Model-Free Cross-Country Term Premia

Ljubica Georgievska*

July 2026

Abstract

I recover cross-country term-premium differentials from two traded prices, the long-tenor cross-currency basis and the swap-spread differential, estimating no expectations model in any country, and a network of seventeen currency pairs holds the bilateral frictions inside dealers' quoted spreads. Ten-year U.S. Treasury term premia sat 28 basis points below German Bund and 37 below Japanese JGB term premia before 2008, and 56 and 83 above them since 2013. The differentials behave as term premia: duration supply prices them at −1.05 on U.S. free float, policy surprises reprice them within days, and survey expectation revisions move the expectations component rather than the differentials.

*CUNEF Universidad & NYU Stern School of Business. Email: lg4086@nyu.edu and ljubica.georgievska@cunef.edu. I thank Max Croce, Francis Longstaff, Emanuel Moench, Angelo Ranaldo, Anthony Saunders, Fabricius Somogyi, and Olav Syrstad for their invaluable comments and suggestions. All errors are my own.

1 Introduction

Does the United States pay more or less than its peer reserve sovereigns for the long-horizon duration risk its debt asks investors to bear? The conventional reserve-currency narrative predicts a yield discount on U.S. Treasuries relative to foreign safe assets, reflecting the non-pecuniary convenience of dollar-denominated safety, liquidity, and reserve services (Krishnamurthy and Vissing-Jorgensen, 2012; Jiang et al., 2021; He et al., 2022). If this convenience is priced at the medium and long maturities where cross-country substitution in sovereign duration is most active, it should appear as a systematically lower U.S. long-horizon term premium than those of peer reserve currencies, so that the question can in principle be answered by comparing term premia across countries.

Answering it in practice has proved difficult, because the standard route to a cross-country term-premium comparison estimates a parametric term-structure model separately for each country and then differences the resulting term-premium estimates, a procedure that requires consistent specifications across countries and is sensitive to model choice within each one (Duffee, 2011; Hamilton and Wu, 2014); term-premium estimates built this way can disagree with one another within a single country by more than the cross-country differences they are asked to measure. This paper develops a route that does not require estimating expectations at all. The long-tenor cross-currency basis swap, an instrument that the literature has read mainly as a barometer of short-horizon dollar funding conditions (Du et al., 2018b; He et al., 2017) and, at longer tenors, as the footprint of corporate issuance flows arbitraging credit-spread differentials across currencies (Liao, 2020), is at the ten- and twenty-year maturities the price at which the market exchanges long-horizon fixed-rate exposure in one currency for long-horizon fixed-rate exposure in another, and the network of such prices, read jointly, allows the cross-country term-premium differential to be recovered directly.

Before describing the identification, I fix the two definitions used throughout the paper. The term premium of a long-term interest rate, whether a government bond yield or a swap rate, is the gap between that rate and the average short-term rate that investors expect to prevail under physical-measure dynamics over the life of the contract. A higher term premium means more compensation demanded for bearing term-structure risk; I am agnostic about the precise composition of the risks bundled into it and adopt the standard view of the term premium as the wedge between the long rate and the physical-measure expected short-rate path (Fama, 1984; Campbell and Shiller, 1991; Cochrane, 2005; Kim and Wright, 2005; Adrian et al., 2013; Bauer and Rudebusch, 2020).

The cross-country term premium of a partner currency relative to the United States is the

differential of the two term premia, measured throughout from the U.S. point of view, so that a positive value means that the partner currency's long-term rates carry more term-structure risk compensation than U.S. long-term rates, and the United States is then the more convenient long-duration borrower at that maturity; this sign convention follows the convenience-yield literature (Krishnamurthy and Vissing-Jorgensen, 2012; He et al., 2022). The cross-country term-premium differential is a long-horizon difference in own-currency term premia, and it is conceptually distinct from, and complementary to, the short-maturity Treasury basis of that literature (Krishnamurthy and Vissing-Jorgensen, 2012; Du et al., 2018b; Jiang et al., 2021; Diamond and Van Tassel, 2026), which is a level wedge between a Treasury yield and a currency-hedged foreign sovereign yield. At the long end the two are built from the same three quoted components and are linked by an exact identity, stated as a displayed equation below, which is why they coincide before 2008 and part ways after.

What I measure is a pair of traded prices rather than an estimate. By a no-arbitrage decomposition, each bilateral cross-currency basis swap price, whose collateralised cash-flow construction bounds the par-swap representation wedge at under two basis points, is the sum of two pieces: the cross-country term-premium differential between the two currencies, which is the quantity I want to recover, and a bilateral pair-specific friction, which I call the pure parity of the pair and which captures the residual deviation from frictionless cross-currency parity that carries no term-premium content. The first piece is the quantity of interest and the second is the error term of the measurement, so what has to be established is not how to estimate the first but how large the second can be, and that is the question the network answers.

It comes from a triangular arbitrage trade. An investor moving long-horizon fixed-rate exposure from U.S. dollars into Japanese yen can trade the dollar-yen basis directly or pass through a third currency, say the euro, and both routes deliver the same final exposure, so the gap between the two costs carries no term premium at all: the dollar and yen premia appear in both routes and drop out on subtraction, while the euro premium enters the two legs of the indirect route with opposite signs and cancels within the route itself. Triangular gaps therefore reflect only the pair-specific pure parities, never the cross-country term-premium differentials that the individual quoted prices embed. A single triangle, however, is one equation in three unknown pair-specific frictions; the seventeen independently quoted pairs among ten currencies contain eight independent triangles[1], and stacking them determines the pair-specific parity

[1]The lower-tenor boundary is set by the availability of cross-currency basis swap prices, which begin at the two-year tenor. The Swiss franc enters the descriptives with its minimum-exchange-rate window of September 2011 through January 2015 excluded and stays outside the estimation panels, since the administered exchange rate makes the cross-currency basis price over that window a function of the policy regime rather

deviations under one explicit cross-sectional restriction, namely that they are organised by pair rather than by currency, which is what allows the error term of the measurement to be bounded from the prices themselves. This restriction is an economic assumption rather than an implication of no-arbitrage; I state it, defend it on institutional grounds, and treat the term-premium reading it delivers as a hypothesis to be tested throughout the paper rather than as a settled label.

What prices identify, and what they do not, is settled here rather than after the results. The construction is expectations-model-free, not assumption-free, and the distinction is exact. Physical expected short-rate differentials cancel algebraically between the swap-rate differential and the long-horizon forward premium, so no expectation of future short rates is formed or estimated anywhere in the paper. The triangles then identify the component of the remaining wedge that is not currency-additive, and that component, the bilateral pair-specific content, is measured from prices alone. They do not identify the currency-additive remainder, and cannot: for any vector of currency-level values, raising every term premium by that vector and lowering every pair friction by the offsetting amount reproduces every quoted basis and every triangular restriction exactly, so no quantity of price data distinguishes the two. Internet Appendix Proposition 3 states that equivalence formally.

Reading the currency-level remainder as cross-country term-premium content, rather than as a currency-level wedge in the forward market, is therefore an identifying restriction and not a finding; I impose the pair-specificity restriction explicitly, and what follows treats the term-premium reading as a hypothesis with testable implications rather than as something the prices have established. One question the prices do settle is the level at which the content is organised, currency rather than pair, which is what the triangles deliver. What prices alone cannot settle, for this measurement or for any model-based term-premium estimate, is the mixture of risk compensation, convenience, and intermediation cost that a term premium bundles; the model-based U.S. term premia of Adrian et al. (2013) and Kim and Wright (2005) bundle the safety and liquidity convenience of U.S. Treasuries in the same way (Krishnamurthy and Vissing-Jorgensen, 2012), and the literature has not required that bundle to be resolved before calling those estimates term premia. For the one direction that no test reaches, a funding wedge common to every dollar pair, I report a benchmark-calibrated sensitivity frontier rather than an identified set, and show that each of the paper's findings survives across it; Table 9 records the epistemic standing of each layer of the paper, what is observed, what is algebraic, what is assumed, and what is interpreted.

The measured quantity is therefore the quoted basis itself, read as the cross-country swap term-

than of the underlying pricing.

premium differential up to the bounded pure parity, and a within-currency bridge identity, which adds the directly observable swap-spread differential, the country-by-country wedge between the local fixed-for-floating swap rate and the local government bond yield (Feldhütter and Lando, 2008; Klingler and Sundaresan, 2019; Jermann, 2020; Boyarchenko et al., 2018; Hanson et al., 2022)[2], converts them into the cross-country government bond term-premium differentials, the policy-relevant counterparts that anchor discussions of safe-asset convenience and reserve-currency status. No expectation of future short rates is formed or estimated anywhere in the construction, so the measurement is expectations-model-free, though not assumption-free. Throughout the paper I report the headline results on the quoted prices themselves, and use the network-extracted series as the check that the pure parity is not carrying them; the two agree to a fraction of a basis point, which is the point rather than a caveat. The identified pair-specific frictions turn out to be small: they average under three percent of the quoted basis and sit inside dealers' quoted bid-ask spreads in most months, so that nearly all economically material variation in long-tenor cross-currency prices is organised at the level of currencies rather than of pairs.

The first contribution is the measurement itself: expectations-model-free cross-country term-premium differentials, of the swap and the government bond kind, for nine currencies against the U.S. dollar, monthly from 2002 to June 2026 and daily where the exercises require it. On the government side the differential shares its three quoted components with the long-maturity government deviation from covered interest parity of Du et al. (2018a), differing by the sign the quoted basis enters with, as the cost of the hedge there and as term-premium content here; the term-premium reading, the network bound, and the cross-country swap term-premium differential, which lies outside that construction, are this paper's. The measurement survives quoted transaction costs, since perturbing every quoted price within its own half-spread and re-solving moves the recovered differentials by 3.6 to 5.1 basis points on average, well below the levels they measure. It also does not depend on the dollar: independently quoted euro crosses, in which no dollar leg appears, reproduce the dollar-implied differentials at correlations of 0.993 to 0.999.

The construction equally delivers the complement of the premium: subtracting the recovered term-premium differential from the observed yield differential gives the expectations component, the expected average short-rate differential purged of term premia, at every analysed tenor and pair, the identification caveats of the differential transferring with the opposite sign; at the long horizons this is a market-implied neutral-rate differential obtained with no

[2]Swap spreads can take either sign; the U.S. swap spread has been persistently negative since 2008, a regularity studied by Klingler and Sundaresan (2019), Jermann (2020), Boyarchenko et al. (2018), and Hanson et al. (2022).

expectations model, and both series, the cross-country term-premium differentials and the neutral-rate differentials at all analysed tenors and countries, are posted as downloadable datasets on my website.

The second contribution is the headline finding, and it is a reversal of the received ranking. Ten-year U.S. Treasury term premia sat below their German Bund and Japanese JGB counterparts before the crisis, by 28 and 37 basis points on average, and have sat above them since, by 56 and 83 basis points over 2013 to June 2026. Formal break estimation dates the onset in late 2007, with the funding crisis, and the consolidation of the new level in October 2011, and every segment mean of the cross-sectional average after the onset is negative under any segmentation the data select, so the configuration has been in place for a decade and a half rather than emerging recently. Du et al. (2026) independently document a post-crisis sign change in the related government deviation, driven by the decline of the United States swap spread, which is the component the two constructions share.

The reversal does not say that the United States has lost its reserve advantage altogether. The short-maturity convenience discount on U.S. Treasury bills survives throughout the sample, and the richer reading is coexistence: a sovereign that retains its convenience at the short and collateral-sensitive end while supplying long-duration safe assets whose duration now commands a term premium rather than a discount.

The third contribution is the economics: the recovered differentials behave as term premia should. Government debt duration supply prices them with the signs preferred-habitat theory predicts (Greenwood and Vayanos, 2014; Vayanos and Vila, 2021), the coefficient on U.S. log free-float supply at -1.05 with a Driscoll-Kraay t-statistic of -10.7 at the Schwert-rule lag and -4.4 under two-way clustering that respects the U.S. variable being common across pairs, and materially stronger once central-bank holdings are stripped from gross debt, consistent with quantitative easing operating through the absorption of duration, and present at the announcement frequency across ninety-nine pre-specified U.S. Treasury refunding surprises (Georgievska et al., 2025). And identified policy surprises reprice them within days: six pre-specified central-bank surprises, the two Bank of Japan yield-curve-control steps, the Bank of England gilt-purchase launch, the European Central Bank's expanded asset purchases, and the two 2008 to 2009 Federal Reserve announcements, move the recovered ten-year differentials by two to eight daily standard deviations within $[-1, +2]$ trading-day windows, all six predicted signs realised, seven of nine on the wider screen disclosed in the text, the reading resting on magnitudes and placebo percentiles rather than on any sign count, while the same windows leave the observable one-year funding basis largely undisturbed and anticipated announcements move the recovered series very little.

The paper complements four literatures, and the specialised strands are taken up where each result is presented. Relative to the model-based term-structure decompositions (Dai and Singleton, 2000; Kim and Wright, 2005; Joslin et al., 2011; Adrian et al., 2013; Bauer and Rudebusch, 2020), which deliver a full decomposition of yields for each country at the cost of an expectations model, whose sensitivity to specification is itself documented (Duffee, 2011; Hamilton and Wu, 2014) and which have been extended across countries under affine structures fitted country by country (Dahlquist and Hasseltoft, 2013; Jotikasthira et al., 2015), I deliver the cross-country term-premium differential alone from the contemporaneous cross-section of quoted prices, with no dynamic specification, factor structure, or survey data, and the measured differentials constitute a set of cross-country moments to which any candidate set of country-level term premia must difference. Relative to the covered-interest-parity and convenience literature (Du et al., 2018b; He et al., 2017; Liao, 2020), which reads the cross-currency basis at short tenors as a funding friction, I show that the long-tenor content of the same instrument is a cross-country term-premium differential, connecting to Lustig et al. (2019), who find in realised returns that carry premia decline with bond maturity as local-currency term premia offset currency risk premia, so that the cross-country term compensation their long-maturity portfolios stop earning is the quantity priced directly here, and the two readings partition by horizon.

The relation to that literature's government-side deviation is exact, and stating it matters because the two are close enough to be conflated and different enough that conflating them reverses a sign. Writing x for the quoted basis and ΔSS for the swap-spread differential, the long-maturity government deviation from covered interest parity, introduced by Du et al. (2018a) and carried through the benchmark reform by Du et al. (2026), is

$$\text{government CIP deviation} \;=\; \Delta SS - x, \qquad \text{whereas} \qquad \widehat{ctp}^{\,gov} \;=\; \Delta SS + x.$$

The two differ by twice the quoted basis and coincide only where it is zero, which is why they agree before 2008 and part ways after. The papers therefore ask different questions of the same three market ingredients: the government deviation asks what a currency-hedged foreign sovereign yields relative to Treasuries, treating the basis as the cost of the hedge, while the cross-country term-premium differential asks by how much the two sovereigns' own-currency long-horizon term compensation differs, treating the basis as that compensation. Diamond and Van Tassel (2026) measure convenience against option-implied benchmarks across the same currencies, similar in spirit, with the swap rate in the role of the box rate and reaching ten years where box rates, available to about three, do not.[3] On the component they share,

[3] A direct numerical comparison is not practical: box rates are estimated from index options to horizons of

Du et al. (2026) document the post-crisis decline of their deviation as driven by the United States swap spread, which is the bridge's swap-spread channel.

