Asset Elasticities and Currency Risk Transfer*

Carol Bertaut

Ester Faia
Goethe University Frankfurt and CEPR

Sebnem Kalemli-Özcan
Brown University, NBER and CEPR

Camilo Marchesini
Brown University

Simon Paetzold
German Bundesbank

Martin Schmitz
European Central Bank

This draft: November 2026. First draft: February 2023.

Abstract

When a government borrows in more than one currency, the elasticity of investors' demand for its debt in each currency determines who bears currency risk and at what price. These elasticities remain unestimated. We extend the Koijen-Yogo asset-demand system to investor sectors disciplined by regulatory mandates and a sovereign that internalizes their demand, and estimate investor-sector-level elasticities for the currency segments of sovereign debt on the universe of euro-area holdings. Demand is more elastic for advanced-economy than for emerging-market issuers. Insurers and pension funds, usually regarded as sticky, respond most to the price of currency risk; nonbanks also absorb it; banks do not respond for emerging markets. In a complementary within-issuer design, U.S. investors, unlike euro-area ones, rebalance away from emerging-market bonds when the currency wedge widens. Emerging-market issuance falls as the currency wedge widens, such that investors' elasticity matters for emerging-market borrowing costs.

Keywords: Currency risk, FX and local currency bonds, banks and nonbanks.
JEL classification: F3, G2.

*The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Board and System, Bundesbank, European Central Bank or the Eurosystem. The paper relies on the confidential administrative datasets Security Holdings Statistics (SHS-Plus) for the euro area and the U.S. Treasury International Capital (TIC) survey, both of which have been accessed through a coauthor. Other proprietary and administrative datasets, used and matched in the analysis, have been accessed through the ECB (centralized returns database) and Goethe University Frankfurt (Consensus Forecasts). We thank Anusha Chari, Hanno Lustig, Federico Mainardi, Motohiro Yogo and Haonan Zhou for insightful comments and participants at seminars and conferences. Jakob Staps provided excellent research assistance.

E-mail addresses: *Corresponding author* Sebnem Kalemli-Özcan (sebnem_kalemli-ozcan@brown.edu). Carol Bertaut (Carol Bertaut passed away on May 25, 2026), Ester Faia (faia@wiwi.uni-frankfurt.de), Camilo Marchesini (camilo_marchesini@brown.edu), Simon Paetzold (simon.paetzold@bundesbank.de), Martin Schmitz (Martin.Schmitz@ecb.europa.eu).

1. Introduction

Sovereign governments issue bonds in more than one currency. A large share of that debt is held across borders by heterogeneous institutional investors: banks, insurers and pension funds, mutual funds and hedge funds. Whenever a bond is denominated in a currency other than the investor's own, someone must bear the currency risk: the investor, if not hedged, or the borrower through higher funding costs. Who bears the currency risk for how much compensation, and whether different currency bonds of the same issuer are easily substitutable, are questions that can be answered if one knows the elasticity of demand for bonds in different currencies. These elasticities are the central object of this paper and, to the best of our knowledge, they have not been estimated before. Using confidential regulatory holdings data from European and U.S. investors, we estimate them separately for emerging-market and advanced-economy sovereign issuers, who issue both in their own currency and in the investor's currency, the euro or the dollar.

There is a large literature that shows the importance of asset demand elasticities in closed and open economies as they are the essential input for a broad class of policy questions: the price impact of large-scale asset purchases and their unwinding, the yield consequences of capital-flow reversals and the change in sovereigns' funding costs when the investor base shifts under new regulations for global banks toward benchmark-driven nonbank holders, are just a few examples. Under infinitely elastic demand, the workhorse assumption of much of asset pricing, none of these shifts would move prices through the demand channel, holding fundamentals fixed. The asset-demand approach (Koijen and Yogo (2019), Gabaix and Koijen (2023), and references therein) has shown that this assumption is counterfactual and that finite elasticities are needed to explain observed prices.

Our main finding is that there is no single marginal holder of currency risk. Who absorbs it depends on the investor's regulatory mandate and on the issuer's endogenous supply response to required compensation by investors in general equilibrium. Demand for assets with higher currency risk is downward sloping for most investors—when the compensation for holding a bond in a given currency rises, investors who are free to take that risk move toward it. For advanced-economy bonds, every euro-area investor type does so. For emerging-market bonds, nonbanks and insurers and pension funds do, whereas banks, which regulation requires to cover currency exposure, show no significant response in our baseline and reduce their portfolio share when we compare bonds of the same issuer. U.S. investors differ from euro-area investors in their demand for emerging-market bonds: when

the currency wedge widens, they rebalance their portfolios away from these bonds. For advanced-economy bonds, the U.S. sample does not identify the response. The composition of the investor base, not a single marginal investor, therefore determines how currency risk is priced. Demand is more elastic for advanced-economy than for emerging-market bonds, and emerging-market sovereigns issue less when the price of currency risk rises.

Estimating these elasticities faces a hard identification problem, as prices and quantities are determined jointly in general equilibrium. We use a shift-share instrument derived from our structural model. From the standpoint of one investor sector, the demand of all other sectors shifts the price of currency risk along that sector's own demand curve. We measure this residual demand with the other sectors' portfolio shares in each issuer-currency segment, fixed at their values in a 2014 base window so that they are predetermined for every observation, and scaled by those sectors' current wealth in sovereign bonds, in the leave-out spirit of Koijen and Yogo (2019). The instrument varies across issuer-currency segments and over time. The elasticity it identifies is therefore that of a sector's demand for a segment, namely the local- or investor-currency debt of a sovereign, and not the elasticity of substitution among the nearly identical bonds inside a segment. Identification rests on the predetermined shares, whose balance on security characteristics we test. The specification conditions on the model's aggregate states rather than on time fixed effects, which a segment-level instrument cannot accommodate. Because we let the issuer's supply respond, the same instrument identifies the supply elasticity, and the estimates are general-equilibrium objects rather than the partial-equilibrium response to a single shock to currency risk.

Three features distinguish our paper from the existing literature. First, the object. Our estimand is the response of a sector's holdings of a sovereign's debt in a given currency to the currency wedge on that issuer-currency segment, estimated on the local-currency debt of each issuer group and also on its local- and investor-currency debt. Existing demand systems aggregate portfolios by country and cannot separate the currency of denomination from the issuer. Second, the data. We use the security-by-security holdings of euro-area investor sectors recorded in the ECB's Securities Holdings Statistics, and the U.S. counterpart from the Treasury International Capital (TIC) survey. Third, the structure. The estimating equations are derived from the optimality and market-clearing conditions of a model with mandate-constrained heterogeneous investors and an optimizing sovereign, such that each estimate maps into a structural parameter and a single instrument

identifies both the investor sectors' demand elasticities and the supply elasticity through the sovereign's issuance response.

The model has three blocks. On the demand side, investors in each sector allocate their portfolio across the issuer-currency segments in which sovereigns borrow. They maximize expected excess returns subject to their mandate: the risk of the portfolio, measured as the sector's regulation measures it, must be covered by the sector's loss-absorbing capital. The constraint is modeled on the sector's actual regulatory requirement, and it makes each sector's demand finitely elastic. Market clearing aggregates the sector demands into a downward-sloping price schedule, whose slope is a holdings-weighted average of the sector frictions, such that it moves with the composition of the investor base without any shock or change in regulation. On the supply side, we extend the sovereign-debt model of Eaton and Gersovitz (1981): the sovereign chooses issuance in both currencies and internalizes this schedule. The resulting uncovered interest parity (UIP) wedge has three terms, namely expected depreciation and the currency risk premium,[1] a term proportional to the outstanding stock of local-currency debt, scaled by demand, and default risk. The empirical design follows from the model. Each sector expands its holdings in a segment until the wedge equals the marginal cost of the position, which is the regulatory weight on the segment times the shadow price of the sector's capital. A first-order approximation of this condition gives our second stage, a regression of a sector's log holdings of a security on that security's currency wedge. Market clearing, which links the residual demand of the other sectors to the wedge, gives the first stage. The sovereign's optimality condition, which equates the marginal revenue of one more unit of local-currency debt, reduced by the slope of demand and by the default premium, to the expected cost of repaying it, gives the supply regression.

Our demand system shares the structure of Koijen and Yogo (2019) and Koijen et al. (2021): heterogeneous price-taking investors whose demands are disciplined by a quadratic risk term and, together with market clearing, determine prices. The difference is where the quadratic term comes from. In their framework, it is the idiosyncratic return variance. Idiosyncratic return variance is a security-level residual that a diversified holder diversifies away, such that at the level of

1 The spot exchange rate is pinned down by the ratio of domestic and foreign stochastic discount factors as in Lustig, Roussanov and Verdelhan (2011). Investors also choose how much of the exposure to hedge, which implies that the covered interest parity (CIP) wedge replaces the currency risk premium with the cost of hedging, documented by De Leo, Keller and Zou (2024). Our focus is on currency risk transfer, and we treat the covered extension in Appendix S6.

an issuer-currency segment it would not, on its own, discipline demand across the currencies of a given issuer. The regulatory measure instead loads on the currency of denomination, which every security in the segment shares. Diversifying across securities removes their idiosyncratic residuals but not the common currency exposure. That non-diversifiable exposure is what makes cross-currency demand finitely elastic, and its slope is what we estimate. It is also consistent with insurers and pension funds, which are sticky in the aggregate demand system of Koijen et al. (2021), displaying a sizeable elasticity at the level of the currency segment in our setting.

Euro-area investors have downward-sloping demand for advanced-economy and emerging-market issuers alike. For advanced-economy issuers, every euro-area investor class absorbs currency risk rather than shedding it. For emerging-market issuers, insurers and pension funds and nonbanks do: euro-area investors hold non-euro (and non-dollar) sovereign bonds, from Brazil to Canada, for the currency risk premium they earn. Elasticities are larger for advanced-economy issuers, with a euro-area aggregate of about 25 against roughly 7 to 12 for emerging markets, consistent with advanced-economy local-currency bonds being close substitutes in deep currency markets.[2] Euro-area banks are the one sector whose demand for emerging-market bonds shows no significant response, probably due to bank regulation requiring currency exposure to be covered.

We complement the segment-level estimates with a within-issuer design, which is also the only design that the U.S. data support. It regresses a sector's portfolio share in a security on that security's currency wedge with issuer-by-time fixed effects, and instruments the wedge with a security-level shift-share. The fixed effects absorb every shock common to an issuer in a quarter, including the aggregate shocks that the segment-level design can only control for, so the comparison is among bonds of the same sovereign, in local and in investor currency. The price is a narrower object, namely substitution among the securities of one issuer, and an instrument whose shares are measured when a security is first held rather than in a common base period.[3] For the euro area, this design confirms that insurers,

2 These magnitudes are in line with the finite, and often low, demand elasticities documented in the asset-demand literature, although the objects differ. Koijen and Yogo (2019) estimate own-price demand elasticities for equities held by U.S. institutions. Closest to our setting, Koijen and Yogo (2026) estimate a global demand system on portfolios aggregated at the country level and report demand elasticities, averaged across years and issuer countries, of 27.9 for short-term debt and 3.2 for long-term debt. Our euro-area aggregates, estimated from security-level holdings by sector on bonds with a residual maturity of 1 to 5.5 years, lie between these two values. All point in the same direction: demand is far more inelastic than the frictionless benchmark.

3 The TIC holdings are annual, recorded at the custodial holder of record rather than by

pension funds and nonbanks raise the share of emerging-market bonds in their portfolio as the currency wedge widens, and it shows that banks reduce theirs. For advanced-economy issuers the within-issuer first stage is weak and the design is uninformative. U.S. investors do the opposite of euro-area insurers and nonbanks: U.S. nonbanks, banks and other investors rebalance away from emerging-market sovereign bonds as the currency wedge widens. Since the dollar is the global funding currency, U.S. investors push emerging-market currency risk back onto the issuer, while euro-area insurers, pension funds and nonbanks hold it for compensation.

To illustrate the importance of our elasticities, we run within the model two policy counterfactuals and show their results in Appendix S8. A prudential loosening of euro-area insurers under Solvency II and a euro-area quantitative tightening, each of which moves emerging-market borrowing costs. The model-consistent magnitudes are obtained with the identified elasticities entering through the composition of the investor base. For instance, applying a single aggregate elasticity uniformly overstates the prudential effect by roughly a factor of eight and misses that the same shock is absorbed very unevenly across sectors.

The rest of the paper is structured as follows. Section 2 reviews the literature. Section 3 presents data and descriptive statistics. Section 4 explains the model. Section 5 presents the empirical strategy based on the model. Section 6 presents the results, including the within-issuer evidence and the comparison with U.S. investors. Section 7 concludes. The Appendix contains the proofs. Sections, tables, figures and equations whose number carries the prefix S are in the Supplemental Material, and those with the prefix D are in the Data Documentation.[4]

2. Related Literature

Our study contributes to the literature on asset demand systems with a new object and an identification strategy grounded in a structural model. The object is the elasticity of demand for the debt of a sovereign in each of the currencies in which

end-investor sector at the issuer-currency segment, and available to us only through 2022. These features rule out the quarterly, currency-segment-level specification used for euro-area investors.

4 This paper supersedes an earlier version under the same NBER Working Paper 34275. Relative to that version, the baseline specification is now derived from the model, the dependent variable is log holdings rather than the log portfolio share, following Koijen et al. (2021), the instrument is built from segment shares fixed in a 2014 base window rather than from security-level shares measured when a security is first held, and the euro-area sample runs through 2025Q4 rather than 2022Q2. We retain the earlier estimates in full in subsection 6.4, where they serve as the within-issuer evidence and as the basis for the comparison with U.S. investors.

it borrows, by investor sector, which is what identifies which investors bear and price the currency risk. The strategy is a shift-share instrument that the model delivers as the residual demand of the other investor sectors.

The asset demand system literature (Koijen and Yogo (2019), Koijen et al. (2021), Gabaix and Koijen (2023), Koijen and Yogo (2026)) starts from the premise that investors' demands have finite elasticities, which can be heterogeneous across investors, and that finite demand combined with market clearing determines asset prices. To overcome the endogeneity problem, a large part of the literature instruments excess demand with the idiosyncratic shocks of large investors—the granular instrumental variable of Gabaix and Koijen (2020). The underlying assumption is that large investors move the market. The literature on exchange rates suggests for international securities things may be different and exploiting variation from the full investors' heterogeneity may provide more information. Barbiero et al. (2026) show, for instance, that the small FX positions of broker-dealers can move currency returns and covered interest parity (CIP) deviations, and Rey et al. (2024) find that equity investors' positions, small relative to those of bond investors can move the exchange rate.

Our approach is data driven. As we cover the universe of euro-area and the U.S. investors, composed of mutual funds, hedge funds, pension funds, insurers, governments, and other corporate and household sectors, at the security level, we do not need to assume apriori whether large or small investors or broker-dealers vs insurance/pension funds or banks are moving the market. We can measure this in the data by estimating the demand of each investor sector for each currency segment. Our shift-share instrument proxies excess demand for bonds in each currency by apportioning to all investors except the one whose demand is being estimated a counterfactual change in their holdings.[5]

A recurring theme of this literature is that the elasticity one estimates depends on the level of aggregation. Chaudhary, Fu and Li (2023) show for corporate bonds that demand is far more elastic among close substitutes than across broad portfolios, and Haddad et al. (2025) state the conditions under which a cross-sectional regression with fixed effects recovers a relative elasticity. We report estimates at two levels. The baseline is identified across issuer-currency segments and over time. It is closest to the elasticities of Koijen and Yogo (2026), who estimate a global demand system on portfolios aggregated at the country level. We disaggregate that object along three margins that country-level data cannot

5 The closest paper to us, in empirical design, is Koijen and Yogo (2019), but they do not focus on currencies.

capture: the investor sector within a country, the currency of denomination within an issuer, and the issuer's own supply response. The second set of estimates is identified among the securities of one issuer in one quarter, and is closer to the security-level elasticities of the corporate-bond literature.

Our instrument belongs to the family of shift-share designs built from the demand of other investors. Koijen and Yogo (2019) instrument prices with the counterfactual demand implied by the investment universes of other investors. Koijen et al. (2021) use Eurosystem purchases apportioned across countries. Flow-induced trading designs (Coval and Stafford (2007), Lou (2012)) apportion fund flows by lagged portfolio weights. Ours apportions the residual sectors' wealth in sovereign bonds by their segment shares, fixed in a base period. Following Goldsmith-Pinkham, Sorkin and Swift (2020) and Borusyak, Hull and Jaravel (2025), identification rests on the shares, and we test their balance on security characteristics.

We provide foundations for the instrument through a model of debt issuance in the tradition of Eaton and Gersovitz (1981), extended to heterogeneous investors with finitely elastic demand, due to regulatory constraints or clientele mandates, for debt denominated in domestic and foreign currency. The object of interest is the elasticity of demand for the bonds of a sovereign in each currency of denomination, and we document in the data that both emerging-market and advanced-economy sovereigns issue in more than one currency. Our theoretical foundations are consistent with Greenwood et al. (2023), Stein and Wallen (2025), and Chaudhary, Fu and Li (2024). On the supply side, we share with Moretti et al. (2024) the idea of embedding finitely elastic demand in an Eaton and Gersovitz (1981) issuance problem. They estimate the slope of active demand from index-rebalancing flows in dollar-denominated emerging-market sovereign bonds and calibrate a sovereign-default model to Argentina. The key difference is that our sovereign issues in more than one currency, consistent with the data, and investors choose which currency to hold together with how much to hedge.[6] Furthermore, our estimates are from the universe of securities from two large areas. On the measurement of the wedge, Du and Schreger (2016) show that for emerging-market issuers the covered wedge computed from bond yields also reflects credit risk that covaries with the currency.

6 We also limit the payoff contamination that Moretti et al. (2024) and Koijen and Yogo (2026) emphasize by restricting estimation to the short and medium maturities over which supply is plausibly predetermined. Our panel leans toward investment grade, such that default risk is remote. A shift in creditor demand is borne mainly by the currency premium rather than by default risk or the quantity of borrowing.

For this reason our baseline uses the uncovered wedge and conditions on credit quality, and we report the covered specifications in subsection S6.1.

Our model is related to the literature in which financial intermediaries with limited risk-bearing capacity price currency risk. In Gabaix and Maggiori (2015) and Itskhoki and Mukhin (2021), intermediaries absorb currency imbalances at a price that rises with their position, which delivers a downward-sloping demand for currency risk governed by a single aggregate friction. In that view, a representative intermediary prices currency risk. In ours, the price of currency risk is set jointly by the mandates of heterogeneous investor sectors and by the issuer, who internalizes their demand when choosing how much to borrow in each currency. We measure the response of each sector in the data and show how the composition of the investor base determines the aggregate friction that this literature takes as a primitive. Maggiori, Neiman and Schreger (2020) document, with security-level holdings of mutual funds, that investors tilt strongly toward bonds denominated in their own currency and in dollars. For sovereign bonds held by euro-area investors, we do not find such a bias. These investors hold about the same share of an issuer's local-currency debt as of its total debt (Figure 1), and they expand their holdings of a currency segment when the compensation for its currency risk rises. In our data, the currency of denomination is a priced margin of demand rather than a fixed preference. On the issuer's side, Ottonello and Perez (2019), Engel and Park (2022), and Du, Pflueger and Schreger (2020) study the sovereign's choice between local- and foreign-currency debt, which trades off the hedging benefit of local-currency debt against the temptation to inflate or depreciate it away. In our model the currency composition of issuance also responds to the price that mandate-constrained investor sectors require for currency risk in each segment, which differs across sectors and moves with the composition of the investor base.

Our study also relates to the growing literature that uses confidential regulatory data on securities holdings to study international asset pricing and its macro consequences, such as Fang, Hardy and Lewis (2022), Faia, Lewis and Zhou (2025), Bertaut, Bruno and Shin (2021). We have a direct antecedent in the earlier portfolio-balance literature of Friedman (1977), which rested on the same premise as the modern asset-demand approach, with assets being imperfect substitutes, such that shifts in their relative supplies or demands move relative yields rather than being absorbed at unchanged prices.

3. Data, Measurement, and Descriptive Statistics

This section describes the institutional setting, the data, and the terms we use throughout the paper. A longer description of the data is in Appendix D9.

We start with less well-known dataset of the European Central Bank's Securities Holdings Statistics (SHS), the ECB's administrative dataset on securities holdings.[7] It records the portfolio holdings of euro-area investors in the rest of the world. The SHS is a harmonized security-by-security collection of the direct portfolio holdings of investors domiciled in the euro area, for all asset classes, such as government bonds, equities, mutual fund shares, covered bonds, and others. The harmonized collection started in the last quarter of 2013 and continues to the present.[8] The data are quarterly and follow the residence principle for both issuer and investor. Our sample is 2014Q1 to 2025Q4. Reporting is mandatory and legally enforced; we describe the administrative details in Appendix D9.

The data record the holder of each security by country of domicile and by investor sector. Table D15 lists the sector coverage of the SHS and our classification of each sector into the four groups we use: monetary financial institutions (MFI), excluding the European Central Bank and the national central banks; insurance corporations and pension funds (ICPF); nonbank financial intermediaries (OFI in the SHS classification, which we call nonbanks), which include mutual funds and hedge funds; and a residual category of other investors (OTH), which consists of households, nonfinancial corporations, and the government sector. In the text we refer to the four groups as banks, insurers and pension funds, nonbanks, and other investors, and we use the acronyms only in tables.

The investors in our data choose their portfolios under different mandates, namely the regulations, institutional arrangements, and clientele preferences that constrain how each sector takes risk: Solvency II for insurers and pension funds, the Basel framework for banks, and lighter prudential rules for investment funds. These mandates are why demand elasticities differ across investor-sector, and why we estimate the demand system sector by sector rather than for a representative investor. Appendix S4 describes them.

Variable Construction and the Sample of Analysis. To compute arbitrage deviations, we match the securities' returns with Bloomberg data on forward, swap, and spot exchange rates, and with quarterly exchange-rate expectations

7 The better-known TIC data is described below.

8 National central banks have collected the data since 2003 at monthly frequency. The harmonized dataset needed for our analysis starts in 2013.

from Consensus Forecasts. We include in the sample only countries for which all of these variables are available. We also require that the foreign exchange market is unified at the end of each quarter, that the exchange rate is floating or a managed float under the classification of Ilzetzki, Reinhart and Rogoff (2022), and that a positive fraction of sovereign bond issuance is denominated in both euro and local currency throughout the period. Finally, we keep only countries whose sovereign securities are held by euro-area investors for an average of at least one billion euros at face value throughout the period. Details are in Appendix D10.

