey's, so we take the vendor's quote on every Turkish session from 2018. On the dates the test selects the curve is inverted, so the substitution runs high.

A second dated window corrects the date a value carries. For 1987 through 2005 the vendor's par ten-year column for Norway reports the previous session's news, and the column self-corrects in 2006 with no duplicated or skipped print; we re-date it one session earlier across the window. Greece's five- and twenty-year columns carry the same defect from January 2004 to March 2008, and there we drop the two columns from the fit grid. No yield is altered in either case.

Daily return. With $y_{i,t}$ the 10-year zero-coupon yield in percent per year and $n = 10$, we compute the daily bond return as the duration approximation

$$r^b_{i,t} \;=\; -\,n\,\Delta y_{i,t}\,/\,100 \qquad \text{(decimal per day)},$$

which omits roughly 0.02% per day of carry as negligible against daily yield volatility. We align the yield to the country's own equity trading calendar and take the change on that calendar, so the bond return pairs with a same-day equity return.

Monthly returns and the splice. We form the annual correlation from monthly returns because the daily zero-coupon panel does not begin until 2000 for the median country. For each country we splice the monthly series from up to three non-overlapping segments in priority order: the assigned daily source resampled to month-end; the Datastream benchmark par curve fitted to zero-coupon; and Global Financial Data redemption yields for the earliest history, with one override, that a layer quoting the ten-year itself outranks a layer above it that is extrapolating one. Within each segment we compute the monthly return by the log-linear approximation of Campbell et al. (1997),

$$r^b_{i,t} \;=\; y_{i,t-1}/12 \;-\; n\,\Delta y_{i,t} \qquad \text{(percent per month)},$$

which retains carry, about 0.4% per month at a 5% yield and material at this frequency. Here n is the maturity of the series everywhere except in the deep-history segments whose instrument has a moving duration of its own, where it is that instrument's duration (below). We splice on returns, not on yield levels: we compute the return inside each segment independently and drop the first observation of each segment, so a level mismatch at a boundary cannot manufacture a return. The monthly yield level panel is spliced across the same segments, and it is the level that enters the fiscal regressor.

The deep-history segment, and its quality. Global Financial Data enters only the third segment, only in the monthly series, and never in the daily panel. It supplies redemption yields, which we put through the same par-consistent fit; on just over half the months we use, it quotes the ten-year alone, and the value we carry is that quote converted from annual to continuous compounding. It reaches 18 of the 44 countries in the estimation sample, supplies 95.2% of the ten-year months the panel carries from the 1960s and 89.6% of those from the 1970s, and its last month anywhere is March 2003 (Figure B.2).

The deep-history source supplies 4,797 country-months inside the estimation sample, spread over 18 countries; those months fall in 405 of the 1,788 country-years the annual panel carries (22.7%), and every figure below is weighted by that usage. Where the ten-year is quoted the two sources agree at 0.997 in levels; where we extrapolate it, which happens only for Korea, they agree at 0.980. Measured on the annual correlation itself the two sources agree at 0.939 on the quoted segment and 0.825 on the extrapolated one.

Figure B.2: Sources supplying the ten-year zero-coupon yield across the estimation sample

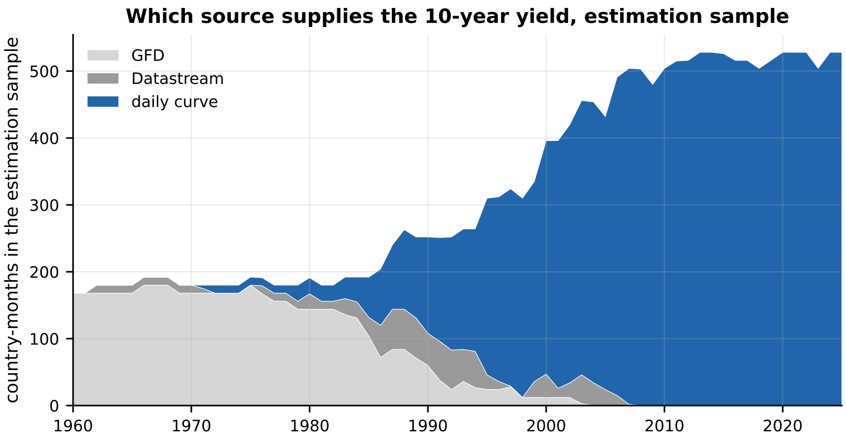

Note: Sources are counted in country-months, by year. Only country-months falling in a country-year that enters the annual regression panel are counted, so the figure and the robustness columns that drop deep-history observations describe the same sample.

In sixteen countries the vendor's instrument notes say the series booked as a ten-year is something else for part of its life, usually another obligation of the same government at a different maturity and, in a few countries for part of their early history, an index that includes other issuers; we keep every era, and the replication package catalogues each in the vendor's words. Seven eras have an instrument whose duration moves inside the year the correlation is computed over. In these seven we replace the grid maturity with the instrument's own duration, measured at the start of each holding month, on 1,508 shipped monthly ten-year returns.

Screens and dated exclusions. We apply four screens, all uniformly across countries, on the daily yields and the returns formed from them. We set to missing any daily bond return implied by a one-day yield move above 5 percentage points, which in a ten-year zero-coupon yield is a restructuring, a redenomination or a data artifact; any return spanning a quote gap longer than 100 days, because differencing across silence books the whole stretch to one session; any run of identical daily yields longer than 15 sessions, since a vendor repeats a benchmark that has stopped updating; and any return that spans a change in how the yield was constructed. Five rules of the same kind screen the monthly returns where they are formed, and further rules act on the fit's inputs and on the month-end levels. The exemption list carries 35 entries, each keeping a real repricing one of the monthly rules would otherwise remove. A few dated exclusions act where a uniform rule cannot, among them long tenors extrapolated before the bond existed, series that start at a currency or instrument switch, and six dated windows recorded in a shipped reference table. The replication package documents every screen, rule and window with its reason, and the monthly rules and the dated windows with their counts.

B.3 Equity returns

The equity leg is the return on a broad national equity index, and as with bonds we build two series; Table B.2 reports each country's daily index, its coverage, and the evidence on which we admit its monthly splice.

Daily panel. The United States is the CRSP value-weighted market return, daily from 1926. Every other country takes the earliest index that varies day to day, between the local Datastream market index and the MSCI country index: local Datastream indices begin between 1986 and 2001 and vary daily from the start; MSCI varies daily only from January 2001, its earlier record being interpolated month-end data. This rule selects local Datastream for 28 countries and MSCI for 15. We infer a daily start from the data, not from vendor metadata, because selecting on declared coverage would treat interpolated month-end values as daily observations. We compute the daily equity return as the log change in the index, taken on the raw price grid before placing the series on the country's trading-day calendar, so a return never spans a market holiday; the CRSP series is already an arithmetic value-weighted return, so we use it directly. We set to missing one Greek return, that of 3 August 2015, the first session after five weeks of closure under capital controls, which is the move across the closure.

Table B.1: 10-year zero-coupon government bond data: coverage and source by country.

Country	Daily	Daily source	Monthly from	Months	GFD share
AT	1991-01–2026-03	Datastream-DL	1960	795	38%
AU	1994-12–2026-02	Bloomberg	1960	762	25%
BE	1991-01–2026-03	Datastream-DL	1960	795	45%
BR	2007-03–2026-02	Bloomberg	2006	242	0%
CA	1987-01–2026-02	BIS	1960	794	39%
CH	1993-05–2026-02	BIS	1960	794	42%
CL	2005-09–2026-02	Bloomberg	2005	246	0%
CN	2006-01–2026-02	Bloomberg	2002	285	0%
CO	2006-04–2026-02	Bloomberg	2002	282	0%
CZ	2001-01–2026-02	Bloomberg	1999	324	0%
DE	1994-12–2026-02	Bloomberg	1960	794	26%
DK	2001-01–2026-03	Datastream-DL	1960	795	35%
ES	1992-02–2026-03	Datastream-DL	1960	795	47%
FI	1991-09–2026-03	Datastream-DL	1960	794	48%
FR	1994-12–2026-02	Bloomberg	1960	794	39%
GB	1986-01–2026-02	BoE GLC	1960	794	29%
GR	1998-05–2026-03	Datastream-DL	1960	636	46%
HK	1994-12–2026-02	Bloomberg	1994	377	0%
HU	2001-03–2026-02	Bloomberg	1999	326	0%
ID	2004-01–2026-02	Bloomberg	2003	274	0%
IE	1988-01–2026-03	Datastream-DL	1960	783	38%
IL	2005-03–2026-02	Bloomberg	2002	282	0%
IN	1998-12–2026-02	Bloomberg	1960	781	53%
IT	1998-01–2026-02	BIS	1960	794	47%
JP	1990-01–2026-02	Bloomberg	1960	794	39%
KR	2002-11–2026-02	Bloomberg	1987	464	28%
LT	2008-06–2026-03	Datastream-direct	1997	302	25%
MX	2003-08–2026-02	Bloomberg	2001	296	0%
MY	2001-09–2026-02	Bloomberg	2001	294	0%
NL	1987-01–2026-03	Datastream-DL	1987	471	0%
NO	1987-01–2026-03	Datastream-DL	1987	467	0%
NZ	2001-01–2026-03	Datastream-DL	1985	493	0%
PE	2006-05–2026-02	Bloomberg	2006	238	0%
PH	2001-01–2026-02	Bloomberg	1996	350	0%
PL	1999-05–2026-02	Bloomberg	1999	322	0%
PT	2001-01–2026-03	Datastream-DL	1994	383	0%
SE	2001-01–2026-02	BIS	1986	478	0%
SG	2001-01–2026-02	Bloomberg	1994	375	0%
SI	2007-04–2026-03	Datastream-DL	2007	226	0%
TH	1999-09–2026-02	Bloomberg	1999	318	0%
TR	2010-01–2026-03	Datastream-DL	2010	195	0%
TW	2001-01–2026-03	Datastream-DL	1999	324	0%
US	1971-08–2025-12	GSW	1962	768	0%
ZA	1995-07–2026-02	Bloomberg	1995	374	0%

Note: Yields are 10-year zero-coupon rates in percent per year, continuously compounded. Each country is assigned a *single* daily source, chosen by longest active history; a bad date is set to missing, not filled from a second provider, because a series stitched across two curve-fitting conventions carries joins that look like price moves. *Monthly from* and *Months* describe the analysis sample, which begins in 1960; the monthly panels hold a short buffer before that date so the first year's rolling windows have inputs, and it is excluded here. *GFD share* is the fraction of those months supplied by Global Financial Data, which reaches 18 countries and is never a primary source and never daily. It supplies redemption yields, put through the same par-consistent fit as the Datastream curve: where the vendor quotes the ten-year alone the fit reduces to that quote converted from annual to continuous compounding, and where it quotes several maturities a curve is fitted across them.

Trading calendars. We difference returns on each country's own trading calendar. We build that calendar from an exchange calendar, frozen once from a calendar library and hashed with the other inputs, and from the variation of the index and of the country's own bond quote. A day the calendar records as open is a session whatever the two series did, since a flat session on a trading day is an observation about co-movement; a day it records as closed is kept only if both the index and the bond quote moved on it; and where no calendar reaches, which is Japan before 1997, a day is kept if either moved. The rule removes 13,814 delivered days across 43 countries, each listed with its date, price and reason in a shipped record, and a second and separate rule removes 1,670 more, Israel's Fridays before the Tel Aviv exchange moved to a Monday-to-Friday week in January 2026. Because the rule conditions on the realized value of one of the two series whose correlation we compute, we price it on a rebuild. Switching it off adds 55 monthly ten-year cells and no country-year and no country, changes the correlation in 6,266 others, three quarters of them by less than 0.01, and moves the coefficient this appendix supports by less than a tenth of a standard error.

Monthly returns and the splice. For the monthly series we take the MSCI country index, or the CRSP return for the United States, resample it to month-end, and splice it back before its start with Global Financial Data, splicing on returns and dropping the first observation of each segment. We admit a splice only on at least 24 months of joint quotation and a return correlation of at least 0.85 over that overlap. 36 of the 44 countries carry an extension; the median admitted overlap correlation is 0.978. We set to missing the 16 monthly returns that are exactly zero, since an exact zero in a log return means a byte-identical index level, which is a stale quote or a closed market, and Colombia's January 1988 return, computed from an index level of exactly 100, the vendor's base value, against a following month quoted on a different scale, so the implied fall of 90.6 percent is a change of units.

B.4 Stock-bond correlation

Our dependent variable is the realized stock-bond correlation, the Pearson correlation between the equity return of Section B.3 and the bond return of Section B.2. Both returns are already defined, so the correlation adds only the window and the minimum number of

Table B.2: Equity return data: coverage, source and splice evidence by country.

Country	Daily	Daily source	Monthly source	History from	Overlap	Corr.
AT	1991-01–2026-03	Datastream local	MSCI	1960	675	0.984
AU	1992-06–2026-03	Datastream local	MSCI	1960	675	0.979
BE	1991-01–2026-03	Datastream local	MSCI	1960	675	0.940
BR	1991-12–2026-03	Datastream local	MSCI	1960	459	0.974
CA	1987-01–2026-03	Datastream local	MSCI	1960	675	0.984
CH	1993-05–2026-03	Datastream local	MSCI	1960	675	0.995
CL	2001-01–2026-03	MSCI	MSCI	1975	459	0.974
CN	2001-01–2026-03	MSCI	MSCI	1993	–	–
CO	2001-01–2026-03	MSCI	MSCI	1988	398	0.974
CZ	2001-01–2026-03	MSCI	MSCI	1995	–	–
DE	1986-01–2026-03	Datastream local	MSCI	1960	675	0.993
DK	2001-01–2026-03	MSCI	MSCI	1960	675	0.983
ES	1992-02–2026-03	Datastream local	MSCI	1960	675	0.967
FI	1991-01–2026-03	Datastream local	MSCI	1960	459	0.977
FR	1988-01–2026-03	Datastream local	MSCI	1960	675	0.996
GB	1986-01–2026-03	Datastream local	MSCI	1960	675	0.992
GR	1998-05–2026-03	Datastream local	MSCI	1988	–	–
HK	1990-01–2026-03	Datastream local	MSCI	1960	675	0.978
HU	1997-10–2026-03	Datastream local	MSCI	1991	375	0.974
ID	1996-07–2026-03	Datastream local	MSCI	1988	–	–
IE	1988-01–2026-03	Datastream local	MSCI	1960	459	0.949
IL	1994-01–2026-03	Datastream local	MSCI	1993	–	–
IN	1998-12–2026-03	Datastream local	MSCI	1960	399	0.976
IT	1998-01–2026-03	Datastream local	MSCI	1960	675	0.987
JP	1986-01–2026-03	Datastream local	MSCI	1960	675	0.986
KR	1994-01–2026-03	Datastream local	MSCI	1962	458	1.000
LT	2008-06–2026-03	MSCI	MSCI	1996	214	0.875
MX	2001-01–2026-03	MSCI	MSCI	1983	459	0.989
MY	1994-01–2026-03	Datastream local	MSCI	1960	459	0.896
NL	1986-01–2026-03	Datastream local	MSCI	1960	675	0.970
NO	1986-01–2026-03	Datastream local	MSCI	1960	675	0.860
NZ	2001-01–2026-03	Datastream local	MSCI	1960	459	0.937
PE	2001-01–2026-03	MSCI	MSCI	1993	–	–
PH	2001-01–2026-03	MSCI	MSCI	1982	458	0.987
PL	1993-01–2026-03	Datastream local	MSCI	1991	399	0.979
PT	2001-01–2026-03	MSCI	MSCI	1960	459	0.950
SE	2001-01–2026-03	MSCI	MSCI	1960	675	0.983
SG	2001-01–2026-03	MSCI	MSCI	1960	675	0.973
SI	2002-06–2026-03	MSCI	MSCI	2002	–	–
TH	1994-11–2026-03	Datastream local	MSCI	1975	456	0.983
TR	2001-01–2026-03	MSCI	MSCI	1986	459	0.991
TW	2001-01–2026-03	MSCI	MSCI	1988	–	–
US	1926-01–2025-12	CRSP	CRSP	1960	1200	0.990
ZA	1995-07–2026-03	Datastream local	MSCI	1960	399	0.977

Note: Daily returns are log differences of index levels on each country's own trading calendar. Iceland and Estonia ship in no panel: Iceland's index breaks across the 2008 collapse and Estonia has no zero-coupon curve to pair with. *Overlap* and *Corr.* describe the backward extension of the monthly series with Global Financial Data, which is admitted only on at least 24 months of joint quotation and a return correlation of at least 0.85, and is spliced on returns and not on price levels. 36 of 44 countries are extended; the median admitted correlation is 0.978 and the lowest is 0.860. Global Financial Data is never a primary source and never enters the daily panel.

return pairs that window must contain.