Relative to the preferred-habitat and safe-asset supply literature (Krishnamurthy and Vissing-Jorgensen, 2012; Vayanos and Vila, 2021; Greenwood and Vayanos, 2014; Georgievska et al., 2025), which predicts that privately held duration prices term premia, the recovered differentials confirm that pricing in a ten-currency panel with announcement-frequency evidence. And relative to the intermediary and swap-spread literatures (He and Krishnamurthy, 2013; He et al., 2017; Brunnermeier and Pedersen, 2009; Adrian et al., 2014; Feldhütter and Lando, 2008; Klingler and Sundaresan, 2019; Jermann, 2020; Boyarchenko et al., 2018; Hanson et al., 2022), which characterise the credit, liquidity, balance-sheet, and duration-demand channels moving the within-country wedge that the bridge identity makes one of the two components of the cross-country government bond term-premium differential, the measurement supplies the cross-country counterpart; the carry-trade literature (Lustig and Verdelhan, 2007; Lustig et al., 2011; Menkhoff et al., 2012) studies differences in interest-rate levels, whereas I study differences in term-structure premia, and a complementary literature attributes dollar variation to primitives in international asset demand (Jiang et al., 2025; Jiang and Richmond, 2025), with Zou (2026) carrying the demand-side reading to government bond auctions across advanced-economy currencies.

The paper proceeds as follows. Section 2 develops the framework and the identification; Section 3 describes the data and the implementation; and Section 4 presents the empirical results, the measurement, its composition, and the reversal, with its identified intervals and the sensitivity frontier, in Section 4.2, and the two families of validation evidence in Section 4.3, the supply evidence first and the far-forward and policy-surprise evidence in Section 4.3.4; Section 5 concludes. An Internet Appendix contains the formal statements, the full identification battery, and every robustness exercise referenced in the text.

2 Framework

This section develops the decomposition that recovers cross-country term premia for each currency from cross-currency basis swap market data under the pair-specificity restriction, without imposing any structure on how term premia evolve, their correlation patterns, or

about two to three years, where a term premium is still a small part of the yield, while the differentials of this paper carry their content at the five- to twenty-year tenors where duration compensation dominates, so the two measurements barely overlap in horizon and are complementary rather than substitutes, the option-implied benchmark most informative at the short tenors where convenience in the level of rates is priced, and the swap-based construction at the long tenors where term premia show up.

their determinants. The observable I exploit is the cross-currency basis swap spread, which the market prices at each date for every liquid bilateral currency pair and at each tenor. The identification strategy separates, for each cross-currency basis pair, a currency-specific term premium differential from a pair-specific friction, by stacking triangular relationships across multiple currencies. Complete derivations and technical details are in Internet Appendix A.

2.1 The Triangular Intuition

I begin with the institutional setting. Throughout the paper, I work with the *fixed-for-fixed* cross-currency swap, in which two counterparties exchange notional principals at inception and re-exchange them at maturity at the inception spot exchange rate, and in between exchange fixed-rate coupon streams in two different currencies over the life of the contract.[4] The basis spread $x_{t,N}^{j/i}$ is the market-determined adjustment to one of the fixed coupon streams that makes the two legs have equal present value at inception. Cross-currency basis swaps are among the most liquid derivatives in global over-the-counter markets and trade on standard tenors out to thirty years (Andersen et al., 2019; Du et al., 2018b).

The cross-currency basis swap over an N-period horizon embeds the term structure information of both currencies plus a residual specific to the dealer book intermediating the bilateral pair. Formally, the basis decomposes additively as

$$x_{t,N}^{j/i} \;=\; \left(tp_{t,N}^{j} - tp_{t,N}^{i}\right) \;+\; \pi_{t,N}^{j/i} \;=\; \underbrace{ctp_{t,N}^{j,i}}_{\text{currency-specific}} \;+\; \underbrace{\pi_{t,N}^{j/i}}_{\text{pair-specific}} \;,$$

into a currency-specific component, the cross-country term premium differential $ctp_{t,N}^{j,i}$ between currencies i and j, and a pair-specific component, the bilateral friction $\pi_{t,N}^{j/i}$ on the dealer book that intermediates pair (i, j). I derive this decomposition formally in Section 2.2.

Now consider an investor at date t who wishes to move N-period fixed-rate interest rate exposure from currency i to currency k in the cross-currency basis swap market. Two routes are available. The *direct route* uses the i/k cross-currency basis swap to exchange currency-i fixed-rate exposure for currency-k fixed-rate exposure over N periods at the prevailing direct

[4] The fixed-for-fixed quotation is economically equivalent to the floating-for-floating cross-currency basis swap that interbank dealers also trade, where the basis is added to one of the floating legs. The market quotes both conventions at the same basis spread, with the equivalence enforced through the local interest rate swaps in each currency that translate floating cash flows into fixed cash flows on the same yield curve. For my purposes the fixed-for-fixed convention is more transparent because the basis spread maps directly into a differential of fixed long-tenor swap rates after adjustment for the FX forward premium. See Du et al. (2018b) for the standard reconciliation between the two conventions and the long-tenor cross-currency basis identity that I use in Equation (3) below.

basis $x^{k/i}_{t,N}$. The *indirect route* goes through an intermediate currency j: first swap currency-i fixed exposure for currency-j fixed exposure at the basis $x^{j/i}_{t,N}$, and then swap currency-j fixed exposure for currency-k fixed exposure at the basis $x^{k/j}_{t,N}$. Since both routes deliver the identical economic exposure, no-arbitrage in the cross-currency swap market implies that the cost of the two routes coincides up to bilateral frictions. Define the *triangular inconsistency* as the gap,

$$\Delta^{i,j,k}_{t,N} \;\equiv\; x^{j/i}_{t,N} \;+\; x^{k/j}_{t,N} \;-\; x^{k/i}_{t,N},$$

which is zero under perfect bilateral arbitrage. The triangular inconsistency is small relative to the level of the bilateral basis itself in my cross-section over the January 2002 to June 2026 sample. I document this empirically in Section 4.1.

The currency-plus-pair structure of the basis is what makes the triangular argument informative. Substituting the decomposition into the three bases in the triangular inconsistency,

$$\begin{aligned}\Delta^{i,j,k}_{t,N} \;=\; &\underbrace{\left(tp^{j} - tp^{i}\right) \;+\; \left(tp^{k} - tp^{j}\right) \;-\; \left(tp^{k} - tp^{i}\right)}_{=\,0} \\ &+\; \pi^{j/i}_{t,N} \;+\; \pi^{k/j}_{t,N} \;-\; \pi^{k/i}_{t,N},\end{aligned}$$

the currency-specific term premia of all three currencies cancel exactly: the intermediate currency j enters with opposite signs in the first two bases and cancels within the partial sum $x^{j/i} + x^{k/j}$, while the residue $-tp^{i} + tp^{k}$ in this partial sum is offset by the subtraction of the third basis $x^{k/i}$. The triangular inconsistency therefore identifies a combination of the three pair-specific frictions alone: triangular inconsistencies see only bilateral frictions, never the underlying currency-specific term premium content.

The full identification strategy follows in two steps. First, the triangular cancellation identifies the bilateral content of the pair-specific frictions $\pi^{j/i}_{t,N}$ from the system of triangular constraints, since these constraints contain no currency-specific term premium content. With ten currencies and seventeen independently quoted pairs, the resulting linear system leaves the level of $\boldsymbol{\pi}$ underdetermined precisely in the space of currency-level shifts, and I resolve it with a structural identifying restriction, that bilateral pair-specific frictions are not currency-level shifts in disguise.

The restriction has direct institutional grounding, and since it carries the identification I state the grounds explicitly: the balance-sheet costs that generate basis frictions, G-SIB capital surcharges, leverage-ratio usage, central-counterparty margining, and bilateral credit-support annexes, are incurred by the dealer book that intermediates each currency pair and apply pair by pair rather than as currency-level factors; the quoting itself is organised pair by pair, with

pair-specific bid-ask bands and pair-specific outage histories that track the entry and exit of liquidity providers in individual crosses; and the restriction is self-enforcing at the margin, since any friction that did operate at the currency level would be collected by the swap-term-premium definition into the currency's swap term premium rather than misallocated to the pair component. Its empirical content is examined throughout and its failure bounded in Section 4.2. Under this restriction, any currency-c-specific compensation embedded in the swap rate is content of the swap term premium $tp^c_{t,N}$ rather than of $\pi^{j/i}_{t,N}$, and the level of $\boldsymbol{\pi}$ is exactly identified under that restriction. I compute the corresponding solution through a minimum-norm projection orthogonal to the space of currency-specific shifts.

Second, subtracting the recovered friction from the observed basis leaves the cross-country term premium differential $ctp^{j,i}_{t,N}$: the recovery is algebraic rather than statistical, purely cross-sectional at each date, with no model of the term structure, no distributional assumptions, and no time-series averaging or filtering; Section 2.3 formalises the construction.

2.2 Definitions and the Cross-Currency Basis Decomposition

Throughout the paper, all interest rates, swap rates, forward rates, and bond yields are expressed in continuously compounded log terms.[5]

Definition 1 (Swap Term Premium)**.** *The swap term premium for currency i at maturity N is*

$$tp^i_{t,N} \;\equiv\; IRS^i_{t,N} \;-\; \frac{1}{N}\mathbb{E}^{\mathbb{P}}_t\left[\sum_{k=1}^{N} r^{(1),i}_{t+k-1}\right], \tag{1}$$

where $IRS^i_{t,N}$ is the fixed rate on an N-period interest rate swap in currency i, $r^{(1),i}_s$ is the one-period reference rate in currency i at date s, and $\mathbb{E}^{\mathbb{P}}_t[\cdot]$ is the conditional expectation under the physical measure. Equivalently, $tp^i_{t,N} = \frac{1}{N}(\mathbb{E}^{\mathbb{Q}^i}_t - \mathbb{E}^{\mathbb{P}}_t)[\sum r^{(1),i}]$, where $\mathbb{Q}^i$ denotes the currency-i risk-neutral measure. The swap term premium measures the compensation investors demand for bearing currency-i-specific term structure risk over N periods.

Definition 2 (Cross-Country Term Premium)**.** *The cross-country term premium of currency j relative to currency i at maturity N is the differential in the swap term premia,*

$$ctp^{j,i}_{t,N} \;\equiv\; tp^j_{t,N} \;-\; tp^i_{t,N}. \tag{2}$$

Throughout the paper, I measure cross-country term premia from the U.S. dollar point of

[5] Internet Appendix A provides the complete derivation, including the par-swap correction, bounded at under two basis points.

view, that is, with $i = \text{USD}$ and j a partner currency. A negative value of $ctp_{t,N}^{j,USD}$ indicates that the USD swap term premium is higher than that of currency j, meaning that investors demand more compensation for bearing term structure risk in USD than in currency j, so dollar duration is held at the less advantageous term premium relative to currency j.

From the observable basis to the term premium differential. I now derive the decomposition that links the cross-currency basis spread $x_{t,N}^{j/i}$ to the cross-country term premium differential $ctp_{t,N}^{j,i}$. Let $\rho_{t,N}^{j/i} \equiv \frac{1}{N}\log(F_{t,N}/S_t)$ denote the FX forward premium over the N-period horizon, with $F_{t,N}$ the forward exchange rate and S_t the spot exchange rate, both quoted as units of currency j per unit of currency i, so that $\rho_{t,N}^{j/i} > 0$ when j-rates exceed i-rates over the horizon. Throughout the paper, the pair superscript on every differential quantity, in the basis x, the pure parity π, the forward premium ρ, and the term-premium differentials, is read as the first-named currency minus the second-named, a/b meaning a minus b.[6] The standard cross-currency basis identity in long-tenor swap markets (Du et al., 2018b) expresses the basis as the cross-country interest rate swap differential adjusted for the FX forward premium,

$$x_{t,N}^{j/i} \;=\; IRS_{t,N}^{j} \;-\; IRS_{t,N}^{i} \;-\; \rho_{t,N}^{j/i}. \tag{3}$$

Two substitutions reduce this identity to the additive decomposition of the basis. First, each fixed swap rate is, by its risk-neutral representation, the average expected one-period rate under that currency's own risk-neutral measure, $IRS_{t,N}^{c} = \frac{1}{N}\mathbb{E}_t^{\mathbb{Q}^c}[\sum_{k=1}^{N} r_{t+k-1}^{(1),c}]$ for $c \in \{i, j\}$, stated as Internet Appendix Equation (IA.1). Second, the FX forward premium decomposes into the cross-country physical-measure expected short-rate path minus a bilateral pair-specific residual that I label $\pi_{t,N}^{j/i}$,

$$\rho_{t,N}^{j/i} \;=\; \frac{1}{N}\mathbb{E}_t^{\mathbb{P}}\left[\sum_{k=1}^{N} r_{t+k-1}^{(1),j} \;-\; \sum_{k=1}^{N} r_{t+k-1}^{(1),i}\right] \;-\; \pi_{t,N}^{j/i}. \tag{4}$$

I do not assume uncovered interest parity, which would force $\pi_{t,N}^{j/i}$ to zero; I recover it empirically. Nor do I impose covered interest parity: the decomposition instead pinpoints where any parity deviation resides, in the currency-level term-premium differential or in the bilateral friction. Structurally, $\pi_{t,N}^{j/i}$ collects whatever bilateral frictions enter the FX forward leg of the arbitrage, such as balance-sheet, regulatory, and institutional constraints on the dealer book that intermediates the (i, j) pair.

[6] This quotation is the reciprocal of the institutional cash-flow description in Internet Appendix Section B.1; the two are equivalent under FX inversion, and the choice fixes only the algebraic sign of $\rho_{t,N}^{j/i}$ in Equations (3)–(4).

Why differentials admit a model-free measurement. The physical measure exits the problem at one point, which is the source of the method's central advantage and of its boundary. The expected short-rate differential $\frac{1}{N}\big(E_t^{\mathbb{P}}[\bar{r}^{\,j}] - E_t^{\mathbb{P}}[\bar{r}^{\,i}]\big)$, writing $\bar{r}^{\,c}$ for the cumulated one-period rate $\sum_{k=1}^{N} r_{t+k-1}^{(1),c}$, enters the composition twice, through two traded prices: once through the interest rate swap differential via Equation (IA.1), and once, with the opposite sign, through the FX forward premium via Equation (4). In the cross-currency basis, which differences the two, the physical-measure expectations therefore cancel identically. The cancellation is algebraic rather than statistical: no expectation of future short rates is formed, estimated, or approximated in any country at any point in the construction. This is the precise sense in which the measurement is expectations-model-free. Estimating $E_t^{\mathbb{P}}[\bar{r}^{\,c}]$ is where auxiliary structure enters every single-country term premium measurement, whether through vector-autoregressive dynamics, affine no-arbitrage restrictions, or survey inputs, and the differencing replaces that estimation with an exact cancellation between market prices. What the construction requires instead is a statement not about the level of expectations but about the cross-sectional residence of the forward-leg residual: any currency-level component of the deviation between the long-horizon forward premium and the expected short-rate differential is, by the term-premium definition, compensation collected in $tp_{t,N}^{c}$, and Assumption 2, stated formally in Internet Appendix Section A.11, restricts the remainder to be pair-specific, a residual that long-horizon parity evidence suggests is small at the tenors I study (Chinn and Meredith, 2004) and that Section 4.1 bounds within quoted transaction costs. The same algebra delimits what the method cannot do: levels do not cancel, so the term premium of any single country is not recoverable without a model of that country's expectations, which is why the paper claims only differentials. The measured differentials in turn discipline any model-based recovery of levels, since a candidate set of country-level term premia must difference exactly to the differentials measured here, a set of over-identifying restrictions I leave to future work.