The sample has 13 emerging-market (EM) issuers: Brazil, Chile, Colombia, India, Israel, Malaysia, Mexico, Poland, Russia (excluding 2014Q4 to 2015Q4 and 2022Q1 onward), South Africa, South Korea, Thailand, and Turkey (until 2018Q2), and 8 advanced-economy (AE) issuers: Australia, Canada, Switzerland (after 2015Q1), the United Kingdom, Japan, Norway, New Zealand, and Sweden.[9]

Denote the currency wedge on security b with residual time to maturity τ at time t by $x_{b,t,t+\tau}$. The wedges are:

$$x^{UIP}_{b,t,t+\tau} = r_{b,t,t+\tau} - r^{h}_{b,t,t+\tau} - \frac{1}{\tau}\left(e^{\exp}_{t,t+\tau} - e_t\right), \quad h \in \{EA, US\} \qquad \text{(UIP)}$$

$$x^{CIP}_{b,t,t+\tau} = r_{b,t,t+\tau} - r^{h}_{b,t,t+\tau} - \rho_{t,t+\tau}, \quad h \in \{EA, US\} \qquad \text{(CIP swaps-based)}$$

where $r_{b,t,t+\tau}$ $\left(r^{h}_{b,t,t+\tau}\right)$ is the return in local (investor home) currency that accrues to the holder of the security at maturity, e_t is the spot exchange rate in units of the foreign currency, $e^{\exp}_{t,t+\tau}$ is the expected spot exchange rate at maturity, and $\rho_{t,t+\tau}$ is the annualized cost of hedging the exchange rate over the same horizon, measured from cross-currency swap contracts. The forwards-based CIP wedge reported in Table S6 replaces it with the annualized forward premium from forward contracts as opposed to swap contracts. In the estimation we suppress the horizon and write $x_{b,t}$ for the wedge of security b in quarter t, and $\tau_{b,t}$ for the residual maturity of the security.

3.1. Facts in International Portfolios from Administrative Data

Figure S2 uses the SHS micro data to show the share of debt securities in the portfolios of each investor type, broken down by issuer type.[10] The issuers are governments, insurance companies, banks, nonfinancial corporations, and others.

9 We also observe security characteristics other than the currency of denomination, such as maturity and rating, which we use as controls.

10 The classification of the issuing sector is the same as the classification of the investor sector.

Most euro-area holdings of EM securities are sovereign bonds (left panel), whereas most holdings of AE securities are bank bonds (right panel).

We focus on sovereign bonds for two reasons. First, they are the largest class of emerging-market securities in euro-area portfolios, and a sizeable one for advanced economies. Second, we want to measure how investors shift across currencies and within issuers, which requires issuers that borrow in more than one currency.

Figure S3 shows the shares of sovereign bonds held by the four investor types: banks, nonbanks, insurance, and other. For both EM and AE issuers, euro-area nonbanks, mainly mutual funds and hedge funds, hold the largest share. Banks come next, overtaken by insurers for emerging-market issuers in the later years of the sample. This heterogeneity is what our identification design needs. Because sovereign holdings are spread across distinct investor sectors rather than concentrated in one, the question of which sector absorbs currency risk is well posed, and each sector can be assigned its own demand elasticity. The dispersion in shares across sectors is also what gives the shift-share instrument its identifying variation. And since a segment's aggregate elasticity is the holdings-weighted average of the sector elasticities, the large share of nonbanks means that they dominate the aggregate response.

Finally, Figure 1 shows the share of total externally held sovereign debt outstanding that euro-area investors hold. Euro-area investors hold roughly 7 to 16 percent of AE externally held sovereign debt and 19 to 26 percent of EM externally held sovereign debt. For externally held debt denominated in the issuer's own currency, the shares are about 7 to 16 and 16 to 22 percent.

4. A Structural Model of Asset Demand and Supply

This section lays out the model of sovereign bond demand and supply that delivers the specifications we estimate, the instrument we use, and the counterfactuals we run in Appendix S8.

The model has three blocks: investor demand, market clearing, and sovereign optimal supply. From now on, a wedge is the excess return on a currency-segment, namely what a unit of the bond is expected to pay, discounted and adjusted for default and, for a local-currency claim, for expected depreciation, minus its price. On the supply side, the issuer chooses issuance internalizing the market-wide

Figure 1
Coverage of SHS dataset over total externally held sovereign debt, by issuer country group and by currency of debt denomination.
See Appendix D11 for detailed notes to this figure.

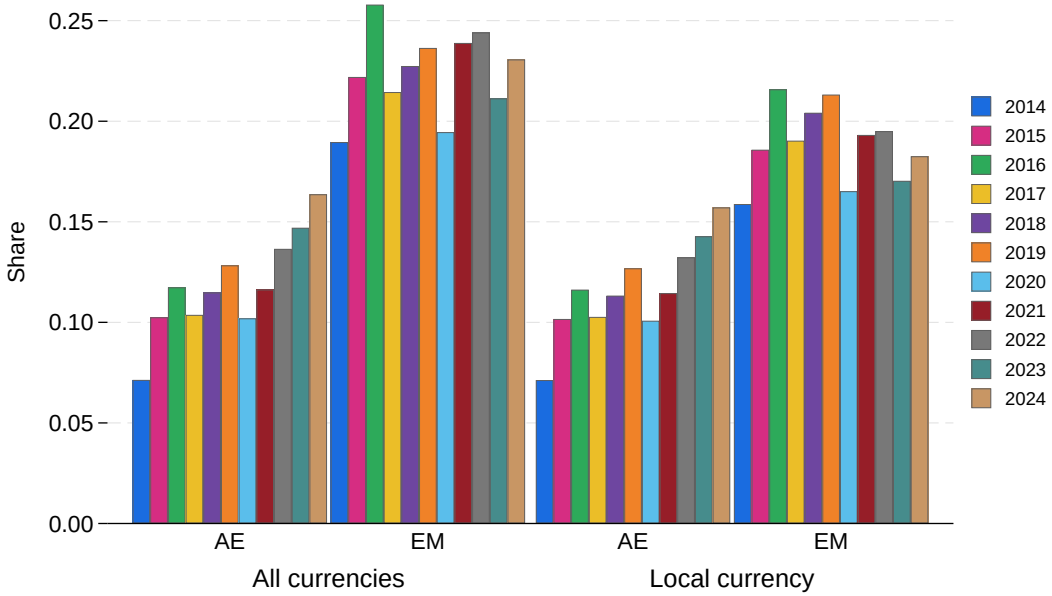

Notes: Share of SHS holding amounts in total externally held amount outstanding, by country sample and currency of denomination.

downward-sloping demand schedule it faces. This is the notion of market discipline we work with, in the tradition of Eaton and Gersovitz (1981) and Bulow and Rogoff (1989), whereby the government is disciplined not only by the default premium investors require, as in the standard formulation, but also by the price impact of its own issuance on demand from different investors. As in the data, a single government issues in more than one currency. The equilibrium has a Stackelberg structure in which the government internalizes investors' demand.

The demand block also differs from the intermediary-friction models used in the literature. First, the aggregate friction on a segment is a composite of sector-level frictions weighted by holdings. Second, each sector-level friction is modeled on the sector-specific regulatory requirement. This adds a role for the composition of the investor base, which gives a structural interpretation to the elasticities we estimate. Each institutional investor has one pool of loss-absorbing equity capital, which has to cover the risk of its positions across all segments.

4.1. The Investors' Optimization Problem

For each segment, the excess return per unit held is the expected discounted payoff, adjusted for default and, in the local-currency case, for expected depreciation, minus the price paid:

$$x_\ell \;\equiv\; \frac{1-\delta}{1+r}\,\frac{e}{E(e')} \;-\; q, \qquad x_* \;\equiv\; \frac{1-\delta}{1+r^*} \;-\; q^*. \tag{1}$$

x_ℓ is the currency wedge we estimate, and by (20) it is the model's counterpart of the measured UIP deviation. Expected profits from the portfolio are linear in positions: $\Pi_i \;=\; \sum_g x_g\, B_{i,g} \;=\; x_\ell\, B_{i,\ell} + x_*\, B_{i,*}$. The sector maximizes profits, subject to its mandate constraint:

$$\max_{\{B_{i,g}\}} \;\sum_g x_g\, B_{i,g} \;-\; \gamma_i^0\, R_i \;-\; \frac{\chi_i}{2}\big(\max\{0,\; R_i - \bar{K}_i\}\big)^2, \qquad R_i \;=\; \tfrac{1}{2}\,\mathbf{B}_i^\top\,\Omega_i\,\mathbf{B}_i, \tag{2}$$

where $\mathbf{B}_i = (B_{i,g})_g$ is the portfolio the sector chooses, Ω_i a positive semi-definite matrix of mandate weights, and $R_i = \frac{1}{2}\mathbf{B}_i^\top\Omega_i\mathbf{B}_i$ its risk-weighted assets, priced at the internal rate $\gamma_i^0 > 0$ and penalized, with intensity $\chi_i \geq 0$, once they exceed the loss-absorbing capital $\bar{K}_i$.[11] The regulatory hard limit $R_i \leq \bar{K}_i$ is the $\chi_i \to \infty$ limit of the penalty, which plays the same role as the variance term in the mean-variance demand of Koijen and Yogo (2019). A change in the mandate tightens the problem either because capital shrinks, lowering $\bar{K}_i$, or because regulation attaches more weight to a segment, raising the corresponding entry of Ω_i; the second is our Solvency II counterfactual from Appendix S8.

As in the actual risk-based regulation, the constraint in (2) is quadratic in holdings. Every regime that binds the sectors in our sample caps a risk measure of the portfolio. We can write the measure more generally as $\rho_i(\mathbf{B}_i)$ and the requirement as $\rho_i(\mathbf{B}_i) \leq \kappa_i$. These measures are all positive multiples of the portfolio standard deviation: $\rho_i(\mathbf{B}_i) \;=\; k_i\,\big(\mathbf{B}_i^\top\Sigma_i\,\mathbf{B}_i\big)^{1/2}$, with k_i being the confidence multiplier of the regime and Σ_i the relevant covariance matrix.[12] The supervisory formulas are usually written in standard-deviation form.

For example, the Solvency II standard formula, which aggregates the capital charges on the currency segments with a correlation matrix and bounds the result by own funds, is of this form. subsection S3.1 gives the mapping, including the diagonal case $\omega_{i,g} = 2\,a_{i,g}^2$ used in the quantitative model, where $a_{i,g}$ is the per-unit regulatory stress on segment g.

11 This is eligible own funds for an insurer, common equity for a bank, the funding surplus for a pension fund, the risk limit in units of net asset value for an asset manager.

12 A cap on the standard deviation is equivalent to a cap on the square. Squaring the requirement gives $\mathbf{B}_i^\top\Sigma_i\mathbf{B}_i \leq (\kappa_i/k_i)^2$, which is exactly (2) with $\Omega_i = 2\Sigma_i$ and $\bar{K}_i = (\kappa_i/k_i)^2$.

The same form applies to the regulatory constraints of the other sectors. Table S8 in Appendix S4 shows the mapping across sectors. In each case the binding rule is a limit on a standard-deviation-based risk measure of the same form, so that squaring it delivers $(\Omega_i, \bar{K}_i)$ as in (2). What differs across sectors is the matrix and the weights, which captures investors' heterogeneity, where Ω_i comes from the regulation and $\bar{K}_i$ from reported capital or coverage data. Although outside the scope of this paper, such constraints can also be microfounded as the choice of a regulator who minimizes a systemic loss function, in the tradition of Koehn and Santomero (1980) and Rochet (1992).

Optimal Demand. Given the problem (2), we can derive the investors' optimal demand. Let $\mu_i \equiv \gamma_i^0 + \chi_i \max\{0,\, R_i - \bar{K}_i\}$ be the marginal cost of risk-weighted assets. The first-order condition for segment g is:

$$x_g \;=\; \mu_i\, \omega_{i,g}\, B_{i,g} \qquad \Longleftrightarrow \qquad B_{i,g} \;=\; \frac{x_g}{\gamma_{i,g}}, \qquad \gamma_{i,g} \;\equiv\; \mu_i\, \omega_{i,g}, \tag{3}$$

or, in vector form, $\mathbf{B}_i = \mu_i^{-1}\Omega_i^{-1}x$, where $x = (x_g)_g$ are the excess returns of (1). Demand is linear in the wedge with slope $1/\gamma_{i,g}$. The friction $\gamma_{i,g}$ has two components. The weight $\omega_{i,g}$ is a primitive capturing the regulatory mandate and varies across segments. The shadow price μ_i is common to all of the sector's positions and is endogenous. It measures the tightness of the regulatory capital constraint, and, as such, it is the channel through which an aggregate shock or a shock in one segment reaches the sector's demand in every other segment.[13] The investor optimization problem, and the related optimal demand, stated here relates to unhedged excess returns. In Appendix S6 we extend it to the case with hedged excess returns, hence a CIP.

The Shadow Price and the Occasionally Binding Constraint. Since the optimal demand depends on the shadow price, μ_i, we now solve for it. Substituting (3) into the constraint gives the sector's risk as a function of prices:

$$R_i = \frac{\mathcal{X}_i}{2\,\mu_i^2}, \qquad \mathcal{X}_i \equiv x^\top \Omega_i^{-1} x = \sum_g \frac{x_g^2}{\omega_{i,g}} \quad \text{(diagonal case)}. \tag{4}$$

The marginal cost of risk-weighted assets, together with (4), pins down the shadow price: μ_i equals the internal rate γ_i^0 while the mandate is slack and rises above it once R_i exceeds $\bar{K}_i$:

$$\mu_i = \gamma_i^0 + \chi_i\, \max\{0,\; R_i - \bar{K}_i\}. \tag{5}$$

The price at which risk-weighted assets equal the ceiling, $R_i = \bar{K}_i$, is $\mu_i^{\text{hard}} =$

13 A model in which each sector pays an unrelated quadratic cost per segment is the special case in which μ_i is held fixed, and it has no such channel.

$(\mathcal{X}_i/2\bar{K}_i)^{1/2}$, which is the limit of a hard wall obtained as $\chi_i \to \infty$. At the other extreme, $\chi_i = 0$ leaves the unconstrained constant friction $\gamma_{i,g} = \gamma_i^0 \omega_{i,g}$.[14]

The kink matters for the estimated elasticities. While the mandate is slack the semi-elasticity is constant at $1/(\gamma_i^0 \omega_{i,g})$; once it binds μ_i rises with the wedge, demand becomes concave, and the semi-elasticity falls as the wedge widens. An estimated slope is therefore local and state-dependent, smaller in states characterized by wide wedges. This implies that a sector's elasticity shifting in response to shocks reflects how binding the regulatory constraint is, rather than a change in preferences.

Market Clearing and Aggregate Demand. The issuer takes the aggregate demand of all sectors as given. Each sector holds a continuum of price-taking investors, and market clearing in segment g reads as follows: $\sum_{i \in \mathcal{I}} m_i \, B_{i,g} = B_g$, with B_g the supply chosen by the government. Substituting the sector demands (3) into the segment market clearing and solving for the wedge gives the inverse aggregate demand:

$$x_g = \Gamma_g \, B_g, \qquad \frac{1}{\Gamma_g} = \sum_{i \in \mathcal{I}} \frac{m_i}{\gamma_{i,g}}. \tag{6}$$

For the local-currency bond this also sets the price $q = \frac{1-\delta}{1+r} \frac{e}{E(e')} - \Gamma_\ell \, B'$, so the UIP wedge the issuer faces, the gap between the domestic bond and the foreign risk-free bond $q^{*,rf}$ brought to local terms at the expected depreciation, is:

$$\begin{aligned} x_\ell^{\text{UIP}} \;\equiv\; \frac{e}{E(e')} \, q^{*,rf} - q &= \frac{e}{E(e')} \frac{1}{1+r^*} - \left(\frac{1-\delta}{1+r} \frac{e}{E(e')} - \Gamma_\ell \, B' \right) \\ &= \frac{e}{E(e')} \left(\frac{1}{1+r^*} - \frac{1-\delta}{1+r} \right) + \Gamma_\ell \, B'. \end{aligned} \tag{7}$$

The friction term $\Gamma_\ell B'$ is the contribution coming from the demand side; the risk-free and default terms price the two bonds. This is the schedule the issuer internalizes and that the empirical section measures, and its econometric specification accounts for the issuer's default risk.

Proposition 1. Aggregation. *Let sector demands be given by (3) and let each segment clear according to segment market clearing condition. Then the inverse demand facing the issuer in segment g is linear in the quantity issued:*

$$x_g \;=\; \Gamma_g \, B_g, \qquad \Gamma_g \;=\; \left(\sum_{i \in \mathcal{I}} \frac{m_i}{\mu_i \, \omega_{i,g}} \right)^{-1} \;=\; \left(\sum_{i \in \mathcal{I}} \frac{m_i}{\gamma_{i,g}} \right)^{-1}, \tag{8}$$

14 Both parameters are observable in the data. For instance, under the Solvency II mapping, with $\bar{K}_i = \text{OF}_i^2$ and $R_i = \text{SCR}_i^2$, the buffer ratio $\gamma_i^0/\mu_i^{\text{hard}} = (\bar{K}_i/R_i)^{1/2} = \text{OF}_i/\text{SCR}_i$ is the solvency coverage ratio, such that $\gamma_i^0 = \mathcal{X}_i^{1/2}/(\sqrt{2}\,\text{SCR}_i)$, pinned down from own funds and the observed coverage ratio.

and the aggregate semi-elasticity of segment g is the holdings-weighted average of the sector semi-elasticities,

$$\beta_{1,g} \;=\; \sum_{i\in\mathcal{I}} h_{i,g}\,\beta_{1i,g}, \qquad h_{i,g} = \frac{m_i\,B_{i,g}}{B_g}, \qquad \beta_{1i,g} = \frac{1}{\gamma_{i,g}\,\bar{D}_{i,g}}, \qquad \sum_{i\in\mathcal{I}} h_{i,g} = 1. \tag{9}$$

The proof is in Appendix A.

Three properties of (8) and (9) describe the channels. First, the aggregate friction is not a primitive. At unchanged mandates it moves with the composition of the investor base $\{m_i\}$, such that a shift toward sectors with high $\gamma_{i,g}$ steepens the schedule the sovereign faces without any regulation having changed. Second, it is segment-specific. A sector that is absent from a segment, or that carries a large $\omega_{i,g}$ there, does not discipline the price of that segment. Hence, the same set of investors can deliver a shallow schedule in one currency and a steep one in the other. Third, with a single investor type, $\mathcal{I} = \{1\}$, equation (8) reduces to $\Gamma_g = \mu\,\omega_{1,g}$. This is the standard intermediary benchmark, in which composition plays no role. In our case elasticity depends on the composition and each investor's elasticity. The weights $h_{i,g}$ in (9) are observed in the holdings data quarter by quarter, such that the right-hand side can be computed directly and compared with the aggregate elasticity estimated on aggregate holdings.

To understand what makes demand sticky, we derive the pass-through of the wedge into the shadow price: zero when the constraint is slack and rising toward one as it binds. This pass-through, which we call θ_i and derive in Proposition 2 below, summarizes how the tightness of the constraint translates into the demand elasticity.

Proposition 2. Own-wedge and relative demand. *Let $\nu_{i,g} = (x_g^2/\omega_{i,g})/\mathcal{X}_i$ be the share of segment g in sector i's mandated risk, and let $\theta_i = 2\phi_i/(1+2\phi_i) \in [0,1]$, where $\phi_i = \chi_i R_i/\mu_i$ if the budget is exceeded and $\phi_i = 0$ otherwise. Then, for any degree of tightness θ_i,*

$$\frac{\partial \log B_{i,g}}{\partial \log x_g} \;=\; 1 - \theta_i\,\nu_{i,g}, \qquad\qquad \log\frac{B_{i,g}}{B_{i,g'}} \;=\; \log\frac{x_g}{x_{g'}} \;-\; \log\frac{\omega_{i,g}}{\omega_{i,g'}}. \tag{10}$$

The first expression is the own-wedge elasticity holding the other wedges fixed, with $\theta_i = 0$ under a constant friction and $\theta_i = 1$ under a hard budget. The second is the relative demand of a sector across two segments, and it is independent of μ_i, hence of $\bar{K}_i$ and χ_i, for every θ_i. The proof is in Appendix A.

Proposition 2 separates two elasticities. The own-wedge elasticity, $1 - \theta_i\nu_{i,g}$, is the response of a sector's holdings to a segment's own wedge. It falls when the regulatory constraint binds, toward zero for a sector concentrated in the segment

($\nu_{i,g}$ near one) and near one for a diversified one. The relative-demand elasticity is a separate object: substitution between the two currency denominations of a given issuer depends only on the ratio of the regulatory risk weights, whatever the budget tightness. A sector can therefore be inelastic in levels yet highly elastic across currencies. This is how a sector such as insurers and pension funds, sticky in the aggregate demand system of Koijen et al. (2021), can in principle still substitute across the currency denominations of the same issuer.

This separation also has important implications for identification, which we exploit in section 5. Since μ_i depends on the whole vector of wedges, any instrument that moves x_g also moves μ_i. Therefore, the own-wedge coefficient conflates the demand slope with the reallocation of the sector's risk budget. Each sector's holding of a segment is $B_{i,g} = x_g/(\mu_i \omega_{i,g})$. For two segments of the same issuer, the ratio is:

$$\frac{B_{i,g}}{B_{i,g'}} = \frac{x_g}{x_{g'}} \cdot \frac{\omega_{i,g'}}{\omega_{i,g}} \tag{11}$$

and μ_i cancels out, leaving only the ratio of wedges and the ratio of mandate weights. The endogenous term that correlates with the instrument disappears, so relative holdings depend only on observed wedges and regulatory charges.

Broader Mandate: Benchmark-tracking sectors. A mandate may also capture clientele requirements, such as tracking an index as in Pavlova and Sikorskaya (2023). The index weight shifts the intercept of demand while the mandate weight continues to set its slope. We refer to this case, derived in Appendix S5, when interpreting the elasticity of nonbanks.

Exchange Rates. To keep the model tractable, we do not add a monetary side. The exchange rate is determined by a no-arbitrage condition on risk-free rates, with stochastic discount factors that are affine in one common global risk factor. subsection S3.2 gives the process.