1. **Non-overlapping monthly, from daily returns.** Within each calendar month, we take the correlation of that month's daily returns on the country's own trading calendar, and use it only if the month carries at least 15 trading days on which both legs are observed. Days the source did not publish are left empty: carrying the last quote forward would imply a bond return of exactly zero and would let a weekly feed reach the daily panel looking daily, and in our daily quality record 0 of 14,219 country-year-series ship more sessions than the source published. This is our dependent variable in the baseline regressions of Section 4, and in the channel-decomposition and fiscal-salience regressions, whose right-hand-side variables are monthly.

2. **Annual, from monthly returns.** Within each calendar year, we take the correlation of that year's monthly equity and bond returns and date it at year-end. We use a year only if at least 10 of its 12 months carry both returns, and we admit a country only if at least 10 years of its history carry a computable cell. This is our dependent variable in the annual regressions, which carry the horse race between interest expense and its components, because the debt ratio moves once a year. We build it from monthly returns because the daily zero-coupon panel does not begin until 2000 for the median country, while the monthly series reach 1960. The difference is 1,788 country-years against 1,225, and it falls in the early, high-inflation decades we need in order to identify anything from within-country variation in the interest burden. The daily-return twin, the calendar-year correlation of daily returns, is built alongside it and ships in the replication package.

3. **Further windows.** For the event study of Section 5 we compute an event-time quarterly correlation of daily returns, binned in three-month windows measured from the 20 October 2009 Greek deficit revision, so that the revelation falls on a bin boundary and no bin straddles it, requiring at least 40 trading days per bin. The package also builds a rolling twelve-month correlation of monthly returns, requiring a full twelve months of paired returns, and calendar quarterly and annual correlations of daily returns, requiring at least 30 and 100 trading days.

We use plain Pearson correlations throughout and apply a Fisher transformation only in the robustness tables. We pair returns on each country's own trading calendar, so a cell never pairs an equity return against a bond return across a market holiday. We use the 10-year bond leg throughout, with a 5-year twin where the regressor is the 5-year CDS spread; because a correlation is invariant to scale, maturity enters only through the shape of yield-curve changes, so the two are nearly identical, and both come from one curve, so their agreement is not two independent measurements of the same market.

B.5 Fiscal and macroeconomic variables

We build the fiscal regressor from annual sources placed on the monthly grid, and we enter the macroeconomic controls as monthly series directly.

The fiscal regressor. Our headline regressor is the forward-looking interest burden, the debt ratio priced at the current long yield,

$$\mathrm{ie}^{\mathrm{fw}}_{i,t} = \frac{D_{i,t}}{Y_{i,t}} \cdot \frac{y_{10,i,t}}{100},$$

with the debt ratio and the 10-year yield both in percent: the interest the sovereign would owe if its whole debt stock were refinanced at today's yield, which moves when the market reprices fiscal risk, before that repricing reaches realized coupons. In the annual panel we take both legs at the end of the year, because a debt stock measured at year end times a yield averaged over the preceding twelve months corresponds to no moment at which a sovereign could have refinanced. In the monthly panel the debt ratio is the last annual figure the source had published, which keeps every monthly observation predetermined, and the yield is that month's own reading. We enter the regressor in levels, because the product $(D/Y)\, y_{10}$ carries the interaction between the quantity of debt and its price that we are studying, and a log specification removes that interaction by construction.

Cascades and their shares. We fill three variables by a cascade, because no single source covers our 44 countries from 1960. We use a layer only where the layer above it has no observation, and we apply the same cascade to every country, with no country-specific overrides. Table B.3 reports how much of the finished panel each layer supplies, because the ordering alone misleads about what the panel rests on: the debt series names the World Economic Outlook first, but the Historical Public Debt Database is effectively the whole of the pre-1990 history. Without it our annual panel would start in 1981 instead of 1961.

Table B.3: Sources supplying the macroeconomic and fiscal controls, by cascade layer.

Variable	Layer	Country-months	Share
Consumer price inflation	OECD	24,600	71.0%
	GFD	6,353	18.3%
	IMF IFS	1,853	5.3%
	GMD	1,265	3.6%
	IMF/DBnomics	595	1.7%
Real activity	OECD industrial production	19,802	57.5%
	GMD real GDP	10,278	29.8%
	IMF/DBnomics	2,205	6.4%
	OECD real GDP	2,157	6.3%
Debt/GDP	IMF WEO	17,408	58.5%
	IMF HPDD	12,204	41.0%
	Eurostat	162	0.5%

Note: Each variable is filled by a cascade that is applied identically to every country: a layer is used only where the layer above it has no observation. Shares are of the country-months the finished monthly panel carries, so they sum to 100% within a variable. The last layer of each cascade is annual, so one figure stands for its twelve months; those cells are recorded in the source grid and are exempt from the build-time repeat check.

Debt and the controls. For the debt-to-GDP ratio we take the IMF World Economic Outlook series for general government gross debt (`GGXWDG_NGDP`), extend it forward with Eurostat quarterly government debt and backward with the IMF Historical Public Debt Database. The vintage we hold measures the debt ratio through 2024 and marks that year as an estimate for 20 of the 44 countries; we keep those values, because across the vintages since 2008 a first estimate is later revised by about as much as a first published outturn, and they reach only the regressor of 2025. We enter it as published and contemporaneous (D_t/Y_t); predeterminedness comes from lagging the whole right-hand side one year, taken on the calendar and not as the preceding row of the estimation panel, which after an interior gap is not a one-year lag.

Three monthly controls follow the same design. Inflation (`cpi`) is year-over-year headline consumer-price inflation, filled by the cascade OECD → IMF International Financial Statistics → IMF/DBnomics → Global Financial Data → Global Macro Database; one Chilean consumer-price print, October 2024, is removed as a vendor error. Real activity (`rgdp`) is filled by the cascade OECD industrial production → IMF/DBnomics → OECD quarterly real GDP → Global Macro Database real GDP; despite the variable's name, this control is industrial-production growth for most of the panel. The policy rate, our monetary control, is filled by the cascade Bank for International Settlements → European Central Bank → Global Financial Data central bank discount rate → Global Financial Data interbank overnight rate; for euro-area members we splice in the ECB rate at euro accession and at nothing else, because several countries pegged to the euro before joining it and admitting one on the strength of its peg would require adjudicating all of them. We use the policy rate and not a short zero-coupon yield for this control because a short yield is cut from the same fitted curve as the ten-year yield inside the fiscal regressor, so conditioning on it removes part of the regressor's own variation. Finally, we control for the global risk environment with two published monthly series, the VIX index sampled at month-end and the excess bond premium of Gilchrist and Zakrajšek (2012), each entered without transformation.

B.6 Sovereign credit default swaps

For the credit channel we use the five-year sovereign CDS spread, the tenor at which sovereign CDS are quoted and traded. Our panel carries 12,142 country-months for 44 countries from January 2001 to July 2026, with a median spread of 49 basis points and a mean of 103. Table B.4 reports coverage and the screen record country by country.

We take these spreads from Markit, which computes the composites, and we read no terminal data for credit default swaps anywhere in the package. A terminal that republishes the composites returns a retired contract's last published quote for every date requested afterwards, so a name that stopped trading arrives as a plausible number repeated for years, and it keys entities by display symbol, so a sovereign and a corporation can resolve to the same string. Every observation is a Government reference entity, senior unsecured tier, five-year tenor, quoted in US dollars for every country except the United States, which is quoted in euros, since a dollar hedge against a dollar default is worth little.

We take the restructuring clause as the vendor gives it. The 2014 ISDA definitions extended coverage to government bail-in, so protection under them is worth more. Of the 44 countries, 24 quote the 2014 clause throughout and 18 switch to it at the vendor's own switch date; we splice at that point because the change is a real price change on the contract the market trades, and the screen record marks which countries carry the discontinuity. Australia and New Zealand come from a separate extract on the modified restructuring clause, the one clause that spans their history.

We apply two screens. A run of identical daily quotes longer than 15 sessions is dropped, with the run length scaled by the level of the spread, since a composite at a few basis points does not move for weeks because there is nothing between the quoting steps for it to move to. A date whose cross-section carries too few distinct values is dropped, because a redistribution fault that pastes one series across many columns leaves every individual column varying and plausible. We also drop zero and negative spreads as impossible; together the screens remove 0.2 percent of daily quotes. We take the last observation of each month with at least 5 quotes, and we leave gaps as gaps and interpolate across none of them, which is why the table reports first and last dates alongside the longest gap. The contract is five-year, so in specifications placing CDS on the right-hand side we use the five-year stock-bond correlation as the dependent variable.

B.7 Real yields and inflation compensation

In the inflation-expectations channel we decompose the nominal yield into a real yield and a breakeven, computing the breakeven as the residual. We take real yields from sovereign inflation-linked bonds, which reach 14 countries and so bound that decomposition. Because the breakeven is a residual, a stale real-yield quote is booked to expected inflation, so we apply the same staleness screen as for nominal yields, requiring at least 5 live quotes a month, and treat a quote departing from both neighbors by more than 5 percentage points and returning immediately as a bad print. South Africa quotes no five-year linker, so its real yield is the midpoint of its four- and six-year, and we remove one Israeli month-end print, in February 2025, that stands apart from flat neighbors while the nominal five-year did not move.

B.8 Sovereign ratings and fiscal-origin classification

Section 6 needs news about the fiscal position that is not itself a market price. Sovereign rating actions are one such source, but an action is not a fiscal event merely because a rating agency took it, so we read what each agency said about its own decision.

We assemble rating histories for Moody's, S&P and Fitch into one file of 1,415 actions covering 43 of the panel countries from April 1974 to April 2026. Moody's and Fitch actions come from a commercial rating history, extended for Fitch with the agency's own

Table B.4: Sovereign CDS: coverage and the effect of the staleness screen.

Country	Coverage	Raw quotes	Screened	Removed	Longest gap	Clause
AT	2001-05–2026-07	6,566	0	0.0%	3 d	switched 2014
AU	2003-04–2026-07	5,972	0	0.0%	59 d	single
BE	2001-01–2026-07	6,647	0	0.0%	4 d	switched 2014
BR	2001-01–2026-07	6,664	0	0.0%	3 d	single
CA	2003-09–2026-07	4,965	0	0.0%	843 d	single
CH	2007-06–2026-07	4,583	0	0.0%	557 d	switched 2014
CL	2002-02–2026-07	6,388	0	0.0%	3 d	single
CN	2001-01–2026-07	6,671	0	0.0%	3 d	single
CO	2001-03–2026-07	6,614	0	0.0%	3 d	single
CZ	2001-03–2026-07	6,524	0	0.0%	105 d	single
DE	2002-07–2026-07	6,217	0	0.0%	81 d	switched 2014
DK	2002-11–2026-07	6,150	0	0.0%	35 d	switched 2014
ES	2001-02–2026-07	6,632	0	0.0%	3 d	switched 2014
FI	2002-07–2026-07	6,263	0	0.0%	11 d	switched 2014
FR	2002-04–2026-07	6,333	0	0.0%	8 d	switched 2014
GB	2006-03–2026-07	5,263	0	0.0%	61 d	switched 2014
GR	2001-01–2026-07	6,315	0	0.0%	456 d	switched 2014
HK	2003-08–2026-07	5,753	0	0.0%	281 d	single
HU	2001-02–2026-07	6,620	0	0.0%	15 d	single
ID	2001-12–2026-07	6,132	104	1.7%	267 d	single
IE	2003-01–2026-07	6,149	0	0.0%	3 d	switched 2014
IL	2001-04–2026-07	6,496	0	0.0%	41 d	single
IN	2003-07–2025-12	3,356	307	9.1%	2,036 d	single
IT	2001-01–2026-07	6,650	0	0.0%	3 d	switched 2014
JP	2001-01–2026-07	6,654	0	0.0%	20 d	switched 2014
KR	2001-03–2026-07	6,609	0	0.0%	3 d	single
LT	2002-04–2026-07	6,278	0	0.0%	41 d	single
MX	2001-01–2026-07	6,669	0	0.0%	5 d	single
MY	2001-04–2026-07	6,592	0	0.0%	3 d	single
NL	2003-07–2026-07	5,588	0	0.0%	572 d	switched 2014
NO	2003-10–2026-07	5,931	0	0.0%	4 d	switched 2014
NZ	2003-07–2026-07	5,713	0	0.0%	273 d	single
PE	2002-02–2026-07	6,388	22	0.3%	3 d	single
PH	2001-03–2026-07	6,614	0	0.0%	3 d	single
PL	2001-01–2026-07	6,646	0	0.0%	16 d	single
PT	2002-02–2026-07	6,384	0	0.0%	3 d	switched 2014
SE	2001-05–2026-07	6,292	0	0.0%	123 d	switched 2014
SG	2003-07–2014-12	1,941	0	0.0%	987 d	single
SI	2002-02–2026-07	6,382	110	1.7%	5 d	single
TH	2001-02–2026-07	6,626	26	0.4%	7 d	single
TR	2001-01–2026-07	6,658	0	0.0%	3 d	single
TW	2006-08–2014-04	642	0	0.0%	1,264 d	single
US	2003-12–2026-07	5,854	0	0.0%	77 d	switched 2014
ZA	2001-01–2026-07	6,652	38	0.6%	16 d	single

Note: Five-year senior sovereign CDS composites taken from the originating vendor, not from a terminal redistribution. *Screened* counts daily quotes removed as a run of identical values longer than the threshold, which is a contract that stopped trading and not a spread that did not move; *Removed* expresses that as a share of raw quotes. *Longest gap* is the longest interval with no quote at all: 7 countries carry a gap of more than a year, and for those the coverage span overstates continuous availability. *Clause* records the restructuring clause: 18 of 44 countries switched convention in 2014 and are spliced across that change, the rest quote one clause throughout.

regulatory disclosure; S&P actions come from the agency's published history, extended with the commercial source where the official record stops short and for the outlook changes it does not publish. An action is an upgrade (502), a downgrade (390) or an outlook change (523); an agency's first rating of a country and a withdrawal are coverage events and not moves, and the 2 such rows in the record score nothing and enter no filter. We keep outlook changes because they are often the first move an agency makes, and we score 464 of them on the same scale at a fraction of a notch, a quarter to three quarters, a fixed score per transition: a half for a move between stable and either outlook, except a quarter for positive to stable, and three quarters for a move between the two outlooks; the 59 moves into or out of Moody's review status score zero. We take the panel's years from the record itself: a fixed window of 1986–2025 would drop 13 actions and 186 country-years, so we impose none.