Substituting Equations (IA.1) and (4) into Equation (3), the physical-measure expected short-rate differential enters twice, once through the two swap rates and once, with the opposite sign, through the forward premium, and cancels term by term; regrouping the surviving expectations currency by currency, the risk-neutral less the physical-measure expected short-rate path of each currency is exactly that currency's term premium by Definition 1, and the basis reduces to the additive decomposition (the line-by-line display is Internet Appendix Equation (IA.2)),

$$x_{t,N}^{j/i} \;=\; \big(tp_{t,N}^{j} - tp_{t,N}^{i}\big) \;+\; \pi_{t,N}^{j/i} \;=\; \underbrace{ctp_{t,N}^{j,i}}_{\text{currency-specific}} \;+\; \underbrace{\pi_{t,N}^{j/i}}_{\text{bilateral pair-specific}} \,. \tag{5}$$

The first term is the cross-country term premium differential I seek, which is the cross-country term premium of Definition 2.

The two currency-specific terms $tp^i_{t,N}$ and $tp^j_{t,N}$ are attributes of currencies i and j individually and enter every basis swap that involves either currency. The second term $\pi^{j/i}_{t,N}$ is the pair-specific residual, an attribute of the bilateral cross-currency pair, for instance, the dealer book that intermediates the (i, j) pair specifically, and that does not enter any other currency pair such as (i, k). The two components of the decomposition are economically and algebraically distinct: cross-country term premium differentials load identically on every basis involving the same two currencies, while pair-specific residuals do not.

Under the normalisation $tp^{USD}_{t,N} = 0$, the triangulation recovers $ctp^{j,USD}_{t,N}$ for every currency j in the cross-section; the choice is arbitrary because only the cross-country differentials are identified, and re-estimating with $tp^{EUR}_{t,N} = 0$ leaves every bilateral differential and every subsequent regression coefficient numerically unchanged. In the empirical analysis I work directly with the differentials measured from the U.S. dollar point of view.

2.3 Triangular Identification and Extraction

For three currencies i, j, k, define the triangular inconsistency at tenor N as

$$\Delta^{i,j,k}_{t,N} \equiv x^{j/i}_{t,N} + x^{k/j}_{t,N} - x^{k/i}_{t,N}, \tag{6}$$

the difference between the cost of the indirect route from currency i to currency k via currency j and the cost of the direct route. Substituting the additive decomposition from Equation (5) into each of the three bases on the right-hand side, the currency-specific term premium components cancel exactly because each currency appears with opposite signs across the three pairs, leaving only the pair-specific frictions:

$$\Delta^{i,j,k}_{t,N} = \pi^{j/i}_{t,N} + \pi^{k/j}_{t,N} - \pi^{k/i}_{t,N}. \tag{7}$$

This is the first step of the identification result. Each observed triangular inconsistency provides one linear equation in the underlying pair-specific frictions, and the equation is free of any currency-specific term premium content.

The procedure then proceeds in two algebraic steps. First, identify the pair-specific frictions from the system of triangular equations. Second, recover the cross-country term premium differentials by going back to the additive decomposition of the basis: subtracting the recovered

pair-specific friction from the observed basis leaves only the currency-specific component, which by Equation (5) is the cross-country term premium differential.

In my study I have ten currencies (USD, EUR, JPY, GBP, CHF, CAD, AUD, NZD, SEK, NOK), with the Swiss franc entering outside its minimum exchange rate window.[7] With ten currencies there are $\binom{10}{2} = 45$ potential bilateral pairs. The cross-currency basis swap market quotes seventeen of these as independent bilateral pairs at long tenors: nine USD-based pairs (USD against EUR, JPY, GBP, CHF, CAD, AUD, NZD, SEK, and NOK) and eight EUR-based pairs (EUR against JPY, GBP, CHF, CAD, AUD, NZD, SEK, and NOK), with the euro-leg quoting histories and outages documented in Internet Appendix Section D.3. The remaining pairs are mechanically constructed from these seventeen through the cancellation of the USD or EUR leg, contain no independent pricing information, and are excluded from the system so that each basis observation contributes independent information to the triangulation. The quoted hub-and-spoke structure, centred on the USD-EUR pair, generates eight independent triangular equations of the form USD-EUR-X for $X \in \{\text{JPY}, \text{GBP}, \text{CHF}, \text{CAD}, \text{AUD}, \text{NZD}, \text{SEK}, \text{NOK}\}$.

The eight triangular equations on the seventeen bilateral frictions can be written compactly in matrix form as

$$\mathbf{D}\,\boldsymbol{\pi}_{t,N} \;=\; \boldsymbol{\Delta}_{t,N}, \tag{8}$$

where $\boldsymbol{\pi}_{t,N}$ is the 17×1 vector of pair-specific frictions across the seventeen independently quoted pairs, $\boldsymbol{\Delta}_{t,N}$ is the 8×1 vector of triangular inconsistencies across the eight available triads, and $\mathbf{D}$ is the 8×17 incidence matrix that encodes which pair-specific frictions enter each triangular equation, with three non-zero entries per row.

The matrix $\mathbf{D}$ has a non-trivial null space. With m independently quoted pairs among n currencies, the cycle space delivers $m - (n-1)$ independent triangular equations for m unknowns, so the null space of $\mathbf{D}$ has dimension $n - 1$; on the full ten-currency network of seventeen quoted pairs the system has seventeen unknowns and eight independent equations,

[7] The Swiss National Bank operated an explicit EUR floor on the CHF from September 2011 through January 2015 and pursued substantial direct FX intervention through much of the rest of the post-2008 sample. The administered exchange rate makes the cross-currency basis swap quote at long tenors a function of the policy regime rather than of the underlying SDF projection, which would distort the recovered term premium differential in ways that the methodology cannot net out. I therefore exclude the floor window from the Swiss series and keep the Swiss pair outside the estimation panels, reporting it for descriptive completeness only. The principle is symmetric across administered markets: I exclude a market over the window in which the administered instrument is the one my quotes price. The Swiss floor administered the exchange rate itself, the instrument underlying the cross-currency basis, whereas Japanese yield-curve control administers the domestic government yield, which enters my construction only through the observable swap spread differential while the yen basis remained market-determined, so the yen pair stays and the Swiss floor window does not; Section 4.3.3 shows that excluding the yield-curve-control window leaves the supply results unchanged.

and the same count applies at each date to the subnetwork of pairs quoted on that date. The null space has a clean economic interpretation: any vector of the form $\pi^{j/i} = \lambda^i - \lambda^j$ for arbitrary currency-specific values $\lambda^1, \ldots, \lambda^n$ produces zero triangular sums by exactly the same algebraic cancellation that eliminates the cross-country term premium differentials. The null space of $\mathbf{D}$ is therefore the space of currency-specific shifts in π. The triangular equations alone leave the level of $\boldsymbol{\pi}$ underdetermined by these $n-1$ dimensions; combined with the structural restriction that $\boldsymbol{\pi}$ contains no currency-level shift, motivated by the institutional pair-specificity of these frictions and made formally explicit through the swap-term-premium definition that absorbs every currency-c-specific compensation in the swap rate into $tp^c_{t,N}$, the level becomes exactly identified.

This structure delivers something precise, and the term-premium interpretation of the recovered cross-country term-premium differential rests on it.

Proposition 1 (Identification and its limits)**.** *Let $\mathbf{x}_{t,N}$ stack the m observed bases among n currencies and let $(\mathbf{tp}_{t,N}, \boldsymbol{\pi}_{t,N})$ be any pair consistent with Equation (5). Then: (i) the cycle-space component of $\boldsymbol{\pi}_{t,N}$, equivalently the vector of triangular inconsistencies and every within-triangle combination of bilateral residuals, is point-identified from observed prices alone, with $\|\boldsymbol{\pi}^{\text{bilat}}_{t,N}\| \leq \sigma^{-1}_{\min}\|\boldsymbol{\Delta}_{t,N}\|$, where $\sigma_{\min}$ is the smallest non-zero singular value of $\mathbf{D}$; (ii) the currency-level component is identified only up to the $(n-1)$-dimensional equivalence class of Internet Appendix Proposition 3; under Assumption 2, which restricts $\boldsymbol{\pi}_{t,N}$ to carry no currency-level component, the node vector $\mathbf{tp}_{t,N}$ is point-identified up to the numeraire normalisation, and it equals the vector of swap term premia if and only if every currency-level component of the long-tenor basis is term-premium content; (iii) the division of the identified currency-level component between physical-measure term premia and any currency-level non-term-premium wedge is not identified by prices within the network, and no reallocation across that division changes any observed basis or triangle.*

Part (i) is what the market data deliver unconditionally, and Section 4.1 shows the identified bilateral component is bounded by executable transaction costs. Part (ii) is where the economics enters: Assumption 2 is a structural restriction on the cross-sectional organisation of frictions, motivated by the pair-specific dealer-book institutions described below and by the swap-term-premium definition that collects currency-specific compensation into $tp^c_{t,N}$, but the definition motivates the restriction rather than enforcing it; the physical-measure expectations themselves cancel identically in the composition of the swap and forward legs and are never estimated, so the restriction concerns only the cross-sectional residence of the forward-leg residual, not the level of expectations. Part (iii) is the residual uncertainty that no norm choice can remove; I treat it by testing every alternative reading of the currency-level compo-

nent that carries measurable implications (Section 4.2) and by bounding the one direction that cannot be tested, so every headline quantity in the paper is to be read as conditional on Assumption 2, with the bounding exercise quantifying its sensitivity to economically calibrated violations.

The complete development, from the system formulation through the minimum-norm solution and the worked numerical example, together with the triangulation standard errors and the formal statement of both assumptions with the catalogue of potential limitations, is in Internet Appendix Sections A.2 through A.9, A.10, and A.11 through A.12.

The bridge to the government-bond differential has one property that belongs in the framework rather than among the results, because it disposes of an entire class of objection before any of them is raised.

Corollary 1 (Swap-rate invariance of the cross-country government bond term premium)**.** *Writing the cross-country government bond term premium as* $x + (SS^{US} - SS^{j})$ *with* $SS^{c} = s^{c} - y^{c}$ *and expanding, the swap rates cancel identically, so the measured differential reduces to* $(y^{j} - y^{US}) - \left[(s^{j} - s^{US}) - x\right]$ *in observables alone. The market forward premium quoted through the swap package at these tenors is* $(s^{j} - s^{US}) + x$*, so the measured differential equals the market-hedged sovereign yield differential plus twice the quoted basis: the hedged differential treats the basis as the cost of the hedge, while the measured differential carries it as term-premium content. Any component residing in the swap rates themselves, including benchmark credit in the floating legs, swap-market funding or collateral effects, and convexity in the fixed rate, cancels exactly in the bridge, whatever its size and dynamics.*

The cancellation is the most direct advantage the construction has over any measure built from swap rates: the swap rates do not merely wash out on average or under a modelling assumption, they cancel line by line, so benchmark credit in the floating legs, swap-market funding and collateral effects, and convexity in the fixed rate are removed exactly, whatever their size, sign, or dynamics, with no estimate, bound, or robustness check needed. What survives is only what lives in the cross-currency basis quote itself, the dollar funding premium of the foreign-exchange swap market, and that single surviving channel is precisely what the sensitivity frontier, the crisis-response ratios, and the overturning thresholds of Section 4 are built to address: the corollary converts an open-ended list of possible contaminations into one named channel with printed bounds on it. Its further consequences are developed in Internet Appendix Section A.14.

2.4 Properties and Inference

The methodology is model-free in a precise sense: no physical-measure expectation of future short rates is estimated in any country, the expectations cancelling identically in the composition of the swap and forward legs (Section 2.2), and the recovery at each date and tenor emerges from no-arbitrage relationships applied to the cross-section of observed bases, requiring no dynamic model of term premia, no cross-currency correlation structure, no distributional or stationarity assumptions, and no market completeness, with covered interest parity nowhere imposed. The identification is exact under, and conditional on, the pair-specificity restriction of Assumption 2, as Proposition 1 states.

Because the identification apparatus will draw the closest scrutiny, Internet Appendix Section E.1 probes it from several directions, and the exercises establish that the recovery does not depend on the apparatus, the split between currency-level and bilateral content, as distinct from what the currency-level component is, which Section 4.2 addresses. The recovered differentials are essentially invariant to the selection rule within the affine solution set, least-absolute-deviations, Chebyshev, and a spread-anchored Bayesian rule included, with correlations at or above 0.9997 and extracted supply headlines within 0.003 of the baseline; they are invariant to excluding any single currency, the euro included, and to bootstrapping the entire extraction. A synthetic-network simulation recovers the true differential at correlation 0.9999 and absorbs a dollar-common friction one for one into the U.S. leg, while a placebo of pure pair-specific funding does not manufacture a currency-level component silently, its implied residual share being sixty-seven percent of quotes against the 2.7 percent observed on the identical convention.

The one solution no criterion selects, calling the entire wedge friction and the term premia zero, fits every triangle by construction but requires frictions of twenty-five to eighty basis points against quoted spreads of one to three and leaves the supply pricing and the event repricings of the joint falsification standard without a source; and random members of the solution set at three contamination scales confirm the validation evidence is not a knife-edge property of the restricted point, the supply coefficient surviving every draw, with discrimination against systematic relabellings supplied by the sensitivity-frontier bounds and the joint falsification standard.

2.5 Stochastic Discount Factor Interpretation

The decomposition has a direct reading through covered interest parity and the change of measure that prices it. Covered parity holds risklessly only for legs that carry no term-horizon risk, so it is a statement about the expectations components of the two swap rates: pricing every leg under a single collateral-consistent measure, the change of measure between the two currencies' money-market accounts leaves the expectations legs matched up to a small residual, the genuine parity deviation. A long-dated swap rate, however, is not risk-free: it carries a term premium, the compensation for bearing duration risk, and two long-dated swaps in different currencies are not the same risk because their term premia differ.

What the cross-currency basis does at the long end is precisely the fair, like-for-like risk adjustment that this difference requires, compensating the leg with the smaller term premium so that the exchange is risk-for-risk fair after the measure change, which is the decomposition $x^{j/i}_{t,N} = ctp^{j,i}_{t,N} + \pi^{j/i}_{t,N}$ with $ctp^{j,i}_{t,N} = tp^{j}_{t,N} - tp^{i}_{t,N}$ and the parity residual $\pi^{j/i}_{t,N}$ bounded within transaction costs by the containment result of Section 4.1. At the short tenors where the covered-interest-parity literature has mostly worked the term-premium differential is negligible, so identifying the whole basis with the parity deviation is innocuous; at long tenors the larger part of the basis is a term-premium adjustment rather than a deviation from parity, and the true deviation is an order of magnitude smaller. Internet Appendix Sections B and B.4 develop the explicit stochastic-discount-factor representation, which delivers the same additive currency-plus-pair structure and locates the restriction precisely, and the covered-interest-parity route in full.