The Issuer's Problem: the Supply Block. The sovereign chooses consumption and next-period debt $\{c, B', B^{*\prime}\}$ to maximize household welfare subject to the resource constraint:

$$\max_{\{c,\, B',\, B^{*\prime}\}} E_0 \sum_{t=0}^{\infty} \varrho^t u(c) \tag{12}$$

$$\text{s.t.} \quad c = y + q\, B' - B + e q^*\, B^{*\prime} - e B^*, \tag{13}$$

where the aggregate state is $s = [y, B, B^*, e]$. The issuer also internalizes the required asset returns that clear the market and are derived in Proposition 1.

Equation (8) gives the downward-sloping schedules for the bonds in each currency that the issuer internalizes:

$$q = \frac{1-\delta(s')}{1+r}\,\frac{e}{E(e')} - \big(\Gamma_\ell\, B' + \xi_\ell\big), \qquad q^* = \frac{1-\delta(s')}{1+r^*} - \Gamma_*\, B^{*\prime} - \xi_*, \tag{14}$$

with the segment frictions $\Gamma_g = \big(\sum_i m_i/\gamma_{i,g}\big)^{-1}$. The currency wedge on segment g therefore has two components, $x_g = \Gamma_g B_g + \xi_g$: the price impact of the quantity issued, $\Gamma_g B_g$, and the required premium ξ_g, which investors demand at any given quantity. The latter is zero in the long-run equilibrium of Proposition 1 and moves when the residual demand of investors shifts following a shock. The sovereign also faces a no-default incentive-compatibility constraint, as it repays whenever $V^R(s) \geq V^D(s)$, which generates an endogenous default probability:

$$\delta(s') = \int_{D(s)} f(y' \mid y) f(e' \mid e) de' dy', \tag{15}$$

where $D(s) = [s : V^R(s) < V^D(s)]$. The equilibrium default makes the schedules (14) shift with issuance as in Eaton and Gersovitz (1981). This part of the model is standard, and the default channel that characterizes it accounts for a small share of the responses we compute: we therefore report its derivation and recursive formulation in subsection S3.3.

The government's first-order conditions for bonds give the supply of bonds in different currencies. To derive it, we first write the problem recursively, so that $V(s) = \max_{B',B^{*\prime}} u(c) + \varrho\, E\big[V(s')\big]$, with c given by (13) and the price schedules (14) substituted in. The envelope condition for a unit of face value repaid next period (only in no-default states, $d' = 0$) is $\frac{\partial V(s')}{\partial B'} = -u_c(c')\,(1-d')$, where d' is next period's default indicator, and u_c is the marginal utility. Differentiating in B' and $B^{*\prime}$ and using $u_c(c)$ as the multiplier on (13) yields the two optimality conditions for debt in local and foreign currency:[15]

$$u_c(c)\Big[\, q + B'\frac{\partial q}{\partial B'}\,\Big] = \varrho\, E\big[u_c(c')(1-d')\big], \tag{16}$$

$$u_c(c)\, e\Big[\, q^* + B^{*\prime}\frac{\partial q^*}{\partial B^{*\prime}}\,\Big] = \varrho\, E\big[u_c(c')(1-d')\, e'\big]. \tag{17}$$

Each condition equates the marginal revenue of issuing one more unit today (left-hand side) to the discounted expected repayment tomorrow (right-hand side). The term $B'\,\partial q/\partial B'$ is what distinguishes the supply condition in our model from the broader class of Eaton and Gersovitz (1981) benchmarks: it provides the impact on bond prices of the wedge that finitely elastic demand imposes on the issuer.

15 For brevity the conditions omit the cross-price terms, through which issuance in one currency moves the price of the other bond via default risk. They operate only through $\partial\delta/\partial B'$.

This is the sense in which it provides the additional discipline device discussed in Bulow and Rogoff (1989). From (14) we obtain:

$$\begin{aligned} \frac{\partial q}{\partial B'} &= -\Gamma_\ell - \frac{e}{E(e')}\frac{1}{1+r}\frac{\partial \delta}{\partial B'}, \\ \frac{\partial q^*}{\partial B^{*\prime}} &= -\Gamma_* - \frac{1}{1+r^*}\frac{\partial \delta}{\partial B^{*\prime}}. \end{aligned} \tag{18}$$

The derivatives above imply that the marginal revenue of issuance is reduced by the aggregate demand friction Γ_g and by the endogenous default premium $\partial\delta/\partial B'$. The first term comes from the aggregate investor demand of bonds in each currency and depends on the composition of the investor base. In Appendix S8 we present a counterfactual exercise, which allows us to quantify those components in a calibrated version of the model that matches the estimated elasticities: there we find that the second component on the right side of $\partial\delta/\partial B'$ is quantitatively small.

Covered parity: a hedging extension. The baseline model derived so far models the demand block in which prices are the unhedged excess returns, the UIP. In Appendix S6 we develop and derive an extension to the case with hedging of the currency exposure, the CIP.

Identified elasticities and their Role for Policy Counterfactuals. In Appendix S8 we show the results of two policy counterfactuals. They illustrate the importance of correctly identifying the elasticities. In the model, simulated so as to match the causally identified slopes of section 5, we run two policy counterfactuals: a change in the insurance and pension funds regulation, Solvency II, and the quantitative monetary tightening. They show the importance of assessing the response to policy while taking into account the composition of the investor base and the general-equilibrium effects that only a structural model delivers.

5. Identification

This section shows how the model's demand and supply lead to our empirical design. Our estimating equation is the own-wedge demand relation of Proposition 2. It identifies the response of a sector's holdings of a sovereign's bonds in a given currency to the currency wedge on that issuer-currency segment. Because the wedge is an equilibrium price, we instrument it with the residual demand of the other investor sectors. The instrument varies across issuer-currency segments and over time, but not across the securities of a segment. The elasticity it identifies is therefore the elasticity of demand for a segment, and not the elasticity of substitution among the nearly identical securities inside a segment, which is the object usually estimated

with security-level variation.[16] The first stage relates the currency wedge to the instrument. The second stage estimates how holdings respond to the instrumented wedge, hence the demand elasticities. A third regression estimates the future response of asset supply or issuance, measured by the amount outstanding, to the same instrumented currency wedge.

Next we describe each stage and then discuss instrument validity and the exclusion restriction.

Demand Specifications. We start from the investors' demand and show how it leads to the instrument and to the two-stage procedure. The first-order condition (3) gives sector i's optimal demand for segment g. For a single issuer borrowing in two currencies it reads as follows:

$$B_{i,\ell} = \frac{1}{\gamma_{i,\ell}} \underbrace{\left(\tfrac{1-\delta}{1+r} \tfrac{e}{E(e')} - q \right)}_{\equiv\, x_\ell} = \frac{1}{\gamma_{i,\ell}} x_\ell, \qquad B_{i,*} = \frac{1}{\gamma_{i,*}} \underbrace{\left(\tfrac{1-\delta}{1+r^*} - q^* \right)}_{\equiv\, x_*} = \frac{1}{\gamma_{i,*}} x_*. \tag{19}$$

The term x_ℓ is the model's currency wedge on the local-currency segment, defined in (1), namely the gap between the return required by arbitrage, adjusted for default and expected depreciation, and the price actually paid. Every security b carries the wedge of its segment $g(b)$, and below we write x_b for the wedge of the segment to which security b belongs. Imposing arbitrage on returns delivers the UIP relation (see also (7)):

$$x_\ell^{\text{UIP}} = \underbrace{\frac{e}{E(e')} \left(\frac{1}{1+r^*} - \frac{1-\delta}{1+r} \right)}_{\text{frictionless parity}} + \underbrace{\Gamma_\ell B'}_{x_\ell}, \qquad \Gamma_\ell = \left(\textstyle\sum_i m_i / \gamma_{i,\ell} \right)^{-1}. \tag{20}$$

The measured deviation $x_{b,t}^{UIP}$ (in the swaps-based version, $x_{b,t}^{CIP}$) is the empirical counterpart of x_ℓ once the frictionless parity term, which the default and macroeconomic controls absorb, is controlled for. Three facts follow from (19), (20), and Proposition 1, and they guide the interpretation of the results. First, the demand is decreasing in the friction: holdings respond to the wedge with slope $1/\gamma_{i,g} = 1/(\mu_i \omega_{i,g})$. A tighter regulatory constraint, through a larger regulatory weight on the segment or a more binding budget, means more inelastic demand. Heterogeneity in $\gamma_{i,g}$ is the model primitive behind the heterogeneity in estimated elasticities, and its segment dimension is the part that regulation pins down. Second, the wedge is an equilibrium price: by (20), x_ℓ arises from market clearing, hence it depends on the aggregate friction Γ_ℓ and on supply B'. This is the simultaneity the instrument shall address. Third, the aggregate is a composite:

16 subsection 6.4 reports complementary estimates that use only variation among the securities of one issuer.

by (9), the elasticity of aggregate demand for a segment is the holdings-weighted average of the sector elasticities.

The relative-demand relation. The second half of equation (10) describes the substitution of a given sector between the local-currency and the hard-currency debt of the same issuer. Let n index issuers, let ℓ and $*$ denote the two currency segments of issuer n, and let $D_{i,g,n}$ be the face value of sector i's holdings of segment g of issuer n (this corresponds to $B_{i,g}$ in the model). Taking logs of (10) and linearizing the log wedge ratio around a reference level $\bar{x}_i$ (possibly the deterministic steady state) delivers:

$$\log \frac{D_{i,\ell,n}}{D_{i,*,n}} = \beta_{0i}^{\text{rel}} + \beta_{1i}^{\text{rel}}\, x_n^{\text{rel}} + \text{FE}_n + \text{FE}_t + \eta_{i,n}, \qquad x_n^{\text{rel}} \equiv x_{\ell,n} - x_{*,n}, \tag{21}$$

with $\beta_{1i}^{\text{rel}} = 1/\bar{x}_i$ and $\beta_{0i}^{\text{rel}} = -\log\left(\omega_{i,\ell}/\omega_{i,*}\right)$, where FE_n and FE_t are issuer and time effects and where η is the residual. In (21) the shadow price has dropped out and the sector intercept is the log ratio of the regulatory weights: hence this relation is independent of whether the constraint binds or not. We do not estimate (21) directly in the baseline, and use it to interpret the own-wedge estimates. Its closest empirical counterpart is the within-issuer design of subsection 6.4, which compares the securities of one issuer in one quarter, in local and in investor currency.

The own-wedge specification. We now derive the specification for a sector's demand slope and its interpretation. Log-linearizing the structural demand (19) around a sector's long-run average (or another reference level) $\bar{x}_i$, which is the first-order expansion of Koijen and Yogo (2019), gives:

$$\log D_{i,b} = \beta_{0i} + \beta_{1i}\, x_b + (\text{controls}) + \eta_{i,b}, \qquad \beta_{1i} = \frac{1 - \theta_i \nu_{i,g}}{\gamma_{i,g}\, \bar{D}_{i,g}}, \tag{22}$$

where the constant absorbs the mandate through $-\log \gamma_{i,g}$. The slope β_{1i} is a decreasing transform of the structural friction $\gamma_{i,g}$, so its dispersion across sectors reflects the model primitive (regulation). The slope is however scaled down by the pass-through θ_i through the factor $1-\theta_i\nu_{i,g} \in [0,1]$. This factor declines the more of the sector's risk-weighted assets is concentrated in the segment. The interpretation therefore depends on whether the constraint binds: when the mandate is slack ($\theta_i = 0$) the slope recovers $1/(\gamma_{i,g}\bar{D}_{i,g})$ exactly, and when it binds it understates it, so the estimate is a lower bound on the structural slope. Differences in estimated slopes across sectors then reflect both their frictions and how concentrated their risk-weighted assets are in the segment being estimated.

As we explain further below, in the data we estimate equations (20) and (22) exploiting both cross-sectional (across investors) and time series variation. For this reason the empirical specifications will also carry an index t.

Residual demand and the first stage. Aggregating (3) over the residual investors $-i = \mathcal{I} \setminus \{i\}$ and imposing market clearing gives the residual aggregate demand for segment g, and hence for a security b in it:

$$\sum_{j \neq i} m_j \, B_{j,g} = \Big(\underbrace{\textstyle\sum_{j \neq i} m_j / \gamma_{j,g}}_{\text{residual demand slope}} \Big) \, x_g \quad \equiv \quad \frac{1}{\Gamma_{-i,g}} \, x_g, \tag{23}$$

which, combined with the sovereign's supply, predetermined in the short run, clears the market for g and determines the equilibrium currency wedge. From the standpoint of sector i, movements in the wealth that the residual sectors bring to segment g, or in the tightness of their mandates, shift x_g along i's own demand curve. This is the shift-share instrument $z_{-i,b,t}$, constructed in subsection 5.1 as the residual sectors' predetermined segment exposure share times their current aggregate wealth, equation (29): as noted above, its data equivalent exploits both cross-sectional and time-series variation. It is the model's residual aggregate demand (23) apportioned to the unit of observation. The first stage is the empirical counterpart of the market-clearing block (23). It regresses, exploiting cross-section and time series variation, the equilibrium wedge on this residual-demand shifter:

$$x_{b,t} = \pi_{0i} + \pi_{1i} \, z_{-i,b,t} + (\text{controls}) + \varepsilon_{b,t}, \qquad |\pi_{1i}| \;\propto\; \Gamma_g, \tag{24}$$

whose relevance ($\pi_{1i} \neq 0$) holds whenever demand for the segment is finitely elastic ($\Gamma_g > 0$). With the supply being predetermined, a shift $u_{-i,g}$ in residual demand at a given wedge changes market clearing to $x_g/\Gamma_g + u_{-i,g} = B_g$ and moves the wedge by Γ_g per unit.

The Supply Curve. Substituting the aggregate demand schedules, (14) and (16), and collecting terms gives optimal issuance as a function of the required wedge and the demand slope:

$$\underbrace{2\,\Gamma_\ell \, B'}_{\text{price-impact}} = \underbrace{\frac{1-\delta}{1+r}\,\frac{e}{E(e')} - \varrho\, E\big[\tfrac{u_c(c')}{u_c(c)}(1-d')\big]}_{\text{net marginal value of funds}} - \underbrace{\xi_\ell}_{\text{required premium}} - \underbrace{\frac{e}{E(e')}\,\frac{1}{(1+r)}\,B'\,\frac{\partial \delta}{\partial B'}}_{\text{default discipline}}, \tag{25}$$

such that:

$$B' = \frac{\dfrac{1-\delta}{1+r}\,\dfrac{e}{E(e')} - \varrho\, E\big[\tfrac{u_c(c')}{u_c(c)}(1-d')\big] - \xi_\ell}{2\,\Gamma_\ell + \dfrac{e}{E(e')}\,\dfrac{1}{(1+r)}\,\dfrac{\partial \delta}{\partial B'}}, \tag{26}$$

and symmetrically for $B^{*\prime}$ from (17), with Γ_* in place of Γ_ℓ. Equation (26) is the model's supply curve.

Lemma 1. Testable Predictions. *Optimal issuance is decreasing in the*

aggregate demand friction Γ_ℓ*. It is also decreasing in the required premium* ξ_ℓ*.*

The first part of Lemma 1 says that the more inelastic investor demand is, the more a given increase in issuance depresses the market-clearing price, so the sovereign issues less. The second part says that a larger required premium leads to a higher marginal cost of funds, to which the issuer responds by issuing less. This is the elastic-supply channel. The sovereign does not place any quantity at the arbitrage price, but it internalizes the aggregate demand semi-elasticity of (9) through Γ_ℓ.

Log-linearizing (26) around the deterministic steady state and mapping B' into the observed amount outstanding $Q_{b,t}$ of security b, so that net issuance is the change in the amount outstanding, gives the supply equation:

$$\Delta \log Q_{b,t} = \alpha_0 + \alpha_1 \, x_{b,t} + (\text{controls}) + u_{b,t}, \qquad \alpha_1 = -\frac{1}{2\,\bar{\Gamma}_\ell \, \bar{B}'} < 0 \tag{27}$$

with $\bar{\Gamma}_\ell$ the steady-state value of $\Gamma_{\ell,t}$ and the wedge $x_{b,t}$ instrumented by the same shift-share $z_{-i,b,t}$ of (29). The instrument is valid here as well. Over the short horizon, when the sovereign's supply is predetermined, it shifts the investor's wedge through residual investor demand. The supply response α_1 is governed by the same structural objects as demand, the aggregate demand friction $\Gamma_{g,t}$ and the default schedule $\partial\delta/\partial B'$, because the issuer internalizes the demand schedule investors generate. Since supply and demand are the two sides of the same market-clearing block, a single instrument identifies both.

5.1. The Instrument

We now describe how we construct each variable in the empirical specifications and discuss the instrument (29) in more detail.

Definitions. There are investor sectors $i \in \mathcal{I}$, sovereign bonds $b = 1, \ldots, N$, and periods $t = 1, \ldots, T$. Denote by $-i = \mathcal{I} \setminus \{i\}$ the set of residual market participants, for each $i \in \mathcal{I}$. Let $D_{j,b,t}$ be the face value of holdings of security b by sector j in period t, and $\mathcal{B}_{-i,t}$ the set of securities held by at least one of the sectors in $-i$ at t. Let $g(b)$ denote the segment of security b, defined as its issuer country $\times$ currency of denomination, and within a given issuer, segments correspond to the currencies in which it borrows. Let $\mathcal{W} = \{2014q1, \ldots, 2014Q4\}$

be the base window. For all $i \in \mathcal{I}$ and every segment g, we compute the residual sectors' portfolio share in segment g over the base window

$$w_{-i,g} = \frac{\sum_{t\in\mathcal{W}} \sum_{b\in\mathcal{B}_{-i,t}:\, g(b)=g} \sum_{j\neq i} D_{j,b,t}}{\sum_{t\in\mathcal{W}} \sum_{b\in\mathcal{B}_{-i,t}} \sum_{j\neq i} D_{j,b,t}}, \tag{28}$$

where the denominator includes, as the outside asset, all holdings of sovereign securities issued by euro-area countries and by all other countries not in our sample. This is an assumption about the relevant outside asset, which is not observed.

It is critical for our approach that the segment shares be predetermined relative to sector i's demand at time t. We achieve this by holding the shares $w_{-i,g}$ constant at their base-window values throughout the estimation sample, such that all time variation in the instrument comes from the residual-investor wealth series and none from the shares. Accordingly, the estimation sample starts after the base window, as for each security we keep only periods $t > \max\{t_{0_i}, 2014q4\}$, where t_{0_i} is the first period in which security b is held by sector i. The instrument for the currency wedge of security b is the predicted residual aggregate demand for that security:

$$z_{-i,b,t} = \log\left(w_{-i,g(b)} \sum_{b'\in\mathcal{B}_{-i,t}} \sum_{j\neq i} D_{j,b',t} \right), \tag{29}$$

where $\sum_{b'\in\mathcal{B}_{-i,t}} \sum_{j\neq i} D_{j,b',t}$ is the face value of all sovereign securities held by residual market participants at t, which we call their nominal wealth in sovereign bonds. The instrument apportions the residual aggregate demand shift to security b through the predetermined segment share $w_{-i,g(b)}$, such that all securities in the same segment receive the same instrument for given (i,t).

Three properties of the instrument shape identification. First, the instrument varies across segments through the predetermined share and over time through common residual wealth, but not within a segment at a point in time. Hence, the identifying variation lies in predetermined cross-segment differences in exposure combined with time variation in residual-investor wealth. Second, and as a consequence, the instrument is identified under issuer- and time-varying controls but is incompatible with issuer-times-currency (segment) fixed effects, which would absorb $w_{-i,g}$ and eliminate the first stage. The covered (CIP) specifications of subsection S6.1 use the same construction. Third, because the instrument is the log of a product, it is additive in the log share and in log residual wealth. A full set of time effects would absorb the wealth term, and issuer effects would absorb the share in the samples restricted to local-currency debt, where each issuer has a single segment. The baseline specification therefore includes neither, and its identifying assumption is correspondingly stronger than in a design with fixed effects:

conditional on the controls, the aggregate sovereign-bond wealth of the residual sectors must be unrelated to sector i's own latent demand for a given security. Shocks that move the wealth of all sectors at once, such as a change in global risk appetite or in euro-area monetary policy, are the threat to this assumption. We address it in two ways. The controls are the model's aggregate states, namely the German yield for the return on the outside asset and issuer GDP for issuer fundamentals. And in subsection 6.4 we report estimates from a security-level design that does include issuer-by-time effects, at the cost of identifying a narrower elasticity.

5.2. A Structural Interpretation of the Instrument

Because the model includes a sector index and a segment index, the instrument (29) has a model counterpart. The share $w_{-i,g}$ is the residual sectors' portfolio share in segment g, which by (3) is a function of the mandate weights $\{\omega_{j,g}\}_{j\neq i}$ and is therefore predetermined with respect to sector i's demand shocks. The shift $\sum_{b'}\sum_{j\neq i} D_{j,b',t}$ is the residual sectors' total wealth, which in the model moves with $\{m_{j,t}\}_{j\neq i}$ and with the tightness of their budgets $\{\bar{K}_{j,t}\}_{j\neq i}$. The exogenous variation the design exploits is thus a shock to the wealth or the mandate of a sector other than the one whose demand is being estimated. That shock shifts residual demand in (23), moves the equilibrium wedge, and traces out sector i's response.

5.3. Instrument Validity

We now discuss instrument relevance and the exclusion restriction.

First, identification rests on the predetermined variation in the exposure shares. Following the exogenous-shares design (see Goldsmith-Pinkham, Sorkin and Swift (2020), and Borusyak, Hull and Jaravel (2025) for discussion of this design too), the identifying assumption is that sector i's demand shocks do not load on the predetermined shares. Specifically, conditional on observable security characteristics, sector i's demand shocks must not move its demand for securities in the same cross-sectional pattern as $-i$'s historical portfolio. Because our shares are aggregated to the issuer-currency segment and held constant at the base window $\mathcal{W}$, the relevant object is the predetermined segment exposure $w_{-i,g}$. Concretely, the concern would be that segments with higher $w_{-i,g}$ are systematically sorted on attributes for which sector i has an unobserved, time-varying demand. For

instance, high-$w_{-i,g}$ segments could be more collateral-eligible than low-$w_{-i,g}$ ones while banks have a time-varying demand for collateral eligibility.