We map letters to a common notch scale, sum signed notches across agencies and divide by the number of agencies covering that country in that year, since without a divisor a country all three agencies follow would look three times as shocked as one followed by a single agency. Coverage comes from the action record: an agency covers a country from its first recorded action or first rating onward, so an agency that rated a country for years without moving stays invisible until its first move and the divisor runs small over the early record. All 98 headline-filter actions the scale signs as a downgrade, 58 downgrades and 40 outlook moves the scale scores in the downgrade direction, fall in country-years all three agencies cover; 11 of the 162 signed as an upgrade do not, so the divisor moves only the upgrade arm of the reported estimate. We assign zero to a country-year in which no agency moved, because no agency moving is a fact about that year.

We have a large language model read the newswire coverage of each action, a bundle of Reuters stories from the five days around it that republish or paraphrase the agency's statement, and return for every event a primary origin and four confound flags (cyclical, banking, external, political). We read each story three times and consolidate by majority, because the model's output is not reproducible to the letter at a fixed temperature: the three readings of one story differ on at least one field for 10.8 percent of stories and, for 57.3 percent of events, on at least one of the event's stories; replacing the three readings with a single one changes an event's consolidated label on at least one field for 10.5 percent of events. Several reports can describe a single event, so in reducing to one row per event we let an agency's own release outrank a paraphrase and a paraphrase outrank a wire flash, take the maximum of the confound flags because a confound named in one report of an event is a confound, and leave a field unclassified where the deciding tier ties or every eligible report declines to state an origin. We recovered and classified a rationale for 1,044 of the 1,415 actions. Of the remaining 371, nothing was read for 161: the corpus holds no story at all for 143, and for 18 only headline-only wire flashes carrying no text, which the classifier never reads; for 210 the stories it holds settle no origin: for 142 no story eligible to vote states one, and for 68 the deciding stories tie, between two origins for 65 and among three for 3. Both survive only the widest filter. The newswire archive begins in January 1996, so no earlier action can be classified and the shifter is zero before then by construction. Three nested filters make the ladder: F_0 keeps every action, 1,415 events; F_B keeps those the classifier reads as fiscal in origin, 776; and F_C, our headline filter, additionally requires that the agency named no confound, 273. Restricting the events to the country-months of Section 6's estimation sample leaves 1,054, 655 and 215 under F_0, F_B and F_C, the counts its tables print. We split the headline filter by which kind of fiscal origin the agency stated, 144 structural and 129 governance, to show that the result does not rest on one of the two.

We ship the action signals at two frequencies. Section 6's rating tables read the monthly panel, whose within-year sums reproduce the annual panel exactly; the annual panel pairs the same signals with the dependent variable and the controls of the annual baseline, so a coefficient there sits on the same scale as Section 4's. A level panel records where each country's rating stood at the end of every year. Both action panels take their dependent variable and controls from the panels every other exhibit reads; the level panel carries no dependent variable and no control.

B.9 Wall Street Journal fiscal-salience index

The rating shifter is sparse, a few hundred actions over five decades, and a rating action can respond to market prices. The newspaper shifter is a monthly share measured from text for the country-months with attributed coverage. For each country-month we have a large language model read the Wall Street Journal (WSJ) articles that a separate attribution pass by a language model assigns to that country, as articles substantively about it, and return the share in which it finds a substantive discussion of the sustainability of the country's public finances; the classifier is asked about the topic and not about tone, so the index measures salience. The corpus holds 200,391 unique articles; an article attributed to several covered countries enters each of their cells, giving 387,066 country-article pairs across 44 countries. The attribution is the model's: 7.2 percent of those pairs assign an article to a country of which the article contains no alias from the fixed table, such as its name, adjective or demonym, its capital or a major city, its central bank, treasury or seat of government, or its sovereign bond nickname, a share highest in the euro area's small members, and the package carries the share computed on alias-confirmed pairs and on the classifier's own list of the countries whose fiscal position a flagged article discusses beside the headline share, for the measurement table of Appendix C.3. The WSJ panel runs 2000:01 through 2026:04. The share averages 0.129 across the 13,142 country-months that carry one, with a standard deviation of 0.195, so the typical country-month has almost no fiscal coverage and the variation we use falls in episodes. We enter the article count as a control in every specification, in logs, because the index would otherwise be partly a measure of attention. Against this shifter we use the non-overlapping monthly correlation of daily returns at the primary maturity, the same object Section 4 aggregates to a year.

Three placebos ship as columns of the same panel. Bad news, not fiscal news: Loughran-McDonald dictionary negativity (Loughran and McDonald, 2011), computed over the identical 387,066 country-article pairs and assigned to countries by the same rule as the index, so a horse race between them is a contest between two readings of an article and never between two ways of attributing one. The classifier keying on country names: we run one classifier, with the prompt and the model identical in both arms, twice over a probe subsample, once on unmodified text and once on text rewritten by a second language model to replace country names, currencies, officials, institutions and dates by placeholders. In both arms articles are assigned to countries by the alias table on the unmodified text, so the contrast is one of flagging alone. The rewriting is itself measured: a name, demonym or city of an attributed country survives in 24.7 percent of the anonymized articles, and some alias the rewriter was told to replace in 35.4 percent. The two arms' country-month shares correlate at 0.88 (Pearson) against the 0.85 we set in advance, a margin we report because it is narrow. Where the anonymized arm returns a placeholder in place of a country we record it under a reserved code that no reported quantity reads. General policy uncertainty: the Economic Policy Uncertainty index (Baker et al., 2016), monthly, lagged on its own dense grid before merging. In a separate panel we replace the stock-bond correlation with the correlation between local equity returns and US equity returns, on the same monthly grid and the same 15-day minimum, giving 13,028 country-months for 43 countries. The 43 are the estimation sample without the United States, whose correlation with US equity is one by construction: if fiscal coverage moved the stock-bond correlation through general attention, it should move this correlation too.

B.10 What our construction does to the data

We remove very little, and where we cannot estimate we carry the vendor's own number. We remove a quote where it cannot be a market price of the sovereign instrument: a level that is arithmetically impossible, a run of identical values, a one-day move above the jump threshold, a value contradicted by an independent measurement, or a segment the vendor's own notes show is another instrument, with a dated exemption list that keeps the real repricings a rule would otherwise catch; and we remove a return where its two ends are not consecutive prices of one construction. On a date the source published nothing, we carry nothing, and where we substitute, the substitute is the vendor's own published number at the maturity the vendor priced. Across the eleven monthly maturities the screens above mask 2,466 returns, and at the ten-year the masked count is 399 of the returns a screen could have reached; the substitutions are the larger count, 64,003 values where the vendor's own quote stands in for a curve the grid could not identify, and 27,921 sessions carrying a converted redemption quote instead of a fitted curve point, most of them Lithuania's, whose grid ends at the ten-year and is never fitted. And the 28 declared rules are individually switchable: the package refuses to ship a build with any of them disabled and reproduces the shipped panels when all of them are on.

Tables C.20 to C.22 re-estimate the monthly specification ladder, the annual horse race and the monthly baseline on later samples on a full rebuild of the data with every switchable rule off and the two estimation floors kept. The coefficient the paper is about keeps its sign and its significance in every column of every panel, and most differences from the shipped estimate are smaller than one standard error. That is why this appendix carries only the summary: the replication package documents each screen, exclusion and rule with the reason for it and, where a record ships, its count, and neither the sign nor the significance of the estimates depends on them.

B.11 Variable definitions

Table B.5 defines every column of the annual estimation panel and reports its moments. We generate the table from the panel; the builder fails if the panel carries a column the table does not define, or defines one the panel does not carry. The dependent variable is present for all 1,788 country-years; the regressor is present for 1,757 of them, and that intersection is what we call the estimation sample. The VIX and the excess bond premium are global series that begin later than the panel does, and year fixed effects absorb both in every specification that carries them, so they bind only where we omit year effects.

Table B.5: Variable definitions and summary statistics, annual panel.

Variable	Units	Definition	N	Mean	SD
$\rho^{sb}_{i,t}$	correlation	Calendar-year correlation of monthly equity log returns with monthly 10-year zero-coupon bond returns. The dependent variable.	1,788	0.13	0.39
$IE^{fw}_{i,t}$	% of GDP	Forward-looking interest expense: general government debt over GDP multiplied by the 10-year zero-coupon yield, $(D/Y)_{i,t} \times y_{i,t}$.	1,734	3.08	2.89
$IE^{fw}_{i,t-1}$	% of GDP	One-year lag of the above. This is the regressor throughout; the lag is what makes it predetermined.	1,757	3.11	2.92
$(D/Y)_{i,t}$	% of GDP	General government gross debt over GDP, contemporaneous in both legs.	1,737	57.12	36.61
$(D/Y)_{i,t-1}$	% of GDP	One-year lag.	1,765	56.97	36.44
$y_{i,t}$	% p.a.	10-year zero-coupon yield, continuously compounded, the December reading of the monthly series.	1,785	5.92	3.77
$y_{i,t-1}$	% p.a.	One-year lag.	1,779	5.98	3.79
$IE/Y_{i,t}$	% of GDP	Backward-looking interest expense: general government interest paid over GDP, as reported. Shown for comparison with the forward-looking measure, not used in the baseline.	1,662	2.85	2.14
Structural balance	% of potential GDP	General government structural balance.	1,270	-2.68	3.20
Expected inflation	% p.a.	Expected consumer price inflation for the year, taken from the previous autumn's World Economic Outlook forecast in the historical vintage database, so it is what forecasters expected at the time, not the later-revised outturn. Constant within the year.	1,290	2.93	3.04
$\pi_{i,t}$	% p.a.	Consumer price inflation, year on year, from the cascade described above, averaged over the twelve monthly readings of the year.	1,788	3.98	4.41
$\pi_{i,t-1}$	% p.a.	One-year lag.	1,774	4.04	4.44
Activity	% p.a.	Real activity growth. Industrial production growth where it exists, then the remaining layers of the activity cascade. Named for real GDP in the code; it is not literally real GDP for most cells.	1,780	3.05	6.61
Activity, lagged	% p.a.	One-year lag.	1,774	3.06	6.63
$i^p_{i,t}$	% p.a.	Central bank policy rate at year end, and the monetary control the paper reports. Filled by the cascade BIS policy rate, then the ECB rate from each country's own euro entry, then a central bank discount rate, then an interbank overnight rate where no policy rate exists.	1,746	4.88	4.33
$i^p_{i,t-1}$	% p.a.	One-year lag.	1,731	4.95	4.32
$i^s_{i,t}$	% p.a.	Retained for comparison only: the two-year zero-coupon yield at year end, superseded as the monetary control by the policy rate above because it is cut from the same fitted curve as the ten-year yield inside IE^{fw}, so conditioning on it removes part of the regressor's own variation.	1,437	4.55	3.98
$i^s_{i,t-1}$	% p.a.	One-year lag.	1,419	4.62	3.98
VIX	index	CBOE implied volatility index, month-end closes averaged over the year. Global, so it is absorbed by year fixed effects and enters only where those are absent.	1,311	19.58	5.77
EBP	% p.a.	Gilchrist-Zakrajšek excess bond premium, annual mean. Global, same caveat as the VIX.	1,592	0.05	0.45
Advanced economy	0/1	One for the 29 advanced economies, zero for the 15 emerging markets.	1,788	0.78	0.41
$\Delta\rho^{sb}_{i,t}$	correlation	First difference of the dependent variable.	1,726	0.00	0.45
$\Delta IE^{fw}_{i,t}$	% of GDP	First difference.	1,710	-0.02	1.56
$\Delta\pi_{i,t}$	% p.a.	First difference.	1,774	-0.04	2.81
Activity, differenced	% p.a.	First difference.	1,766	-0.04	8.71

Note: Definitions and moments for every column of the annual estimation panel, which carries 1,788 country-years for 44 countries over 1960–2025. Country and date identifiers are omitted. Moments are computed over the non-missing cells of each column, so N differs across rows and a regression using several of them runs on their intersection. The table is generated from the panel: the builder fails if the panel carries a column this table does not define, or defines one the panel does not carry.

C Robustness and additional results

This appendix is organized by paper section. Appendix C.1 collects the robustness for Section 4: the construction of the fiscal measure, the stability of the estimate across countries, and alternative dependent variables. Appendix C.2 re-runs the euro-area design of Section 5 under alternative treatments, windows, samples and dependent variables. Appendix C.3 collects the robustness for Section 6: the two shifters beside the interest-expense measure, the rating shifter by bin of the lagged fiscal burden, the newspaper index against its placebos, by country group and by how it is measured, and leave-one-out over all 44 countries. Appendix C.4 re-estimates the credit decomposition of Section 7 under tighter distress thresholds. Appendix C.5 re-runs the paper's main tables on a dataset rebuilt without the discretionary screens that Appendix B documents.

C.1 Robustness for Section 4

Four figures and six tables support Section 4. Figure C.1 varies the construction of the fiscal measure itself: panel (a) moves the regressor from realized, backward-looking interest payments to the forward-looking ratio, and panel (b) rebuilds $\mathrm{ie}^{\mathrm{fw}}$ with the yield at every maturity the panel carries. The estimate is indistinguishable from zero at the backward-looking end, rises monotonically toward the forward-looking measure, and across maturities peaks at the ten-year maturity of the baseline. Figure C.2 re-estimates the baseline forty-four times, dropping one country at a time; no exclusion changes the sign on either panel. Table C.1 replaces the dependent variable with its Fisher transformation and with the volatility-scaled correlation, the object on the left-hand side of the model's regression form. Table C.2 compares the two constructions of the monthly debt leg, the reported December step and a linear interpolation. Tables C.3 and C.4 report the regression country by country at both frequencies, so the countries the main text discusses can be placed in the distribution they come from; these regressions carry no time effects, so they retain the common global movement in correlations and are not comparable with the panel estimates. Table C.5 races inflation against the fiscal variable on the country-months with inflation at most 10 percent. In Figure C.2 the estimates run from 0.0122 to 0.0147 on the monthly panel and from 0.0125 to 0.0157 on the annual, against 0.0138 and 0.0145 on the full samples. We report the step construction of the debt leg in Table C.2 because a linear interpolation between year-end figures places part of the following December's value inside every earlier month of the year, which the one-month lag does not remove.

Inflation is the leading account of the stock-bond correlation in United States data (Campbell et al., 2017, 2020; Li et al., 2022): the correlation is positive when inflation shocks dominate bond risk and negative when growth shocks do, a regime the whole cross-section shares in a given month. The panel can ask a question the United States time series cannot, whether a country's own inflation explains its own correlation once that common regime is absorbed. Table C.5 asks it. On the full sample, pooled, inflation and the fiscal variable both enter positively; with country and month effects the inflation coefficient turns negative, and the sign comes from the country-months of a few emerging markets with inflation above 10 percent, outside the range the United States evidence covers. Columns (3) to (6) therefore keep the country-months with inflation at most 10 percent, 97 percent of the sample. There the inflation coefficient is zero, alone or in the race, while ie^{fw} enters at 0.015 with a t-statistic of 5 alone and beside it, close to its full-sample value, and the annual panel gives the same ranking. The two accounts are complements. Inflation describes the regime the cross-section moves through together and the level differences across countries; the fiscal variable describes which country's correlation moves when its own position changes, which is the variation a two-way fixed-effects panel isolates.