3 Data and Implementation

3.1 Data Sources

The primary dataset comprises monthly mid-market prices for cross-currency basis swaps from Bloomberg for seventeen independently quoted bilateral pairs among ten currencies (USD, EUR, JPY, GBP, CHF, CAD, AUD, NZD, SEK, NOK): nine dollar-based and eight euro-based pairs. The remaining cross-rates among non-dollar, non-euro currencies are not independently quoted at long tenors and are mechanically implied by the seventeen independent pairs, so I exclude them so that each observation contributes independent information to the triangulation. The sample spans January 2002 through June 2026 at headline tenors of two, five, and ten years; the one-year tenor enters as the observable short-end funding benchmark,

and the twenty-year tenor, quoted on every dollar pair, enters the tenor-slope, far-forward, and return-forecasting constructions and the robustness exercises. The currency node recovered from the swap market is euro-area-wide, since euro interest rate swaps reference the area-wide benchmark, while the government bond leg of the euro pair uses German Bunds as the euro-area reference sovereign; I accordingly write the pair as USD-EUR when the reference is the currency node and refer to Bunds when the statement concerns the government bond leg.

Government bond yields come from Bloomberg sovereign zero-coupon curves and swap rates come from Bloomberg swap zero-coupon curves for each of the ten currencies, enabling construction of the conventional swap spread $SS^k_{t,N} \equiv IRS^k_{t,N} - r^{gov,k}_{t,N}$ in each currency. The cross-country swap spread differential $\Delta SS^{USD,j}_N \equiv SS^{USD}_N - SS^j_N$ that enters the bridge from the swap to the cross-country government bond term premium is computed for every currency j in the cross-section and every tenor used in the analysis.

The macro-financial dataset carries the intermediary capital ratio of He et al. (2017), the broker-dealer leverage factor of Adrian et al. (2014), the VIX, the MOVE index, a corporate credit spread, the broad nominal dollar index, and country-specific quantitative-easing indicators; the supply-side dataset carries gross debt, central-bank holdings, the free float they imply, and par-weighted average maturity for each sovereign. Internet Appendix Section D gives every series, its source, its frequency, and its construction.

3.2 Benchmark Consistency of the Swap-Spread Differential

Every leg of every series is chain-reconciled to a single benchmark convention across the IBOR-to-RFR transitions, using overlap means so that no series carries a splice discontinuity; throughout the paper, "benchmark" attached to a series name refers to this construction. Figure 5 settles what is at stake more directly than the construction can: the reconciled and the quoted-legacy series differ by 0.2 basis points through the cessation window, and the central reversal is present in full on a sample that ends in December 2021, before any cessation date, at +27.9 to −51.3 basis points for the euro pair and +37.4 to −75.7 for the yen, so the transition cannot be what produces the headline. The full-depth seventeen-pair network through December 2021 is accordingly the primary identification sample, every headline result present in full upon it, and the benchmark-consistent series through June 2026 is an extension, whose thinner post-transition network is documented where the quoting histories are recorded. Internet Appendix Section D.2 reports the construction leg by leg, the transition dates, the size of the discontinuities the treatment removes, the pre-2002 extension and its per-exhibit use, and the residual reconciliation diagnostics by currency.

3.3 Implementation

Internet Appendix Section B.1 develops the institutional cash flows of the fixed-for-fixed contract, the collateral and discounting conventions by era, the treatment of the notional exchanges and of mark-to-market as against constant-notional quotation, and the exact construction of synthetic fixed-for-fixed prices, with the par-swap representation wedge bounded at under two basis points; Section 3.2 chain-reconciles every benchmark transition leg by leg on identical conventions and dates; and quoted prices are validated against transaction activity in the swap data repository in Internet Appendix Table IA.14.

For each date and tenor I form the available independent triangles, stack them as in Equation (8), and solve for the bilateral pair-specific frictions by the minimum-norm least-squares solution via the Moore-Penrose pseudoinverse; subtracting the recovered frictions from the observed bases gives the cross-country swap term premia, and adding the directly observable swap-spread differential gives the cross-country government bond term premium through the bridge, which holds exactly by construction at every date and tenor, so no individual-country term-structure model is estimated at any step. Standard errors are computed analytically by linear variance propagation and by cross-sectional bootstrap of the within-date quotes. Internet Appendix Section A develops every step, with the solution and recovery in Internet Appendix Equations (IA.5) and (IA.6) and the bridge in Internet Appendix Equation (IA.7).

4 Empirical Results

I proceed in three steps. Section 4.1 establishes from quoted prices that the bilateral, pair-specific component of the basis is quantitatively negligible, so the observed network is close to transitive; because the minimum-norm construction selects the smallest friction vector consistent with the observed triangles, I present this as a property of the data rather than as validation of the economic interpretation. Section 4.2 documents the post-crisis reversal in the recovered government term premia and stress-tests it over the unidentified currency-node space, an Andrews single-break test dating the consolidation at October 2011, read as a location estimate rather than a significance statement under the series' persistence, with the multiple-break analysis in Internet Appendix Section E.2 and the margin-by-margin decomposition in Internet Appendix Table IA.15. Section 4.3 then asks whether the remaining currency-level component displays the supply, maturity, and policy-shock behaviour of term compensation; the exploratory dollar-prediction exercise is in Internet Appendix Section I.1. The paper's four claims and their epistemic standing are restated at the head of Internet

Appendix Section E and in Table 9; Internet Appendix Table IA.13 consolidates the sample, frequency, and sources behind every headline exhibit, and Internet Appendix Table IA.16 documents the free-float construction country by country.

4.1 Empirical Composition of the Cross-Currency Basis

The decomposition in Equation (5) writes the bilateral cross-currency basis additively into the cross-country term premium differential $ctp^{j,i}_{t,N}$ and the pair-specific component $\pi^{j/i}_{t,N}$, and the relative empirical content of the two establishes the near-transitivity of the network.

The quoted dealer spreads sharpen this reading into a bound. The mean absolute residual ranges from 0.2 to 0.9 basis points across the eight pairs, against mean quoted half-spreads of 1.2 to 3.0 basis points, and between 80 and 99 percent of monthly observations lie inside the half-spread (Internet Appendix Figure IA.1). The pair-specific component is therefore bounded by transaction costs, the configuration triangle arbitrage at executable prices enforces, while the currency-level component, an order of magnitude outside the band, is a price rather than a friction; and because the band is the market's rather than the estimator's, the bound also answers the concern that the minimum-norm construction makes the residuals small by design. Reported transactions carry the same reading at the level of prices rather than spreads: over 2013 to 2020 the monthly median traded cross-currency basis sits 1.2 basis points from my quoted series at the median of 1,005 pair-month cells, inside the pair's own quoted half-spread in 60 percent of them, and correlates with it between 0.82 and 0.99 across the five pairs the repository covers deeply (Internet Appendix Table IA.2).

A network simulation confirms both properties of the extraction, recovering the true currency-level differential at a correlation of 0.9999 and, under a placebo of pure pair-specific funding, not manufacturing a currency-level component silently (Internet Appendix Table IA.12; the behavioural validation of the bilateral frictions against funding-stress and quarter-end indicators, and the swap-data-repository transaction evidence, are in Internet Appendix Sections H.2 and A.13).

The cross-country term premium differential is accordingly the first-order empirical content of the bilateral basis: across the twenty-four pair-tenor combinations the share $|\pi|/|x|$ averages 2.3 percent, and 0.8 percent outside two near-zero-basis cells whose near-zero denominators inflate the ratio, while the two components are approximately uncorrelated in the time series (Internet Appendix Table IA.12 and Figure IA.2, with the complete pair-and-tenor statistics in Internet Appendix Section E.4). The pair-specific and currency-specific components are therefore both quantitatively distinct in magnitude and statistically nearly orthogonal in their

variation, consistent with the structural reading developed in Internet Appendix Section B.

4.2 A Reversal in Relative Sovereign Duration Premia

The recovered ten-year government term premia are persistent and wide relative to their measurement error, with first-order autocorrelations from 0.90 to 0.99, standard deviations from 21 to 57 basis points against quoted half-spreads of one to three basis points, and a common component tied to the shared U.S. leg that carries 68 percent of the cross-sectional variation (Internet Appendix Section E reports the principal-component structure and leave-own-out loadings). The reversal is not a restatement of observable swap spreads: for the yen the post-GFC swap-spread differential averages −21.8 basis points against a cross-country government bond term premium of −83.2, the balance of −61.4 residing in the triangulated component, which therefore contributes just under three quarters of the reversal for that pair (Internet Appendix Table IA.15).

Table 2 reports the full-sample mean, standard deviation, and observation count of the three series, the government differential, its swap component, and its swap-spread component, by pair at the two-, five-, and ten-year tenors; sub-period means on the calendar windows fixed across the descriptive and reversal exhibits, pre-crisis January 2002 to December 2007, a crisis window January 2008 to December 2012 bracketing the September 2008 Lehman collapse (Jermann, 2020; Klingler and Sundaresan, 2019), and post-crisis January 2013 to June 2026, are reported with their persistence-robust inference in Table 3, the split corroborated by formal estimation dating the onset in late 2007 with every subsequent segment mean below zero (Internet Appendix Sections E.2 and A.15).

Table 2 about here

Figures 1 and 2 about here

Two features anchor what follows. First, under my sign convention $\widehat{ctp}_N^{gov,j,USD} = tp_N^{gov,j} - tp_N^{gov,USD}$, a positive value indicates that U.S. Treasuries carry the lower term premium of the two and command a term-premium advantage at horizon N, while a negative value indicates the opposite. The average is positive in the pre-GFC subsample (2002 to 2007) for seven of the eight pairs with a pre-crisis identified history, indicating that U.S. Treasuries commanded a cross-country term premium against the major reserve-currency sovereigns before 2008. The pair-level mean flips sign during the GFC subsample for six of the eight pairs, the New Zealand pair flipping in the opposite direction, and is negative in the post-GFC subsample for seven

of the eight, sharply so for the euro, the yen, the Canadian dollar, and the two Scandinavian currencies, with means declining monotonically across the three sub-periods for five of the eight pairs. The New Zealand pair is the one significant counterexample: its shift of +33.7 basis points, with a Newey-West t-statistic of +3.60 and lying outside its no-break band, runs in the opposite direction, and I report it as such rather than explain it away.

TABLE 3 ABOUT HERE

The ten-year term premium of German Bunds over U.S. Treasuries averaged +28 basis points pre-GFC, −33 basis points during the GFC, and −56 basis points post-GFC; for Japanese government bonds the corresponding averages are +37, −28, and −83 basis points (Table 3, which reports the sub-period means pair by pair together with persistence-robust inference on the pre-GFC to post-GFC shift). The reserve-currency finding that motivates the paper is therefore a post-crisis phenomenon rather than a sample-wide average.

Second, the cross-country swap spread differential ΔSS switches sign across the sub-periods for five of the nine pairs, reflecting the emergence of persistently negative U.S. swap spreads after the crisis (Klingler and Sundaresan, 2019; Jermann, 2020; Hanson et al., 2022), and under the bridge identity of Internet Appendix Equation (IA.7) it amplifies the swap-term-premium component on the six pairs where the reserve-currency reversal lives while offsetting it for the Australian and New Zealand pairs (Internet Appendix Section E).

Figures 1 and 2 display the bridge-identity decomposition pair by pair. The post-2008 downward drift of $\widehat{ctp}_{10}^{gov,j,USD}$ against the euro and Canadian pairs is driven primarily by the persistently negative cross-country swap spread differential, while the swap-term-premium component is more stable, and for the yen pair the two components contribute roughly equally (Internet Appendix Section E reads the trajectories pair by pair). Figure 3 displays the complementary split at ten years for the euro and yen pairs: the observed government yield differential, its expectations component, the neutral-rate differential released with the paper, and the recovered term-premium differential, the wedge between the two. The expectations component follows the yield differential closely throughout, so the movement of the wedge across the crisis is the reversal of Table 3 seen from the identity's other side.

FIGURE 3 ABOUT HERE

Under the bridge identity $\widehat{ctp}^{gov,j,USD} = \widehat{ctp}^{swap,j,USD} + \Delta SS^{USD,j}$, the sign reversal in the cross-country government bond term premium for EUR, GBP, and CAD is the algebraic consequence of the post-2008 swing of the U.S. swap spread to persistently negative territory

while peer-country swap spreads remained near zero or positive (Klingler and Sundaresan, 2019; Jermann, 2020; Boyarchenko et al., 2018; Hanson et al., 2022).[8]

What the cross-currency component adds. Because the post-crisis behaviour of United States swap spreads is separately documented, the question a reader will press is whether the cross-currency basis is doing any work at all. Table 4 answers it by putting the basis, the swap-spread differential, and their sum through identical tests, and it places each component in the literature it belongs to: the swap-spread differential is the within-country wedge whose post-crisis United States behaviour the swap-spread literature documents (Feldhütter and Lando, 2008; Klingler and Sundaresan, 2019; Jermann, 2020; Boyarchenko et al., 2018; Hanson et al., 2022), the long-tenor basis is the instrument the covered-interest-parity literature reads as a funding friction at short tenors (Du et al., 2018b; Liao, 2020; Du et al., 2026), and the composite is the government term-premium differential that neither literature measures. What was knowable from the swap-spread differential alone is the familiar post-crisis decline of United States swap spreads; what the basis adds is the cross-country swap-side component, nearly half of the Japanese reversal; and only the composite carries the government term-premium reading, the largest supply loading, and both purchase-factor responses. Three things follow. Duration supply prices the composite more strongly than either part, at -1.088 against -0.825 for the swap-spread differential and -0.257 for the basis, so the sum is not a diluted version of its larger component. Purchase surprises point the same way on both legs: the composite responds to the Federal Reserve asset-purchase factor at $+0.68$ basis points per standard deviation against $+0.56$ for the swap-spread differential and $+0.12$ for the basis, and to the European Central Bank purchase factor at -0.97 against -0.61 and -0.37, so on the policy margin too the composite carries more than either part. And the composite is the series with the largest reversal, -50.3 basis points on average against -41.1 and -9.2. The covered-parity deviation of the government CIP literature is carried as a fourth numerical column, being the swap-spread differential less the basis rather than plus it, and the evidence selects the sum: the difference shows a smaller mean reversal at -31.8 basis points, a supply loading of -0.562 against the composite's -1.088, purchase-factor responses that are smaller on both legs and insignificant on the euro leg, and affine-model agreement of four of six pairs against five, so the sign combination that defines the government term-premium differential,

[8]The persistent post-2008 negativity of U.S. swap spreads is conventionally attributed in this literature to dealer balance-sheet capacity, the post-Basel III supplementary leverage ratio, and pension and insurance-company receive-fixed demand at long tenors. Under my bridge identity, any such dealer-balance-sheet, regulatory, or institutional-demand effect that loads on SS_N^{USD} enters $\widehat{ctp}_N^{gov,j,USD}$ algebraically through $\Delta SS_N^{USD,j}$ rather than through $\widehat{ctp}_N^{swap,j,USD}$, and the channel decomposition in Section 4.3.2 traces each driver to the bridge component that carries it.

not the one that defines the CIP deviation, is the one the term-premium predictions single out.