This identifying assumption exploits differences in mandates, regulatory constraints, and risk-management rules between i and the sectors in $-i$. In Table S7 we show that only one observable characteristic, the log of amount outstanding (in euros), a proxy for liquidity, is robustly correlated with the exposure shares, irrespective of the composition of $-i$ and both between and within issuers and time periods. Seniority and collateral eligibility show no imbalance within issuers and time periods. The relevance condition hinges on the intensive margin of portfolio adjustment, whereby in the cross-section, high-exposure segments tend to be associated with high currency wedges. The exclusion restriction requires that predicted holdings of investors $-i$ affect the demand of sector i only through their price impact, that is, only through the currency wedge. Figure S4 illustrates the strategy. We estimate the elasticity of a given investor's demand, or portfolio share, to a shock that induces excess demand from the remaining investors.

Second, we aggregate the exposure shares to the issuer-currency segment rather than measuring them at the security level, which improves both the relevance and the validity of the instrument, as the segment share, averaged over the base window, is substantially less noisy than a security-level share measured off holdings of a security at any specific date, and it is predetermined at a common base period. Identification then draws on predetermined differences in $-i$'s exposure across segments, interacted with the time variation of residual-investor wealth that is common across all securities. These cross-segment differences are the exogenous shares of the design. Conditional on the controls, the balance tests indicate that they are as good as randomly assigned, and identification comes from historical exposure patterns that predate, and are orthogonal to, sector i's contemporaneous demand.

Empirical Specification for Demand Elasticities. The empirical specification follows the asset demand system of Koijen et al. (2021). Beyond our shift-share instrument, we add a larger set of security-level controls. The two-stage regression is

$$x_{b,t} = \pi_{0i} + \pi_{1i} z_{-i,b,t} + \pi_{2i} \log\left(D_{i,b,t_{0_i}}\right) + \pi_{3i} y_t^{DE} + \pi_{4i}' X_{b,t} + \varepsilon_{b,t}, \tag{30}$$

$$\log(D_{i,b,t}) = \beta_{0i} + \beta_{1i} x_{b,t} + \beta_{2i} \log\left(D_{i,b,t_{0_i}}\right) + \beta_{3i} y_t^{DE} + \beta_{4i}' X_{b,t} + \eta_{i,b,t}, \tag{31}$$

where $\log\left(D_{i,b,t_{0_i}}\right)$ controls for omitted fixed security attributes and the yield on the 10-year German bond y_t^{DE} controls for the attractiveness of alternative investments inside the euro area. The vector $X_{b,t}$ contains the log of real GDP in the issuer

country of security b at time t, which controls for macroeconomic conditions in the issuer country, and time-varying security-level attributes, namely the lag of the log of amount outstanding, residual maturity dummies, and credit quality dummies. The term $\varepsilon_{b,t}$ is a security-level disturbance to the currency wedge at t, and $\eta_{i,b,t}$ is the latent demand of sector i for security b at t. The coefficient of interest is β_{1i} for each investor i, which we call the demand slope.

Mapping between Empirical Controls and Model Elements. The controls map into elements of the structural model. $\log D_{i,b,t_{0_i}}$ absorbs the fixed security attributes and the sector-level term $-\log \gamma_{i,g}$, the fixed effect of the demand system. y_t^{DE}, the 10-year German yield, is the return on the outside euro-area asset, the reference leg r^* in the wedge (20). The vector $X_{b,t}$ (issuer real GDP, amount outstanding, residual-maturity and credit-quality dummies) proxies for the arguments of the default probability $\delta(s')$ and the issuer-risk terms that enter (19). Conditioning on them leaves the currency-risk variation that identifies β_{1i}.

Computation of the Demand Elasticities. We follow Koijen and Yogo (2019) and Koijen et al. (2021) and compute the price elasticity of demand implied by our estimates for each sector, security, and quarter as

$$-\frac{\partial \log\left(D_{i,b,t}\right)}{\partial \log\left(P_{b,t}\right)} = 100\frac{\beta_{1i}}{\tau_{b,t}}\left(1 - w_{i,b,t}\right), \tag{32}$$

where β_{1i} is the semi-elasticity of demand from (31), $\tau_{b,t}$ is the maturity of security b at t, and $w_{i,b,t}$ is sector i's portfolio share in security b in quarter t. Recall that $D_{i,b,t}$ is measured at face value.[17] We then compute the median sectoral elasticity by pooling across securities and quarters, and aggregate to a euro-area total as a weighted average of the sectoral medians, with weights equal to sector i's 2014Q1 holdings in the relevant currency-country sample over all sectors' holdings in that sample. Combining (32) with the expression for β_{1i} in (22), the model-implied elasticity is proportional to $1/\gamma_{i,g}$, which links the estimated object directly to the structural primitive, such that the least elastic sectors are those with the tighter mandate in that segment.

17 The portfolio share of security b in sector i's portfolio at t is

$$w_{i,b,t} = \frac{D_{i,b,t}}{\sum_{b' \in \mathcal{B}_{i,t}} D_{i,b',t}}, \tag{33}$$

where $\mathcal{B}_{i,t}$ is the set of securities held by sector i at t.

5.4. The Supply Specification

We derived the supply specification from the model in equation (27). Here we describe how the variables are measured. The dependent variable is the log change in the amount outstanding of security b between t and $t+4$, $\Delta \log(Q_{b,t+4})$. The two stages are:

$$x_{b,t} = \pi_0 + \pi_1 z_{-i,b,t} + \pi_2 \log(Q_{b,t-1}) + \pi_3 y_t^{US} + \pi_4 \left(\frac{debt}{gdp}\right)_t + u_{b,t}, \quad (34)$$

$$\Delta \log(Q_{b,t+4}) = \alpha_0 + \alpha_1 x_{b,t} + \alpha_2 \log(Q_{b,t-1}) + \alpha_3 y_t^{US} + \alpha_4 \left(\frac{debt}{gdp}\right)_t + v_{b,t}. \quad (35)$$

We select as controls the variables that matter for the issuer's decision. The 10-year U.S. Treasury yield captures global conditions and the opportunity cost of funds. The ratio of debt (in local and foreign currency) to GDP captures the issuer's fiscal position. The lag of the log of the amount outstanding captures the liquidity of the security. The last control matters because the issuer may issue more of exactly the securities that investors prefer to hold, namely the more liquid ones, as our balance tests show, which would bias the estimates. Because the shares are predetermined, we use only data after 2014Q4, as in the demand estimation. To maximize statistical power, we use the set of securities held by euro-area nonbanks, the largest and most representative high-frequency sample we have, such that $-i$ in the instrument consists of banks, insurers and pension funds, and other investors.

6. Empirical Results

We present the results in order, describing the summary statistics and balance tests first, and subsequently the demand and supply estimates.

6.1. Preliminaries

Summary Statistics. Table S6 reports summary statistics for the wedges in the estimation sample of securities with 1 to 5.5 years of residual maturity. Emerging-market local-currency UIP wedges average 2.75 percentage points, with a 5th to 95th percentile range of about −2.0 to +7.7. Advanced-economy local-currency wedges average 0.83, with a range of −3.1 to +4.4. For emerging-market local-currency bonds the swaps-based CIP wedge averages 0.085, such that almost all of the wedge is the uncovered component. The endogenous regressor therefore has large, two-sided, currency-driven dispersion.

Balance Tests and Exclusion Restriction. We next examine the validity of the instrument. In our design, identification relies on the predetermined leave-out exposure shares, not on the exogeneity of the shifts. In the baseline, the instrument uses the segment share $w_{-i,g}$, namely the leave-out exposure of the residual sectors to the issuer-currency segment $g(b)$, averaged over the 2014 base window and held fixed thereafter. Predetermination comes from holding the shares constant at the base window. As-good-as-random assignment requires that these shares not be systematically sorted on security characteristics that some sectors prefer.

We assess sorting with a balance regression. To capture the preference for a particular security, we compute the portfolio share that the sectors in $-i$ assign to security b the first time ($t_{0_{-i}}$) they hold it, and denote it $w_{-i,b,t_{0_{-i}}}$. The segment share is a within-segment average of these finer security-level leave-out shares. We run the diagnostic on the security-level shares, which is the more demanding test, as any sorting of $w_{-i,g}$ on security attributes must originate in sorting of the underlying $w_{-i,b,t_{0_{-i}}}$, and averaging over the securities in a segment can only attenuate such sorting, but not exacerbate it.[18] Balance of the security-level shares within issuer-times-time cells therefore implies at least as good a balance of the segment share we construct our instrument with. Following Borusyak, Hull and Jaravel (2022), we regress $w_{-i,b,t_{0_{-i}}}$ (in basis points) on the full set of characteristics under three progressively more demanding fixed-effect structures, none, issuer plus time, and issuer times time, separately for each own-sector. Table S7 reports the results.

Without fixed effects, exposure tilts toward large, collateral-eligible bonds, but almost all of this imbalance is between issuers and between periods. Within issuer-times-time cells, the only characteristic that remains correlated with the shares is the log amount outstanding, a proxy for liquidity. A characteristic biases the design only if it is omitted, and therefore, following Borusyak, Hull and Jaravel (2022), we include the log amount outstanding among the controls. Appendix S2 reports the table and discusses it column by column.

18 The portfolio share of security b in the portfolio of $-i$ in quarter t is:

$$w_{-i,b,t} = \frac{\sum_{j \neq i} D_{j,b,t}}{\sum_{b' \in \mathcal{B}_{-i,t}} \sum_{j \neq i} D_{j,b',t}} \tag{36}$$

and we evaluate it at the first quarter in which at least one sector in $-i$ holds security b, defined as $t_{0_{-i}} = \min\{t \in 1, \ldots, T : \sum_{j \neq i} D_{j,b,t} > 0\}$, which yields $w_{-i,b,t_{0_{-i}}}$.

6.2. Demand for Local- vs. Investor-Currency Bonds

Table 1 reports the second-stage estimates of the demand slope β_{1i} in specification (31). Each cell is the 2SLS coefficient of log holdings on the instrumented UIP wedge, by investor sector (panels) and currency-country sample (columns), with the first-stage F reported beneath each estimate. Through the lens of the model, $\beta_{1i} = (1 - \theta_i \nu_{i,g})/(\gamma_{i,g} \bar{D}_{i,g})$ by (22). A positive coefficient is the model-consistent, downward-sloping demand, whereby holdings rise as the wedge widens, that is, as the compensation for bearing currency risk increases. Its magnitude is inversely proportional to the sector's friction $\gamma_{i,g}$. By contrast, a negative coefficient signals the opposite behavior, indicating that the sector cuts exposure precisely when the wedge is wide, transmitting rather than absorbing currency risk. Table 2 translates the slopes into price elasticities of demand, rescaling each slope by the security's maturity and own-weight and aggregating to a euro-area total with fixed 2014Q1 weights.

The first stage is strong in most cells. The F-statistic is above 10 everywhere except for banks and other investors in the emerging-market local-currency-only column ($F = 6.4$ and $F = 1.3$). We flag these as weak-instrument cells and do not interpret them. Relevance is strongest for advanced-economy issuers, where the first-stage F ranges from roughly 90 to 270. Building the exposure share at the issuer-times-currency level concentrates identification on the variation the instrument is designed to trace, and delivers a relevant first stage.

The advanced-economy estimates are the cleanest in the table and show broad absorption of currency risk. All slopes are positive. For insurers and pension funds, nonbanks, and other investors, they are highly significant, between 0.56 and 1.45. Banks are positive and significant in local currency (0.49), though not once euro securities are added. The implied elasticities are correspondingly large, with around 49 for insurance corporations and pension funds, 21 for nonbanks, 20 for the other sector, and 17 for banks in the local-currency sample. The euro-area aggregate is near 25. For comparison, Koijen and Yogo (2026) report average elasticities of 27.9 for short-term debt and 3.2 for long-term debt in a demand system estimated on country-level portfolios. Every advanced-economy investor class has a downward-sloping, finitely elastic demand for currency risk, consistent with these bonds being close substitutes that carry modest and well-priced currency risk.

For emerging-market issuers, the two well-identified sectors, insurers and pension funds and nonbanks, both have a positive and statistically significant demand slope,

such that both absorb currency risk rather than shed it. Insurance corporations and pension funds have a positive, significant slope in both columns (1.10 and 0.51), with implied elasticities of about 40 in the local-currency sample and 18 once euro securities are added. Nonbanks, which dominate the emerging-market investor base with an aggregation weight of 0.71 to 0.80, also have a positive, well-identified slope throughout (0.33 and 0.27, with $F = 70$ and $F = 152$, respectively), with implied elasticities near 12 and 9. This is the central finding: the largest euro-area holder of emerging-market currency risk has an ordinary, downward-sloping demand, and absorbs rather than transmits that risk. Banks and the other sector are not cleanly identified in the specification that isolates currency risk. The bank slope is insignificant in the local-currency-only column, where the first stage is weak. It is negative but insignificant once euro securities are included (-0.165). We read it with caution, since that sample includes euro-denominated emerging-market bonds that carry no currency risk for euro-area investors. The other sector is imprecisely signed in local currency and carries negligible weight.

The euro-area aggregates reflect this composition. Because nonbanks absorb most of the weight, the emerging-market total is positive, about 7 in the local-plus-euro sample, where every underlying cell clears the weak-instrument threshold. The local-currency-only aggregate is larger, around 20, but this figure rests on the weak-instrument bank cell, which alone contributes nearly half. Restricting attention to the two well-identified sectors results in a local-currency emerging-market elasticity closer to 12. Either way, the aggregate response of emerging-market demand to the currency wedge is positive, set by the sectors that absorb currency risk.

Comparing across issuer groups, advanced-economy demand is more elastic than emerging-market demand, with a euro-area total near 25 against 7 to 12. The gap is narrower than a stark currency-risk ranking would suggest, and for insurance corporations and pension funds the two are of similar size (about 49 against 40 in local currency). The difference between advanced- and emerging-market elasticities is therefore a feature of the level of the wedge and of the composition of demand, not of a categorical difference in how steeply any given sector's demand slopes.

6.3. Supply Elasticities

Table 3 reports the second-stage supply coefficient α_1 from specification (35) under the uncovered instrument. The coefficient measures how the amount a government issues over the following year, $\Delta \log Q_{b,t+4}$, responds to the instrumented currency wedge $x_{b,t}$. A negative α_1 is a disciplining response, whereby when investors require

more compensation to bear the security's currency risk, the sovereign issues less of it. This is the supply-side counterpart of the finitely elastic demand estimated above, and it is the mechanism through which investor demand disciplines issuance.

The uncovered specification shows a disciplining response where it is identified. The coefficient for emerging-market bonds in the local-plus-euro sample is -0.088 and significant at the 10 percent level. The first stages are weak elsewhere, and the remaining cells are insignificant, including a local-currency emerging-market cell whose first stage collapses entirely ($F = 0.13$). We therefore read the uncovered supply estimates with caution. The covered supply estimates in subsection S6.2, where the first stage is strong, show a negative and significant response for emerging-market issuers in both samples.

The controls confirm that the estimated response is disciplining behavior rather than a mechanical correlation. Conditioning on the lagged log amount outstanding is essential. Without it, the sovereign's tendency to issue more of the liquid securities that investors prefer to hold (the one characteristic our balance tests find correlated with the exposure shares) would induce co-movement between issuance and the wedge that has nothing to do with discipline. Netting out this liquidity channel leaves the supply response to the price of currency risk. The 10-year U.S. Treasury yield and the debt-to-GDP ratio absorb the global cost of funds and the fiscal-space conditions that also shift issuance and the wedge. Hence, α_1 captures the sovereign's issuance response to the price of currency risk alone.

6.4. Within-Issuer Evidence and the U.S. Comparison

The baseline identifies demand across issuer-currency segments and over time, and it does so without fixed effects. This section reports a second design that is its mirror image. It absorbs all issuer-by-time variation and identifies demand from differences among the securities of one issuer in one quarter. It is the specification that the U.S. data support: the confidential TIC holdings are annual, recorded at the custodial holder of record rather than by end-investor sector, and available to us only through 2022. We report its estimates in full for euro-area and U.S. investors, together with the corresponding supply estimates.

What the design identifies. With issuer-by-time fixed effects, the currency wedge of a security is compared with that of the other securities of the same sovereign in the same quarter. In the samples that pool local- and investor-currency debt, this comparison includes bonds of the same issuer in different currencies, which is the relative-demand relation (21) of the model. In the local-currency

Table 1
UIP: Slope of the Demand Function

Dependent variable: Log Holdings	EM + Local & EUR	EM + Local	AE + Local & EUR	AE + Local
Panel A: Insurance				
$x_{b,t}^{UIP}$	0.508**	1.098**	1.258***	1.454***
	(0.225)	(0.513)	(0.274)	(0.361)
Observations	5929	5085	6275	6190
F-statistic	49.200	32.494	100.803	92.724
Panel B: Nonbanks				
$x_{b,t}^{UIP}$	0.273***	0.327**	0.556***	0.630***
	(0.062)	(0.134)	(0.129)	(0.136)
Observations	6903	6054	7328	7229
F-statistic	151.581	70.135	169.472	185.239
Panel C: Banks				
$x_{b,t}^{UIP}$	-0.165	1.725	0.141	0.486***
	(0.183)	(1.194)	(0.105)	(0.111)
Observations	3389	2539	4256	4166
F-statistic	37.377	6.392	229.758	270.872
Panel D: Other				
$x_{b,t}^{UIP}$	0.161	-1.272	0.587***	0.579***
	(0.228)	(1.852)	(0.132)	(0.154)
Observations	3915	3060	5232	5150
F-statistic	14.781	1.349	139.176	135.784

Notes: The coefficients are the slopes of demand estimated from the two-stage regression (31). Each panel reports the slopes for one investor type. The last row in each panel reports the Cragg and Donald (1993) F-statistic. Robust standard errors clustered at the quarter level are in parentheses. One, two, and three stars denote significance at the 10, 5, and 1 percent levels.

Table 2

Implied Elasticities

	N	Median	SD	5th pct.	95th pct.	Weight
Panel A: EM bonds denominated in local currency						
Insurance	5085	39.545	25.756	21.410	105.411	0.046
Nonbanks	6054	11.603	7.669	6.335	31.015	0.802
Banks	2539	67.143	43.676	34.006	177.410	0.145
Other	3060	-43.553	29.172	-119.048	-24.598	0.008
Euro area total		20.479				1.000
Panel B: EM bonds denominated in local currency and EUR						
Insurance	5929	17.890	11.789	9.864	48.143	0.064
Nonbanks	6903	9.480	6.337	5.277	25.530	0.712
Banks	3389	-6.049	4.056	-16.426	-3.199	0.202
Other	3915	5.363	3.643	3.105	15.072	0.021
Euro area total		6.791				1.000
Panel C: AE bonds denominated in local currency						
Insurance	6190	48.903	34.562	27.824	137.073	0.180
Nonbanks	7229	21.224	15.268	12.064	60.854	0.532
Banks	4166	17.150	12.232	9.296	49.925	0.186
Other	5150	19.643	13.481	11.079	53.217	0.102
Euro area total		25.290				1.000
Panel D: AE bonds denominated in local currency and EUR						
Insurance	6275	42.318	29.871	24.077	118.582	0.184
Nonbanks	7328	18.787	13.436	10.633	53.358	0.519
Banks	4256	4.997	3.534	2.694	14.456	0.190
Other	5232	19.954	13.649	11.228	53.932	0.108
Euro area total		20.622				1.000

Notes: The table reports the implied elasticities computed from formula (32) using the UIP demand slopes in Table 1. Panel A is for EM bonds in local currency, Panel B for EM bonds in local currency and euro, Panel C for AE bonds in local currency, and Panel D for AE bonds in local currency and euro.

Table 3

UIP: Slope of the Supply Function

Dependent variable: $\Delta \log(Q_{b,t+4})$				
	EM + Local & EUR	EM + Local	AE + Local & EUR	AE + Local
$x_{b,t}^{UIP}$	-0.088*	2.473	-0.068	-0.084
	(0.048)	(25.113)	(0.072)	(0.088)
Observations	6068	5320	6138	6045
F-statistic	24.001	0.128	14.082	11.189

Notes: The coefficients are the slopes of supply from the two-stage regression (35). The last row reports the Cragg and Donald (1993) F-statistic. Robust standard errors clustered at the quarter level are in parentheses. One, two, and three stars denote significance at the 10, 5, and 1 percent levels.

samples, it is a comparison across issues and maturities in one currency. The estimated object is therefore an elasticity of substitution among the securities of one issuer, which are closer substitutes than the segments of the baseline. The fixed effects absorb every shock common to an issuer in a quarter, including the aggregate shocks that the baseline can only control for.

The design has two limitations, which are the reason it is not our baseline. First, the shares in the instrument are measured in the first quarter in which a security is held, and the estimation sample is not restricted to later quarters, such that the shares are not predetermined for every observation. The baseline removes this concern by fixing the shares in a common base window that precedes the estimation sample. Second, the dependent variable is the log portfolio share rather than log holdings, and the sample ends in 2022Q2. For these reasons we read the estimates of this section for their sign and significance, and we do not compare their magnitudes with those of Table 1.

Specification. The dependent variable is the log of sector i's portfolio share in security b, $\log(w_{i,b,t})$. The instrument is built at the security level. It is the excess weight of security b in the portfolio of the residual sectors in the first quarter in which they hold it, relative to an equal-weighted portfolio, scaled by the growth of their nominal wealth in sovereign bonds since that quarter. Unlike the baseline instrument (29), it is a product in levels and varies across the securities of an issuer, such that it is not absorbed by issuer-by-time fixed effects. Both stages include issuer-by-time fixed effects and residual-maturity and rating controls. The supply counterpart uses as dependent variable the growth of the amount outstanding of the security over the following quarter, on the securities held by euro-area nonbanks. The sample is 2014Q1 to 2022Q2 and includes securities with a residual maturity of 1 to 5.5 years. Because the instrument has few time clusters, the tables report 95 percent Anderson and Rubin (1949) confidence intervals, which are robust to weak instruments, numerically computed with the wild restricted efficient bootstrap of

Davidson and MacKinnon (2010), together with the effective F-statistic of Olea and Pflueger (2013). Appendix S7 gives the equations.