Two figures give the slower-moving evidence behind the eras paragraph of Section 4. Figure C.3 re-sorts the countries every month on their lagged interest expense and plots the average monthly correlation of the top and bottom thirds: the top third sits above the bottom third in all but one of the months plotted, by 0.20 on average, and the gap is widest in the 2010s. It is the pooled relationship of the first column of Table 2 drawn over time, a comparison across countries at each date and not the within-country estimate the paper reports. Figure C.4 traces the United States alone since 1962: the correlation is positive in most years through the four decades in which the interest burden rose toward six percent of GDP, negative in 14 of the 22 years from 2000 to 2021 with the burden averaging 2.6 percent, and positive again in 2022 to 2024 as the burden returned above five. The two annual series correlate at 0.60.

Table C.6 asks which way the association runs, on the annual panel, where the debt leg of the regressor moves. Lagged interest expense predicts the correlation beside the correlation's own lag, at 0.012 with a p-value of 0.014, while the lagged correlation does

Table C.1: Alternative dependent variables

	(1)	(2)	(3)	(4)	(5)	(6)
	Fisher z Monthly 10Y	Fisher z Monthly 5Y	Fisher z Annual 10Y	$\rho \times \sigma_b$ Monthly 10Y	$\rho \times \sigma_b$ Monthly 5Y	$\rho \times \sigma_b$ Annual 10Y
$IE^{fw}_{i,t-1}$	0.0153***	0.0150***	0.0178***	0.0003***	0.0002***	0.2420***
	(0.0034)	(0.0045)	(0.0056)	(0.0000)	(0.0000)	(0.0904)
Observations	14,414	13,903	1,757	14,414	13,903	1,757
Baseline	0.0138***	0.0133***	0.0145***	0.0138***	0.0133***	0.0145***

Note: This table reports the baseline specification with the dependent variable transformed. Columns (1) to (3) replace the correlation with its Fisher transformation, $z = \text{artanh}(\rho^{sb})$; columns (4) to (6) multiply it by the standard deviation of bond returns within the same period, the left-hand side of the model's regression form in Section 2. The regressor, the controls, the fixed effects and the standard errors are those of Table 2, and the baseline row repeats the estimate on the untransformed correlation for the same panel. The volatility-scaled columns carry the units of the return standard deviation they multiply, daily in the monthly columns and monthly in the annual one, so each column is read against its own baseline row. Standard errors, in parentheses, are two-way clustered by country and by period. *** p<0.01, ** p<0.05, * p<0.1.

not predict interest expense: 0.04 percentage points of GDP per unit of correlation, with a p-value of 0.83.

Table C.2: The debt leg, December step against interpolated

	Monthly 10Y	Monthly 5Y
December (reported)	0.0134***	0.0127***
(SE)	(0.0031)	(0.0040)
Interpolated	0.0136***	0.0112***
(SE)	(0.0029)	(0.0038)
Observations	14,118	13,607

Note: This table reports the baseline specification of Table 2 with the debt leg of IE^{fw} constructed two ways on identical observations: the reported series, the last annual debt-to-GDP figure the source had published, held through the year; and a linear interpolation of the annual figures onto the monthly grid. The two agree to within 0.0002 at ten years and 0.0015 at five. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1.

Table C.3: The regression country by country, monthly

Country	Group	No controls	(SE)	R^2	With controls	(SE)	R^2	N	Sample
HK	AE	0.5081	(0.5870)	0.011	0.5186	(0.5936)	0.021	276	2002-02–2025-01
AU	AE	0.2339***	(0.0306)	0.229	0.2419***	(0.0334)	0.235	370	1995-01–2025-10
US	AE	0.2292***	(0.0205)	0.414	0.2108***	(0.0217)	0.434	651	1971-09–2025-12
GB	AE	0.2100***	(0.0280)	0.280	0.1792***	(0.0314)	0.341	481	1986-01–2026-01
KR	AE	0.1667*	(0.0913)	0.033	0.1740*	(0.1017)	0.047	277	2002-12–2026-01
ES	AE	0.1363***	(0.0181)	0.244	0.1629***	(0.0189)	0.294	408	1992-02–2026-01
MY	EM	0.1170**	(0.0479)	0.030	0.1308**	(0.0591)	0.032	280	2001-10–2025-01
NL	AE	0.1107***	(0.0138)	0.266	0.1116***	(0.0144)	0.277	468	1987-02–2026-01
NO	AE	0.1121***	(0.0234)	0.131	0.1091***	(0.0240)	0.133	450	1987-09–2026-01
AT	AE	0.1092***	(0.0171)	0.176	0.1070***	(0.0174)	0.185	421	1991-01–2026-01
FI	AE	0.1120***	(0.0193)	0.166	0.1068***	(0.0195)	0.172	411	1991-09–2026-01
PL	EM	0.0859***	(0.0286)	0.064	0.1046***	(0.0278)	0.075	320	1999-06–2026-01
SG	AE	0.0972***	(0.0277)	0.083	0.0980***	(0.0281)	0.084	289	2001-01–2025-01
NZ	AE	0.1130**	(0.0516)	0.036	0.0917*	(0.0557)	0.061	298	2001-01–2025-10
CA	AE	0.0859***	(0.0080)	0.270	0.0859***	(0.0093)	0.270	466	1987-01–2026-01
HU	EM	0.0627***	(0.0116)	0.116	0.0746***	(0.0109)	0.137	298	2001-04–2026-01
BR	EM	0.0438*	(0.0232)	0.051	0.0574**	(0.0244)	0.077	226	2007-04–2026-01
IL	AE	0.0741***	(0.0168)	0.097	0.0548**	(0.0247)	0.126	216	2005-04–2026-01
FR	AE	0.0589	(0.0361)	0.032	0.0546	(0.0369)	0.050	373	1995-01–2026-01
DE	AE	0.0489	(0.0364)	0.022	0.0514	(0.0392)	0.046	373	1995-01–2026-01
BE	AE	0.0483***	(0.0067)	0.182	0.0474***	(0.0074)	0.191	421	1991-01–2026-01
IE	AE	0.0348***	(0.0075)	0.127	0.0357***	(0.0073)	0.140	441	1988-01–2026-01
PT	AE	0.0311***	(0.0095)	0.062	0.0350***	(0.0086)	0.120	301	2001-01–2026-01
SI	AE	0.0304**	(0.0141)	0.015	0.0343**	(0.0138)	0.017	213	2007-05–2026-01
ZA	EM	0.0309***	(0.0104)	0.051	0.0297***	(0.0104)	0.059	364	1995-07–2025-10
PE	EM	0.0336	(0.0318)	0.006	0.0231	(0.0443)	0.008	222	2006-06–2025-01
JP	AE	0.0292	(0.0209)	0.016	0.0198	(0.0193)	0.061	433	1990-01–2026-01
CZ	AE	0.0241	(0.0336)	0.002	0.0149	(0.0391)	0.006	301	2001-01–2026-01
TH	EM	0.0191	(0.0172)	0.005	0.0134	(0.0177)	0.014	304	1999-10–2025-01
GR	AE	0.0120	(0.0073)	0.076	0.0118*	(0.0066)	0.121	331	1998-05–2026-01
LT	AE	0.0086	(0.0275)	0.001	0.0088	(0.0303)	0.002	139	2008-06–2026-01
MX	EM	0.0007	(0.0242)	0.000	0.0015	(0.0263)	0.001	269	2003-08–2026-01
TR	EM	-0.0403***	(0.0144)	0.063	0.0011	(0.0219)	0.125	192	2010-02–2026-01
CL	EM	0.0002	(0.0255)	0.000	0.0010	(0.0277)	0.001	243	2005-10–2026-01
IN	EM	-0.0013	(0.0139)	0.000	-0.0028	(0.0145)	0.018	326	1998-12–2026-01
PH	EM	-0.0020	(0.0048)	0.001	-0.0109***	(0.0042)	0.044	289	2001-01–2025-01
IT	AE	-0.0248	(0.0235)	0.013	-0.0196	(0.0237)	0.022	336	1998-01–2026-01
CO	EM	-0.0066	(0.0067)	0.001	-0.0296**	(0.0127)	0.032	237	2006-05–2026-01
CH	AE	-0.0399**	(0.0200)	0.012	-0.0482**	(0.0216)	0.018	335	1998-01–2026-01
SE	AE	-0.0619*	(0.0322)	0.024	-0.0613*	(0.0319)	0.078	300	2001-01–2026-01
ID	EM	-0.0478***	(0.0086)	0.082	-0.0659***	(0.0091)	0.107	261	2004-02–2026-01
DK	AE	-0.0784**	(0.0327)	0.042	-0.0686**	(0.0349)	0.084	301	2001-01–2026-01
CN	EM	-0.1252***	(0.0470)	0.040	-0.0954**	(0.0388)	0.062	229	2006-01–2025-01
TW	AE	-0.2579***	(0.0635)	0.055	-0.2560***	(0.0614)	0.057	274	2001-01–2025-01

Note: This table reports one time-series regression per country of the non-overlapping monthly correlation of daily equity and daily 10-year zero-coupon bond returns on the forward-looking interest expense ratio, lagged one month, without controls and then with CPI inflation and industrial production growth, both lagged one month, on the same observations. Countries with fewer than 120 monthly observations are omitted, which removes none of the 44. Standard errors, in parentheses beside each estimate, are Newey-West with a twelve-month bandwidth. Rows are ordered by the controlled estimate. These regressions carry no time fixed effects, so their coefficients are not comparable with the panel estimates of Table 2. *** p<0.01, ** p<0.05, * p<0.1.

Table C.4: The regression country by country, annual

Country	Group	No controls	(SE)	R^2	With controls	(SE)	R^2	N	Sample
HK	AE	0.5060	(0.9587)	0.012	0.7985	(0.8060)	0.031	24	2002–2025
KR	AE	0.4055***	(0.1467)	0.178	0.3092*	(0.1590)	0.237	37	1988–2025
AU	AE	0.2191***	(0.0469)	0.187	0.2370***	(0.0528)	0.196	59	1961–2025
PL	EM	0.2437***	(0.0411)	0.376	0.2303***	(0.0401)	0.483	26	2000–2025
CL	EM	0.1137**	(0.0548)	0.069	0.1417**	(0.0692)	0.214	20	2006–2025
JP	AE	0.1126**	(0.0495)	0.124	0.1315***	(0.0499)	0.195	65	1961–2025
NO	AE	0.1699***	(0.0309)	0.193	0.1289**	(0.0529)	0.222	38	1988–2025
US	AE	0.1649***	(0.0314)	0.181	0.1145***	(0.0333)	0.349	63	1963–2025
MY	EM	0.1343	(0.1768)	0.028	0.1119	(0.1789)	0.071	24	2002–2025
SG	AE	0.0996	(0.0790)	0.066	0.0913	(0.0667)	0.142	31	1995–2025
CZ	AE	0.0983	(0.0818)	0.029	0.0892	(0.1173)	0.037	26	2000–2025
HU	EM	0.0967***	(0.0270)	0.236	0.0811***	(0.0274)	0.279	26	2000–2025
BR	EM	0.0627**	(0.0312)	0.125	0.0751***	(0.0284)	0.156	19	2007–2025
CA	AE	0.0942***	(0.0223)	0.176	0.0733***	(0.0226)	0.285	65	1961–2025
AT	AE	0.0675	(0.0438)	0.049	0.0689	(0.0423)	0.107	65	1961–2025
DK	AE	0.0635***	(0.0188)	0.248	0.0632***	(0.0171)	0.284	65	1961–2025
FI	AE	0.0843*	(0.0491)	0.060	0.0605	(0.0392)	0.167	60	1966–2025
PE	EM	-0.0100	(0.0531)	0.000	0.0570	(0.0743)	0.028	19	2007–2025
ES	AE	0.0676*	(0.0383)	0.065	0.0566	(0.0404)	0.096	65	1961–2025
NL	AE	0.0617*	(0.0339)	0.085	0.0527	(0.0348)	0.150	38	1988–2025
FR	AE	0.0496	(0.0517)	0.019	0.0479	(0.0543)	0.096	65	1961–2025
PT	AE	0.0382*	(0.0204)	0.072	0.0428**	(0.0176)	0.166	31	1995–2025
TR	EM	-0.0498	(0.0319)	0.052	0.0383	(0.0321)	0.533	15	2011–2025
SE	AE	0.0606**	(0.0305)	0.090	0.0382	(0.0307)	0.200	39	1987–2025
IE	AE	0.0177	(0.0141)	0.032	0.0330**	(0.0136)	0.117	62	1961–2025
BE	AE	0.0327**	(0.0131)	0.096	0.0330**	(0.0134)	0.100	65	1961–2025
MX	EM	0.0424	(0.0688)	0.013	0.0294	(0.0936)	0.098	24	2002–2025
NZ	AE	-0.0041	(0.0219)	0.001	0.0269	(0.0462)	0.029	39	1986–2025
CH	AE	0.0556	(0.0673)	0.019	0.0265	(0.0772)	0.032	65	1961–2025
CO	EM	-0.0259**	(0.0114)	0.037	0.0211	(0.0191)	0.221	23	2003–2025
IT	AE	0.0160	(0.0123)	0.024	0.0194	(0.0128)	0.027	65	1961–2025
ZA	EM	0.0174	(0.0201)	0.013	0.0164	(0.0211)	0.020	30	1996–2025
GR	AE	0.0121***	(0.0046)	0.103	0.0163**	(0.0080)	0.214	34	1988–2025
IN	EM	0.0074	(0.0201)	0.003	0.0044	(0.0188)	0.108	57	1961–2025
GB	AE	0.0424**	(0.0211)	0.052	-0.0045	(0.0245)	0.139	65	1961–2025
DE	AE	0.0050	(0.0711)	0.000	-0.0107	(0.0515)	0.269	65	1961–2025
PH	EM	-0.0102	(0.0072)	0.024	-0.0135	(0.0102)	0.045	28	1997–2025
ID	EM	0.0058	(0.0205)	0.001	-0.0182	(0.0361)	0.049	22	2004–2025
IL	AE	0.0180	(0.0334)	0.013	-0.0221	(0.0376)	0.217	22	2003–2025
SI	AE	-0.0134	(0.0341)	0.004	-0.0280	(0.0285)	0.103	15	2011–2025
TH	EM	-0.0291	(0.0362)	0.006	-0.0944*	(0.0505)	0.179	26	2000–2025
LT	AE	-0.0856**	(0.0414)	0.117	-0.0966**	(0.0444)	0.278	20	1999–2025
TW	AE	-0.3122**	(0.1569)	0.076	-0.2799*	(0.1455)	0.096	22	2000–2025
CN	EM	-0.3650***	(0.1098)	0.249	-0.5061***	(0.1429)	0.335	23	2003–2025

Note: This table reports the annual twin of Table C.3: one time-series regression per country of the calendar-year correlation of monthly equity and monthly 10-year zero-coupon bond returns on the forward-looking interest expense ratio, lagged one year, without controls and then with CPI inflation and real activity growth, both lagged, on the same observations. Countries with fewer than 15 annual observations are omitted, which removes none of the 44. Standard errors, in parentheses beside each estimate, are Newey-West with 3 lags. Rows are ordered by the controlled estimate. These regressions carry no time fixed effects, so their coefficients are not comparable with the panel estimates. *** p<0.01, ** p<0.05, * p<0.1.