What the measurement adds beyond the observable wedge. The measurement adds four things to what the observable wedge already delivers. The first is that it is a cross-country differential rather than a U.S. level, so that every component common to swap markets globally, bank credit quality and common regulatory costs among them, nets out. The second is that U.S. duration supply prices both bridge components with the same sign, the basis carrying roughly a fifth of the loading, an over-identifying consistency across two distinct markets that a pure balance-sheet reading has no reason to deliver. The third is that the tenor and ceiling battery of Section 4.3.4 tests the term-premium interpretation instead of taking it as given. And the fourth is that the containment result bounds the pair-specific residual within quoted transaction costs, which is what licenses reading $x + \Delta SS$ as a term-premium differential at all. Nor is either component alone the differential: reading the swap-spread differential alone as the premium imposes identical swap term premia across countries, which the recovered swap-side series rejects at -61 basis points for the yen after 2013, the division of variance between the components is itself time-varying with the crisis as the hinge, and for the Australian and New Zealand pairs the two components take opposite signs in level, a configuration no single-component reading can produce. Internet Appendix Table IA.17 reports the channel split and the time-varying division.

The bound over the unidentified direction. The direction no test can reach is bounded by a calibrated sensitivity frontier. The logical standing of the four steps, the unconditional network property, the pair-specificity assumption, the tested term-premium interpretation, and the sensitivity frontier for the unreachable direction, is set out in Internet Appendix Section H.4; nothing of economic substance rests on the label alone.

Identified intervals for every headline quantity. The unidentified direction's effect on each headline result quantifies the conditionality of Proposition 1. The unidentified direction is a dollar-wide level wedge λ_t added to every USD-pair differential. Table 7 reports, for each headline quantity, the smallest wedge path that would overturn it, against two verified calibration anchors from the observed funding benchmark, the cross-pair average one-year USD basis, which is the funding benchmark of record throughout the paper, the own-pair one-year basis entering only where a pair-specific comparator is the more informative one and each use stating which series it employs: its sustained post-2009 level of -11.6 basis points,

and its worst twelve-month episode of 26.3 basis points in absolute value. Three features stand out. First, the reversal is immune to any constant wedge, which cancels in the pre-post difference; overturning it requires the wedge itself to have moved by 84 to 120 basis points at the break, seven to ten times the sustained level of any observed funding basis. Second, overturning the supply coefficient requires a wedge that loads on log U.S. free-float debt at $+1.05$, whereas the observed funding benchmark loads at -0.233 ($t = -5.24$) on the eight-pair funding panel, the supply coefficient itself estimated on the seven-pair free-float panel as the table's caption records, the opposite sign, so no rescaling of observed funding behaviour can generate the overturning path. Third, the December 2022 event response requires the wedge to jump 12.5 basis points inside a two-day window in which the observed funding benchmark moved 1.2 basis points, both measured over the same $[-1, +2]$ window. The carry results are rank-based, and a uniform wedge shifts all USD-pair signals equally, leaving the cross-country ranking unchanged. Every headline sign survives throughout the economically calibrated set.

An executable band built from perturbing every quoted basis within its half-spread averages 3.6 to 5.1 basis points across the dollar pairs, an order of magnitude below the post-2013 means, so the reversal survives the entire band and the monthly interval excludes zero in 96 to 100 percent of post-2013 months for the German, Japanese, Canadian, and Scandinavian pairs (Internet Appendix Table IA.27).

TABLE 7 ABOUT HERE

The identifying restriction itself is bounded, and the bound must run over the whole space the restriction selects from rather than along one direction of it. The unidentified space has dimension $n - 1$, one free level per currency: a wedge specific to the yen, or to the euro, survives every triangle and reproduces in dollar-free crosses exactly as a dollar-wide wedge does, so an exercise that reallocates only a dollar-common component leaves those directions untouched. I therefore bound all of them at once.

Each currency is allowed a long-horizon wedge equal to κ_c times its own node in the THREE-MONTH network, solved on the same seventeen pairs and the same incidence structure as the long-tenor system, with each κ_c free in $[-\bar{\kappa}, +\bar{\kappa}]$ and free to take whichever sign is most adverse, independently across currencies. The short-end node is the right scale because any funding wedge must appear where funding lives and where a term premium is negligible, and κ_c is then the pass-through from the short node to the ten-year node, measured at about 0.05 in the data. Table 5 reports the worst case over that whole family. The family indexes each wedge's time path to its own short-end funding node, so the frontier covers every cross-sectional combination of capitalised funding wedges; a long-horizon currency-level wedge with

no short-end footprint lies outside it, and that direction is confronted by the joint falsification standard and the external validation rather than by this frontier.

At $\bar{\kappa} = 1$, which is full capitalisation and the ceiling a permanent funding state attains, the reversal keeps its sign for Germany, Japan, Canada, Sweden and Norway, the German interval running from -126 to -42 basis points and the Japanese from -161 to -80; and the supply coefficient moves from -0.953 to -0.526 with a t-statistic of -2.28 under the adversarial combination of signs across all currencies (Table 5). Those conclusions therefore cannot be overturned by any capitalised funding wedge whatsoever, in any combination. Three pairs do not clear that bar: sterling, the Australian dollar and the New Zealand dollar can have their sign flipped at $\bar{\kappa}$ near one, and I report them as bounded rather than established. They are the three pairs whose observed movement is smallest relative to the short-end variation of their own currency, which is the expected place for a bound to bite and not a separate finding.

No regime partition removes the reversal. The break date is a location estimate rather than a significance result, the single-break test being oversized under this much persistence, so the question worth settling is whether any reasonable partition removes the result: sweeping the start of the post-window month by month across 2010 to 2014 with the pre-window fixed, every one of the eight partners keeps its sign at every start, New Zealand included and positive throughout, and a moving-block bootstrap of twenty-four months on the post-window observations keeps the sign at the adverse tail for every pair as well (Figure 6). Nor is the euro reversal a Bund-specific artefact: re-benchmarking the euro leg to French or Dutch ten-year sovereign yields preserves the sign change, at -40 and -67 basis points against -84 on Bunds (Internet Appendix Section E.2).

One multiplicity treatment over the whole family. Correcting a preferred subset and disclosing the failures separately is transparent but weak, so a single Holm-Bonferroni adjustment is applied under one membership rule: the family is every t-statistic this paper reports, in a table or in prose, each test counted once, spanning the supply battery, the policy-factor cells, the composition-news cell, and the twelve cells of the survey-revisions orthogonality exercise, failures included. The component and covered-parity columns of Table 4 are linear recombinations of the composite's two observable parts, reported for comparison rather than as additional hypotheses, and are outside the family for that reason, as is the composite's own supply cell there, which restates the supply battery's design rather than adding a test. That gives thirty-five tests, of which twenty-six reach five percent uncorrected and twenty survive the correction: the supply battery in its entirety, all eight tenor cells and all five low-frequency

cells, the twenty-year European Central Bank purchase cell at $t = -3.1$ with a Holm-adjusted p of 0.037, and six cells of the survey-revisions exercise, the five expectations-component loadings and the sterling premium loading whose sign is opposite to the contamination direction. No other policy cell survives, the nearest being the twenty-year Federal Reserve purchase cell at a Holm-adjusted p of 0.109, and I report that rather than shrink the family until it does. That ordering is the paper's own: the supply evidence survives family-wise correction comfortably and the policy-factor evidence at its longest tenor, which is why the two families are presented in the order they are. The event-window sign universe is not in this family and cannot be: binomial probabilities were removed from the paper, that evidence rests on magnitudes and placebo percentiles, and there are no p-values there to correct.

What the reversal does and does not claim. The reversal involves claims of increasing strength. That the observable configuration of yields, swap rates, and bases changed after the crisis is directly observable in prices. That the currency-additive component of the long-tenor basis, and with it the measured cross-country term premium, changed sign is the paper's measurement, exact under Assumption 2 and bounded against its failure in Table 7. That relative duration compensation changed is supported by two families of evidence on the measured cross-country term-premium differential: it is priced by duration supply with the preferred-habitat signs across two decades, and its far-forward component behaves as persistent compensation and not as any stationary funding process, repricing within days of identified policy surprises. The paper claims no more than this: I do not claim that relative convenience services in the sense of the short-maturity literature changed sign, which is a different quantity, and the reserve-currency implication is asserted only in the dimension measured here, the long-horizon term compensation at which the two sovereigns' duration trades. Within the conditioning the panel supports, the reversal survives the five-column controls and holds at the two-, five-, and ten-year tenors with the monotone gradient the duration mechanism implies; components the panel cannot separate are absorbed into the bounding exercise rather than claimed away.

The reversal against independent survey expectations. Professional forecasters' projections carry no swap-market information of any kind, so replacing the expectations content of the yield differential with what forecasters said they expected provides a check a funding account cannot pass by construction; the informative question is incremental, whether subtracting the recovered differential moves the implied expectations closer to the survey than

assuming the differential is zero.[9] It does so in both populated one-year cells, the euro root mean squared distance falling from 96.3 to 85.9 basis points over 95 quarterly rounds and sterling from 36.3 to 24.9 over 59 with the correlation rising from 0.945 to 0.974, while the sixteen-round sterling two-year cell is marginally the other way and reported as such. On the reversal itself the euro cell is the one that can speak: its recovered differential falls by 33.5 basis points between the pre-2008 and post-2013 windows against a survey-implied fall of 42.9, while the sterling one-year cell runs the other way, the pair whose recovered post-2009 mean on the model-comparison dates lies within four basis points of zero and the single miss in the model-based comparison of Panel C of Table 6, recorded rather than set aside. At ten years, the one public long-horizon source, the Consensus panel underlying Wright (2011), ending April 2009, sits on the same side as the recovered measure pre-2008 in five of eight pairs with the euro cell clean, the Japanese miss reflecting long-horizon survey rates held far above realised JGB yields, the stickiness that motivates estimating no expectations model in the first place. The surveys also separate the construction's two components at the revisions margin: survey expectation revisions load on the expectations component with coefficients rising to 0.87 at one year, close to one for one, and on the recovered differential at zero, so independent expectation revisions move the component built to carry expectations and not the recovered premium, the direct answer to the concern that the differential is mismeasured expectations (Internet Appendix Table IA.6). The same margin validates the released expectations component directly: its revisions track the independent survey revisions at correlations rising with the revision horizon, from 0.36 to 0.60 on the euro leg and from 0.63 to 0.93 on the sterling one-year leg between three and twelve months, so the neutral-rate differential co-moves with genuine survey expectations on the margin that no construction choice can manufacture. Panels A and B of Table 6 report both exercises, and Internet Appendix Section H.3 reports every cell, the survey-leg constructions, the long-horizon exercise, the revisions-margin exercise, and the forecast-accuracy diagnostics.

The reversal against conventional term-structure models. Surveys reach one and two years and the headline differential is at ten, so the second external check comes from the methods this paper sets out to avoid: a Gaussian affine term-structure model estimated country by country in the manner of Joslin et al. (2011) with the small-sample bias correction

[9]The comparison is confined to genuine surveys of forecasters. No practical alternative exists. Broader country coverage is available in licensed forecaster panels, but their long-horizon collection reaches only one and two years ahead and the longer-horizon series began only recently, so no survey anywhere supplies the ten-year expected-rate path that a direct test of the headline differential would require. The public evidence used here is therefore the frontier the data permit rather than a convenient subset of something better. The headline horizon is therefore validated against a conventional yields-only term-structure model instead, in Panel C.

of Bauer et al. (2012), sharing no data construction and no identifying restriction with the measurement here, and a dynamic Nelson-Siegel factor model in the manner of Diebold and Li (2006), outside the affine class, with vector-autoregressive factor dynamics and no no-arbitrage restriction. In Panel C of Table 6 the affine model places five of the six pairs on the same side of zero as the recovered measure in the post-2009 window, the euro at -30 against -52 basis points and the yen at -16 against -75, sterling the miss and the pair whose recovered mean on the comparison dates lies within four basis points of zero; the factor model places all six on the same side on the identical dates, sterling included at -5.3 against -3.4, and all eight when Sweden and Norway are added on their full samples. The agreement extends beyond signs to the ordering of the cross-section: the Spearman rank correlation between the recovered post-2009 means and the model means across the six pairs is 0.94 for the factor model and 0.49 for the affine model, so the independent methods rank the post-crisis differentials much as the recovered measure does. Monthly correlations between either model and the recovered series are low, between -0.55 and $+0.60$, and I read the exercise accordingly: methods built on entirely different principles do not reproduce the high-frequency variation but independently reproduce the sign of the economically important post-crisis configuration, corroboration of the configuration rather than of the series.

The counterexample, and why it matters. The reversal is not a mechanical United-States-against-everyone result, and the most direct evidence for that is that one pair goes the other way. The New Zealand differential moves $+33.7$ basis points between the same two windows while the other seven move negative, and it is the only partner whose quoted basis and swap-spread differential both rise. A construction that manufactured a dollar-side reversal by arithmetic could not produce that. New Zealand is also the pair the sensitivity frontier of Table 5 cannot pin down, its interval widening to $[-73, +141]$ at full capitalisation because its own short-end node moves further than any other currency's. I therefore treat it as a genuine counterexample rather than as noise: what distinguishes the pairs that reverse from the one that does not is a question about sovereign supply, domestic pension and insurance demand for long duration, and the size of the local swap market relative to the bond market, and it is one the measurement makes askable rather than one it answers here.

Heterogeneity across the reserve sovereigns. The pairs are not homogeneous reserve assets: Internet Appendix Section E.1 records the pair-specific configurations, Bund scarcity, collateral rules, and the redenomination tail on the German side, yield-curve control and the compressed free float on the Japanese, the yen exhibits essentially invariant to excluding the yield-curve-control window. The honest description of the government-side series is relative

long-horizon sovereign duration pricing, the sovereign leg carrying each market's own scarcity, collateral, and convenience content alongside duration compensation, while the swap-side series, with no sovereign leg, carries the identification.

4.3 The Recovered Differentials Behave as Term Premia

That the recovered differentials are term premia rather than an accounting residual is established by two families of evidence, developed in this section and the next: they are priced by government debt duration supply with the signs preferred-habitat theory predicts, and they reprice within days of identified policy surprises. This section presents the first, the supply evidence, after fixing the currency-level attribution that the differentials permit.

Each recovered differential is the difference of two currency-level term premia, identified up to a component common to all currencies, the global term-premium factor of the international term-structure literature (Dahlquist and Hasseltoft, 2013; Jotikasthira et al., 2015), to which the differentials are invariant by construction. Decomposing the monthly variance of each dollar-pair differential into its two sides attributes on average 37 percent to the dollar's own term premium and 63 percent to the partner's, the dollar share having been 17 percent before 2008; because this covariance-splitting attribution is descriptive rather than structural, the statements that carry evidential weight are made on quotes rather than shares, and Internet Appendix Section E.3 reports the pair-by-pair shares by subsample with the full node decomposition.

The partner attribution is verifiable in quotes that contain no dollar leg at all, and the verification must run on raw quoted prices, since the recovered differentials satisfy additivity by construction. For any two partner currencies, the change in the difference of their raw dollar-pair bases, which carries dollar legs, loads on the change in the corresponding raw euro-cross difference, which carries none, with a pooled slope of 0.90 and an R^2 of 0.88 across twenty-one partner pairs and 5,697 monthly observations, and per-pair slopes between 0.80 and 0.96; the shortfall from one is what classical measurement error in the euro-cross quotes produces by attenuation, so 0.90 reads as a lower bound on the shared content. Figure 4 makes the verification visible pair by pair, overlaying each raw quoted euro cross on its dollar-implied counterpart from the two dollar pairs: the series coincide at level correlations of 0.993 to 0.999 with mean absolute differences of 0.5 to 2.4 basis points, and the pooled monthly changes lie along the forty-five-degree line.