Euro-area investors. Table 4 reports the estimates for euro-area investors. For emerging-market issuers, in the sample that pools local-currency and euro bonds, the first stage is strong for insurers and pension funds, nonbanks, and banks, with effective F-statistics of 89.5, 24.7, and 31.8. A wider wedge raises the portfolio share of insurers and pension funds (0.58) and of nonbanks (2.07), and lowers that of banks (−1.42). All three are significant at the 5 percent level.[19] In the local-currency sample, the response of nonbanks remains positive and significant (1.39), that of insurers and pension funds is small and not different from zero (0.13), and the first stages for banks and other investors are weak (2.1 and 4.5). For advanced-economy issuers, the effective F-statistic is below 3 in every cell and the confidence intervals are wide or unbounded. The within-issuer design is uninformative for these issuers, consistent with there being little dispersion in the wedge across the bonds of an advanced-economy sovereign once issuer-by-time effects are absorbed.

Two results carry over from the baseline. Insurers and pension funds and nonbanks move toward emerging-market bonds when the compensation for currency risk rises, in both designs. And banks do not: their response is insignificant in the baseline and negative here, such that among the bonds of the same emerging-market sovereign, banks shift away from those whose wedge widens. One result does not carry over. Within issuer, nonbanks respond more than insurers and pension funds, whereas the baseline slopes are larger for insurers and pension funds. The two designs measure different objects with different dependent variables, as discussed above, and we do not attempt to reconcile the ranking.

U.S. investors. Table 5 reports the same specification on the annual TIC surveys that fall in the 2014Q1 to 2022Q2 window, with year in place of quarter effects. The data are described in subsection D9.1. In the sample that pools the local-currency and dollar bonds of emerging-market issuers, a wider wedge lowers the portfolio share of nonbanks (−2.16), banks (−1.59), and other investors (−2.11). All three are significant at the 5 percent level, with effective F-statistics between 13 and 24. The coefficient for insurers is also negative (−1.89), but its confidence interval includes zero. In the local-currency sample and for advanced-economy issuers the first stages are weak, with effective F-statistics of 3.1 or less, and we do not interpret those cells.

19 The wedge is in percentage points. A coefficient of 0.583 means that a 1 basis point increase in the wedge raises the portfolio share by $e^{0.00583} - 1 = 0.58$ percent.

U.S. investors therefore rebalance away from emerging-market sovereign bonds when the currency wedge widens. This is what euro-area banks do, and the opposite of what euro-area insurers, pension funds and nonbanks do. The similarity between U.S. and euro-area banks is consistent with both being subject to the Basel requirements described in section 3. For the other sectors, one interpretation of the contrast is that the dollar is the global funding currency, such that U.S. investors push emerging-market currency risk back onto the issuer, whereas euro-area investors hold it for the premium. The U.S. sample ends in 2022 and cannot be updated, and the contrast we draw is one of the direction of rebalancing, not of magnitude.

Supply. Table S12 reports the within-issuer supply estimates. For emerging-market issuers, a wider wedge lowers the growth of the amount outstanding over the following quarter, with coefficients of -0.154 in the sample that pools local-currency and euro bonds and -0.103 in the local-currency sample. Both are significant at the 5 percent level and have strong first stages (effective F-statistics of 23.5 and 41.1). For advanced-economy issuers the first stage is weak and we do not interpret the estimates. The sign agrees with the baseline supply estimates of Table 3: under both designs, emerging-market sovereigns issue less of a bond when the compensation that investors require for its currency risk rises. In the working-paper version we used these estimates, together with their covered counterparts, to gauge magnitudes. An 8 basis point widening of the wedge, which is the response of the spread of an investment-grade emerging-market sovereign to a shock to the VIX estimated by Gilchrist et al. (2022), lowers the net issuance of emerging-market local-currency debt by about 0.26 percent of GDP on average over 2014Q1 to 2022Q2.

7. Conclusion

This paper measures who bears currency risk in sovereign debt markets and at what price. We build a shift-share instrument from the predetermined, different currency bond exposures of residual investor sectors and ground it in a sovereign-issuance model with heterogeneous, finitely elastic investors. With it we estimate, for each investor sector, the elasticity of demand for the debt of a sovereign in each currency of denomination, for advanced and emerging issuers alike.

Some investors absorb, others transfer and yet some others shy away from taking any currency risk at all. Thus, there is no single investor pricing the currency

Table 4
Within-Issuer Design, Euro-Area Investors: UIP Wedge and Portfolio Shares

Dependent variable: Portfolio share, $\log(w_{i,b,t})$	EM + Local & EUR	EM + Local	AE + Local & EUR	AE + Local
Panel A: Insurance				
$x_{b,t}^{UIP}$	0.583**	0.128	-1.682	-0.959
	[0.272,0.623]	[-0.119,0.190]	[-18.347,-0.856]	[-Inf,+Inf]
Observations	3050	2428	3643	3541
Effective F-statistic	89.526	81.182	2.265	0.632
Panel B: Nonbanks				
$x_{b,t}^{UIP}$	2.074**	1.393**	4.969	6.421
	[0.973,2.736]	[0.716,1.652]	[0.123,14.867]	[0.211,20.952]
Observations	6209	5579	5756	5639
Effective F-statistic	24.676	43.654	1.719	1.227
Panel C: Banks				
$x_{b,t}^{UIP}$	-1.417**	1.313	44.521	-128.256
	[-2.175,-0.945]	[-Inf,+Inf]	[-Inf,+Inf]	[-Inf,+Inf]
Observations	1987	1364	2414	2312
Effective F-statistic	31.804	2.064	0.357	0.052
Panel D: Other				
$x_{b,t}^{UIP}$	5.408	1.070	-0.664	-0.279
	[-Inf,+Inf]	[0.259,2.127]	[-17.317,1.472]	[-14.168,1.811]
Observations	1678	1054	3856	3766
Effective F-statistic	0.404	4.495	2.976	2.701
Quarter FE	Yes	Yes	Yes	Yes
Issuer Country FE	Yes	Yes	Yes	Yes
Quarter x Issuer Country FE	Yes	Yes	Yes	Yes
Rating controls	Yes	Yes	Yes	Yes
Maturity controls	Yes	Yes	Yes	Yes

Notes: Estimates of β_{1i}^{w} in (S7.5). The sample is 2014Q1 to 2022Q2. The 95 percent Anderson and Rubin (1949) confidence intervals are reported in square brackets. They are computed using 999 wild restricted efficient bootstrap (WRE) replications following Davidson and MacKinnon (2010), clustering at the time level. An interval reported as $[-\text{Inf}, +\text{Inf}]$ is unbounded: the instrument is too weak to bound the coefficient. The Olea and Pflueger (2013) critical value for the effective F-statistic at a worst-case bias of 30 percent is 12.039. A value of the F-statistic above this threshold denotes a strong instrument. Two stars next to the coefficient estimate denote significance at the 5 percent level.

Table 5
Within-Issuer Design, U.S. Investors: UIP Wedge and Portfolio Shares

Dependent variable: Portfolio share, $\log(w_{i,b,t})$	EM + Local & USD	EM + Local	AE + Local & USD	AE + Local
Panel A: Insurance				
$x_{b,t}^{UIP}$	-1.888	6.137	-2.300	-1.714
	[-39.992,0.996]	[1.660,16.447]	[-Inf,+Inf]	[-Inf,+Inf]
Observations	1820	1077	1259	1202
Effective F-statistic	12.793	2.581	0.416	1.207
Panel B: Nonbanks				
$x_{b,t}^{UIP}$	-2.163**	5.692	-0.668	0.350
	[-6.708,-1.336]	[0.939,11.868]	[-Inf,+Inf]	[-Inf,+Inf]
Observations	3280	1894	1491	1407
Effective F-statistic	23.886	0.889	0.623	3.058
Panel C: Banks				
$x_{b,t}^{UIP}$	-1.592**	3.776	-44.126	2.109
	[-3.502,-1.196]	[-Inf,+Inf]	[-Inf,+Inf]	[-Inf,+Inf]
Observations	759	309	765	703
Effective F-statistic	13.163	0.190	0.025	1.278
Panel D: Other				
$x_{b,t}^{UIP}$	-2.105**	3.920	-93.099	-4.848
	[-7.584,-1.191]	[-Inf,+Inf]	[-Inf,+Inf]	[-Inf,+Inf]
Observations	2099	731	1043	968
Effective F-statistic	13.630	0.948	0.000	0.658
Year FE	Yes	Yes	Yes	Yes
Issuer Country FE	Yes	Yes	Yes	Yes
Year x Issuer Country FE	Yes	Yes	Yes	Yes
Rating controls	Yes	Yes	Yes	Yes
Maturity controls	Yes	Yes	Yes	Yes

Notes: Estimates of β_{1i}^{w} in (S7.5) on the annual TIC surveys in the 2014Q1 to 2022Q2 window. The 95 percent Anderson and Rubin (1949) confidence intervals are reported in square brackets. They are computed using 512 wild restricted efficient bootstrap (WRE) replications following Davidson and MacKinnon (2010), clustering at the time level. An interval reported as $[-\text{Inf}, +\text{Inf}]$ is unbounded: the instrument is too weak to bound the coefficient. The Olea and Pflueger (2013) critical value for the effective F-statistic at a worst-case bias of 30 percent is 12.039. A value of the F-statistic above this threshold denotes a strong instrument. Two stars next to the coefficient estimate denote significance at the 5 percent level.

wedge and no single marginal holder of currency risk. This changes by investor type and by the issuer's endogenous supply response.

For emerging markets, insurers and pension funds and nonbanks all absorb currency risk. Banks show no significant response, and when we compare bonds of the same issuer they reduce their portfolio share as the wedge widens. U.S. investors, in the same within-issuer comparison, rebalance away from emerging-market bonds. For advanced economies, every investor type absorbs currency risk, with implied elasticities between 17 and 49. Demand is more elastic for advanced-economy currency risk, which is cheap and substitutable, than for emerging-market risk. On the supply side, a wider currency wedge restrains subsequent emerging-market issuance, weakly when the wedge is measured on an uncovered basis (given expectational errors) and significantly on a covered basis, such that the shape of investor demand feeds back onto the sovereign in terms of the high compensation investors require to hold the currency risk.

A. Proofs of Propositions 1 and 2

Proof of Proposition 1. Substitute the sector demand $B_{i,g} = x_g/\gamma_{i,g}$ from (3) into the segment market-clearing condition $\sum_i m_i B_{i,g} = B_g$,

$$x_g \sum_i \frac{m_i}{\gamma_{i,g}} = B_g \quad \Longrightarrow \quad x_g = \Big(\sum_i \tfrac{m_i}{\gamma_{i,g}} \Big)^{-1} B_g = \Gamma_g \, B_g,$$

which is (8). For the semi-elasticity, differentiate $B_g = \sum_i m_i B_{i,g}$ using $\partial B_{i,g}/\partial x_g = 1/\gamma_{i,g}$,

$$\frac{\partial \log B_g}{\partial x_g} = \frac{1}{B_g} \sum_i \frac{m_i}{\gamma_{i,g}} = \sum_i \frac{m_i B_{i,g}}{B_g} \frac{1}{\gamma_{i,g} B_{i,g}} = \sum_i h_{i,g} \, \beta_{1i,g},$$

evaluated at the sector-segment average level $\bar{D}_{i,g} = B_{i,g}$, which is (9). □

Proof of Proposition 2. From (3), $\log B_{i,g} = \log x_g - \log \mu_i - \log \omega_{i,g}$, so

$$\frac{\partial \log B_{i,g}}{\partial \log x_g} = 1 - \frac{\partial \log \mu_i}{\partial \log x_g}.$$

When the mandate is slack, $\mu_i = \gamma_i^0$ is constant and the last term is zero. When it binds, $\mu_i = \gamma_i^0 + \chi_i(R_i - \bar{K}_i)$ with $R_i = \mathcal{X}_i/(2\mu_i^2)$ from (4). Differentiating this relation and using $\partial \mathcal{X}_i / \partial \log x_g = 2x_g^2/\omega_{i,g}$,

$$\Big(1 + \frac{\chi_i \mathcal{X}_i}{\mu_i^3}\Big) \frac{\partial \log \mu_i}{\partial \log x_g} = \frac{\chi_i}{\mu_i^3} \frac{x_g^2}{\omega_{i,g}}.$$

Since $\phi_i = \chi_i R_i/\mu_i$ and $\mathcal{X}_i = 2\mu_i^2 R_i$, one has $\chi_i \mathcal{X}_i/\mu_i^3 = 2\phi_i$ and $\chi_i(x_g^2/\omega_{i,g})/\mu_i^3 = 2\phi_i\nu_{i,g}$, so

$$\frac{\partial \log \mu_i}{\partial \log x_g} = \frac{2\phi_i}{1+2\phi_i}\,\nu_{i,g} = \theta_i\,\nu_{i,g},$$

which gives the first expression, with $\theta_i = 0$ when $\phi_i = 0$ and $\theta_i \to 1$ as $\chi_i \to \infty$. For the second, taking the ratio of (3) across two segments of the same sector cancels μ_i, $B_{i,g}/B_{i,g'} = (x_g/x_{g'})(\omega_{i,g'}/\omega_{i,g})$, and in logs this is the relative-demand relation, independent of μ_i and hence of $\bar{K}_i$ and χ_i. □

References

Anderson, Theodore W, and Herman Rubin. 1949. "Estimation of the parameters of a single equation in a complete system of stochastic equations." *The Annals of mathematical statistics*, 20(1): 46–63.

Arslanalp, Serkan, and Takahiro Tsuda. 2014. "Tracking Global Demand for Advanced Economy Sovereign Debt." International Monetary Fund 3.

Barbiero, Omar, Falk Bräuning, Gustavo Joaquim, and Hillary Stein. 2026. "Dealer Risk Limits and Currency Returns." Social Science Research Network SSRN Working Paper Series 4728116.

Bertaut, Carol C., Valentina Bruno, and Hyun Song Shin. 2021. "Original Sin Redux." Working Paper.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2022. "Quasi-Experimental Shift-Share Research Designs." *The Review of Economic Studies*, 89(1): 181–213.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2025. "A Practical Guide to Shift-Share Instruments." *Journal of Economic Perspectives*, 39(1): 181–204.

Bulow, Jeremy, and Kenneth Rogoff. 1989. "A Constant Recontracting Model of Sovereign Debt." *Journal of Political Economy*, 97(1): 155–178.

Chaudhary, Manav, Zhiyu Fu, and Jian Li. 2023. "Corporate Bond Multipliers: Substitutes Matter." Working paper, Columbia Business School.

Chaudhary, Manav, Zhiyu Fu, and Jian Li. 2024. "Corporate Bond Elasticities: Substitutes Matter." *Working Paper*.

Coval, Joshua, and Erik Stafford. 2007. "Asset Fire Sales (and Purchases) in Equity Markets." *Journal of Financial Economics*, 86(2): 479–512.

Cox, John C., Jonathan E. Ingersoll, and Stephen A. Ross. 1985. "A Theory of the Term Structure of Interest Rates." *Econometrica*, 53(2): 385–407.

Cragg, John G., and Stephen G. Donald. 1993. "Testing Identifiability and Specification in Instrumental Variable Models." *Econometric Theory*, 9(2): 222–240.

Davidson, Russell, and James G MacKinnon. 2010. "Wild bootstrap tests for IV regression." *Journal of Business & Economic Statistics*, 28(1): 128–144.

De Leo, Pierre, Lorena Keller, and Dongchen Zou. 2024. "Speculation, forward exchange demand, and cip deviations in emerging economies." *The Wharton School Research Paper*.

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

Du, Wenxin, Carolin E. Pflueger, and Jesse Schreger. 2020. "Sovereign Debt Portfolios, Bond Risks, and the Credibility of Monetary Policy." *The Journal of Finance*, 75(6): 3097–3138.

Eaton, Jonathan, and Mark Gersovitz. 1981. "Debt with Potential Repudiation: Theoretical and Empirical Analysis." *The Review of Economic Studies*, 48(2): 289–309.

Engel, Charles, and JungJae Park. 2022. "Debauchery and Original Sin: The Currency Composition of Sovereign Debt." *Journal of the European Economic Association*, 20(3): 1095–1144.

Faia, Ester, Karen K Lewis, and Haonan Zhou. 2025. "Do investor differences impact monetary policy spillovers to emerging markets?" *Journal of International Economics*, 104100.

Fang, Xiang, Bryan Hardy, and Karen K. Lewis. 2022. "Who Holds Sovereign Debt and Why It Matters." National Bureau of Economic Research NBER Working Paper 30087.

Friedman, Benjamin M. 1977. "Financial Flow Variables and the Short-Run Determination of Long-Term Interest Rates." *Journal of Political Economy*, 85(4): 661–689.

Gabaix, Xavier, and Matteo Maggiori. 2015. "International Liquidity and Exchange Rate Determination." *The Quarterly Journal of Economics*, 130(3): 1369–1420.

Gabaix, Xavier, and Ralph S. J. Koijen. 2020. "Granular Instrumental Variables." National Bureau of Economic Research NBER Working Paper 28204.

Gabaix, Xavier, and Ralph S. J. Koijen. 2023. "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis." *Working Paper*.

Gilchrist, Simon, Bin Wei, Vivian Z. Yue, and Egon Zakrajšek. 2022. "Sovereign Risk and Financial Risk." *Journal of International Economics*, 136: 103603.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik Instruments: What, When, Why, and How." *American Economic Review*, 110(8): 2586–2624.

Greenwood, Robin, Samuel G. Hanson, Jeremy C. Stein, and Adi Sunderam. 2023. "A Quantity-Driven Theory of Term Premia and Exchange Rates." *The Quarterly Journal of Economics*, 138(4): 2327–2389.

Haddad, Valentin, Zhiguo He, Paul Huebner, Péter Kondor, and Erik Loualiche. 2025. "Causal Inference for Asset Pricing." Available at SSRN 5187305.

Ilzetzki, Ethan, Carmen M. Reinhart, and Kenneth S. Rogoff. 2022. "Rethinking Exchange Rate Regimes." In *Handbook of International Economics*. Vol. 6. Elsevier.

Itskhoki, Oleg, and Dmitry Mukhin. 2021. "Exchange Rate Disconnect in General Equilibrium." *Journal of Political Economy*, 129(8): 2183–2232.

Koehn, Michael, and Anthony M. Santomero. 1980. "Regulation of Bank Capital and Portfolio Risk." *The Journal of Finance*, 35(5): 1235–1244.

Koijen, Ralph S. J., and Motohiro Yogo. 2019. "A Demand System Approach to Asset Pricing." *Journal of Political Economy*, 127(4): 1475–1515.

Koijen, Ralph S. J., and Motohiro Yogo. 2026. "Exchange Rates and Asset Prices in a Global Demand System." *Journal of Political Economy*, 134(8): 2441–2487.

Koijen, Ralph S. J., François Koulischer, Benoît Nguyen, and Motohiro Yogo. 2021. "Inspecting the Mechanism of Quantitative Easing in the Euro Area." *Journal of Financial Economics*, 140(1): 1–20.

Lou, Dong. 2012. "A Flow-Based Explanation for Return Predictability." *The Review of Financial Studies*, 25(12): 3457–3489.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan. 2011. "Common Risk Factors in Currency Markets." *The Review of Financial Studies*, 24(11): 3731–3777.

Maggiori, Matteo, Brent Neiman, and Jesse Schreger. 2020. "International Currencies and Capital Allocation." *Journal of Political Economy*, 128(6): 2019–2066.

Moretti, Matías, Lorenzo Pandolfi, Sergio L. Schmukler, Germán Villegas-Bauer, and Tomás Williams. 2024. "Inelastic Demand Meets Optimal Supply of Risky Sovereign Bonds." International Monetary Fund IMF Working Paper 2024/227.

Olea, José Luis Montiel, and Carolin Pflueger. 2013. "A Robust Test for Weak Instruments." *Journal of Business & Economic Statistics*, 31(3): 358–369.

Ottonello, Pablo, and Diego J. Perez. 2019. "The Currency Composition of Sovereign Debt." *American Economic Journal: Macroeconomics*, 11(3): 174–208.

Pavlova, Anna, and Taisiya Sikorskaya. 2023. "Benchmarking Intensity." *The Review of Financial Studies*, 36(3): 859–903.

Rey, Hélène, Adrien Rousset Planat, Vania Stavrakeva, and Jenny Tang. 2024. "Elephants in Equity Markets." National Bureau of Economic Research Working Paper 32756.

Rochet, Jean-Charles. 1992. "Capital Requirements and the Behaviour of Commercial Banks." *European Economic Review*, 36(5): 1137–1170.

Stein, Jeremy C, and Jonathan Wallen. 2025. "The Imperfect Intermediation of Money-Like Assets." *The Journal of Finance*, 80(6): 3185–3221.

Supplement to "Asset Elasticities and Currency Risk Transfer"

Carol Bertaut Ester Faia Sebnem Kalemli-Özcan

Camilo Marchesini Simon Paetzold Martin Schmitz

November 2026

This supplement contains additional descriptive evidence, the balance tests, additional details on the model, the covered (CIP) specifications, the equations of the within-issuer design, and the model counterfactuals. Sections, tables, figures and equations of the supplement carry the prefix S. Items with the prefix D are in the Data Documentation, and items without a prefix are in the main paper.

S2. Additional Descriptive Evidence, Identification Scheme, and Balance Tests

Figure S2 shows the share of debt securities in the portfolios of each investor type, by issuer type, discussed in subsection 3.1. Table S7 reports the balance tests discussed in subsection 6.1.

Summary Statistics. Table S6 reports summary statistics for the wedges. Values are annualized wedges in percentage points, and the sample is the 1 to 5.5 year securities used in estimation. Six facts stand out. First, by construction, for

*The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Board and System, Bundesbank, European Central Bank or the Eurosystem. The paper relies on the confidential administrative datasets Security Holdings Statistics (SHS-Plus) for the euro area and the U.S. Treasury International Capital (TIC) survey, both of which have been accessed through a coauthor. Other proprietary and administrative datasets, used and matched in the analysis, have been accessed through the ECB (centralized returns database) and Goethe University Frankfurt (Consensus Forecasts). We thank Anusha Chari, Hanno Lustig, Federico Mainardi, Motohiro Yogo and Haonan Zhou for insightful comments and participants at seminars and conferences. Jakob Staps provided excellent research assistance. E-mail addresses: *Corresponding author* Sebnem Kalemli-Özcan (sebnem_kalemli-ozcan@brown.edu). Carol Bertaut (Carol Bertaut passed away on May 25, 2026), Ester Faia (faia@wiwi.uni-frankfurt.de), Camilo Marchesini (camilo_marchesini@brown.edu), Simon Paetzold (simon.paetzold@bundesbank.de), Martin Schmitz (Martin.Schmitz@ecb.europa.eu).