Table C.5: Inflation and interest expense on the low-inflation sample

	(1)	(2)	(3)	(4)	(5)	(6)
	$\rho^{sb,10Y}_{i,t}$	$\rho^{sb,10Y}_{i,t}$	$\rho^{sb,10Y}_{i,t}$	$\rho^{sb,10Y}_{i,t}$	$\rho^{sb,10Y}_{i,t}$	$\rho^{sb,10Y}_{i,y}$
Sample	Full	Full	≤10%	≤10%	≤10%	≤10%
$IE^{fw}_{i,t-1}$	0.0373***	0.0137***	—	0.0147***	0.0147***	0.0167***
	(0.0088)	(0.0030)		(0.0029)	(0.0029)	(0.0044)
CPI inflation$_{i,t-1}$	0.0073**	-0.0062***	-0.0001	—	0.0002	-0.0072
	(0.0030)	(0.0015)	(0.0041)		(0.0039)	(0.0070)
Real activity growth$_{i,t-1}$	0.0013	0.0001	0.0000	0.0001	0.0001	-0.0013
	(0.0008)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0010)
Country FE	No	Yes	Yes	Yes	Yes	Yes
Time FE	No	Yes	Yes	Yes	Yes	Yes
Observations	14,282	14,282	13,878	13,878	13,878	1,634
R^2	0.099	0.442	0.439	0.443	0.443	0.400

Note: This table reports panel regressions of the stock-bond correlation on consumer price inflation and on forward-looking interest expense. Columns (1) and (2) use the common sample of Table 2; columns (3) to (5) keep the country-months of that sample whose lagged inflation is at most 10 percent in absolute value, marked ≤10% in the Sample row, and column (6) is the annual panel of Table 3 under the same restriction. Inflation is year-over-year consumer price inflation and real activity is industrial production growth for most of the panel, both lagged one period. A dash marks a regressor the column does not enter. Standard errors, in parentheses, are two-way clustered by country and by period. *** p<0.01, ** p<0.05, * p<0.1.

Table C.6: Which way the association runs

VARIABLES	(1) $\rho^{sb}_{i,t}$	(2) $IE^{fw}_{i,t}$
$\rho^{sb}_{i,t-1}$	0.1399***	0.0383
	(0.0361)	(0.1786)
$IE^{fw}_{i,t-1}$	0.0118**	0.6818***
	(0.0048)	(0.1160)
Macro controls	Yes	Yes
Country FE	Yes	Yes
Year FE	Yes	Yes
Observations	1,719	1,672
R^2	0.396	0.782

Note: This table asks which way the association of Table 2 runs, on the annual panel. Column (1) regresses the stock-bond correlation on its own lag and on lagged forward-looking interest expense; column (2) regresses forward-looking interest expense on its own lag and on the lagged correlation. Both columns carry the controls and the country and year effects of the baseline, and the lag of the correlation is taken on the calendar, so a country-year after a gap has no lag. The coefficient on the lagged correlation in column (2) is in percentage points of GDP per unit of correlation. A lagged dependent variable beside country effects is biased by a term of order one over the years per country, about forty here. Standard errors, in parentheses, are two-way clustered by country and by year. *** p<0.01, ** p<0.05, * p<0.1.

Figure C.1: Which measure, and which maturity

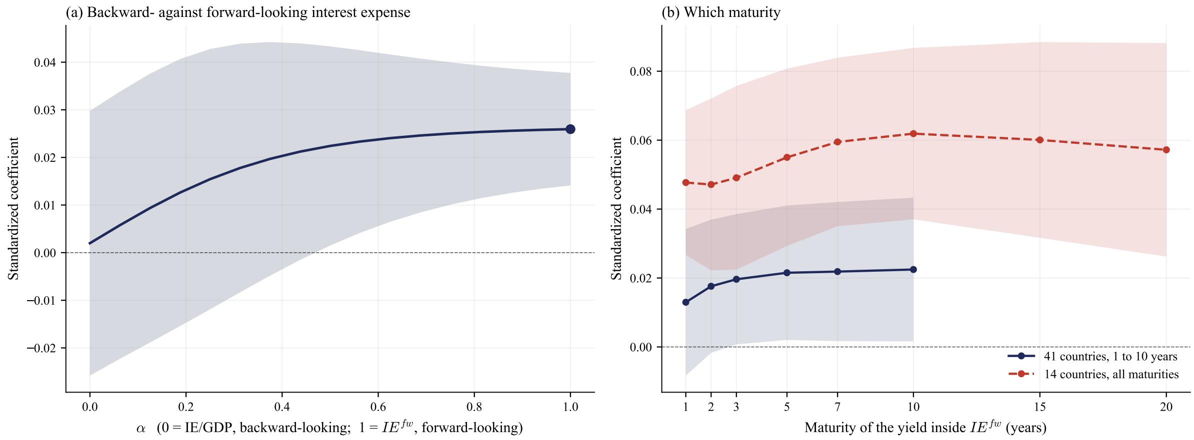

Note: Two sweeps over the construction of the fiscal measure, both estimated on the monthly panel with the baseline specification of Table 2. Panel (a) replaces ie^{fw} with the blend $\alpha \cdot \text{ie}^{\text{fw}} + (1-\alpha) \cdot IE/Y$, where IE/Y is realized interest payments over GDP, and traces the estimate as α moves from the backward-looking measure to the forward-looking one. Panel (b) rebuilds ie^{fw} with the yield at each maturity the panel carries; the wide line uses the 41 countries with every maturity from one to ten years and the narrow line the 14 countries that also have fifteen- and twenty-year yields. Both panels plot the standardized coefficient, the estimate multiplied by the within-country standard deviation of its own regressor, because the blend and the maturity substitution change the scale of the regressor. Shaded bands are 95 percent confidence intervals from standard errors two-way clustered by country and month.

Figure C.2: Leave-one-out

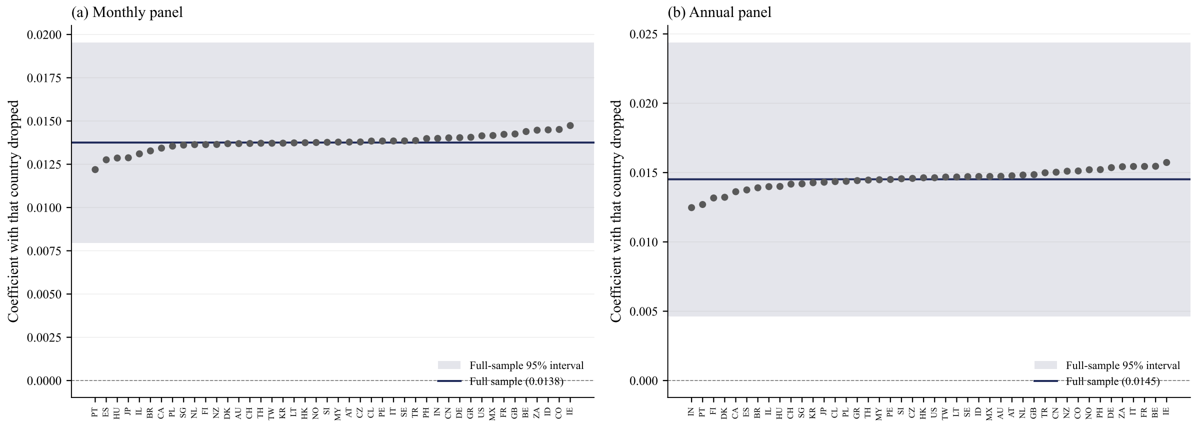

Note: The baseline estimate re-computed forty-four times, each time dropping one country. Panel (a) is the monthly panel and panel (b) the annual. Each point is the coefficient with that country excluded; the horizontal line is the full-sample estimate and the shaded band its 95 percent confidence interval. Countries are ordered by the estimate. The baseline is the specification of column (5) of Table 2, estimated on every row it reaches, so the full-sample estimate reads 0.0138 where that table's common-sample column prints 0.0137.

Figure C.3: Countries sorted on their interest expense, month by month

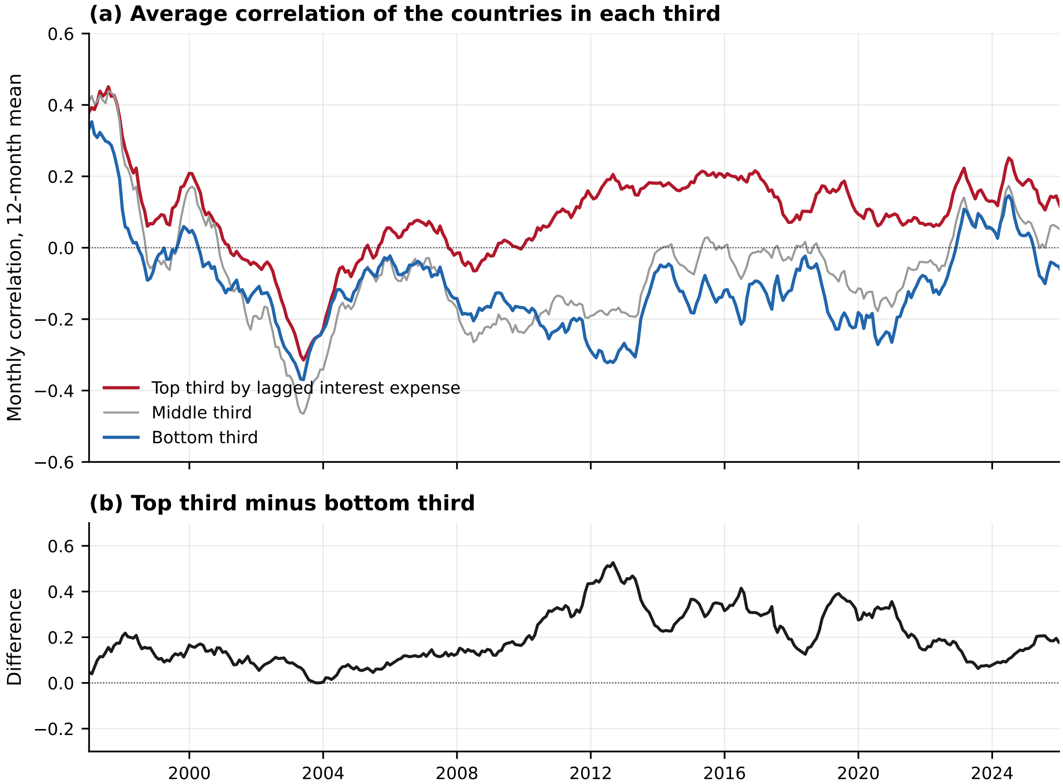

Note: Each month, the countries with a monthly stock-bond correlation are ranked on the previous month's forward-looking interest expense and split into thirds. Panel (a) plots the average correlation of each third, smoothed with a twelve-month trailing mean; panel (b) the top third less the bottom third. Membership is re-drawn every month. The series begin in 1997 and rest on at least fifteen countries.

Figure C.4: The United States since 1962

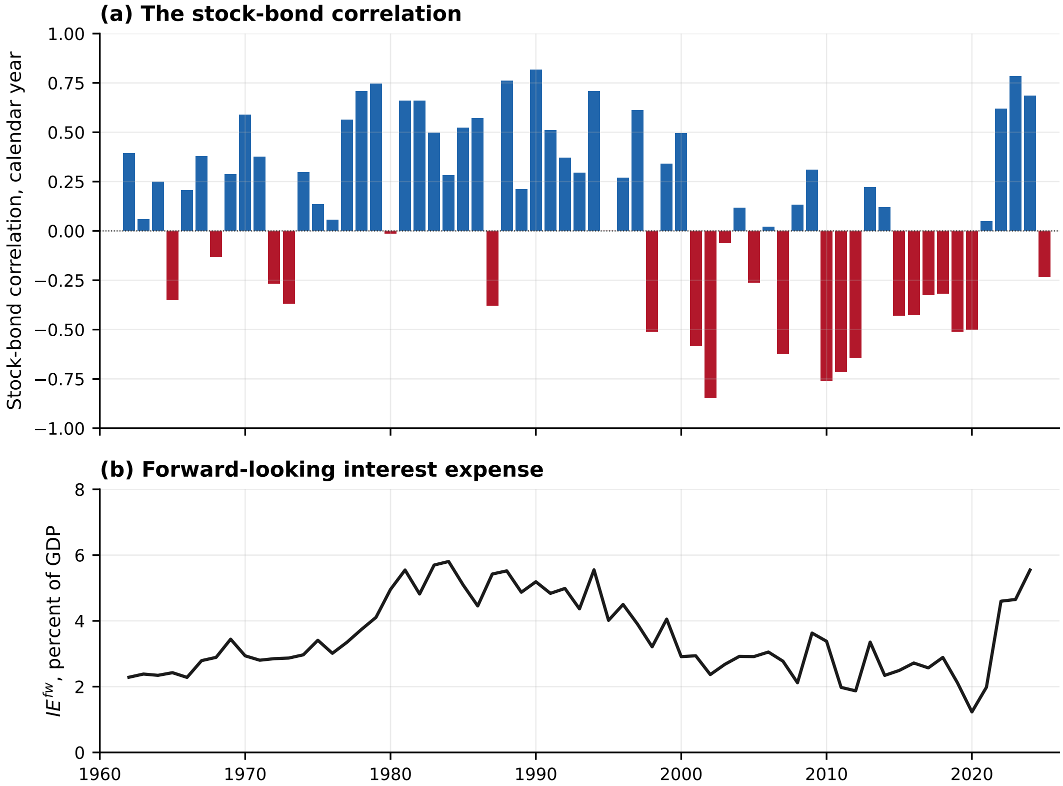

Note: Panel (a) is the calendar-year correlation of monthly equity and ten-year zero-coupon bond returns for the United States, the annual panel's dependent variable, positive years in blue and negative in red. Panel (b) is the United States' forward-looking interest expense, general government debt over GDP times the ten-year zero-coupon yield, in percent of GDP.

C.2 Robustness for Section 5

One figure supports Section 5. Figure C.5 plots the reported euro-area difference-in-differences estimate beside sixteen variations of the design: the treatment definition, the window length, the country sample, each periphery member dropped in turn, two alternative dependent variables, and three placebo boundaries inside the pre-announcement period. Every row estimates the same object as the reported specification, so the estimates are directly comparable down the column. The placebo rows are the exception by construction: a placebo that read as large as the reported estimate would be evidence against the design. Across the thirteen variations of treatment, window, sample and dependent variable the estimate lies between 0.37 and 0.49, against 0.44 in the reported specification; the largest departure is 0.7 standard errors of the reported estimate, every variation keeps its sign, and every interval excludes zero. Dropping Greece gives 0.40. The three placebo rows read 0.04 to 0.11, and none is distinguishable from zero.