A dollar-funding account of the differentials lives on the dollar leg by definition and cannot generate this pattern. Even a reading prepared to relabel the entire dollar-side component

as funding claims at most the 37 percent that the dollar side contributes; the remaining 63 percent is reproduced nearly one for one in quotes with no dollar leg, any funding component common to all currencies is absent from the differentials mechanically along with the rest of the global factor, and the funding content within the dollar-side share is itself the subject of the horizon test and the sensitivity-frontier bounds of Section 4.3.4. The panel regressions accordingly use pair fixed effects, absorbing time-invariant cross-sectional heterogeneity, and pooled slopes estimating the average within-pair response to each driver.

FIGURE 4 ABOUT HERE

4.3.1 First: Duration Supply Prices the Differentials

Under the preferred-habitat framework (Vayanos and Vila, 2021; Greenwood and Vayanos, 2014), the quantity that matters for term premium pricing is privately held duration: larger stocks of privately held government debt and longer weighted-average maturities of the privately held free float both elevate the compensation marginal holders demand for bearing term risk. Motivated by this framework, I enter U.S. and partner-country supply conditions separately in a panel specification of the ten-year cross-country government bond term premium differential $\widehat{ctp}_{10}^{gov,j,USD}$ on log-transformed privately held debt-to-GDP and the par-weighted average maturity of the privately held free float, on the seven dollar pairs with constructed free-float series over January 2002 to June 2026, a panel of 2,043 pair-months over 294 months in the main-supply specification of Table 1; the Swiss pair is excluded throughout, and the Swedish pair enters the gross-debt robustness, which requires no free-float series.

The privately held free-float supply $\widetilde{d}_{k,t} \equiv d_{k,t} - c_{k,t}$ is gross sovereign debt-to-GDP net of central bank holdings, and $\widetilde{w}_{k,t}$ is the par-weighted maturity of that free float; netting central bank holdings recovers the supply the marginal yield-sensitive holder faces, since those holdings do not enter the demand-supply equilibrium that determines the term premium (construction in Internet Appendix Section D, gross-debt robustness in Section F.4.1). A three-currency preferred-habitat model rationalising the two-coefficient structure is in Internet Appendix Section G.1. Frontier and mean-group estimators preserve the headline signs, while a twelve-month first-difference specification returns a coefficient near zero, an informative null consistent with a low-frequency relation between levels of supply and levels of the premium rather than a flow relation (Internet Appendix Section G.2).

Under my sign convention, the preferred-habitat prediction is that the coefficient on $\log(\widetilde{d})_{\text{US}}$ is negative (rising privately held U.S. supply raises $tp^{gov,USD}$ and compresses the bilateral differential), the coefficient on $\log(\widetilde{d})_j$ is positive (rising privately held partner supply raises

$tp^{gov,j}$ and widens the bilateral), and the same ordering holds for $\widetilde{\text{WAM}}_{\text{US}}$ and $\widetilde{\text{WAM}}_j$ through the duration channel. Pair fixed effects are included in the panel specifications. Throughout the paper's panels I report Driscoll and Kraay (1998) standard errors with a Bartlett kernel and Schwert-rule lag $L = \lfloor 12(T/100)^{1/4} \rfloor$, which account jointly for cross-sectional correlation across pairs and serial correlation within each pair; table captions state any departure from this convention. Table 1 reports the results.

TABLE 1 ABOUT HERE

Column (1) reports the main preferred-habitat panel regression with currency-pair fixed effects and the four free-float supply variables alone. All four coefficients carry the predicted sign. The two quantity coefficients, which carry the more specific prediction, are precisely estimated: the U.S. log free-float coefficient is -1.051 ($t = -10.72$) and the partner-country coefficient is $+0.521$ ($t = +6.87$), both significant at the one-percent level. The two maturity coefficients are not distinguishable from zero, at -0.060 ($t = -1.26$) and $+0.023$ ($t = +0.99$). They also lack the monotone scaling across tenors that the quantity coefficients display, as Internet Appendix Section F.4.2 records. I therefore read the supply result off the quantity coefficients, and treat the duration margin as directionally consistent rather than independently informative. The pair fixed effects absorb between-pair differences in chronic free-float levels, so that the within-pair time-series variation is what identifies the preferred-habitat mechanism across the full cross-section.

Columns (2) through (5) progressively add the country-specific quantitative-easing indicators, the intermediary capital ratio and its interaction with U.S. asset purchases, and the macro-financial controls; the two quantity coefficients retain their predicted signs and significance throughout, the maturity coefficients keep the predicted sign in every column but one and reach conventional significance only on the U.S. side, and the quantitative-easing indicators enter with contrasting signs across the two legs. The supply loading itself is state-dependent in the direction a preferred-habitat account with intermediary absorbers predicts, attenuated when intermediary capital is abundant and steepest when it is scarce, at a t-statistic of 4.2 on the interaction, with the loading unchanged in sign and significance throughout; the quantitative-easing interaction carries the opposite state dependence and is reported as returned. The asymmetry between the two quantity coefficients, the U.S. coefficient roughly twice the partner coefficient in absolute magnitude, is treated as descriptive, consistent with more inelastic U.S. issuance over the sample (Georgievska et al., 2025) but equally open to measurement and variance accounts. Internet Appendix Sections F.1, G.2 and F.2 discuss each column, the interaction battery, and the elasticity interpretation in turn.

Because the U.S. regressors are common across pairs, the U.S. supply coefficient could in principle reflect a common trend rather than within-pair identification, and three pieces of evidence weigh against that reading. The partner-country supply variables enter as pair-specific time series and load with the predicted positive signs, so the supply mechanism operates symmetrically across the two sides of the bilateral, which a common trend would not produce; the coefficient survives purging U.S. supply of the global supply factor essentially unchanged at -1.05 ($t = -4.7$, Internet Appendix Table IA.18); and the loading localises predominantly in the swap-spread component of the bridge, the incidence the preferred-habitat reading predicts. The pre-2020 to post-2020 sign flip that these panels display is resolved on the extended sample as pandemic-window state dependence, the relationship suspended over the pandemic window, March 2020 to December 2021, and holding outside it at a slope of -1.13 (Section 4.3.3, Internet Appendix Table IA.19; the full three-part identification argument is in Internet Appendix Section F.4.3).

4.3.2 Channel Decomposition Through the Bridge Identity

By the within-currency identity $\widehat{ctp}_N^{gov,j,USD} = \widehat{ctp}_N^{swap,j,USD} + \Delta SS_N^{USD,j}$, every coefficient on the cross-country government bond term premium decomposes additively into a swap-term-premium component and a swap-spread-differential component, and the economic content lies in which component carries each loading: the supply and asset-purchase loadings have their direct incidence in the government market and are carried predominantly by the swap-spread component, while the intermediary-capital interaction and the macro-financial factors load mainly on the swap component, the pattern the equilibrium framework predicts. Internet Appendix Table IA.8 re-estimates the five-column specification with each component in turn as the dependent variable on an identical regressor set and sample, so that $\widehat{\beta}^{\text{gov}} = \widehat{\beta}^{x} + \widehat{\beta}^{\Delta SS}$ to machine precision, Internet Appendix Table IA.17 extracts the column (5) coefficients onto a single page, and Internet Appendix Section F.3 discusses each channel in turn.

Internet Appendix Section F.4 reports complementary exercises on the supply panel: pre-pandemic subsample stability and the post-pandemic break resolved as pandemic-window state dependence, gross debt in place of the free float, the two- and five-year tenors where the U.S. supply coefficient scales monotonically with maturity, expected inflation in place of realised, and the He et al. (2017) capital ratio against the Adrian et al. (2014) leverage factor.

4.3.3 Extraction Check, Reserve Demand, and State Dependence

Taking the prediction to the data requires separating two official holders that sit on opposite sides of the market. The domestic central bank operates on the supply side: when it purchases its own government's bonds under an asset-purchase programme, it withdraws them from the float that price-sensitive private investors must absorb, so the supply measure relevant for pricing is the privately held free float, gross debt net of domestic central bank holdings, $\widetilde{d}_{c,t} \equiv d_{c,t} - c_{c,t}$, consistent with the treatment of central bank absorption in the quantitative easing literature (Wright, 2022; Diez de los Rios, 2024). Foreign official institutions operate on the demand side: foreign central banks and reserve managers hold Treasuries for reserve and policy purposes and are relatively insensitive to yields, as Fang et al. (2026) document for the foreign official investor base, so their holdings act as a shifter of inelastic demand rather than a component of the supply that private investors price. I therefore estimate the supply relationship on the free float while holding the foreign official reserve demand fixed, which reports the supply association with the demand side held fixed, and which has the attraction that the demand control coincides with the reserve-currency channel that is the subject of the paper.

Replacing the quoted basis with the triangle-extracted premium changes essentially nothing: on the triangle-identified panel of 1,888 pair-months the supply-only coefficient is -1.066 ($t = -9.89$) against -1.051 on the quoted construction, the fully conditioned coefficient -1.216 ($t = -5.78$) against -1.134, and the channel allocation is preserved, with the component split at -0.187 and -0.879 against -0.194 and -0.857 (Internet Appendix Tables IA.1, IA.9 and IA.20; the comparison-country supply term, the fiscal-balance flows, and the free-float maturity construction are in Internet Appendix Sections G.2 and D.6).

The extracted result is insensitive to the recovered residuals: rescaling every residual by up to a factor of ten leaves the supply coefficient within 0.07 of its value, and the quoted construction is invariant to them by definition (Internet Appendix Section F.4).

Holding foreign official reserve demand fixed tightens the interpretation without overturning it. With the foreign official Treasury stock included as a demand control the supply-only coefficient is -0.930 ($t = -4.63$), and the official-demand control itself enters with the sign the reserve-currency channel implies. Under the full conditioning set the supply coefficient attenuates to -0.438 ($t = -2.05$), so part of the conditioned loading is shared with official demand. An effective float netting out both price-insensitive holders leaves the coefficient at -1.052 ($t = -10.71$). The demand side also carries a flow margin: on the valuation-adjusted foreign official Treasury flows of Bertaut and Judson (2014), whose final release ends

in December 2020, the twelve-month official flow relative to GDP enters the panel at +3.309 ($t = 2.30$) with the sign inelastic official absorption implies, a descriptive result whose design, pre-registered sign, and failed stock margins are reported in Internet Appendix Section F.4.6. I therefore read the supply and demand coefficients jointly (Internet Appendix Section F.4). The long-horizon intermediation alternative is confronted on observable proxies as well: with the BIS swap-market notional by currency and the BIS international debt securities outstanding by currency of denomination added to the panel, the free-float loading survives essentially unchanged, which was the pre-specified prediction, and the battery is in Internet Appendix Section G.2.

The result is robust along the dimensions a single-country supply series would leave open, and is not a repackaging of short-end dollar funding: orthogonalising U.S. supply against the global supply factor and a common trend leaves the coefficient economically unchanged, dropping each pair in turn leaves it in a narrow band, and the loading on the short-end dollar funding factor declines toward zero as the tenor lengthens and, at the ten-year tenor, sits far below the excess-sensitivity bound the measured funding persistence permits, the discrimination resting on that exceedance comparison rather than on the decline itself, which a decaying funding pass-through also produces (Internet Appendix Table IA.18).

Duration supply rather than debt supply. Preferred habitat names the duration supplied to private investors, not the quantity of debt, and the paper's maturity coefficients are not distinguishable from zero, so the natural test forms the theory's quantity directly. Multiplying each free float by the modified duration of a par bond at its own par-weighted average maturity and its own yield gives a duration-weighted float, which in logs is the existing log free float plus log modified duration and so needs no new data; implied durations average 4.8 years for the United States, 6.2 for Germany and 11.4 for the United Kingdom. The measure is an approximation: it captures the level and the time variation of supplied duration, not the convexity of the true maturity distribution, which would need security-level holdings. The informative test is not significance, which the debt measure already delivers, but whether the loading scales with tenor, because a duration mechanism must price the twenty-year differential more than the two-year and a quantity-of-safe-debt mechanism has no reason to. It does: the duration-weighted United States loading runs -0.64, -0.87, -0.89 and -1.18 at two, five, ten and twenty years, strictly steepening with maturity. The debt measure delivers the same profile, -0.70 to -1.29, so the tenor scaling is a property of the supply relation itself rather than of the duration weighting.

Composition news at a held financing requirement. The most direct supply experiment the public record permits holds the total fixed and moves the maturity mix. From every United States auction announcement since 2002, two thousand and one coupon auctions on five hundred and eighty-one announcement dates, the news in each announcement is the change in offering size against the previous auction of the same tenor and the same new-issue-or-reopening status, weighted by the modified duration of the announced tenor; announcements aggregate into long-end news, ten years and beyond, short-end news, and the unweighted total, and the regression holds the total fixed so the long-end coefficient is the composition margin the theory names. Auction sizes are partly telegraphed at the preceding refunding, so measured news understates true surprise and attenuates the estimate toward zero, which makes the test conservative; no instrument is involved. At the headline ten-year tenor the differential falls 0.33 basis points per standard deviation of long-end news ($t = -2.15$) within the $[-1, +2]$ window, while short-end news, the total, and the one-year funding placebo all sit at zero: more long-duration issuance at a held financing requirement lowers the partner-minus-United-States differential, which is the sign the supply panel carries, arriving within days of the announcement. The response splits across the bridge exactly as the panel does: the swap-spread differential carries -0.274 ($t = -1.86$) and the basis -0.055 ($t = -1.34$), summing to the composite by the identity, and the basis share of about a sixth sits beside the roughly one-fifth channel split the levels panel reports, an agreement across two very different frequencies. Neither component clears conventional significance alone; the composite does. The estimate is descriptive support for the duration mechanism on the announcement margin, read alongside the panel rather than in place of it. The auction window also prices the differential from the demand side: on ten-year auctions across the same advanced economies, demand surprises revealed by the bid-to-cover ratio lower domestic long-maturity yields with the maturity-localised signature preferred-habitat theory predicts and move the covered sovereign differential within days (Zou, 2026), quantity-side event evidence complementary to the composition margin measured here.

Because the U.S. supply variable is common across pairs within a month, the effective time-series information is the 294 months rather than the 2,043 pair-months, and because it is also highly persistent, the ordinary panel standard error is not the relevant discipline. Two-way clustering by month and pair leaves the coefficient at -1.051 ($t = -4.39$); collapsing the panel to a single time series gives -0.61 (Newey-West $t = -5.4$); and the test that speaks directly to persistence, a low-frequency cosine regression in the manner of Müller and Watson (2018), retains the U.S. coefficient at every bandwidth tried, running from -0.93 to -0.97 with t-statistics of -3.6 to -8.3, rejecting even when only periods of eight years and longer are retained (Internet Appendix Table IA.21).

Three qualifications belong beside it, reported rather than left to be discovered. The partner-country coefficient does not survive the same low-frequency test, its t-statistics running between 0.8 and 1.4, so the bilateral symmetry of the levels panel is a panel result and not a low-frequency one. An Ibragimov-Muller group-t on eight non-overlapping time blocks does not reject, a design that splits a near-unit-root regressor into windows over which it barely moves, while grouping by currency pair rejects at -2.86, so the negative loading is common across currencies rather than carried by one or two. And the twelve-month first-difference specification returns $+0.09$ ($t = 0.24$): the relation is a low-frequency stock relation, as a preferred-habitat mechanism implies, and not a monthly flow elasticity. It is insensitive to excluding the yen during yield-curve control and the benchmark-cessation windows and to perturbing every quoted basis within its half-spread, and I rest the interpretation on the state-dependence, demand-control, and event evidence rather than an estimated equilibrium relation; Internet Appendix Section G.2 and Table IA.21 report the complete battery.