Figure S2

Issuer share in the total portfolio holdings of investors, by country sample See section D3 for detailed notes to this figure. SHS holdings by investor type: left panel holdings in EME, right panel holdings in AE

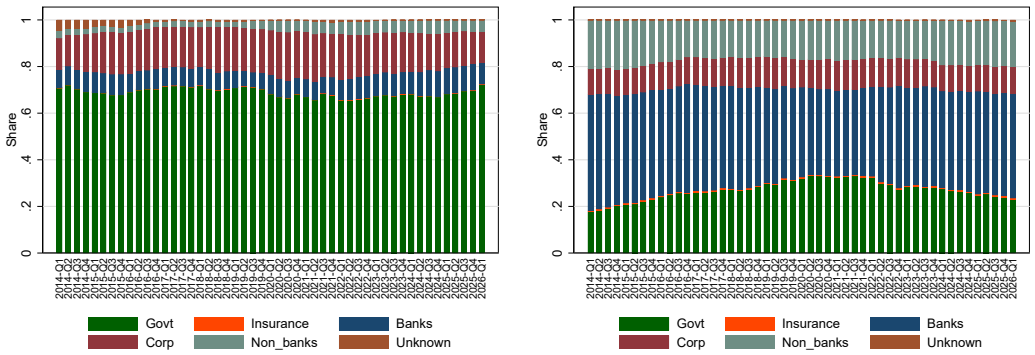

Figure S3

Investor shares in sovereign bonds, by country sample

See section D3 for detailed notes to this figure. SHS holdings by investor types: left panel EM securities, right panel AE.

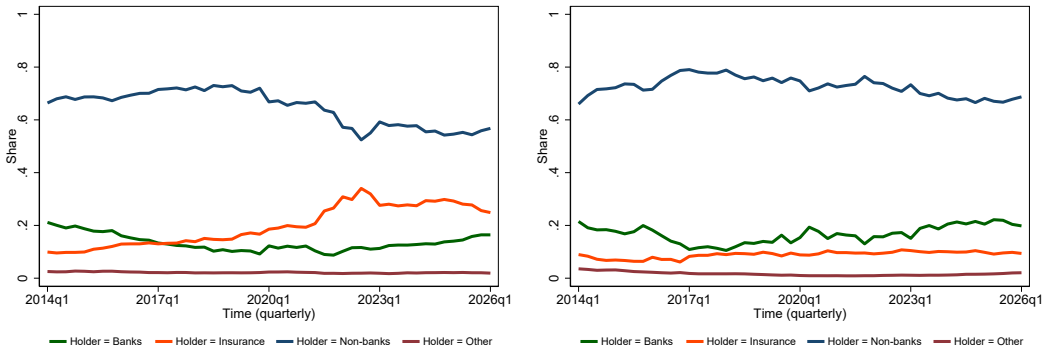

EM and AE securities issued in euros, the UIP, forwards-based CIP, and swaps-based CIP wedges are identical. Second, the local-currency wedges are large and widely dispersed, which provides the identifying variation. EM local-currency UIP averages 2.75 with a 5th to 95th percentile range of about -2.0 to $+7.7$, a spread of almost 10 percentage points. AE local-currency UIP averages 0.83 with a range of -3.1 to $+4.4$. The endogenous regressor has large two-sided cross-sectional dispersion, such that there is plenty of variation in the price of currency risk for the instrument to trace the demand curve. Third, EM wedges vary much more than AE wedges, both in level (2.75 against 0.83) and in the tails. Since elasticities are identified off movements in the wedge, our design is most powerful for EM local-currency bonds. Fourth, for EM local currency, the UIP wedge of 2.75 far exceeds the swaps-based CIP wedge of 0.085, such that almost the entire wedge, about 2.67 percentage points, is the uncovered component, namely expected depreciation plus the currency risk premium, rather than a covered basis. What we call the currency wedge is therefore, at least unconditionally, overwhelmingly currency risk and not other premia. Fifth, the sample is built around local-currency bonds with euro bonds as the single benchmark, and we deliberately exclude dollar bonds to keep one clean benchmark. The sample is larger for local-currency bonds (6,798 EM and 7,973 AE observations), which gives us statistical power exactly where currency risk lies. Sixth, EM and AE securities have similar residual maturities of around 3 years, which rules out duration as an explanation for differences in EM and AE elasticities we documented in the main text.

In sum, Table S6 validates the wedge construction and shows that the endogenous variable has large, two-sided, currency-driven dispersion.

Figure S4 illustrates the identification strategy described in subsection 5.3.

Figure S4

Identification Scheme

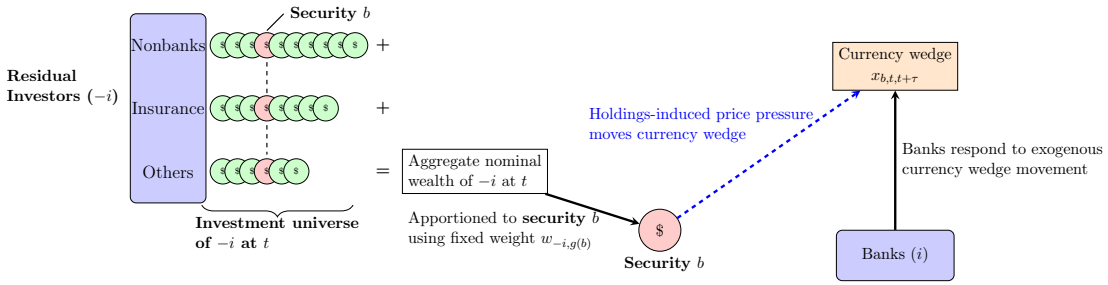

Balance tests. Without fixed effects (columns 1 to 4), some characteristics (log amount outstanding, collateral eligibility, euro denomination, and some seniority categories) are significantly related to the shares, as, unconditionally, exposure

Table S6

Summary Statistics

Variable	Mean	Median	5th pct	95th pct	N
Panel A: Emerging Markets					
Subpanel (a): Securities denominated in euros					
$x_{b,t}^{UIP} = x_{b,t}^{CIP}$ (swaps-based) $= x_{b,t}^{CIP}$ (forwards-based)	1.039	0.821	0.281	2.585	817
Residual maturity	3.272	3.310	1.148	5.279	817
Subpanel (b): Securities denominated in local currency					
$x_{b,t}^{UIP}$	2.752	2.755	-2.025	7.686	6798
$x_{b,t}^{CIP}$ (swaps-based)	0.085	0.306	-3.311	2.227	6798
$x_{b,t}^{CIP}$ (forwards-based)	1.138	1.107	-2.343	4.256	6798
Residual maturity	3.045	2.956	1.068	5.197	6798
Panel B: Advanced Economies					
Subpanel (a): Securities denominated in euros					
$x_{b,t}^{UIP} = x_{b,t}^{CIP}$ (swaps-based) $= x_{b,t}^{CIP}$ (forwards-based)	0.214	0.180	-0.347	0.896	111
Residual maturity	2.934	2.992	1.085	4.773	111
Subpanel (b): Securities denominated in local currency					
$x_{b,t}^{UIP}$	0.825	1.005	-3.076	4.404	7973
$x_{b,t}^{CIP}$ (swaps-based)	0.399	0.358	-0.438	1.411	7973
$x_{b,t}^{CIP}$ (forwards-based)	0.254	0.322	-0.602	1.026	7973
Residual maturity	2.989	2.926	1.019	5.225	7973

Notes: Summary statistics for emerging-market and advanced-economy arbitrage deviations (UIP and CIP). Values are annualized wedges in percentage points; the sample includes the 1 to 5.5 year securities used in estimation.

tilts toward large, collateral-eligible bonds. However, almost all of this imbalance is between issuers and between periods. When we add issuer and time fixed effects (columns 5 to 8), the collateral-eligibility and euro-denomination correlations become insignificant, and with issuer-times-time fixed effects (columns 9 to 12) they remain so. Callable, coupon, maturity bucket, seniority-3, and credit quality are essentially never imbalanced within issuer-times-time cells. The attributes one would worry about most, collateral eligibility first of all, do not vary with the shares once we compare within issuer-times-time cells. Log amount outstanding is the only characteristic that remains correlated with the shares after fixed effects: for all four sectors under issuer-plus-time, and for insurers and pension funds (at 10%) and other investors (at 5%) under issuer-times-time. Larger and more liquid issues carry larger exposure shares. However, a characteristic predicts the shares only if it is omitted, and therefore, following Borusyak, Hull and Jaravel (2022), we include log amount outstanding among the controls in our regressions.

More broadly, the conditional share exogeneity required by Goldsmith-Pinkham, Sorkin and Swift (2020) and Borusyak, Hull and Jaravel (2025) holds. Under issuer-times-time fixed effects, the joint F-tests fail to reject balance for three of the four sectors. The exception is the other-investors group, which has a very small aggregation weight. For the sectors that bear most weight in the aggregate, the predetermined shares are jointly uncorrelated with observables even in the most demanding comparison.

To sum up, the exclusion restriction fails if $\mathrm{Cov}(w_{-i,g}, \eta_{i,b,t}) \neq 0$. Table S7 shows that the underlying security-level shares are balanced on the observable determinants of demand, with the single exception of a liquidity proxy, and by the averaging argument above the segment shares inherit this balance. There is therefore a credible case for the exogeneity of the shares, our identifying assumption, once we control for liquidity. The only imbalance in amount outstanding that survives the most demanding comparison is in the other-investors group, which lumps together households, government, and the like and whose size is negligible in our sample.

S3. Additional Model Details

This section provides additional details on the model: the Solvency II example, the exchange-rate process, the sovereign's default decision, the regulatory constraints

Table S7

Balance Tests

Dependent variable: $w_{b,-i,t_{0_{-i}}}$ (bps)	No fixed effects				Issuer FE, Time FE				Issuer x Time FE			
	ICPF	OFI	MFI	OTH	ICPF	OFI	MFI	OTH	ICPF	OFI	MFI	OTH
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Callable	-0.967	-0.761	-0.575	-0.153	-0.953	-0.688	-0.747	-0.464	-0.468	-0.247	-0.717	-0.168
	(0.792)	(0.712)	(0.471)	(0.523)	(0.702)	(0.595)	(0.471)	(0.560)	(0.608)	(0.531)	(0.668)	(0.358)
Log amount outstanding (EUR)	0.261**	0.202*	0.187*	0.228**	0.690**	0.649**	0.460**	0.289**	0.495*	0.510	0.349	0.224**
	(0.121)	(0.107)	(0.094)	(0.085)	(0.271)	(0.255)	(0.209)	(0.116)	(0.293)	(0.328)	(0.244)	(0.104)
Coupon	-0.042	-0.041	-0.007	-0.030	0.002	0.007	0.017	-0.013	0.024	0.031	0.019	-0.005
	(0.042)	(0.048)	(0.030)	(0.022)	(0.011)	(0.016)	(0.018)	(0.012)	(0.026)	(0.024)	(0.018)	(0.010)
Maturity bucket	-0.175	-0.167	-0.066	-0.054	-0.095	-0.127	0.003	-0.079	-0.077	-0.067	-0.016	-0.035
	(0.171)	(0.179)	(0.117)	(0.112)	(0.140)	(0.145)	(0.100)	(0.109)	(0.113)	(0.114)	(0.070)	(0.065)
Euro denomination	-1.310*	-1.247*	-0.548	-1.494***	1.583	1.789	1.774	0.305	3.506	6.190	5.725	-1.121
	(0.761)	(0.674)	(0.452)	(0.492)	(1.095)	(1.149)	(1.131)	(0.197)	(2.690)	(4.118)	(3.822)	(1.016)
Seniority = 2	-0.899*	-0.808*	-0.560	-0.521**	-0.499	-0.370	-0.343	-0.462*	-0.469	-0.372	-0.379	-0.322*
	(0.508)	(0.475)	(0.371)	(0.251)	(0.495)	(0.414)	(0.412)	(0.269)	(0.288)	(0.239)	(0.228)	(0.168)
Seniority = 3	-0.619*	-0.438	-0.179	-0.115	-0.067	-0.605*	0.055	-0.124	-0.407	-1.068	-0.234	0.075
	(0.355)	(0.404)	(0.125)	(0.201)	(0.572)	(0.336)	(0.340)	(0.272)	(0.562)	(0.648)	(0.433)	(0.282)
Collateral eligibility	-3.820***	-3.422***	-2.170***	-2.949***	1.559	2.097	2.060	-1.000	4.834	6.099	6.334	-2.481
	(0.743)	(0.636)	(0.566)	(0.668)	(1.794)	(2.001)	(2.075)	(0.618)	(3.584)	(4.454)	(4.538)	(1.558)
Credit quality	0.187	0.168	-0.023	0.211*	0.210	0.180	0.201	0.128	0.034	0.025	0.054	0.117*
	(0.273)	(0.252)	(0.223)	(0.113)	(0.209)	(0.197)	(0.145)	(0.084)	(0.126)	(0.131)	(0.113)	(0.065)
Observations	1239	1197	1267	1154	1231	1189	1260	1147	909	891	937	842
R-squared	0.279	0.234	0.159	0.408	0.513	0.470	0.374	0.591	0.743	0.721	0.649	0.787
Adjusted R-squared	0.273	0.228	0.153	0.404	0.474	0.425	0.325	0.555	0.640	0.608	0.515	0.692
Joint test (p)	0.000	0.000	0.000	0.000	0.000	0.000	0.002	0.000	0.259	0.669	0.796	0.000
Issuer FE	NO	NO	NO	NO	YES	YES	YES	YES	NO	NO	NO	NO
Time FE	NO	NO	NO	NO	YES	YES	YES	YES	NO	NO	NO	NO
Issuer x Time FE	NO	NO	NO	NO	NO	NO	NO	NO	YES	YES	YES	YES

Notes: The column name denotes sector i, with ICPF being 'insurance', OFI being 'nonbanks', MFI being 'banks', and OTH being 'other'. Seniority consists of three categories: (1) Unguaranteed/Senior/Unsecured, (2) Government-Treasury-guaranteed/Senior/Unsecured, (3) Other. The reference category is (1). Credit quality consists of 7 categories: AAA, AA, A, BBB, BB, B, and C.

of each sector, benchmark-centered mandates, and a variant in which investors choose how much currency exposure to hedge, which delivers the CIP wedge.

S3.1. The Solvency II Example

To fix ideas, let us see how the constraint nests the ones imposed by the Solvency II regulation. For each currency segment g the capital charge is $\mathrm{SCR}_{i,g} = a_{i,g}\, B_{i,g}$, where $B_{i,g}$ is the sector's exposure to that segment, the holdings of (2), and $a_{i,g}$ is the per-unit stress the module applies to it. Aggregating the charges across modules with the correlation matrix gives: $\mathrm{SCR}_i = \left(\sum_g \sum_{g'} \mathrm{Corr}_{gg'}\, \mathrm{SCR}_{i,g}\, \mathrm{SCR}_{i,g'} \right)^{1/2} \leq \mathrm{OF}_i$. Hence, the solvency condition is the square root of a quadratic form in exposures, bounded by own funds. Squaring this last expression gives (2) where: $\left[\Omega_i\right]_{gg'} = 2\,\mathrm{Corr}_{gg'}\, a_{i,g}\, a_{i,g'}, \qquad \bar{K}_i = \mathrm{OF}_i^2$. Setting the off-diagonal correlations to zero gives the diagonal case used in the quantitative model solved numerically below, $\omega_{i,g} = 2\, a_{i,g}^2$.

S3.2. Exchange Rates

To keep the model tractable, we do not add a monetary side. To determine the exchange rate we rely on a no-arbitrage condition on risk-free rates under an affine structure, $\frac{e'}{e} = \frac{M'^*}{M'}$. This requires a process for the stochastic discount factors M and M^*, which we model as affine in one common global risk factor $\tilde{\omega}$. We mark the parameters of this block with a tilde to keep them apart from the demand-side objects:[1] $-\log(M') = \tilde{\delta} + \left(\tilde{\psi} + \frac{\tilde{\lambda}^2}{2}\right)\tilde{\omega} + \tilde{\lambda}\,\tilde{\omega}^{\frac{1}{2}}\tilde{\epsilon}'$, $-\log(M'^*) = \left(1 + \frac{\tilde{\lambda}^{*2}}{2}\right)\tilde{\omega} + \tilde{\lambda}^*\tilde{\omega}^{\frac{1}{2}}\tilde{\epsilon}'$. We normalize the foreign block, $\tilde{\delta}^* = 0$ and $\tilde{\psi}^* = 1$, without loss of generality: only the domestic-minus-foreign differences $\tilde{\delta}$, $\tilde{\psi} - 1$, and $\tilde{\lambda} - \tilde{\lambda}^*$ enter the exchange rate $e'/e = M'^*/M'$ and are identified.

S3.3. Government Strategic Default

Strategic Default. The government can default, and if it does so, it defaults on all of its debt. It repays if the maximized value from repayment, V^R, is bigger than that of default, V^D, conditional on the aggregate states, $s = [y, B, B^*, e]$, hence

$$d = \begin{cases} 0, & \text{if } V^R(s) \geq V^D(s) \\ 1, & \text{otherwise.} \end{cases}$$

1 $\tilde{\omega}$ follows a square-root process as in Cox, Ingersoll and Ross (1985): $\tilde{\omega}' = (1-\tilde{\phi})\tilde{\kappa} + \tilde{\phi}\,\tilde{\omega} + \tilde{\sigma}\,\tilde{\omega}^{\frac{1}{2}}\tilde{\epsilon}'$, where $\tilde{\epsilon}$ is independently and normally distributed with mean zero and variance one.

The values of repayment and default can be derived as follows. If the government decides to repay, the economy consumes its endowment and borrowed resources, less the repayment of existing borrowing, and obtains the discounted future value. If the government decides to default, it faces a default cost and can only consume y^d. Additionally, it is excluded from financial markets and re-enters only with probability ϑ in the next period with a zero debt position. Hence:

$$V^R(y, B, B^*, e) = \max_{B', B^{*\prime}} u(y + qB' - B + \tag{S3.1}$$
$$eq^* B^{*\prime} - eB^*) + \varrho \int_{y', e'} V^0(y', B', B^{*\prime}, e') f(y' \mid y) f(e' \mid e) de' dy',$$

where $V^0 = \max\{V^R, V^D\}$. The value of default is:

$$V^D(y; y^d) = u(y^d) + \varrho \int_{y', e'} [\vartheta V^0(y', 0) + (1 - \vartheta) V^D(y', 0)] f(y' \mid y) f(e' \mid e) de' dy'. \tag{S3.2}$$

The default probability is given by δ over the default set $D(s) = [s : V^R(s) < V^D(s)]$. In this default probability, the government takes into account the expected future income realization. Given this relationship, the probability of default is given by:

$$\delta(s') = \int_{D(s)} f(y' \mid y) f(e' \mid e) de' dy'. \tag{S3.3}$$

S4. Other Regulatory Constraints

In the euro area, insurers and pension funds are subject to the Solvency II Directive, a risk-based solvency-capital regime that attaches capital charges to market and credit risk. These charges shape how much currency and country risk these investors bear. On the other hand, investment funds operate under lighter prudential constraints. Solvency II has evolved over our sample. Since the 2024 to 2026 review it has moved toward lower effective risk sensitivity. Banks are governed by the Basel framework, revised in three waves (Basel I, II, and III). Basel III combines a Tier 1 minimum[2] with surcharges that national regulators may impose, some common across currency areas and some jurisdiction-specific.[3]

The friction in the model is denoted by $\gamma_{i,g} = \mu_i \omega_{i,g}$. The segment weight $\omega_{i,g}$ is determined in each sector by the relevant regulator. We discuss here in more detail three specific regulations that apply to the investors in our sample. For

2 This consists of Common Equity Tier 1 (CET1) and Additional Tier 1 (AT1), implying a common-equity target no lower than 6 percent of risk-weighted assets.

3 Examples include bank-specific stress capital buffers, the Pillar 2 requirement assigned by the ECB in the euro area, and the capital surcharge for global systemically important banks (G-SIBs) imposed by the Federal Reserve.

insurance corporations and pension funds, Solvency II adds a regulatory charge on any position held in a currency other than that of the insurer's liabilities and left unhedged. This is on top of the credit charge on the bond itself. Formally $\omega_{i,\ell} > \omega_{i,*}$. For banks, Basel and the liquidity rules require little or no equity capital when a sovereign bond is in the sovereign's own currency and is funded in that currency, and impose a higher charge otherwise. Hence, even in this case the local-currency weight exceeds the hard-currency one. Nonbanks face little to no regulation, and hence, the weight is a penalty for deviating from the benchmark: $\omega_{i,g}$ follows index membership and weights. In each case, $\omega_{i,g}$ is a sector-by-segment component, whose primitive parameters can be inferred from the regulatory laws. The ratio $\omega_{i,\ell}/\omega_{i,*}$, the relative cost of holding local- rather than hard-currency debt of the same issuer, is steepest for insurers and pension funds, whose regulation is the most currency-contingent.

Table S8 summarizes the cases. The Basel credit-risk requirement is linear in exposures, $\sum_g a_g^{\text{cr}} B_{i,g}$ with a_g^{cr} being the rating-based weight. A bank faces both this linear constraint and the quadratic risk-based limits. Which one binds at the margin is an empirical question. Since euro-area banks held capital well above their requirements throughout our sample, the linear constraint is slack, and the risk-based limits shape the allocation across segments.

Table S8

The risk budget across sectors

Sector	Binding rule	Ω_i	$\bar{K}_i$
Insurers and pension funds (ICPF)	Solvency II standard formula, SCR $\leq$ OF, spread and currency modules	$2\,\text{Corr}_{gg'} a_{i,g} a_{i,g'}$, charges from the modules	OF_i^2, from published coverage ratios
Banks (MFI)	Market-risk capital on the trading book, expected shortfall under FRTB	$2k_i^2\Sigma_i$, Σ_i the return covariance of the segments	squared market-risk capital allocated to the book
Investment funds (OFI)	UCITS global exposure, value at risk, or a tracking-error mandate	$2k_i^2\Sigma_i$, or the covariance of active positions for tracking error	squared risk limit in units of net asset value
Other (OTH)	Internal economic capital	internal covariance matrix	squared internal limit

S5. Benchmark-Centered Mandates

For a sector whose mandate is a tracking-error limit relative to an index, as in Pavlova and Sikorskaya (2023), the risk measure is centered on the benchmark portfolio λ_i rather than on zero, $\frac{1}{2}(\mathbf{B}_i - \lambda_i)^\top \Omega_i (\mathbf{B}_i - \lambda_i) \leq \bar{K}_i$. The first-order condition (3) becomes $\mathbf{B}_i = \lambda_i + \mu_i^{-1}\Omega_i^{-1}x$, that is, $B_{i,g} = \lambda_{i,g} + x_g/\gamma_{i,g}$. The index weight shifts the intercept of demand while the mandate metric continues to govern its slope, such that every result in the main text holds with holdings measured as deviations from the index. In particular, the relative-demand relation of Proposition 2 holds for $B_{i,g} - \lambda_{i,g}$, and the sector intercept of (21) combines the mandate ratio with the index ratio whenever the latter is not zero. Since index weights are observable, the two components can be separated, which is how we read the nonbank sector in the empirical section.