Figure C.5: The euro-area estimate under alternative designs

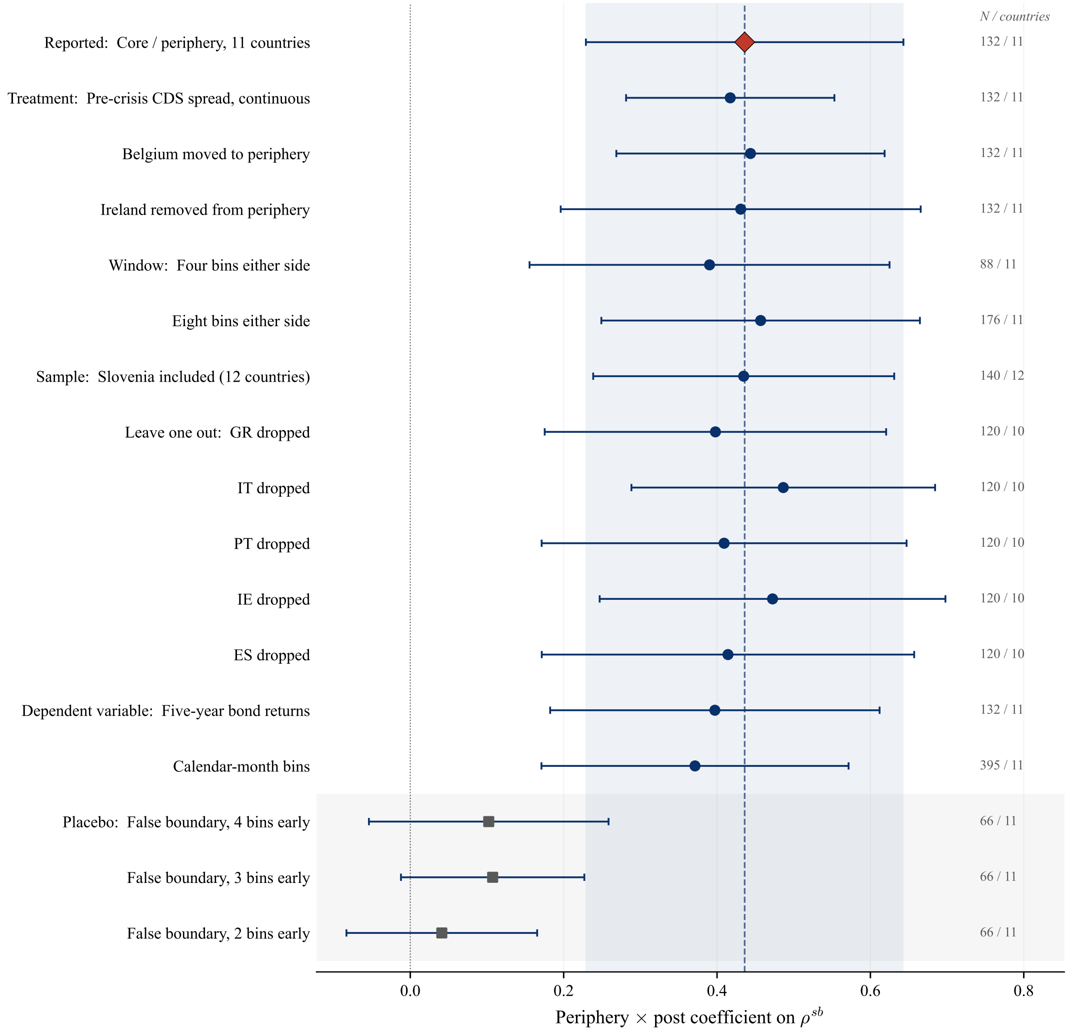

Note: Each row reports the coefficient on the interaction of the periphery indicator with an indicator for the period after 20 October 2009, from the specification of Figure 3. The diamond is the reported estimate and the shaded band its 95 percent confidence interval; bars are the corresponding intervals for each variation, from country-clustered standard errors with the small-sample correction of Figure 3, at each row's own degrees of freedom. The right margin reports the number of country-bin observations and of countries. The treatment rows replace the indicator with the continuous pre-crisis credit spread, move Belgium into the periphery, and remove Ireland; the window rows shorten and lengthen the event window; the sample row restores Slovenia; the leave-one-out rows drop each periphery member in turn; two rows change the dependent variable, to five-year bond returns and to calendar months; and the placebo rows anchor a false boundary two, three and four bins before the announcement, estimated on the pre-announcement sample alone.

C.3 Robustness for Section 6

Table C.7 re-estimates the reported column of each shifter table with the lagged interest-expense measure added, on the identical rows. The newspaper index keeps a positive association beside the measure, 0.081 against 0.101 alone, while the rating shifter falls from 0.158 to 0.064 and is indistinguishable from zero. The lagged measure carries the previous month's sovereign yield, which already prices the fiscal position an agency cites, so we read the conditioning as removing the part of the rating association the yield already carries. Table C.8 repeats the two panels of Table 6 with the bins cut once on the whole sample instead of on an expanding window. The estimate increases across all four bins on both sorting variables. Under two-way clustering with a normal reference distribution, the third bin is positive at the five percent level and the top bin at the one percent level, 0.209 on the interest-expense sort and 0.167 on the debt sort; the wild cluster bootstrap on the debt sort's third bin, which holds 29 countries, gives a p-value of 0.066. Equal slopes are rejected at the one percent level in both panels. Table C.9 holds the reported specification of Table 7 fixed and moves only the rationale filter, through three nested filters: F_0 admits every rating action, F_B keeps those whose stated rationale reaches the public finances, and F_C, the filter the paper reports, adds the requirement that the agency named no cyclical, banking, external or political cause beside it. With every rating action admitted the coefficient is 0.120, significant at the one percent level, and the symmetry restriction is rejected; under F_C it is 0.158 and the restriction is no longer rejected. The downgrade arm is about 0.2 under each of the three filters, and the upgrade arm is near zero under the first two and negative and imprecise under F_C. The last two columns split the F_C actions by the kind of fiscal rationale the agency gave. The 116 actions citing the structural fiscal position (the deficit, the debt level, its trajectory or the interest burden) give 0.189, significant at the five percent level, with a within-country randomization p-value of 0.003; the 99 citing fiscal governance give 0.101, with a randomization p-value of 0.31 and a 95 percent interval that covers both zero and the structural estimate. Table C.10 sorts the rating sample into four bins of the lagged fiscal burden, as Tables 6 and C.8 do for the newspaper index, with the bins cut on an expanding window in Panels A and B and once on the whole sample in Panel C. On either expanding sort the top bin is the only one distinguishable from zero under two-way clustering; the wild cluster bootstrap on that bin's own countries separates it from zero on the debt sort and does not on the interest-expense sort. With 205 action months, 195 of them moving the shifter, spread over four bins the lower bins are imprecise, and the test of equal slopes does not reject in any panel.

Tables C.11 through C.16 ask what else the newspaper index could be measuring, and how much its measurement matters. Five of them hold the specification of the reported column of Table 5 fixed; Table C.14 uses the specification of column (2), without the macro controls, because the placebo panel carries none. Table C.11 puts the index beside the mean Loughran-McDonald negativity of the same articles, lagged the same way. The negativity score predicts the correlation on its own, 0.710 (0.263); beside the index it keeps 0.445 (0.215) and the index keeps 0.093, so the two readings of one set of articles carry separate information and a general bad-news score does not absorb the index. Table C.12 adds the logarithm of the Economic Policy Uncertainty index on the nineteen countries it covers, and the index moves from 0.260 to 0.256. The higher slope on this subsample comes from the countries it covers: the reported specification estimated on the same nineteen countries over all of their months gives 0.247 on 5,488 country-months. Table C.13 reclassifies a probe subsample of 4,584 country-months on text rewritten by a second language model to replace country names, currencies, officials, institutions and dates by placeholders. In the index column articles are assigned to countries by a separate attribution pass of a language model; in the probe columns they are assigned by the alias table on the unmodified text, identically in both arms, because the anonymized arm's model output names countries by placeholder tokens that cannot be mapped to a country code. The rewriting is measured: a name, demonym or city of an attributed country survives in 24.7 percent of the anonymized articles and some alias the rewriter was told to replace in 35.4 percent. The slope is 0.075 on the original text and 0.049 on the anonymized text, significant at the one percent and at the ten percent level; the point estimate falls by a third and keeps its sign, and dropping every article-country pair whose alias survives gives 0.043 (0.025) on 4,424 country-months. The two arms agree at the article level, with a Cohen's kappa of 0.86, and the country-month shares from the two arms correlate at 0.88. Table C.14 runs the same specification on two dependent variables, the stock-bond correlation and the correlation of the country's equity return with the United States equity return. The index gives 0.105 on the first and 0.003 on the second, so a general attention channel, which would move both, does not carry the association.

Table C.15 estimates the reported column separately on the 29 advanced economies and the 15 emerging markets. Both groups are positive, 0.133 and 0.085, and each survives a wild cluster bootstrap on its own countries. Table C.16 holds the reported specification fixed and changes how the share is measured. The model attributes 7.2 percent of the article-country pairs to a country of which the article contains no alias from the fixed table, such as its name, adjective or demonym, its capital or a major city, its central bank, treasury or seat of government, or its sovereign bond nickname, a share highest in the euro area's small members; on alias-confirmed pairs alone, with the attention control on the same count, the estimate is 0.095 (0.029) on 11,473 country-months, a diagnostic bound and not a replacement. Counting in the numerator only the articles whose fiscal discussion the classifier tied to that country, over the reported denominator, gives 0.351 (0.062). Restricting to the country-months whose lagged article count is at least two, three and five gives 0.151, 0.188 and 0.283 on 10,794, 9,784 and 8,090 country-months: the association is larger where the share is measured on more articles, which is what attenuation from a share built on one or two articles would produce, and the reported column is estimated on every country-month.

Table C.17 re-estimates each reported column 44 times, leaving one country out in turn. The rating estimate runs from 0.090 without Greece to 0.179 without Indonesia and the newspaper estimate from 0.084 without Spain to 0.108 without South Africa, and no drop changes either sign. The range of the point estimates is the informative statistic; with 44 clusters, whether one drop crosses a significance threshold is substantially a property of the variance estimator.

Table C.7: The two shifters beside the interest-expense measure

VARIABLES Shifter	(1) $\rho_{i,t}^{sb,10Y}$ Newspaper index	(2) $\rho_{i,t}^{sb,10Y}$ Newspaper index	(3) $\rho_{i,t}^{sb,10Y}$ Rating actions	(4) $\rho_{i,t}^{sb,10Y}$ Rating actions
$FSI_{i,t-1}$	0.1014***	0.0809**		
	(0.0342)	(0.0319)		
$Action_{i,t}^{F_C}$			0.1580**	0.0643
			(0.0635)	(0.0585)
$IE_{i,t-1}^{fw}$		0.0139***		0.0134***
		(0.0039)		(0.0031)
Attention control	Yes	Yes	No	No
Country FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Macro controls	Yes	Yes	Yes	Yes
IE^{fw} control	No	Yes	No	Yes
Observations	11,802	11,802	14,385	14,385
Countries	44	44	44	44
R^2	0.398	0.402	0.438	0.442

Note: This table re-estimates the reported column of each shifter table with the lagged interest-expense measure added, on the identical rows. The dependent variable is the monthly stock-bond correlation, the non-overlapping correlation of daily equity and daily 10-year zero-coupon bond returns within a calendar month. Columns (1) and (2) regress it on the newspaper fiscal-salience index lagged one month, $FSI_{i,t-1}$, on 11,802 country-months over 44 countries and 2000-02–2026-01; columns (3) and (4) regress it on the net sovereign rating action under the F_C filter, $Action_{i,t}^{F_C}$, on 14,385 country-months over 44 countries and 1974-02–2026-01. Columns (1) and (3) are the reported columns of Tables 5 and 7; columns (2) and (4) add $IE_{i,t-1}^{fw}$, the product of debt to GDP and the 10-year sovereign yield, lagged one month. The newspaper index keeps a positive association beside the measure; the rating shifter's estimate falls and is no longer distinguishable from zero. Each pair is estimated on one sample fixed before its first column, over every regressor in either column, so a change within a pair is the added measure and never the sample. The macro controls are contemporaneous on the newspaper panel and lagged one month on the rating panel, as in the two source tables; the newspaper columns carry the attention control, the logarithm of one plus the article count. The global series (`vix`, `ebp`) take one value per month, so the month effects absorb them and they are carried in the declared control set alone. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1. P-values behind the stars use the standard normal reference for the clustered t-statistic in the Python build and a t-distribution with the smaller cluster count less one degrees of freedom in the Stata build.

Table C.8: Newspaper fiscal salience by bin of the lagged fiscal burden, whole-sample cuts

VARIABLES	(1) $\rho_{i,t}^{sb,10Y}$	(2) $\rho_{i,t}^{sb,10Y}$	(3) $\rho_{i,t}^{sb,10Y}$	(4) $\rho_{i,t}^{sb,10Y}$	(5) $\rho_{i,t}^{sb,10Y}$
Panel A. $IE_{i,t-1}^{fw}$*, whole-sample cut*					
Observed range	[-0.63,1.07]	[1.07,1.98]	[1.98,3.19]	[3.19,54.59]	
Mean $IE_{i,t-1}^{fw}$	0.47	1.51	2.55	5.60	2.53
$FSI_{i,t-1}$	-0.0220	-0.0127	0.0992**	0.2092***	-0.0591*
	(0.0249)	(0.0400)	(0.0399)	(0.0602)	(0.0307)
Wild cluster p				[0.002]	
× Bin 2					0.0456 (0.0376)
× Bin 3					0.1598*** (0.0418)
× Bin 4					0.3118*** (0.0695)
Equal slopes, p					<0.0001
Observations	2,951	2,950	2,950	2,951	11,802
Countries	35	41	40	29	44
Panel B. Debt/GDP$_{i,t-1}$*, whole-sample cut*					
Observed range	[0.05,38.01]	[38.05,53.82]	[53.90,82.01]	[82.17,258.37]	
Mean Debt/GDP$_{i,t-1}$	26.26	44.73	66.09	119.16	64.00
$FSI_{i,t-1}$	-0.0180	0.0042	0.1075**	0.1666***	0.0033
	(0.0374)	(0.0429)	(0.0534)	(0.0505)	(0.0271)
Wild cluster p	[0.628]		[0.066]	[0.001]	
× Bin 2					0.0081 (0.0448)
× Bin 3					0.0797 (0.0696)
× Bin 4					0.2937*** (0.0636)
Equal slopes, p					<0.0001
Observations	2,953	2,957	2,950	2,942	11,802
Countries	28	30	29	18	44
Attention control	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes
Macro controls	Yes	Yes	Yes	Yes	Yes

Note: Each panel re-estimates the headline column of Table 5 inside bins of the sorting variable named in the panel heading, with the cut points computed once on the whole estimation sample; Table 6 cuts the same bins on an expanding window. Column (5) estimates all bins jointly with the index's slope free to differ by bin. The interaction rows report the difference from Bin 1, with the standard error in parentheses beside the coefficient, and Equal slopes reports the joint test that those differences are zero. Bins are cut across the pooled estimation sample and do not partition countries: a country moves between bins as its sorting variable moves. Intervals are half open and closed on the right, so a value sitting exactly on a cut enters the lower bin alone. The debt series carries the preceding December's annual value through the following calendar year, so 0 of 1,188 country-years take more than one monthly value and ties on a cut point are the ordinary case in Panel B. A bin whose country count falls below 30 carries a wild cluster restricted bootstrap-t p in brackets, 9,999 draws with Rademacher weights clustered on country. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1. P-values behind the stars use the standard normal reference for the clustered t-statistic in the Python build and a t-distribution with the smaller cluster count less one degrees of freedom in the Stata build.

Table C.9: The rating shifter under successively tighter rationale filters

VARIABLES	(1) F_0	(2) F_B	(3) F_C	(4) F_C structural	(5) F_C governance
Panel A. Net action					
$Action^F_{i,t}$	0.1198***	0.1432***	0.1580**	0.1893**	0.1007
	(0.0316)	(0.0361)	(0.0635)	(0.0872)	(0.1055)
Panel B. Two arms					
$Down^F_{i,t}$	0.1963***	0.2249***	0.1878***	0.2268***	0.1485
	(0.0360)	(0.0382)	(0.0592)	(0.0790)	(0.1427)
$Up^F_{i,t}$	0.0046	-0.0183	-0.1127	-0.1452	0.0159
	(0.0304)	(0.0355)	(0.1031)	(0.1021)	(0.1711)
Symmetry p	0.000	0.000	0.505	0.414	0.494
Randomization p				[0.0032]	[0.3135]
Country FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes
Macro controls	Yes	Yes	Yes	Yes	Yes
Filter-passing events	1,054	655	215	116	99
Months with an action	926	594	205	113	93
Nonzero country-months	895	576	195	109	87
Observations	14,385	14,385	14,385	14,385	14,385
R^2	0.439	0.439	0.438	0.438	0.438

Note: This table holds Table 7's reported specification fixed and moves only the rationale filter. F_0 keeps every rating action, F_B keeps the actions whose stated rationale reaches the public finances, and F_C adds the requirement that the agency named no cyclical, banking, external or political cause beside it. The last two columns cut F_C into the structural and governance rationale categories it partitions into, and their coefficients are estimated on the same rows as the rest of the table. The filters and the classification behind them are described in Appendix B. *Filter-passing events* counts the rating actions the filter keeps whose country-month is a row of the estimation sample, taken from the event file's own membership flags; *Months with an action* counts the country-months in which at least one of them falls; *Nonzero country-months* counts the months in which those actions moved the shifter by at least a fraction of a notch, which is fewer, because a move into or out of review status, which the scale scores at zero, passes a filter and shifts nothing. The sample is 44 countries over 1974-02–2026-01, 14,385 country-months. Panel A uses the net action $Action = Down - Up$; Panel B frees the arms and reports the p-value of the restriction Panel A imposes, that $b_{Down} + b_{Up} = 0$. The last row of Panel B reports, for the two sub-type columns, a within-country randomization p: each of 9,999 draws reorders that column's shifter across the months of each country, so every country keeps its own actions and only the months they land in change, and the p is the share of draws whose coefficient is at least as large in absolute value as the one reported. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1. P-values behind the stars use the standard normal reference for the clustered t-statistic in the Python build and a t-distribution with the smaller cluster count less one degrees of freedom in the Stata build.