The supply relationship is state dependent on the extended sample: it is suspended during the pandemic window of 2020 and 2021, the slope outside the window at -1.134 with a window interaction of $+0.921$ ($t = +7.98$), and racing the candidate accounts retains the pandemic window while the deficit-pace and quantitative-easing alternatives do not once the window is controlled, the deficit-pace interaction being significant only on the sample through 2021. Part of the state dependence operates through the maturity of the additional supply, the pandemic shock having been financed at the short end so that the long-end free float rose by less than the headline deficit suggests (Georgievska et al., 2025), and I read the maturity of issuance and the pace of the deficit as complementary components of a single fiscal-surge channel (Greenwood et al., 2015): the mechanism is not broken but conditional, holding across two decades outside the exceptional issuance environment of the pandemic window (Internet Appendix Table IA.19).

Refunding-announcement surprises. From the complete record of Treasury marketable-borrowing announcements I construct ninety-nine borrowing-estimate surprises over 2000 to 2026, each the estimate for a quarter announced at its start minus the estimate for the same quarter three months earlier; announcement-window responses of the quoted-construction differential carry the predicted negative sign for the core pairs over the following week, and excluding the March 2020 announcement the euro-pair response is significant at -0.38 basis points per one hundred billion dollars of surprise borrowing ($t = -2.2$), the yen and Norwegian responses of the same sign at -0.57 and -0.48, consistent with the stock-level relationship above.

I attempted two routes to an instrument for the supply channel and report both outcomes. Treasury auction surprises and fiscal-news shocks were examined in two forms, each as a reduced-form event study at announcement frequency, and neither yielded a usable design. Interest-rate surprises were not attempted. No instrumental-variables estimation was run on the supply channel at any point, so the panel is descriptive throughout and there are no first stages, weak or otherwise, to report.

Three further checks close the battery and are reported in Internet Appendix Section F.4: a reverse-timing placebo, in which the current stock retains its coefficient (-0.65, $t = -3.0$) while its twelve-month-ahead value is insignificant, against the pattern reverse causality would produce; a moving-block bootstrap placing the ninety-five percent interval at $[-0.85, -0.08]$; and the confirmation that the bilateral construction carries information beyond the familiar within-U.S. swap-spread relation, since partner-country swap spreads load on their own free floats weakly or not at all while the triangulated differential loads significantly.

4.3.4 Second: Far-Forward Behaviour and Policy Surprises Discriminate a Term Premium

What prices can and cannot decide. Three questions can be asked of the recovered differential, and prices settle two: premium rather than expectations is decided without an expectations model, by construction, and currency-level rather than pair-level is decided by the triangular containment within quoted spreads. What mixture the premium bundles, risk compensation, convenience, or intermediary shadow cost, is undecidable from prices, here as for every model-based estimate, whose term premia bundle the same Treasury convenience without the literature requiring that bundle resolved (Adrian et al., 2013; Kim and Wright, 2005; Krishnamurthy and Vissing-Jorgensen, 2012); nor is demand pressure an alternative to a term premium, being in the preferred-habitat framework the term premium's source (Vayanos and Vila, 2021), and when the marginal holder is a global currency-hedging investor (He et al., 2017; Du et al., 2018b; Liao, 2020), persistent priced compensation for currency-attached duration risk is a constituent of that currency's term premium rather than a wedge outside it. The dollar-side attribution bounds even the surviving alternative reading to at most the one third of variation the dollar's own movement contributes (Internet Appendix Section J).

A battery of tests dispenses with the allocation altogether, exploiting the property that funding components are short-concentrated or tenor-flat while a term premium accumulates with duration: the slope constructions computed from raw quotes are immune to any tenor-flat

currency-level contamination, while the five-year-five-year forward construction passes a tenor-flat component through one for one and is disciplined instead by the loading and volatility bounds below. The tenor slope of the raw benchmark construction exhibits the central reversal, the euro moving from +21 to −19 basis points and the yen from +27 to −40; supply prices that slope at −0.75 ($t = -6.6$), the loading rising monotonically from −0.67 at two years to −1.30 at twenty, the signature of priced duration risk; and identified policy shocks discriminate in the same direction, the two genuine Bank of Japan yield-curve-control surprises moving the daily yen benchmark by +12.5 and +9.7 basis points, 3.2 and 2.5 daily standard deviations, with five of seven asset-purchase announcements carrying the predicted sign, headlined by the March 2009 gilt launch at eight (Internet Appendix Section H.4).

Funding is separated quantitatively rather than categorically: in the March 2020 shock the ten-year bases moved at 4 and 32 percent of the short end for the euro and yen, the government differential carries no incremental forecast content for future funding, and the pass-through of three-month funding changes decays monotonically along the curve for every pair, with cross-pair medians of 0.48 at six months, 0.26 at one year, 0.08 at five years, and 0.05 at ten (Internet Appendix Table IA.22). Internet Appendix Section H.4 reports the complete battery. For the cross-country government bond term premium, the bridge delivers a stronger, exact statement.

The measurement without the dollar. One structural property removes the dollar from most of the measurement: dropping all nine dollar pairs leaves a euro-hub tree in which every differential among the non-dollar currencies is read directly from euro-cross quotes, and those dollar-free readings coincide with the full-network recovery at correlations above 0.999 and differences below 0.2 basis points, so no dollar-specific wedge can shape the recovered structure among the other currencies; the dollar node's own level is the sole place a dollar wedge could reside, exactly where the bounds and the excess-sensitivity restriction operate. The one node-space alternative, a forward dollar exchange-rate premium residing in the basis itself, is bounded and tested by four pieces of evidence, headlined, as Figure 4 already made visible in raw quotes, by the December 2022 surprise appearing undiminished in dollar-free form: +15.0 basis points in the network-implied yen-euro benchmark differential, in which the dollar node cancels, against +12.5 for the dollar pair, and +8.5 in the directly quoted euro-yen basis, a swap-side quote with no swap-spread leg, against the +8.2 swap-side component of the dollar-pair response, its like-for-like comparator (Internet Appendix Sections E.1 and H.4, the constructions collected in Table IA.23). The quoted crosses are moreover independently discovered prices on their dynamics as well as their levels: in a bivariate error-correction bat-

tery, each cross with a stored daily history and its dollar-implied counterpart adjust toward one another two-sidedly, the implied series correcting toward the quote as well as the quote toward the implied series, with the quote's Gonzalo-Granger permanent-component share between 0.37 and 0.76 across the six crosses where the adjustment is measured with precision (Internet Appendix Table IA.24). Internet Appendix Table IA.25 summarises the disposition of every alternative reading, algebraic cancellation, rejection by a testable restriction, quantitative bounding, or direct control.

The basis package's own duration structure supplies a further argument: the cross-country government bond term premium scales with the leg-duration mismatch ($t = 4.5$) and the swap term premium with the package's average duration ($t = 7.4$), while a funding component scales with neither (Internet Appendix Section H.5).

Proposition 4, stated formally in Internet Appendix Section H.4, bounds what any stationary funding process can deliver: if the currency-level component at tenor τ equalled the expected average of a stationary funding process with monthly autocorrelation ρ plus a constant risk premium, the five-year-five-year forward component implied between tenors, $2\,x_{10Y} - x_{5Y}$, could load on the current funding state by at most $\frac{1}{60}\sum_{k=60}^{119}\rho^{k}$ and would be strictly less volatile than the funding state, for any $\rho < 1$. Its parameter-free corollary, obtained in the limit of a permanent funding state, is tighter still and requires no persistence estimate at all: the ten-year component cannot exceed the short-horizon funding cost itself in absolute size, since that ceiling is attained only when the funding state never mean-reverts and any mean reversion lowers it.

The short-horizon funding cost is directly observable. Reconstructing the three-month cross-currency basis from spot and forward quotes against overnight index swap rates, in the manner of the covered-interest-parity literature, it averages -43.2 basis points for the yen and -23.1 for the euro over the post-2008 sample, against observed ten-year bases of -54.6 and -20.1. For the yen the ceiling is therefore violated outright, since no funding process, however persistent, capitalises a 43 basis point short-horizon cost into a 55 basis point ten-year one; the same holds for the franc, the Australian dollar, the New Zealand dollar, and the krona. At the measured persistence of the three-month basis the ceiling binds far more tightly, the implied ten-year footprint being -1.42 basis points for the yen and -0.88 for the euro, between 0.4 and 7.1 per cent of the observed ten-year levels. Rationalising the euro level as capitalised funding would require a monthly persistence of 0.998, a half-life near twenty-four years (Internet Appendix Section H.4).

The far-forward component also carries the paper's economics: U.S. supply prices the government five-year-five-year forward at -0.76 with a t-statistic of -6.4 on the eight-pair panel

(Internet Appendix Section H.4), which persistent government debt duration supply can do and transitory funding cannot, and the December 2022 policy surprise moved the yen forward component by +8.5 basis points in two days (Internet Appendix Table IA.23), repricing compensation five to ten years ahead instantly. Auction evidence from outside the swap market points the same way. Zou (2026) finds that demand surprises revealed by bid-to-cover ratios in ten-year government bond auctions, across eight of the economies studied here, move covered sovereign yield differentials of the Du et al. (2018a) form nearly uniformly across maturities, by 0.75 to 1.3 basis points per unit of surprise from the one-year to the ten-year tenor, while the yield response itself rises monotonically with maturity: liquidity demand enters the covered differential as a level across tenors, consistent with the ceiling's premise that funding and liquidity content is anchored at the short end, while maturity-increasing variation of the kind the supply loadings display carries term-structure content.

The same construction can be read at the short end, where by its own logic it cannot be measuring term risk, and what it recovers there is the short-horizon counterpart of the convenience yield the money-market literature measures: the swap-spread component of the three-month differential, which contains no basis at all, stood at +51.7 basis points for the euro pair and +23.9 for the yen pair before 2008 and at −11.6 and −5.1 after 2013, so the richness of the United States bill against its own curve fell relative to both partners across the crisis, complementing the decline in the short-horizon Treasury premium that Du et al. (2018b) document. The full short-end row, the benchmark-route disclosure, and the wider forward half-spreads that make it less precisely measured than the long-tenor quotes are in Internet Appendix Section H.4.

The ceiling rejects the reading of the long-tenor currency-level component as a capitalised funding cost, at any persistence and without recourse to an estimated one, and it does so without leaning on the allocation restriction at all, since it takes the observed ten-year component as given, whatever mixture of term premium and dollar-wide wedge that component contains, and caps the total; reallocating between the two therefore cannot relax it; the only surviving version of the funding hypothesis is a funding risk premium, which is priced compensation for long-horizon risk rather than a funding cost. The allocation restriction is therefore needed only for the level attribution of the currency component, whose failure the thresholds cap, while the reversal, the supply pricing, the maturity profile, the policy-event responses, and the forecast asymmetry stand without it.

Table 8 about here

Identified policy-surprise factors. The systematic policy evidence comes first, because it does not depend on which windows are selected. Identified announcement-surprise factors on both legs deliver it as regressions over every announcement in the sample, pre-specified with their signs in the replication materials. On the United States leg, using the Swanson (2021) decomposition of Federal Reserve announcement surprises into target, forward-guidance, and asset-purchase factors, in the tradition of Gürkaynak et al. (2005), the ten-year benchmark differentials load on the asset-purchase factor at +0.68 basis points per standard deviation across the 163 in-sample announcements ($t = 2.1$; the euro pair at +0.62, $t = 2.3$, the yen at +1.18, $t = 3.3$), while the target factor is largest and most precisely estimated at the two-year tenor (+0.98, $t = 2.2$), its longer-tenor profile non-monotone with a smaller, marginal ten-year loading, and the forward-guidance factor enters marginally at the two- and five-year tenors and moves nothing beyond five years, as expectations content that differences out of a premium differential should not. On the euro leg, the four factors of Altavilla et al. (2019) deliver the mirror image: the euro-pair differential falls on European Central Bank purchase surprises with a magnitude rising monotonically in tenor, −0.5, −0.7, −1.0, and −1.4 basis points per standard deviation at two, five, ten, and twenty years ($t = -3.1$ at twenty), while the euro funding placebo moves −0.1 ($t = -0.3$). Duration news therefore reprices the differentials on whichever leg it lands, graded monotonically by duration on the euro leg and concentrated at the five- to twenty-year tenors on the United States leg, while short-rate news is largest at the short end or enters flat; that is the announcement-frequency behaviour of a term-premium differential, not of a funding spread. Three honest entries accompany it in Internet Appendix Section H.6: the pre-specified prediction that the asset-purchase response exceeds the target response at ten years fails on point estimates; the Federal Reserve-side funding placebo is distinguishable from zero at roughly half the size of the differential response; and the European Central Bank target factor enters significantly with a roughly flat tenor profile, each reported as returned.

The six headline episodes. Table 8 consolidates the identified policy windows into one pre-specified exhibit, each row a genuine surprise with its direction stated ex ante: the two Bank of Japan yield-curve-control surprises, the Bank of England gilt-purchase launch, the European Central Bank's expanded asset-purchase announcement, and the two 2008 to 2009 Federal Reserve announcements read cross-pair. All six predicted signs are realised, and the responses run from two to eight daily standard deviations, at the ninety-fifth to one-hundredth percentile of all same-length windows of the same series, while the pair's own one-year basis moves within ordinary daily variation. These episodes are economic illustration rather than the statistical case, which the factor regressions above carry, and the reason is that any selected

universe invites the question of why those windows. The wider screen is therefore disclosed rather than left implicit: two of the seven major asset-purchase windows miss, giving seven of nine, and pooling every disjoint pre-specified sign universe in the paper gives fifteen of twenty-seven, which is indistinguishable from chance. I attach no formal probability to any sign count and rest the reading on magnitudes and placebo percentiles. The full universe, the anticipated-announcement placebos, and the leverage actions are in Internet Appendix Section H.4.

An identified duration-demand event. The best-identified duration-demand shock in the sample is the Bank of England's announcement of temporary, targeted purchases of long-dated gilts on 28 September 2022, in the middle of the liability-driven-investment episode, which carries almost no policy-path news because the Bank simultaneously signalled further tightening and postponed gilt sales. The sterling ten-year benchmark differential fell 7.6 basis points over the $[-1, +2]$ window, at the 93rd percentile, while the funding placebo moved 0.3; the twenty-year differential fell 32.6 basis points, 6.8 daily standard deviations at the 99.8th percentile, concentrated precisely at the tenors the operation targeted. A response that scales with proximity to the purchased habitat is preferred-habitat pricing read from the demand side, on the one margin where announcement-frequency variation is otherwise scarce. The episode's other window belongs beside it and cuts the other way: over the preceding forced-selling week the ten-year differential moved 22.7 basis points in the opposite direction, that window pricing dollar funding stress in the basis while the gilt-specific movement sat at twenty years, so the measure separates a funding episode from a demand shock rather than confirming every window. Across the full pre-specified demand screen of eight windows three signs are realised, and I rest nothing on that count.