S6. Deriving CIP Through Investors' Optimal Hedging

Our empirical strategy estimates the elasticity of holdings to instrumented arbitrage deviations, hedged and unhedged. To give the covered specification the same structural interpretation as the uncovered one, we extend the investor problem with an endogenous hedging choice. We define the covered wedge like the uncovered wedge of (1) with the forward in place of the expected spot. Different sectors are subject to different regulatory requirements for currency risk, but for all sectors covering a share c_i of the local-currency holding lowers the regulatory weight $\omega_{i,\ell}(c_i)$ attached to that currency segment.

Hedging extension. Sector i covers a share $c_i \in [0,1]$ of its local-currency position by solving:

$$\max_{B_{i,\ell},\, B_{i,*},\, c_i} \; \left[c_i\, x_\ell^{\text{CIP}} + (1-c_i)\, x_\ell\right] B_{i,\ell} + x_*\, B_{i,*} \qquad \text{s.t.} \qquad \tfrac{1}{2}\,\omega_{i,\ell}(c_i)\, B_{i,\ell}^2 + \tfrac{1}{2}\,\omega_{i,*}\, B_{i,*}^2 \;\leq\; \bar{K}_i, \tag{S6.1}$$

where $\omega_{i,\ell}(c_i) = \omega_{i,*} + a_{2,i}(1-c_i)^2$ is the regulatory weight on the local-currency segment, decreasing in the covered share, with $a_{2,i} > 0$ being the currency charge, and $x_\ell^{\text{CIP}} = \frac{1-\delta}{1+r}\frac{e}{f} - q$ being the covered wedge. The latter is basically the uncovered wedge x_ℓ of (1) with the market forward f in place of $E(e')$. Let us define the

multiplier μ_i and $\Delta_f \equiv x_\ell^{\text{CIP}} - x_\ell$. The first-order conditions for both bonds and for the currency coverage are:

$$B_{i,\ell} = \frac{x_\ell + c_i \Delta_f}{\mu_i \, \omega_{i,\ell}(c_i)}, \qquad B_{i,*} = \frac{x_*}{\mu_i \, \omega_{i,*}}, \qquad -\Delta_f = \mu_i \, a_{2,i} \, (1 - c_i) \, B_{i,\ell}. \tag{S6.2}$$

The first two are the demands of (3) with the local-currency weight at the covered share, so $\gamma_{i,\ell}(c_i) = \mu_i \omega_{i,\ell}(c_i)$ falls in c_i. The third determines the covered share: the sector covers until the premium it pays, $-\Delta_f$, equals the equity capital that covering frees.

Solving the covered-share condition in (S6.2) gives:

$$1 - c_i = \frac{-\Delta_f}{\mu_i \, a_{2,i} \, B_{i,\ell}}, \tag{S6.3}$$

so the covered share rises with the tightness of the mandate μ_i, with the currency charge $a_{2,i}$, and with the size of the position, and falls with the insurance premium paid. If a sector is required to cover currency risk (for instance, banks), this corresponds to the limit $a_{2,i} \to \infty$, where $c_i \to 1$.

Aggregation. Summing the local-currency demand in (S6.2) across sectors and imposing market clearing gives:

$$B_\ell = x_\ell \sum_i \frac{m_i}{\gamma_{i,\ell}(c_i)} + \sum_i \frac{m_i \, c_i \, \Delta_f}{\gamma_{i,\ell}(c_i)}, \qquad \Gamma_\ell(\mathbf{c}) = \Big(\sum_i \frac{m_i}{\gamma_{i,\ell}(c_i)} \Big)^{-1}, \tag{S6.4}$$

which is (8) where the frictions are computed at the optimal covered shares. Hedging flattens the schedule, since each $\gamma_{i,\ell}(c_i)$ falls, and shifts it through the second term.

S6.1. Demand for Currency Risk: the Covered (CIP) Slopes

The construction of the instrument at the segment level allows the covered (CIP) specifications to be identifiable under the same design. The covered basis, whose microfoundation we give in Appendix S6 by extending the investor's problem with an endogenous hedging choice, is, to a first-order approximation, a property of the issuer-currency segment rather than of the individual security. Most of the cross-sectional variation in $x_{b,t}^{CIP}$ is between segments, as opposed to within segments, and hence by varying at the same level as the covered basis, the segment-share instrument is relevant. We therefore use the same instrument for the covered and uncovered specifications, without using segment fixed effects, because they would absorb the predetermined share $w_{-i,g}$ on which identification rests.

Table S9 reports the second-stage slopes when the currency wedge is measured on

a covered, swaps-based basis, $x_{b,t}^{CIP}$. Table S10 reports the corresponding implied elasticities. The specification, instrument, and control variables are identical to the case of the uncovered wedge presented in the main text, with the same predetermined segment-share instrument $z_{-i,b,t}$ of (29) and no additional fixed effects. Since the covered wedge nets out expected depreciation and the currency risk premium, leaving the near-arbitrage basis, comparing its slopes with the uncovered ones isolates the extent to which each sector responds to the hedged, rather than to the currency-bearing, component of the excess return.

The first stage is strong almost everywhere across the table, with the first-stage F ranging between about 4 and 264 and clearing the weak-instrument threshold in every cell except one, namely the residual sector in the emerging-market local-currency column ($F = 3.8$), which we do not interpret. Nonbanks, which carry the bulk of the aggregation weight, display a positive and precisely estimated slope in every sample—0.44 and 0.34 for emerging markets, 1.64 and 1.78 for advanced economies, with first-stage F between 92 and 264. Once currency exposure is hedged away, the dominant sector behaves as an ordinary downward-sloping demander, buying more of a segment when its covered excess return rises. This is the same sign the sector displays on the uncovered wedge, such that nonbanks absorb the covered and uncovered components alike rather than reversing across them.

The remaining sectors carry less weight. Insurers and pension funds are strongly positive for advanced economies (3.69 and 3.72) and so is coefficient for the residual sector, while banks are negative for advanced economies, significantly so in local currency (-5.49). For emerging markets, insurers display negative and significant covered slopes (-0.46 and -1.22), as do banks in local currency (-0.96). We do not extrapolate much from those estimates primarily since the scaling of the elasticities may be mechanically driven by the size of the covered currency wedge which is small (albeit positive). The elasticity rescales the slope by the covered wedge, and since the latter is close to zero, it automatically inflates the elasticities. Therefore, large elasticities do not necessarily mean that investors are very responsive. We therefore interpret the covered specification only when it is precisely estimated and carries some sensible results, not mechanically driven by statistical biases.

S6.2. The Covered (CIP) Supply Slope

Table S11 reports the supply coefficient α_1 of (35) when the currency wedge is measured on the covered, swaps-based basis. The covered specification delivers a clean disciplining response, as the coefficient is negative throughout and significant

Table S9
CIP (swaps-based): Slope of the Demand Function

Dependent variable: Log Holdings				
	EM + Local & EUR	EM + Local	AE + Local & EUR	AE + Local
Panel A: Insurance				
$x_{b,t}^{CIP}$	-0.461***	-1.223***	3.687***	3.723***
	(0.116)	(0.355)	(1.008)	(0.991)
Observations	5929	5085	6275	6190
F-statistic	98.627	42.478	108.978	131.673
Panel B: Nonbanks				
$x_{b,t}^{CIP}$	0.436***	0.335***	1.638***	1.783***
	(0.109)	(0.077)	(0.196)	(0.249)
Observations	6903	6054	7328	7229
F-statistic	91.555	104.794	221.148	263.501
Panel C: Banks				
$x_{b,t}^{CIP}$	0.104	-0.955***	-1.705	-5.494**
	(0.104)	(0.267)	(1.569)	(2.794)
Observations	3389	2539	4256	4166
F-statistic	144.322	31.866	21.806	29.403
Panel D: Other				
$x_{b,t}^{CIP}$	-0.140	-0.910*	3.644***	3.340***
	(0.170)	(0.552)	(1.068)	(1.059)
Observations	3915	3060	5232	5150
F-statistic	29.787	3.786	36.514	41.314

Notes: The coefficients are the slopes of demand estimated from the two-stage regression (31), with the wedge measured on a covered, swaps-based basis. Each panel reports the slopes for one investor type. The last row in each panel reports the Cragg and Donald (1993) F-statistic. Robust standard errors clustered at the quarter level are in parentheses. One, two, and three stars denote significance at the 10, 5, and 1 percent levels.

Table S10
Implied Elasticities: CIP (swaps-based)

	N	Median	SD	5th pct.	95th pct.	Weight
Panel A: EM bonds denominated in local currency						
Insurance	5085	-44.064	28.699	-117.457	-23.857	0.046
Nonbanks	6054	11.868	7.844	6.479	31.722	0.802
Banks	2539	-37.165	24.175	-98.199	-18.823	0.145
Other	3060	-31.169	20.877	-85.197	-17.604	0.008
Euro area total		1.893				1.000
Panel B: EM bonds denominated in local currency and EUR						
Insurance	5929	-16.246	10.706	-43.719	-8.958	0.064
Nonbanks	6903	15.135	10.118	8.424	40.760	0.712
Banks	3389	3.826	2.565	2.023	10.389	0.202
Other	3915	-4.676	3.176	-13.141	-2.707	0.021
Euro area total		10.407				1.000
Panel C: AE bonds denominated in local currency						
Insurance	6190	125.241	88.514	71.257	351.047	0.180
Nonbanks	7229	60.052	43.201	34.135	172.180	0.532
Banks	4166	-194.004	138.370	-564.749	-105.158	0.186
Other	5150	113.347	77.791	63.928	307.081	0.102
Euro area total		30.034				1.000
Panel D: AE bonds denominated in local currency and EUR						
Insurance	6275	124.016	87.537	70.560	347.510	0.184
Nonbanks	7328	55.408	39.625	31.358	157.365	0.519
Banks	4256	-60.574	42.836	-175.237	-32.653	0.190
Other	5232	123.939	84.780	69.738	334.990	0.108
Euro area total		53.381				1.000

Notes: The table reports the implied elasticities computed from formula (32) using the covered (CIP) demand slopes in Table S9. Panel A is for EM bonds in local currency, Panel B for EM bonds in local currency and euro, Panel C for AE bonds in local currency, and Panel D for AE bonds in local currency and euro.

for emerging-market issuers (-0.034 and -0.053), with first stages far above any conventional threshold (F between 211 and 485). A one-percentage-point widening of the covered currency wedge lowers subsequent issuance growth by three to five percent. The advanced-economy coefficients are also negative but smaller and imprecisely estimated, consistent with these issuers facing a weaker disciplining margin on the part of investors.

We note that, since the currency-segment-share instrument traces variation at the issuer-currency level and the covered wedge is itself an issuer-currency object, the first stage is unsurprisingly very strong. On the contrary, the uncovered wedge carries a large security-specific and expectational component, from expected depreciation and security-level returns, that the segment instrument does not move, which weakens its relevance.

Table S11
CIP (swaps-based): Slope of the Supply Function

Dependent variable: $\Delta \log (Q_{b,t+4})$				
	EM + Local & EUR	EM + Local	AE + Local & EUR	AE + Local
$x_{b,t}^{CIP}$	-0.034**	-0.053***	-0.044	-0.046
	(0.014)	(0.015)	(0.053)	(0.053)
Observations	6068	5320	6138	6045
F-statistic	211.132	370.197	443.336	485.054

Notes: The coefficients are the slopes of supply from the two-stage regression (35), with the wedge measured on a covered, swaps-based basis. The last row reports the Cragg and Donald (1993) F-statistic. Robust standard errors clustered at the quarter level are in parentheses. One, two, and three stars denote significance at the 10, 5, and 1 percent levels.

S7. Within-Issuer Design: Specification and Supply Estimates

This section gives the equations of the within-issuer design of subsection 6.4 and reports its supply estimates. Given the definitions in subsection 5.1, sector i's portfolio share in security b at time t is

$$w_{i,b,t} = \frac{D_{i,b,t}}{\sum_{b' \in \mathcal{B}_{i,t}} D_{i,b',t}}, \tag{S7.1}$$

where $\mathcal{B}_{i,t}$ is the set of securities sector i holds at time t. The instrument is built at the security level. Let $t_{0_{-i}}$ be the first quarter in which security b is held by at least one of the sectors in $-i$, and let $w_{-i,b,t_{0_{-i}}}$ be the share of security b in the portfolio of $-i$ in that quarter, as in the balance tests of subsection 6.1. The share in the instrument is the excess of this portfolio share over the weight that

security b would receive in an equal-weighted portfolio of the $|\mathcal{B}_{-i}|$ securities in the investment universe of $-i$:

$$w^{\text{ex}}_{-i,b} = w_{-i,b,t_{0_{-i}}} - \frac{1}{|\mathcal{B}_{-i}|}. \tag{S7.2}$$

The instrument scales this time-invariant excess share by the growth of the residual sectors' nominal wealth in sovereign bonds since the same quarter:

$$z^{\text{sec}}_{-i,b,t} = w^{\text{ex}}_{-i,b} \frac{\sum_{b' \in \mathcal{B}_{-i,t}} \sum_{j \neq i} D_{j,b',t}}{\sum_{b' \in \mathcal{B}_{-i,t_{0_{-i}}}} \sum_{j \neq i} D_{j,b',t_{0_{-i}}}}. \tag{S7.3}$$

Unlike the baseline instrument (29), $z^{\text{sec}}_{-i,b,t}$ is a product in levels and varies across the securities of an issuer, such that it is not absorbed by issuer-by-time fixed effects. The two stages are

$$x_{b,t} = \pi^{w}_{0i} + \pi^{w}_{1i}\, z^{\text{sec}}_{-i,b,t} + \text{FE}_{n(b),t} + \pi^{w\prime}_{2i} X_{b,t} + \varepsilon_{b,t}, \tag{S7.4}$$

$$\log\left(w_{i,b,t}\right) = \beta^{w}_{0i} + \beta^{w}_{1i}\, x_{b,t} + \text{FE}_{n(b),t} + \beta^{w\prime}_{2i} X_{b,t} + \eta_{i,b,t}, \tag{S7.5}$$

where $n(b)$ is the issuer of security b, $\text{FE}_{n(b),t}$ are issuer-by-time fixed effects, which subsume issuer and time effects, and $X_{b,t}$ contains residual-maturity and rating controls. The coefficient β^{w}_{1i} is the semi-elasticity of sector i's portfolio share with respect to the currency wedge. The supply counterpart replaces the dependent variable of (S7.5) with the growth of the amount outstanding of the security over the following quarter, $\Delta \log(Q_{b,t+1})$, and is estimated on the securities held by euro-area nonbanks. The sample is 2014Q1 to 2022Q2 and includes securities with a residual maturity of 1 to 5 years. Because the instrument has few time clusters, the tables report 95 percent Anderson and Rubin (1949) confidence intervals, which are robust to weak instruments, computed with the wild restricted efficient bootstrap of Davidson and MacKinnon (2010), together with the effective F-statistic of Olea and Pflueger (2013).

Table S12
Within-Issuer Design: UIP Wedge and Supply

Dependent variable: Change in amount outstanding, $\Delta \log(Q_{b,t+1})$				
	EM + Local & EUR	EM + Local	AE + Local & EUR	AE + Local
$x^{UIP}_{b,t}$	-0.154**	-0.103**	0.578	0.706
	[-0.248,-0.067]	[-0.150,-0.046]	[0.166,1.681]	[0.187,2.378]
Observations	6025	5419	5511	5396
Effective F-statistic	23.507	41.096	2.282	1.828
Quarter FE	Yes	Yes	Yes	Yes
Issuer Country FE	Yes	Yes	Yes	Yes
Quarter x Issuer Country FE	Yes	Yes	Yes	Yes
Rating controls	Yes	Yes	Yes	Yes
Maturity controls	Yes	Yes	Yes	Yes

Notes: Within-issuer supply estimates on the securities held by euro-area nonbanks, 2014Q1 to 2022Q2. The dependent variable is the growth of the amount outstanding of the security over the following quarter. The 95 percent Anderson and Rubin (1949) confidence intervals are reported in square brackets. They are computed using 999 wild restricted efficient bootstrap (WRE) replications following Davidson and MacKinnon (2010), clustering at the time level. The Olea and Pflueger (2013) critical value for the effective F-statistic at a worst-case bias of 30 percent is 12.039. Two stars next to the coefficient estimate denote significance at the 5 percent level.

S8. Model Counterfactuals

S8.1. Counterfactual I: Prudential Regulation of Euro-Area Creditors

With the demand elasticities in hand, we calibrate the model to match them and use it to examine counterfactual policy changes. The first question is how emerging-market sovereign borrowing conditions respond when the prudential regime of their euro-area creditors changes. The natural example is the reform of Solvency II, which pertains to euro-area insurers and pension funds. Our calibration window (2014 to 2022) covers the pre-review regime. The 2024 to 2026 review, adopted as an amending Directive in 2025 with national transposition due by end-January 2027, loosens it. It cuts the cost-of-capital rate underlying the risk margin from 6% to 4.75%, reducing the risk margin by roughly one-fifth, introduces a time-dependent risk-tapering factor, and broadens access to reduced capital charges on long-term equity. We ask how much such a relaxation changes the demand of insurers and pension funds for emerging-market bonds, and what that does to emerging-market borrowing costs.

The mandate structure of section 4 determines where such a reform enters. Solvency II sets the capital charges $\omega_{\mathrm{ICPF},g}$ and, through eligible own funds, the loss-absorbing capital $\bar{K}_{\mathrm{ICPF}}$, and therefore changes one sector's friction $\gamma_{\mathrm{ICPF},g} = \mu_{\mathrm{ICPF}}\,\omega_{\mathrm{ICPF},g}$. By (8), the friction affects aggregate demand through the composite, weighted by the share of the segment the sector holds, and the issuer internalizes the change through that composition. The review targets the risk margin and long-term-equity charges rather than sovereign risk weights directly, such that we take from it the direction of the change and discipline its size as follows.

Disciplining the size of the friction change. The direction of the friction change comes from the regulation, its size from our elasticity estimates. Through the demand-system mapping $\beta_{1i} = 1/(\gamma_{i,g}\bar{D}_{i,g})$, the friction is the inverse of the sector's estimated semi-elasticity, such that the perturbation can be mapped onto the range that the semi-elasticity spans in the data. For insurers and pension funds, the estimated slope varies by about a factor of two across specifications, from 0.51 in the local-plus-euro sample to 1.10 in the local-currency sample. We use estimates whose first stages are well above the weak-instrument threshold, namely the insurance and pension sector ($\hat{\beta}_1 = 1.098$, $F = 32.5$) and the investment-fund sector ($\hat{\beta}_1 = 0.327$, $F = 70.1$). With holdings weights renormalized over these

two sectors, $h_{\text{ICPF}} = 0.054$ and $h_{\text{OFI}} = 0.946$, the composite (9) gives an aggregate semi-elasticity of 0.369.

We consider a tightening that scales the mandate of insurers and pension funds by $\varsigma = 1.30$, well inside the factor-of-two dispersion that their slope itself spans across specifications. Applied to that sector alone, the tightening moves the composite from 0.369 to 0.355, a fall of 3.7 percent, and it therefore scales the composite friction by 1.039, from $\Gamma_\ell = 0.125$ to $\Gamma_\ell = 0.130$. The contrast with a rescaling of the aggregate is the quantitative content of the sector structure. Applied to the sector for which it binds, the reform raises the composite friction by about 4 percent (from $\Gamma_\ell = 0.125$ to 0.130); applied to every sector at once it would raise it by 30 percent (to 0.163), about eight times as much, because it would implicitly assign the insurance and pension sector the whole segment instead of the 5.4 percent it holds.

Table S13

Prudential regulation of euro-area creditors: stationary equilibria

Moment	Frictionless $\Gamma_\ell = 0$	Baseline $\Gamma_\ell = 0.125$	Tightening $\Gamma_\ell = 0.130$	Data
LC debt to output $\bar{D}$	0.663	0.144	0.143	0.070
Mean UIP deviation (pp)	0.743	3.400	3.507	3.310
Std. dev. of UIP deviation (pp)	1.327	1.781	1.750	4.226
Default frequency (% per year)	0.875	0.768	0.786	0.1 to 1
Friction wedge $E[\Gamma_{\ell,t} D_t]$ (pp)	0.000	2.991	3.092	3.030
β_1	n.a.	0.334	0.323	0.330
Mass at the debt floor	0.000	0.017	0.017	
Change against baseline				
Mean UIP deviation (pp)			+0.107	
LC debt to output			−0.001	
β_1			−0.011	

Notes: Stationary equilibria at the frozen calibration, varying only the aggregate friction. The tightening column applies $\varsigma = 1.30$ to the mandate of the insurance and pension sector alone, which by (9) scales the composite friction by 1.039. The frictionless column is reported as a bound on the cumulative importance of the channel. Moments are unconditional ergodic means over both friction states and non-default observations, in the yield convention of the empirical estimates. β_1 is undefined at $\Gamma_\ell = 0$, where demand is perfectly elastic. Only β_1, and through it the friction wedge, is targeted.

From the model UIP to the empirical wedge. The model's UIP deviation is the same object as in the empirical analysis, namely the emerging-market local-currency borrowing cost over the risk-free rate, net of expected depreciation, which consists of a currency-risk premium, a default premium, and the endogenous friction

wedge. Only the semi-elasticity β_1, and hence the implied wedge, is targeted in the calibration. The mean model UIP is nevertheless 3.40 pp, close to the raw-yield data counterpart of 3.31 pp, with the friction wedge accounting for 2.99 pp of it. The model understates the volatility of the wedge, 1.78 pp against 4.23 pp in the data, because the simulated series omits the higher-frequency variation in raw yields (risk-premium fluctuations, liquidity, and measurement noise), and we read the model volatility as a lower bound. The local-currency debt-to-output ratio, 0.144, is roughly twice the pooled Arslanalp and Tsuda (2014) mean of 0.070 but within the range of country means, which runs from 0.010 for India to 0.164 for South Africa.