Table C.10: The rating shifter by bin of the lagged fiscal burden

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Bin 1	Bin 2	Bin 3	Bin 4	Pooled
Panel A					
$Action^{F_C}_{i,t}$	-0.0430	0.2715	-0.0329	0.1515**	0.0069
	(0.1616)	(0.2508)	(0.1774)	(0.0722)	(0.1535)
× Bin 2					0.1168 (0.3044)
× Bin 3					-0.1128 (0.1935)
× Bin 4					0.2089 (0.1653)
				[0.1507]	
State, bin mean	1.16	2.45	3.63	6.87	2.79
State, range	−0.6–3.7	1.1–4.8	2.2–6.2	3.4–54.6	−0.6–54.6
Interactions jointly zero, p					0.065
Observations	6,340	3,376	2,249	2,380	14,345
Countries	39	38	35	25	44
Panel B					
$Action^{F_C}_{i,t}$	0.0020	0.1084	-0.0853	0.2199***	-0.0244
	(0.1426)	(0.1355)	(0.2254)	(0.0667)	(0.1641)
× Bin 2					0.1676 (0.1944)
× Bin 3					-0.0438 (0.2344)
× Bin 4					0.2480 (0.1698)
	[0.9868]		[0.6438]	[0.0249]	
State, bin mean	28.09	46.69	64.76	108.94	63.75
State, range	0.1–44.1	37.4–61.4	46.0–78.4	44.8–258.4	0.1–258.4
Interactions jointly zero, p					0.318
Observations	3,722	3,627	2,750	4,246	14,345
Countries	29	33	28	22	44
Panel C					
$Action^{F_C}_{i,t}$	0.2123*	-0.0329	0.2338	0.1214*	0.2265
	(0.1256)	(0.2133)	(0.1795)	(0.0689)	(0.1398)
× Bin 2					-0.3524 (0.2356)
× Bin 3					-0.1602 (0.1900)
× Bin 4					-0.0361 (0.1495)
State, bin mean	0.55	1.70	2.87	6.06	2.80
State, range	−0.6–1.2	1.2–2.2	2.2–3.7	3.7–54.6	−0.6–54.6
Interactions jointly zero, p					0.458
Observations	3,597	3,596	3,596	3,596	14,385
Countries	36	41	38	30	44
Country FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes
Macro controls	Yes	Yes	Yes	Yes	Yes

Note: Each panel re-estimates Table 7's reported specification inside each bin of the lagged fiscal burden, on one estimation sample of 44 countries over 1974-02–2026-01, 14,385 country-months. Panels A and B cut 4 bins with an EXPANDING WINDOW: a country-month is assigned under cut points computed on every observation dated strictly before its own month, and the 40 earliest country-months carry no bin because the history before them is shorter than the 40 observations the cut needs. Panel C repeats Panel A with one set of cut points taken from the whole estimation sample. Bins are half-open and closed on the right, so a value sitting exactly on a cut enters one bin; in Panel B that is load-bearing, because the monthly debt series takes one value a year and a value landing on a cut would otherwise place a whole country-year in two bins. Cut points are taken across the pooled sample and not within each country, so a bin is not a group of countries and a country moves between bins over time; the expanding cuts also move, which is why the level ranges of adjacent bins overlap. The pooled column estimates one regression over every binned row with the shifter's slope free to differ by bin, and *Interactions jointly zero* reports the p-value of the test that the bin slopes are equal; the indented rows are the interaction coefficients, each the difference from the lowest bin's slope. Brackets report a wild cluster restricted bootstrap-t p-value with Rademacher weights, clustered on country over 9,999 draws, on every bin column whose country count is below 30, because clustered asymptotics overstate precision at that many clusters; the pooled column carries none, its restriction being the joint test. Cut points in the units of the sorting variable, all computed at build time: Panel A's cuts move over 584 months, from 3.152, 3.274 and 3.333 in the first month it can cut to 1.191, 2.246 and 3.670 in the last; Panel B's cuts move over 584 months, from 41.443, 42.592 and 44.603 in the first month it can cut to 39.019, 55.288 and 79.804 in the last; Panel C cuts once on the whole estimation sample, at 1.193, 2.247 and 3.673. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1. P-values behind the stars use the standard normal reference for the clustered t-statistic in the Python build and a t-distribution with the smaller cluster count less one degrees of freedom in the Stata build.

Table C.11: Newspaper fiscal salience against a dictionary negativity placebo

VARIABLES Regressors	(1) $\rho^{sb,10Y}_{i,t}$ Index	(2) $\rho^{sb,10Y}_{i,t}$ Lexicon	(3) $\rho^{sb,10Y}_{i,t}$ Horse race
$FSI_{i,t-1}$	0.1014***		0.0931***
	(0.0342)		(0.0328)
$LM_{i,t-1}$		0.7102***	0.4451**
		(0.2628)	(0.2151)
Attention control	Yes	Yes	Yes
Country FE	Yes	Yes	Yes
Month FE	Yes	Yes	Yes
Macro controls	Yes	Yes	Yes
IE^{fw} control	No	No	No
Observations	11,802	11,802	11,802
Countries	44	44	44
R^2	0.398	0.397	0.399

Note: This table puts the newspaper fiscal-salience index against a dictionary placebo. The dependent variable is the monthly stock-bond correlation, the non-overlapping correlation of daily equity and daily 10-year zero-coupon bond returns within a calendar month; a country-month enters when it carries at least 15 trading days. The placebo is the mean Loughran-McDonald negativity of the same country-month's articles, lagged one month, and it is what the index would look like if it were measuring bad news about a country instead of fiscal news. Articles are assigned to countries by a separate attribution pass of a language model, on both sides of the share, under attribution rule attr_v1.0. The two measures are computed over one set of country-article pairs and share one denominator, which the build asserts on each of the 13,142 country-months that carry an article count. All three columns are estimated on one sample of 11,802 country-months over 44 countries and 2000-02–2026-01, fixed before the first column, so the coefficients are comparable across them and the observation counts are equal by construction. Every column carries the attention control, the logarithm of one plus the article count, which is the denominator both measures divide by. The global series (`vix`, `ebp`) take one value per month, so the month effects absorb them and they are carried in the declared control set alone. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1. P-values behind the stars use the standard normal reference for the clustered t-statistic in the Python build and a t-distribution with the smaller cluster count less one degrees of freedom in the Stata build.

Table C.12: Newspaper fiscal salience against economic policy uncertainty

VARIABLES Regressors	(1) $\rho^{sb,10Y}_{i,t}$ Index	(2) $\rho^{sb,10Y}_{i,t}$ Index and EPU
$FSI_{i,t-1}$	0.2603***	0.2564***
	(0.0907)	(0.0907)
	[0.005]	[0.005]
$\log EPU_{i,t-1}$		0.0304
		(0.0237)
Attention control	Yes	Yes
Country FE	Yes	Yes
Month FE	Yes	Yes
Macro controls	Yes	Yes
IE^{fw} control	No	No
Observations	5,384	5,384
Countries	19	19
R^2	0.462	0.463

Note: This table adds the Economic Policy Uncertainty index to the specification of the main-text salience table's reported column. The dependent variable is the monthly stock-bond correlation, the non-overlapping correlation of daily equity and daily 10-year zero-coupon bond returns within a calendar month; a country-month enters when it carries at least 15 trading days. The uncertainty index enters in logarithms and lagged one month, and it is lagged once, on its own contiguous month grid inside the data build, so a month the uncertainty index does not publish gives a missing lag and never a stale one. Articles are assigned to countries by a separate attribution pass of a language model, on both sides of the share, under attribution rule attr_v1.0. The uncertainty index covers 19 of the paper's countries, and those 19 are the sample here: 5,384 country-months over 2000-02–2025-12. The slope on this subsample is not the full-panel slope, and the covered countries are why. The reported column on the whole panel is 0.1014 over 44 countries, and the same column estimated on these 19 countries over all of their months is 0.2471 on 5,488 country-months, so the difference is the country mix and not the uncertainty index's coverage window. The global series (`vix`, `ebp`) take one value per month, so the month effects absorb them and they are carried in the declared control set alone. Because the sample has fewer country clusters than the clustered asymptotics assume, the figure in square brackets is a wild cluster restricted bootstrap-t p-value, Rademacher weights, clustered on country, 9,999 draws. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1. P-values behind the stars use the standard normal reference for the clustered t-statistic in the Python build and a t-distribution with the smaller cluster count less one degrees of freedom in the Stata build.

Table C.13: Newspaper fiscal salience under anonymized article text

VARIABLES Text	(1) $\rho^{sb,10Y}_{i,t}$ Full panel	(2) $\rho^{sb,10Y}_{i,t}$ Probe, raw	(3) $\rho^{sb,10Y}_{i,t}$ Probe, anonymized	(4) $\rho^{sb,10Y}_{i,t}$ Anonymized, residuals dropped
$FSI_{i,t-1}$	0.1014***	0.0752***	0.0492*	0.0432*
	(0.0342)	(0.0268)	(0.0261)	(0.0254)
Attention control	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Macro controls	Yes	Yes	Yes	Yes
IE^{fw} control	No	No	No	No
Observations	11,802	4,584	4,584	4,424
Countries	44	44	44	44
R^2	0.398	0.441	0.440	0.439

Note: This table asks whether the newspaper fiscal-salience index survives removing every country name from the article text. Column (1) is the specification of the main-text salience table's reported column on the whole estimation sample. Columns (2) and (3) are the probe subsample, classified once on unmodified text and once on text with country names, currencies, officials, institutions and dates replaced by placeholders. The dependent variable throughout is the monthly stock-bond correlation, the non-overlapping correlation of daily equity and daily 10-year zero-coupon bond returns within a calendar month; a country-month enters when it carries at least 15 trading days. In column (1) articles are assigned to countries by a separate attribution pass of a language model, on both sides of the share, under attribution rule attr_v1.0. In columns (2) to (4) articles are assigned to countries by the alias table run on the unmodified text, on both sides of the share and identically in both arms, because the anonymized arm's model output names countries by placeholder tokens that cannot be mapped to a country code; holding attribution fixed is what makes the contrast a test of flagging. The anonymizer is itself a language model, and its output is measured: the name, adjective, demonym or a city of an attributed country survives in the anonymized text of 2,463 of the 9,987 articles, 24.7 percent, and an alias of any kind the anonymizer is told to replace, a central bank, treasury, bond or index name or seat of government among them, in 3,539, 35.4 percent; column (4) drops every pair of the wider kind from both sides of the share. Each probe column carries its own attention control, the logarithm of one plus that arm's article count, because the two arms count different articles and the whole panel's count is neither; the macro controls and the fixed effects are the reported column's, with the global series (`vix`, `ebp`) absorbed by the month effects as there; columns (2) and (3) share one sample of 4,584 country-months and column (4) loses the country-months whose articles all carry a residual alias. Against thresholds fixed before the probe ran, 0.85 for the correlation of the two arms' country-month shares and 0.70 for the article-level Cohen's kappa, the arms correlate at 0.88 over 4,801 populated country-months and agree with a kappa of 0.86 over 9,987 article pairs. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1. P-values behind the stars use the standard normal reference for the clustered t-statistic in the Python build and a t-distribution with the smaller cluster count less one degrees of freedom in the Stata build.

Table C.14: Cross-asset placebo for newspaper fiscal salience

DEPENDENT VARIABLE	(1) $FSI_{i,t-1}$	(2) Observations	(3) Countries	(4) R^2
$\rho_{i,t}^{sb,10Y}$	0.1053***	11,802	44	0.396
	(0.0346)			
$\rho_{i,t}^{eq,US}$	0.0030	11,773	43	0.578
	(0.0124)			

Note: This table runs one regression twice with two dependent variables. The benchmark row is the monthly stock-bond correlation, the non-overlapping correlation of daily equity and daily 10-year zero-coupon bond returns within a calendar month. The placebo row is the non-overlapping correlation of the same country's daily equity returns with United States daily equity returns. A country-month enters either panel when it carries at least 15 trading days. The regressor is the newspaper fiscal-salience index lagged one month, and both rows carry the attention control, the logarithm of one plus the article count. Articles are assigned to countries by a separate attribution pass of a language model, on both sides of the share, under attribution rule attr_v1.0. The two rows share the index, the attention control, the country effects and the month effects, so the only difference between them is which correlation is on the left-hand side. If fiscal coverage moved the stock-bond correlation through a general attention channel and not through the fiscal position, it would move the equity-to-equity correlation as well. The placebo panel covers 43 countries against the benchmark's 44: the United States has no equity correlation against its own market, so the object is undefined there and the country is absent by construction. On the 11,452 country-months both panels share, the two estimates are 0.1036 and 0.0030, so the contrast is not the samples differing. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1. P-values behind the stars use the standard normal reference for the clustered t-statistic in the Python build and a t-distribution with the smaller cluster count less one degrees of freedom in the Stata build.

Table C.15: Newspaper fiscal salience by country group

VARIABLES	(1) $\rho^{sb,10Y}_{i,t}$	(2) $\rho^{sb,10Y}_{i,t}$	(3) $\rho^{sb,10Y}_{i,t}$
Sample	Full	Advanced	Emerging
$FSI_{i,t-1}$	0.1014***	0.1327***	0.0853**
	(0.0342)	(0.0408)	(0.0382)
		[0.001]	[0.034]
Attention control	Yes	Yes	Yes
Country FE	Yes	Yes	Yes
Month FE	Yes	Yes	Yes
Macro controls	Yes	Yes	Yes
IE^{fw} control	No	No	No
Countries	44	29	15
Observations	11,802	8,081	3,721

Note: This table re-estimates the specification of the main-text salience table's reported column on the whole estimation sample and separately on each country group. The dependent variable is the monthly stock-bond correlation, the non-overlapping correlation of daily equity and daily 10-year zero-coupon bond returns within a calendar month; a country-month enters when it carries at least 15 trading days. The regressor is the newspaper fiscal-salience index lagged one month, and every column carries the attention control, the logarithm of one plus the article count, because the index is a share whose denominator is that count. Articles are assigned to countries by a separate attribution pass of a language model, on both sides of the share, under attribution rule attr_v1.0. The estimation sample is fixed once, over 44 countries and 2000-02–2026-01, and the two group columns partition it: 29 advanced economies and 15 emerging markets. Group membership follows the International Monetary Fund's World Economic Outlook and is held fixed at each country's current group over the whole sample, so no country changes column inside the panel. The global series (`vix`, `ebp`) take one value per month, so the month effects absorb them and they are carried in the declared control set alone. Because each group column has fewer country clusters than the clustered asymptotics assume, the figure in square brackets is a wild cluster restricted bootstrap-t p-value, Rademacher weights, clustered on country, 9,999 draws. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1. P-values behind the stars use the standard normal reference for the clustered t-statistic in the Python build and a t-distribution with the smaller cluster count less one degrees of freedom in the Stata build.