The role of currency pricing. Close to the calibrated regime, a change in the creditor friction is absorbed almost entirely by the currency premium rather than by quantities. The tightening raises the mean UIP deviation from 3.40 to 3.51 pp, an increase of 0.107 pp, while local-currency debt-to-output falls from 0.1444 to 0.1431, a decrease of less than one percent. By symmetry, a comparable loosening of the mandate of insurers and pension funds compresses the emerging-market premium and raises borrowing modestly. Emerging-market sovereign supply is inelastic in this region, consistent with the supply estimates, which show that issuance responds only weakly to the currency wedge.

Because holdings fall when the friction rises, the friction wedge rises less than one-for-one with Γ_ℓ, and the implied semi-elasticity falls from 0.334 to 0.323. A prudential reform of this size, applied to insurers and pensions alone, moves emerging-market borrowing costs by about a tenth of a percentage point. This is roughly one-eighth of what the same reform appears to deliver when the prudential change is applied to the investor base as a whole, and the gap is the holdings weight of the regulated sector in the segment. The first column of Table S13 presents the frictionless limit, showing that removing the friction entirely would raise sustainable debt-to-output more than fourfold, from 0.144 to 0.663, and lower the mean premium to 0.74 pp. Marginal regulatory changes are therefore modest and fall on prices, while the level of the friction is a first-order determinant of how much emerging markets can borrow.

The discipline operates through demand, not default. The default frequency is low throughout. It is 0.77 percent per year at the baseline, within the range implied by the ratings of our sample countries and an order of magnitude below serial defaulters such as Argentina, and it barely moves under the perturbation, to 0.79 percent. Essentially the entire issuance and pricing response therefore runs

through the demand-friction channel, isolated from the Eaton and Gersovitz (1981) default force.

Welfare. The welfare implications follow the same logic. A more accommodating investor base lowers the emerging-market borrower's cost of funds, such that a Solvency II loosening benefits it, while the loss of the disciplining role that downward-sloping demand plays works in the opposite direction. Which force dominates depends on the issuer's credibility. The sign is the opposite of Moretti et al. (2024), where the government prefers inelastic investors: their model is calibrated to Argentina, where a powerful default-cost channel makes downward-sloping demand valuable because it disciplines dilution and lowers default risk. Our panel is closer to the investment-grade end, the default channel is largely inoperative, and mainly the higher cost of funds from a steeper demand curve remains. Where default risk is high, inelastic demand is a disciplining asset, while where it is remote, it is mainly a cost.

S8.2. Counterfactual II: a Euro-Area Quantitative Tightening

Our second counterfactual examines the transmission to emerging markets of a withdrawal of official-sector demand for sovereign bonds, following the exit from quantitative easing. While the prudential experiment changed the slope of private demand, quantitative tightening changes its level: as the Eurosystem runs off its euro-area holdings, the sectors that held alongside it rebalance away from the emerging-market segment, and the residual demand facing the emerging sovereign declines. We model this as a downward shift $\Delta^{QT} \geq 0$ in the intercept of the local-currency demand schedule,

$$q \;=\; \frac{1-\delta}{1+r}\,\frac{e}{E(e')} \;-\; \Delta^{QT} \;-\; \Gamma_\ell\, B', \tag{S8.1}$$

and, since the currency wedge is $x_b = \frac{1-\delta}{1+r}\,\frac{e}{E(e')} - q$, the withdrawal raises the wedge one-for-one, $\Delta x_b = \Delta^{QT}$. Throughout we state the wedge, and hence Δ^{QT}, in the same yield units as in our empirical estimates.

Calibrating the shock. We obtain the size of the shift by inverting a demand curve. Since sector i's log holdings respond to the wedge as $\Delta \log B_i = \beta_{1i}\,\Delta x$, the wedge change that rationalizes a withdrawal s^{QT} of that sector's holdings is

$$\Delta^{QT} \;=\; \frac{s^{QT}}{\beta_{1i}} \qquad \text{percentage points of yield.} \tag{S8.2}$$

The denominator must be the semi-elasticity of the same sector whose holdings response is in the numerator. We use the nonbank (OFI) sector, which carries the bulk of the aggregation weight, with the baseline shift-share estimate $\hat{\beta}_{1,\mathrm{OFI}} = 0.327$, the specification that identifies that sector most strongly.

We recover the numerator from the Eurosystem's realized runoff in three steps. First, we estimate the response of nonbank holdings of each bond to quantitative easing,

$$\log D_{\mathrm{OFI},b,t} = \alpha + \beta^{QE} z^{QE}_{-i,b,t} + (\text{controls}) + u^{QE}_{b,t}, \tag{S8.3}$$

where $z^{QE}_{-i,b,t} = w_{-i,g(b)}\, QE_t$ uses the same predetermined shares as (28) and (29), QE_t is the Eurosystem's net public-sector purchase flow in quarter t,[4] and the controls and sample match the demand specification. The estimated slope is $\hat{\beta}^{QE} = 1.181$ (s.e. 0.235), clustered by quarter. Second, we aggregate to the portfolio: to first order the change in log total holdings is the holdings-weighted average of the security-level changes, so the relevant exposure is

$$\bar{w}_{\mathrm{OFI}} = \sum_{b \in \mathcal{B}_{\mathrm{OFI},t^\star}} \frac{D_{\mathrm{OFI},b,t^\star}}{\sum_{b' \in \mathcal{B}_{\mathrm{OFI},t^\star}} D_{\mathrm{OFI},b',t^\star}}\, w_{-i,g(b)}, \tag{S8.4}$$

evaluated at $t^\star = 2022\mathrm{Q2}$, the last quarter before the runoff, which gives $\bar{w}_{\mathrm{OFI}} = 6.07 \times 10^{-4}$. Third, we size the withdrawal from the realized policy: the runoff, dated from the peak of cumulated net purchases in 2022Q2 through the end of our sample in 2025Q4, spans 14 quarters averaging 59.9 EUR billion, a cumulative unwind of $\sum_t |QE_t| \approx 839$ EUR billion. Combining the three,

$$\hat{s}^{QT} = \hat{\beta}^{QE}\, \bar{w}_{\mathrm{OFI}} \sum_{t \in \{2022Q3,\ldots,2025Q4\}} |QE_t| = 0.601 \quad \text{log units}, \tag{S8.5}$$

so (S8.2) at $\hat{\beta}_{1,\mathrm{OFI}} = 0.327$ gives $\hat{\Delta}^{QT} = 0.601/0.327 = 1.84$ percentage points of yield. Table S14 reports the stationary equilibria.

Incidence across prices and quantities. The tightening raises the mean UIP deviation from 3.40 to 4.25 pp, an increase of 0.85 pp, and lowers local-currency debt-to-output from 0.144 to 0.100. Both the price and the quantity margins are affected. While a marginal change in the slope of demand, as in the prudential exercise, is absorbed almost entirely by the premium, a shift in the level of demand of the size of the Eurosystem runoff is large enough to move the quantity of

4 We measure Eurosystem purchases from the ECB's published history files for the Public Sector Purchase Programme (PSPP) and the public-sector part of the Pandemic Emergency Purchase Programme (PEPP), which report monthly net purchases by issuer. We take total public-sector net purchases minus purchases of supranational issuers, sum them within each quarter, and convert to billions of euros. The measure is net of redemptions and is therefore negative once redemptions exceed reinvestment. We use the sign switch to date the balance-sheet runoff.

Table S14

Euro-area quantitative tightening: stationary equilibria

Moment	Baseline $\Delta^{QT} = 0$	Tightening $\Delta^{QT} = 1.84$ pp	Data
LC debt to output $\bar{D}$	0.144	0.100	0.070
Mean UIP deviation (pp)	3.400	4.250	3.310
Std. dev. of UIP deviation (pp)	1.781	1.693	4.226
Default frequency (% per year)	0.768	0.763	0.1 to 1
Friction wedge $E[\Gamma_{\ell,t} D_t]$ (pp)	2.991	2.083	3.030
β_1	0.334	0.480	0.330
Mass at the debt floor	0.017	0.101	
Change against baseline			
Mean UIP deviation (pp)		+0.850	
LC debt to output		−0.044	
β_1		+0.146	

Notes: $\Delta^{QT} = s^{QT}/\hat{\beta}_{1,\mathrm{OFI}}$, with $s^{QT} = 0.601$ log units the withdrawal implied by the cumulative 2022Q3 to 2025Q4 Eurosystem runoff estimated on the nonbank panel. Moments are unconditional ergodic means over both friction states and non-default observations, in the yield convention. Only β_1 is targeted in the calibration.

borrowing as well. The incidence on prices alone is therefore a local property around the calibrated regime.

The response confirms that quantitative tightening works through a different margin than the prudential shock. The implied semi-elasticity rises from 0.334 to 0.480. The intercept shift leaves the slope of demand unchanged, but by lowering equilibrium holdings it raises $\beta_1 = 1/(\Gamma_\ell \bar{D})$ through the long-run debt level rather than through the friction. This is the mirror image of the prudential experiment, where tightening the mandate lowered β_1 from 0.334 to 0.323.

Finally, since the model has heterogeneous investors, we can trace the incidence of the withdrawal across the investor base. At the common wedge change Δ^{QT}, each sector adjusts according to its own semi-elasticity:

Sector	$\hat{\beta}_{1i}$	$\Delta \log B_i = \hat{\beta}_{1i}\, \Delta^{QT}$
Insurance corporations and pension funds	1.098	2.019
Investment funds	0.327	0.601
Aggregate (holdings weighted)	0.369	0.678

Insurers respond more than three times as much in proportional terms as nonbanks, and the aggregate response is 12.8 percent larger than the nonbank

figure on which the shock is calibrated. A model with a single creditor friction could deliver only the change in aggregate demand.

Data Documentation for "Asset Elasticities and Currency Risk Transfer"

Carol Bertaut Ester Faia Sebnem Kalemli-Özcan

Camilo Marchesini Simon Paetzold Martin Schmitz

November 2026

This document accompanies the replication materials. It describes the confidential holdings data for the euro area and the United States, the selection of the country sample, and the construction of the descriptive figures. Sections, tables and figures of this document carry the prefix D. Items with the prefix S are in the Supplemental Material, and items without a prefix are in the main paper.

D9. Securities Holdings Statistics from the Euro Area

The Securities Holdings Statistics by sector (SHS) are collected on a security-by-security basis and provide information on securities held by all institutional sectors of euro area investors (with the exception of monetary authorities), broken down by instrument type, issuer country and further security- and sector-specific classifications. These data can be combined with other security-specific information, e.g. on the currency denomination, rating and maturity. Those are particularly

*The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Board and System, Bundesbank, European Central Bank or the Eurosystem. The paper relies on the confidential administrative datasets Security Holdings Statistics (SHS-Plus) for the euro area and the U.S. Treasury International Capital (TIC) survey, both of which have been accessed through a coauthor. Other proprietary and administrative datasets, used and matched in the analysis, have been accessed through the ECB (centralized returns database) and Goethe University Frankfurt (Consensus Forecasts). We thank Anusha Chari, Hanno Lustig, Federico Mainardi, Motohiro Yogo and Haonan Zhou for insightful comments and participants at seminars and conferences. Jakob Staps provided excellent research assistance. E-mail addresses: *Corresponding author* Sebnem Kalemli-Özcan (sebnem_kalemli-ozcan@brown.edu). Carol Bertaut (Carol Bertaut passed away on May 25, 2026), Ester Faia (faia@wiwi.uni-frankfurt.de), Camilo Marchesini (camilo_marchesini@brown.edu), Simon Paetzold (simon.paetzold@bundesbank.de), Martin Schmitz (Martin.Schmitz@ecb.europa.eu).

relevant when focusing on emerging market economies as the selection of duration and rating is typically different compared to advanced economies. Moreover, the holdings data can be matched with data on prices and returns allowing us to also decompose the changes in the shares into their various components: valuation, exchange rate changes and transactions.

In what follows we provide a general description of the framework underlying the securities data.

Securities Holdings Statistics. General Framework The legal basis for collecting SHS data is laid down in Regulation ECB/2012/24, which is further complemented by Guideline ECB/2013/7. Note that the dataset is gathered through a harmonized collection across the central banks of the euro area. To this purpose the second regulation sets out the procedures to be followed by national central banks when reporting to the ECB. The harmonized data have been collected in full since the fourth quarter of 2013.[1] The data are at the security level and contain the breakdown by asset type (debt versus equity, including fund shares; debt is further broken down by type of issuer, such as government, non-financial corporations, etc.).

The main breakdown of the data is by type/sector of investors (see Table D1). There are 22 investor types in total. We focus in particular on debt securities held by banks, insurance companies and pension funds, nonbank financial intermediaries (such as investment funds) and other investors (government, nonfinancial corporations, households). A key aspect lies in the fact that the data contain both investment in securities issued domestically and abroad: this feature, rarely present in holdings data, allows us to compute the portfolio shares of each investor over its total investment. This is crucial for gauging the importance of a given type of investment, such as emerging-market securities, relative to the rest of the portfolio. Finally, note that the securities include holdings by investors residing in the euro area and non-resident investors' holdings of euro area securities that are deposited with a euro area custodian.

The holdings information is complemented by the Centralised Securities Database (CSDB) that contains information such as price, issuer name and outstanding amount, precise debt type and issuer information for over six million outstanding debt securities, equities and investment fund shares.

To ensure good data quality, SHS data are regularly checked against comparable data sources, which include the integrated euro area financial and non-financial

1 Between the first quarter of 2009 and the fourth quarter of 2013, reporting agencies were not obliged to report the data, but many did.

accounts (EAA), balance of payments and international investment position statistics (BoP/IIP), Monetary and Financial Institutions (MFI) balance sheet statistics, insurance corporations and pension fund statistics, investment fund statistics and securities issues statistics, as well as consolidated banking data.

Table D15

Coverage of Holder Sectors and Underlying Subsectors in the SHS

Holder Sector	**Subsectors (identifier code)**
ICPF	- Insurance corporations (S_128). - Pension funds (S_129). - Other insurance corporations and pension funds (S_12QU).
NBFI (= OFI)	- Other financial intermediaries plus financial auxiliaries plus captive financial institutions and money lenders excluding financial vehicle corporations (S_125W). - Financial vehicle corporations (S_125A). - Non-MMF investment funds (S_124).
MFI	- Deposit taking corporations except central banks (S_122). - Money market funds (MMF) (S_123).
OTH	- Non-financial corporations (S_11). - Households excluding non-profit institutions serving households (S_14). - Non-profit institutions serving households (S_15). - Other households and non-profit institutions serving households (S_1MU). - Central government (S_1311). - State government (S_1312) . - Local government (S_1313). - Social security funds (S_1314). - Other general government (S_13U).

D9.1. The U.S. Data: Treasury International Capital (TIC) Survey

For U.S. investors we rely on the securities holdings dataset assembled and maintained by the U.S. Department of the Treasury, the Treasury International Capital (TIC) survey. The TIC collects annual detailed position data on holdings of long-term and short-term securities by U.S.-resident custodians, broker-dealers, and institutional end-investors, who hold foreign securities on behalf of U.S.-resident clients. Reporting and filing of the positions data are mandatory and compliance

is enforced by the U.S. government. The data from the annual survey are available since 2003. The dataset contains information on the security issuer, the currency of denomination, maturity, and duration.

Two features of the TIC shape how we use it. First, our access to the confidential TIC ended in 2022, such that the U.S. estimates are frozen at the 2014Q1 to 2022Q2 window and cannot be refreshed. Second, the survey is annual and custodial, recording the holder of record rather than the end investor at the issuer-currency segment granularity we use for the euro area. Neither the annual frequency nor the custodial recording supports the quarterly, segment-level, levels specification of the main text, which is why the U.S. evidence in subsection 6.4 uses the within-issuer portfolio-share specification and is read as signs rather than calibrated magnitudes.

The U.S. holdings these data record are summarized in Figure D1 and Figure D2, the U.S. counterparts of the euro-area descriptive facts in Figure S1 and Figure S2.

Figure D5
TIC: U.S. holdings of debt securities, by issuer type

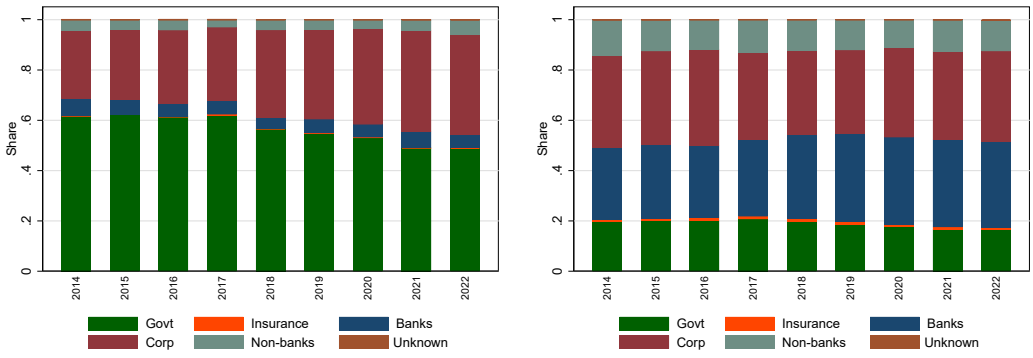

(a) TIC: U.S. Holdings in EM Bonds, by Issuer Type

(b) TIC: U.S. Holdings in AE Bonds, by Issuer Type

D10. Country Sample Selection

The empirical analysis and descriptive statistics use narrow samples of emerging-market and advanced economies. The samples contain countries that are relevant investment destinations for euro-area investors, for which UIP and CIP can be computed, and for which UIP and CIP deviations are not expected to be large because their exchange rates float and their exchange markets are unified at the

Figure D6
TIC: U.S. holdings of sovereign bonds, by investor type

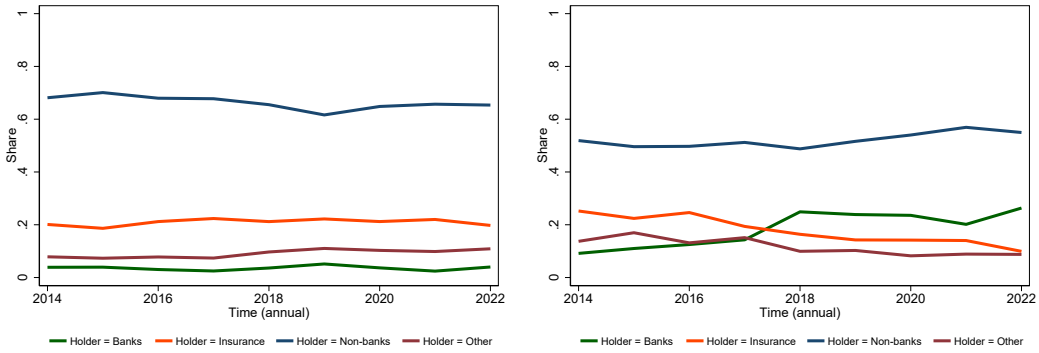

(a) **(b)**

TIC: U.S. Holdings in EM Sovereign Bonds, by Investor Type TIC: U.S. Holdings in AE Sovereign Bonds, by Investor Type

end of each quarter. We arrive at the final sample (Table D2) in the following steps.

- Compute the average holdings of government bonds issued by countries, by euro area investors across all quarters.
- Remove all countries below an average holding of 1bn EUR and keep all countries with both exchange rate forwards and exchange rate expectations.
- **AE sample:** Remove countries with fixed ex-rate (DK), countries that borrow mostly in USD (SG, HK), and countries for which there is no benchmark 1Y local currency bond in Refinitiv or Bloomberg (IS) and the U.S.
- **EM sample:** Keep all countries with floating or managed floats (category 3 & 4 of the classification in Ilzetzki, Reinhart and Rogoff (2022)) and unified exchange markets at the end of each quarter by extending the latest value into the future. Russia is excluded for 2022Q1 and onwards even though data for the classification of the exchange rate regime and unified exchange market were not yet available.

D11. Notes to Figures

Figure 1. We use data from Arslanalp and Tsuda (2014) to obtain outstanding amounts at the end of each year of foreign-held sovereign debt securities denominated in any currency and in local currency. We divide the holding amounts recorded in SHS for euro area investors by the externally held outstanding amounts to assess the coverage of our datasets in the outstanding amount of total externally

Table D16

Country sample for descriptive statistics and UIP/CIP regressions. Countries in gray were removed due to the step-wise sample creation described above.

	Data on forwards, expectations	Final sample
EME	AR, BR, CL, CN, CO, CZ, EG, HR, HU, ID, IL, IN, KE, KR, KZ, LK, MA, MX, MY, NG, PE, PH, PK, SA, PL, RO, RU, TH, TR, UA, ZA	BR, CL, CO, IN, IL, KR, MX, MY, PL, RU (no 14q4-15q4 & 22q1-), TH, TR (until 18q2), ZA
AE	AU, CA, CH, DK, GB, HK, IS, JP, NO NZ, SE, SG, US	AU, CA, CH (from 15q1), GB JP, NO, NZ, SE

held sovereign debt for emerging market economy and advanced economy government debt securities (see section D2 for the country samples). For Israel, we obtain debt securities in any currency held abroad from the World Bank's Quarterly External Debt Statistics and the total amount of debt securities in all currencies from the World Bank Quarterly Public Debt Statistics to calculate the amount of sovereign debt securities in local currencies held abroad. For New Zealand, we subtract the amount outstanding of foreign currency sovereign debt securities from the BIS international debt securities statistics from total debt securities held abroad to obtain the amount of local currency sovereign debt securities held abroad. *Source*: European Central Bank's Securities Holdings Statistics (SHS) and Arslanalp and Tsuda (2014).

Figure S1. The portfolio shares are computed from security holdings of investors in EA by issuer type. For each quarter, we sum the holdings of all EA investors at market value and compute the share of securities issued by each issuer type. We compute the shares in the portfolio of all euro area investors separately for issuers from EMs and AEs. We do the same for U.S. investors using TIC holdings data. *Source*: European Central Bank's Securities Holdings Statistics (SHS).

Figure S2. The investor shares are computed from security holdings of sovereign bonds issued by emerging market and advanced economy governments. For each quarter, we sum the holdings of all euro area investors at market value and compute the share of sovereign debt held by each investor type. *Source*: European Central Bank's Securities Holdings Statistics (SHS).