Table C.16: How the newspaper fiscal-salience index is measured

VARIABLES Share	(1) $\rho_{i,t}^{sb,10Y}$ Reported	(2) $\rho_{i,t}^{sb,10Y}$ Alias-confirmed	(3) $\rho_{i,t}^{sb,10Y}$ Fiscal-bound numerator	(4) $\rho_{i,t}^{sb,10Y}$ ≥ 2 articles	(5) $\rho_{i,t}^{sb,10Y}$ ≥ 3 articles	(6) $\rho_{i,t}^{sb,10Y}$ ≥ 5 articles
$FSI_{i,t-1}$	0.1014***	0.0951***	0.3508***	0.1511***	0.1879***	0.2831***
	(0.0342)	(0.0286)	(0.0625)	(0.0470)	(0.0566)	(0.0745)
Attention control	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Macro controls	Yes	Yes	Yes	Yes	Yes	Yes
IE^{fw} control	No	No	No	No	No	No
Observations	11,802	11,473	11,802	10,794	9,784	8,090
Countries	44	44	44	44	44	44
R^2	0.398	0.403	0.404	0.412	0.425	0.447

Note: This table holds the specification of the main-text salience table's reported column fixed and changes only how the share is measured or which country-months enter. The dependent variable is the monthly stock-bond correlation, the non-overlapping correlation of daily equity and daily 10-year zero-coupon bond returns within a calendar month; a country-month enters when it carries at least 15 trading days. Articles are assigned to countries by a separate attribution pass of a language model, on both sides of the share, under attribution rule attr_v1.0. Column (1) is the reported column, on 11,802 country-months over 44 countries and 2000-02–2026-01. Column (2) keeps an article-country pair only where the alias table finds an alias of the country from its fixed table in the text, such as its name, adjective or demonym, its capital or a major city, its central bank, treasury or seat of government, or its sovereign bond nickname; 7.2 percent of the attributed pairs fail that test, and the column drops them from both sides of the share and from its attention control. Column (3) replaces the numerator with the classifier's own list of the countries whose fiscal position a flagged article discusses, over the reported denominator. Columns (4) to (6) keep the country-months whose lagged article count is at least 2, 3 and 5, 10,794, 9,784 and 8,090 country-months. Every column carries the attention control, the logarithm of one plus the article count its share divides by. The global series (`vix`, `ebp`) take one value per month, so the month effects absorb them and they are carried in the declared control set alone. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1. P-values behind the stars use the standard normal reference for the clustered t-statistic in the Python build and a t-distribution with the smaller cluster count less one degrees of freedom in the Stata build.

Table C.17: Leaving each country out in turn

	Estimate	Std. error	t	Countries
Panel A. Rating shifter, $Action^{F_C}_{i,t}$				
Full sample	0.1580**	(0.0635)	2.49	44
Smallest drop (leaving out GR)	0.0901	(0.0791)	1.14	43
Largest drop (leaving out ID)	0.1792***	(0.0565)	3.17	43
Drops losing significance at 5 percent	1 of 44			
Drops changing the sign	0 of 44			
Panel B. Newspaper fiscal salience, $FSI_{i,t-1}$				
Full sample	0.1014***	(0.0342)	2.97	44
Smallest drop (leaving out ES)	0.0844***	(0.0316)	2.68	43
Largest drop (leaving out ZA)	0.1082***	(0.0347)	3.12	43
Drops losing significance at 5 percent	0 of 44			
Drops changing the sign	0 of 44			

Note: The dependent variable is the monthly stock-bond correlation, the non-overlapping correlation of daily equity and daily 10-year zero-coupon bond returns within a calendar month. Each panel re-estimates one reported column of Section 6 on the sample with one country removed, once for every country the estimation frame carries, and reports the range those 44 estimates span. Panel A is the rating column of Table 7, 14,385 country-months over 1974-02–2026-01; Panel B is the salience column of Table 5, 11,802 country-months over 2000-02–2026-01. The specification, the controls and the fixed effects are the reported column's in each case, and no further restriction is applied. The counts of drops losing significance are fragile and the range of the point estimates is the informative half. With 44 clusters a two-way clustered standard error is estimated from few groups, so whether a single drop leaves a coefficient above or below a threshold is substantially a property of the variance estimator and not a statement about the country's data. The point estimates are what the exhibit is for: Panel A's 44 drops span 0.0892, which is 1.40 full-sample standard errors, and Panel B's span 0.0238, or 0.70. The country-by-country estimates are written to the build's cross-check directory. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1. P-values behind the stars use the standard normal reference for the clustered t-statistic in the Python build and a t-distribution with the smaller cluster count less one degrees of freedom in the Stata build.

C.4 Robustness for Section 7

Tables C.18 and C.19 re-estimate the credit decomposition of Table 8 with the distress exclusion tightened. The baseline excludes country-months whose lagged five-year spread exceeds 1,000 basis points, the level at which a sovereign contract prices a high probability of restructuring; that threshold is a choice, and these tables show what it does by tightening it to 500 and to 200 basis points, removing progressively more of the crisis-level spread observations. The credit coefficient rises from 0.055 at the baseline to 0.097 and then 0.143, and the difference between the credit and non-default shares clears one percent in every column of both tables.

Table C.18: Credit and non-default interest expense, spreads above 500 basis points dropped

VARIABLES	(1) $\rho_{i,t}^{sb,5Y}$	(2) $\rho_{i,t}^{sb,5Y}$	(3) $\rho_{i,t}^{sb,5Y}$	(4) $\rho_{i,t}^{sb,5Y}$	(5) $\rho_{i,t}^{sb,5Y}$	(6) $\rho_{i,t}^{sb,5Y}$
Credit	0.1213***	0.1215***	0.1267***	0.1192***	0.0903***	0.0973***
	(0.0192)	(0.0197)	(0.0206)	(0.0223)	(0.0155)	(0.0177)
Non-default	0.0344***	0.0257***	0.0236**	0.0326***	-0.0016	-0.0181
	(0.0079)	(0.0092)	(0.0095)	(0.0090)	(0.0121)	(0.0118)
Credit − Non-default [p]	0.0001	0.0000	0.0000	0.0010	0.0000	0.0000
Macro controls	No	Yes	Yes	Yes	Yes	Yes
VIX control	No	No	Yes	No	No	No
Country FE	No	No	No	No	Yes	Yes
Month FE	No	No	No	Yes	No	Yes
Observations	10,716	10,716	10,716	10,716	10,716	10,716
R^2	0.108	0.118	0.135	0.307	0.215	0.408

Note: Table 8's specification with the distress exclusion tightened: country-months whose lagged five-year spread exceeds 500 basis points are dropped, 132 further estimation rows against the 1,000 basis-point baseline, with all 44 countries retained. Layout, controls, fixed effects and the standard-error convention are those of Table 8, with one addition: the Credit − Non-default row reports the p-value of the reparameterized difference in every column. The credit coefficient rises as the cut tightens and the difference clears one percent in every column, so the channel ranking is not made by the threshold and the baseline is conservative for the comparison. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1.

Table C.19: Credit and non-default interest expense, spreads above 200 basis points dropped

VARIABLES	(1) $\rho_{i,t}^{sb,5Y}$	(2) $\rho_{i,t}^{sb,5Y}$	(3) $\rho_{i,t}^{sb,5Y}$	(4) $\rho_{i,t}^{sb,5Y}$	(5) $\rho_{i,t}^{sb,5Y}$	(6) $\rho_{i,t}^{sb,5Y}$
Credit	0.1690***	0.1688***	0.1825***	0.1727***	0.1273***	0.1430***
	(0.0324)	(0.0328)	(0.0322)	(0.0331)	(0.0349)	(0.0262)
Non-default	0.0332***	0.0229**	0.0186*	0.0247**	0.0068	-0.0084
	(0.0093)	(0.0096)	(0.0096)	(0.0098)	(0.0134)	(0.0134)
Credit − Non-default [p]	0.0003	0.0001	0.0000	0.0001	0.0010	0.0000
Macro controls	No	Yes	Yes	Yes	Yes	Yes
VIX control	No	No	Yes	No	No	No
Country FE	No	No	No	No	Yes	Yes
Month FE	No	No	No	Yes	No	Yes
Observations	9,843	9,843	9,843	9,843	9,843	9,843
R^2	0.073	0.086	0.116	0.317	0.176	0.405

Note: Table 8's specification with the distress exclusion tightened: country-months whose lagged five-year spread exceeds 200 basis points are dropped, 1,005 further estimation rows against the 1,000 basis-point baseline, with all 44 countries retained. Layout, controls, fixed effects and the standard-error convention are those of Table 8, with one addition: the Credit − Non-default row reports the p-value of the reparameterized difference in every column. The credit coefficient rises as the cut tightens and the difference clears one percent in every column, so the channel ranking is not made by the threshold and the baseline is conservative for the comparison. Standard errors, in parentheses, are two-way clustered by country and by month. *** p<0.01, ** p<0.05, * p<0.1.

C.5 The construction choices of Appendix B

Appendix B documents a set of discretionary screens, masks, exclusions and corrections, each declared and individually switchable. We test whether the results depend on them by rebuilding the dataset from the same raw inputs with the same code, with every declared screen, mask, exclusion and correction switched off and only the estimation floors kept, and re-estimating the paper's main specifications on the result.

Two features of the exercise matter for reading the tables. First, each column is a full rebuild, not a filter applied to the baseline panel. Switching off a screen changes which country-months clear the fifteen-day minimum, which changes which country-years clear the ten-month minimum, and the year-end yield inside $\mathrm{ie}^{\mathrm{fw}}$ moves with them. A specification estimated on a filtered panel would hold all of that fixed and would not be the counterfactual it appears to be. Second, no column here is a build without construction. A zero-coupon yield has to be derived from something, and every column still converts par quotes, fits curves, splices sources and aligns calendars. What the rebuild removes are the discretionary rules the appendix enumerates.

Tables C.20 through C.22 report the results. Differences from the baseline estimate are expressed in standard errors, not in percentages, because a percentage change in a coefficient of this size does not measure economic distance.

Table C.20: The monthly panel rebuilt without the discretionary interventions.

	(1) Pooled	(2) Country FE	(3) Month FE	(4) TWFE	(5) + Macro	(6) + Policy	(7) 5Y DV
A. Baseline							
$IE^{fw}_{i,t-1}$	0.0398***	0.0369***	0.0334***	0.0130***	0.0137***	0.0137***	0.0133***
(SE)	(0.0095)	(0.0119)	(0.0076)	(0.0030)	(0.0030)	(0.0029)	(0.0041)
Observations	14,282	14,282	14,282	14,282	14,282	14,282	13,903
B. Screens off							
$IE^{fw}_{i,t-1}$	0.0406***	0.0394***	0.0340***	0.0149***	0.0160***	0.0161***	0.0100***
(SE)	(0.0082)	(0.0104)	(0.0066)	(0.0037)	(0.0036)	(0.0036)	(0.0024)
Difference, in SE	+0.09	+0.21	+0.08	+0.64	+0.79	+0.85	-0.81
Observations	14,529	14,529	14,529	14,529	14,529	14,529	14,168

Note: The dependent variable is the non-overlapping monthly correlation of daily equity and daily ten-year zero-coupon bond returns. Columns (1) to (6) are estimated on a common sample within each panel, so a change across columns is the specification and never the sample; column (7) replaces the dependent variable with its five-year twin and is estimated on the rows that variable reaches. Standard errors are two-way clustered by country and by month. Every column is a full rebuild of Stage 1 under the named preset, and the difference row is the counterfactual estimate less the baseline, in standard errors of the baseline. Standard errors are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table C.21: The annual horse race rebuilt without the discretionary interventions.

	(1) IE^{fw}	(2) Debt/GDP	(3) Yield	(4) Saturated
A. Baseline				
$IE^{fw}_{i,t-1}$	0.0145***	—	—	0.0131*
	(0.0050)			(0.0074)
$(D/Y)_{i,t-1}$	—	0.0003	—	-0.0002
		(0.0006)		(0.0006)
$y^{10}_{i,t-1}$	—	—	0.0133**	0.0038
			(0.0062)	(0.0070)
Observations	1,757	1,757	1,757	1,757
B. Screens off				
$IE^{fw}_{i,t-1}$	0.0158**	—	—	0.0217**
	(0.0069)			(0.0087)
Difference, in SE	+0.26			+1.15
$(D/Y)_{i,t-1}$	—	0.0001	—	-0.0007
		(0.0007)		(0.0006)
Difference, in SE		-0.36		-0.88
$y^{10}_{i,t-1}$	—	—	0.0095	-0.0038
			(0.0071)	(0.0063)
Difference, in SE			-0.62	-1.09
Observations	1,801	1,801	1,801	1,801

Note: The dependent variable is the calendar-year correlation of monthly equity and monthly ten-year zero-coupon bond returns. IE^{fw} is the product of debt to GDP and the ten-year yield. All four columns within a panel are estimated on the common sample the saturated column defines, so the comparison across columns is the specification and not the rows. Every column carries country and year effects and the two macroeconomic controls; standard errors are two-way clustered by country and by year. A dash marks a regressor the column does not enter, which is not the same as an estimate that rounds to zero. Every column is a full rebuild of Stage 1 under the named preset, and the difference row is the counterfactual estimate less the baseline, in standard errors of the baseline. Standard errors are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table C.22: The monthly baseline on later samples, rebuilt without the discretionary interventions.

	Full	1990+	2000+	2010+	2015+
A. Baseline					
Pooled OLS: $IE^{fw}_{i,t-1}$	0.0399***	0.0387***	0.0330***	0.0354***	0.0504***
	(0.0096)	(0.0092)	(0.0082)	(0.0107)	(0.0061)
Country and month FE, macro controls: $IE^{fw}_{i,t-1}$	0.0138***	0.0145***	0.0171***	0.0156**	0.0317***
	(0.0030)	(0.0029)	(0.0037)	(0.0064)	(0.0065)
Observations	14,414	14,036	12,463	8,279	5,666
Months	653	433	313	193	133
B. Screens off					
Pooled OLS: $IE^{fw}_{i,t-1}$	0.0408***	0.0396***	0.0338***	0.0378***	0.0403***
	(0.0082)	(0.0078)	(0.0070)	(0.0098)	(0.0099)
Difference, in SE	+0.09	+0.09	+0.10	+0.23	-1.65
Country and month FE, macro controls: $IE^{fw}_{i,t-1}$	0.0161***	0.0170***	0.0181***	0.0176***	0.0182***
	(0.0036)	(0.0036)	(0.0038)	(0.0067)	(0.0052)
Difference, in SE	+0.79	+0.84	+0.28	+0.31	-2.10
Observations	14,661	14,281	12,641	8,332	5,692
Months	653	433	313	193	133

Note: The samples are cumulative and nested, so the table answers whether the result survives dropping early data and not when the relationship lives. The cut is on the dependent variable's date and the boundary is inclusive, so the 1990 column begins in January 1990 and uses the December 1989 regressor on its first row. Both specifications within a column are estimated on the same rows, so the two differ by model and not by sample. Every column carries all 44 countries. Standard errors are two-way clustered by country and by month. Every column is a full rebuild of Stage 1 under the named preset, and the difference row is the counterfactual estimate less the baseline, in standard errors of the baseline. Standard errors are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.