The joint falsification standard. The individual tests above can each be argued with; their conjunction is the identification argument. A currency-level mechanism that is not a term premium would have to generate, simultaneously, a ten-year component larger in absolute size than the three-month funding cost itself, which no capitalised funding process delivers at any persistence; pricing of the far-forward component by government debt duration supply with the preferred-habitat sign on the United States side; a 12.5 basis point repricing of the benchmark differential within two days of an identified policy surprise whose dollar-wide funding benchmark moved 1.2 basis points; and magnitudes that survive the breakeven wedges of Table 7. Expected funding costs fail the first by inequality; balance-sheet mechanisms divide by where they reside, anything priced through the swap rates cancelling identically in the bridge by Corollary 1, transitory short-concentrated costs failing the first and the tenor

battery, and a persistent priced component surviving in the government leg being bounded by the sensitivity frontier and, under the global-marginal-investor reading, a constituent of the premium rather than a wedge outside it; and any mechanism confined to the dollar fails the dollar-free replication of the December 2022 event. The burden has a stated boundary, and it is a limit rather than a victory: what these tests reject is the interpretation of the currency component as a capitalised expectation of short-horizon funding cost, what they leave standing is a persistent, priced, currency-attached component compensating a long-horizon exposure, and calling every such component a term premium by definition would make the tests unfalsifiable, so I do not; it is the supply, maturity, policy-shock and expectations evidence, not the definition, that supports reading the bulk of it as term-premium compensation, and the battery does not adjudicate the microfoundation of the persistent branch (Internet Appendix Section J). Two margins are controlled directly rather than left as gaps: subtracting quoted sovereign credit-default-swap differentials in full leaves the post-2013 configuration intact for every pair with quoted protection, where large they are signed so that subtraction enlarges the configuration rather than explaining it (Internet Appendix Section G.3); and the one clean leverage-regulation event, the July 2013 enhanced supplementary leverage ratio proposal, loads on the swap-spread leg at -6.1 basis points with the ten-year basis inside its ordinary range, as the bridge implies, the remaining actions disclosed at chance in Internet Appendix Section G.4, descriptive support on the allocation margin and no more.

5 Conclusion

The long-tenor cross-currency basis and the swap-spread differential are traded prices, and read jointly across a network of seventeen pairs they deliver the comparison of sovereign term premia that a term-structure model is normally estimated to produce, with no expectations model formed in any country and, under the stated pair-specificity restriction whose failure the sensitivity frontier bounds, with the residual bilateral friction held inside dealers' quoted spreads. What the prices then say is that the relative price of sovereign duration has changed sides. Ten-year U.S. Treasury term premia sat 28 basis points below German and 37 below Japanese term premia before the global financial crisis and have sat 56 and 83 above them since 2013, a pattern shared by seven of the eight pairs with identified histories, firmly for five of them, with the sterling and Australian shifts marginal under persistence-robust inference and the New Zealand pair shifting the other way. The reading that resolves the change is not that dollar safety has become less convenient, since the short-maturity convenience of Treasury bills survives throughout the sample, but that the United States now pays for the

duration its debt asks investors to bear: the differentials load on the privately held free float at -1.05, and they reprice within days of the central-bank announcements that change how much duration the private sector must hold.

I view the recovered cross-country term premium series as a building block for future work on monetary policy transmission, sovereign-supply-and-demand co-movement, intermediary asset pricing in the cross-section of currencies, and global imbalance dynamics.

References

Adrian, Tobias, Richard K Crump, and Emanuel Moench, 2013, Pricing the term structure with linear regressions, *Journal of Financial Economics* 110, 110–138.

Adrian, Tobias, Erkko Etula, and Tyler Muir, 2014, Financial intermediaries and the cross-section of asset returns, *Journal of Finance* 69, 2557–2596.

Altavilla, Carlo, Luca Brugnolini, Refet S. Gürkaynak, Roberto Motto, and Giuseppe Ragusa, 2019, Measuring euro area monetary policy, *Journal of Monetary Economics* 108, 162–179.

Amihud, Yakov, and Clifford M. Hurvich, 2004, Predictive regressions: A reduced-bias estimation method, *Journal of Financial and Quantitative Analysis* 39, 813–841.

Andersen, Leif, Darrell Duffie, and Yang Song, 2019, Funding value adjustments, *The Journal of Finance* 74, 145–192.

Avdjiev, Stefan, Wenxin Du, Catherine Koch, and Hyun Song Shin, 2019, The dollar, bank leverage, and deviations from covered interest parity, *American Economic Review: Insights* 1, 193–208.

Bauer, Michael D., and Glenn D. Rudebusch, 2020, Interest rates under falling stars, *American Economic Review* 110, 1316–1354.

Bauer, Michael D., Glenn D. Rudebusch, and Jing Cynthia Wu, 2012, Correcting estimation bias in dynamic term structure models, *Journal of Business & Economic Statistics* 30, 454–467.

Bertaut, Carol C., and Ruth Judson, 2014, Estimating U.S. cross-border securities positions: New data and new methods, International Finance Discussion Papers 1113, Board of Governors of the Federal Reserve System.

Borio, Claudio, Robert McCauley, Patrick McGuire, and Vladyslav Sushko, 2016, Covered interest parity lost: Understanding the cross-currency basis, *BIS Quarterly Review* 45–64.

Boyarchenko, Nina, Pooja Gupta, Nick Steele, and Jacqueline Yen, 2018, Negative swap spreads, *Federal Reserve Bank of New York Economic Policy Review* 24, 1–14.

Brunnermeier, Markus K., and Lasse Heje Pedersen, 2009, Market liquidity and funding liquidity, *Review of Financial Studies* 22, 2201–2238.

Campbell, John Y., and Robert J. Shiller, 1991, Yield spreads and interest rate movements: A bird's eye view, *Review of Economic Studies* 58, 495–514.

Cenedese, Gino, Pasquale Della Corte, and Tianyu Wang, 2021, Currency mispricing and dealer balance sheets, *The Journal of Finance* 76, 2763–2803.

Chinn, Menzie D., and Guy Meredith, 2004, Monetary policy and long-horizon uncovered interest parity, *IMF Staff Papers* 51, 409–430.

Cochrane, John H., 2005, *Asset Pricing*, revised edition (Princeton University Press, Princeton, NJ).

Correa, Ricardo, Wenxin Du, and Gordon Y. Liao, 2025, U.S. banks and global liquidity, *Journal of Political Economy Macroeconomics* 3, 574–620.

Dahlquist, Magnus, and Henrik Hasseltoft, 2013, International bond risk premia, *Journal of International Economics* 90, 17–32.

Dai, Qiang, and Kenneth J. Singleton, 2000, Specification analysis of affine term structure models, *Journal of Finance* 55, 1943–1978.

Diamond, William, and Peter Van Tassel, 2026, Risk-free rates and convenience yields around the world, *Journal of Finance* Forthcoming.

Diebold, Francis X., and Canlin Li, 2006, Forecasting the term structure of government bond yields, *Journal of Econometrics* 130, 337–364.

Diez de los Rios, Antonio, 2024, Estimating the portfolio-balance effects of the Bank of Canada's Government of Canada bond purchase program, Staff Working Paper 2024-34, Bank of Canada.

Driscoll, John C., and Aart C. Kraay, 1998, Consistent covariance matrix estimation with spatially dependent panel data, *Review of Economics and Statistics* 80, 549–560.

Du, Wenxin, Joanne Im, and Jesse Schreger, 2018a, The U.S. Treasury premium, *Journal of International Economics* 112, 167–181.

Du, Wenxin, Ritt Keerati, and Jesse Schreger, 2026, Decoupling dollar and Treasury privilege, NBER Working Paper No. 35000.

Du, Wenxin, Alexander Tepper, and Adrien Verdelhan, 2018b, Deviations from covered interest rate parity, *Journal of Finance* 73, 915–957.

Duffee, Gregory R., 2011, Forecasting with the term structure: The role of no-arbitrage restrictions, Working Paper 576, The Johns Hopkins University, Department of Economics.

Duffie, Darrell, and Kenneth J. Singleton, 1997, An econometric model of the term structure of interest-rate swap yields, *Journal of Finance* 52, 1287–1321.

Fama, Eugene F., 1984, Forward and spot exchange rates, *Journal of Monetary Economics* 14, 319–338.

Fang, Xiang, Bryan Hardy, and Karen K. Lewis, 2026, Who holds sovereign debt and why it matters, *Review of Financial Studies* forthcoming.

Feldhütter, Peter, and David Lando, 2008, Decomposing swap spreads, *Journal of Financial Economics* 88, 375–405.

Geman, Hélyette, Nicole El Karoui, and Jean-Charles Rochet, 1995, Changes of numéraire, changes of probability measure and option pricing, *Journal of Applied Probability* 32, 443–458.

Georgievska, Ljubica, 2019, The persistent element of the no-arbitrage framework: Discount rates, bases, and opportunity cost of collateral, Working paper, University of California, Los Angeles.

Georgievska, Ljubica, Stefano Pegoraro, and Anthony Saunders, 2025, The US Treasury's biggest short: Duration in the shadows, Working Paper, BI Norwegian Business School and University of Notre Dame.

Greenwood, Robin, Samuel G. Hanson, and Jeremy C. Stein, 2015, A comparative-advantage approach to government debt maturity, *Journal of Finance* 70, 1683–1722.

Greenwood, Robin, and Dimitri Vayanos, 2014, Bond supply and excess bond returns, *Review of Financial Studies* 27, 663–713.

Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson, 2005, Do actions speak louder than words? the response of asset prices to monetary policy actions and statements, *International Journal of Central Banking* 1, 55–93.

Hamilton, James D., and Jing Cynthia Wu, 2014, Testable implications of affine term structure models, *Journal of Econometrics* 178, 231–242.

Hansen, Lars Peter, and Scott F. Richard, 1987, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models, *Econometrica* 55, 587–613.

Hanson, Samuel G., Aytek Malkhozov, and Gyuri Venter, 2022, Demand-supply imbalance risk and long-term swap spreads, *Working Paper* Systemic Risk Centre Discussion Paper No. 118.

Haubrich, Joseph, George Pennacchi, and Peter Ritchken, 2012, Inflation expectations, real rates, and risk premia: Evidence from inflation swaps, *Review of Financial Studies* 25, 1588–1629.

He, Zhiguo, Bryan Kelly, and Asaf Manela, 2017, Intermediary asset pricing: New evidence from many asset classes, *Journal of Financial Economics* 126, 1–35.

He, Zhiguo, and Arvind Krishnamurthy, 2013, Intermediary asset pricing, *American Economic Review* 103, 732–770.

He, Zhiguo, Stefan Nagel, and Zhaogang Song, 2022, Treasury inconvenience yields during the COVID-19 crisis, *Journal of Financial Economics* 143, 57–79.

Hodrick, Robert J., 1992, Dividend yields and expected stock returns: Alternative procedures for inference and measurement, *Review of Financial Studies* 5, 357–386.

Huang, Wenqian, Angelo Ranaldo, Andreas Schrimpf, and Fabricius Somogyi, 2024, Constrained liquidity provision in currency markets, *Journal of Financial Economics* Forthcoming.

Jermann, Urban J., 2020, Negative swap spreads and limited arbitrage, *Review of Financial Studies* 33, 212–238.

Jiang, Zhengyang, Arvind Krishnamurthy, and Hanno Lustig, 2021, Foreign safe asset demand and the dollar exchange rate, *Journal of Finance* 76, 1049–1089.

Jiang, Zhengyang, and Robert J. Richmond, 2025, Reserve asset competition and the global fiscal cycle, Working Paper.

Jiang, Zhengyang, Robert J. Richmond, and Tony Zhang, 2025, Understanding the strength of the dollar, *Journal of Financial Economics* 167, 104024.

Joslin, Scott, Kenneth J. Singleton, and Haoxiang Zhu, 2011, A new perspective on Gaussian dynamic term structure models, *Review of Financial Studies* 24, 926–970.

Jotikasthira, Chotibhak, Anh Le, and Christian Lundblad, 2015, Why do term structures in different currencies co-move?, *Journal of Financial Economics* 115, 58–83.

Kim, Don H., and Jonathan H. Wright, 2005, An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates, *Finance and Economics Discussion Series* .

Klingler, Sven, and Suresh Sundaresan, 2019, An explanation of negative swap spreads: Demand for duration from underfunded pension plans, *Journal of Finance* 74, 675–710.

Kloks, Pēteris, Edouard Mattille, and Angelo Ranaldo, 2025, Hunting for dollars, Research paper series, Swiss Finance Institute.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen, 2012, The aggregate demand for Treasury debt, *Journal of Political Economy* 120, 233–267.

Liao, Gordon Y., 2020, Credit migration and covered interest rate parity, *Journal of Financial Economics* 138, 504–525.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan, 2011, Common risk factors in currency markets, *Review of Financial Studies* 24, 3731–3777.

Lustig, Hanno, Andreas Stathopoulos, and Adrien Verdelhan, 2019, The term structure of currency carry trade risk premia, *American Economic Review* 109, 4142–4177.

Lustig, Hanno, and Adrien Verdelhan, 2007, The cross section of foreign currency risk premia and consumption growth risk, *American Economic Review* 97, 89–117.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2012, Carry trades and global foreign exchange volatility, *Journal of Finance* 67, 681–718.

Müller, Ulrich K., and Mark W. Watson, 2018, Long-run covariability, *Econometrica* 86, 775–804.

Rime, Dagfinn, Andreas Schrimpf, and Olav Syrstad, 2022, Covered interest parity arbitrage, *Review of Financial Studies* 35, 5185–5227.

Somogyi, Fabricius, 2025, Dollar dominance in FX trading, Working Paper.

Swanson, Eric T., 2021, Measuring the effects of federal reserve forward guidance and asset purchases on financial markets, *Journal of Monetary Economics* 118, 32–53.

Vayanos, Dimitri, and Jean-Luc Vila, 2021, A preferred-habitat model of the term structure of interest rates, *Econometrica* 89, 77–112.

Wong, Alfred, and Jiayue Zhang, 2018, Breakdown of covered interest parity: Mystery or myth?, *BIS Papers* .

Wright, Jonathan H., 2011, Term premia and inflation uncertainty: Empirical estimates from an international panel dataset, *American Economic Review* 101, 1514–1534.

Wright, Jonathan H., 2022, The extent and consequences of Federal Reserve balance sheet shrinkage, *Brookings Papers on Economic Activity* 53, 259–275.

Zou, Dongchen, 2026, Bond demand and the yield-exchange rate nexus: Risk premium vs. convenience yield, Working Paper, Indiana University.



Figure 1: **Decomposition of the Ten-Year Cross-Country Government Bond Term Premium into its Bridge-Identity Components.** Each panel plots, for one of the nine USD bilateral pairs, the monthly time series of the ten-year cross-country government bond term premium $\widehat{ctp}_{10}^{gov,j,USD}$ (black line), constructed from directly quoted series as $x_{10}^{j/USD} + \Delta SS_{10}^{USD,j}$, together with its two quoted components drawn as stacked bands: the bilateral cross-currency basis $x_{10}^{j/USD}$, representing the cross-country swap term premium up to the contained pair-specific residual, as a band from zero (blue), and the cross-country swap spread differential $\Delta SS_{10}^{USD,j}$ as a band stacked on it (green), all in basis points. By construction the outer envelope of the stacked bands coincides with the government line exactly at every date, on the full quoted sample with no extraction anywhere, so each band's thickness is its component's contribution. Robustness at the five- and two-year tenors is reported in Internet Appendix Figures IA.4 and IA.5. Sample: January 2002 to June 2026, monthly frequency. Series are drawn continuously across data gaps of at most one year; gaps exceeding one year appear as breaks, the only such gap being the Swiss franc's exclusion during the minimum

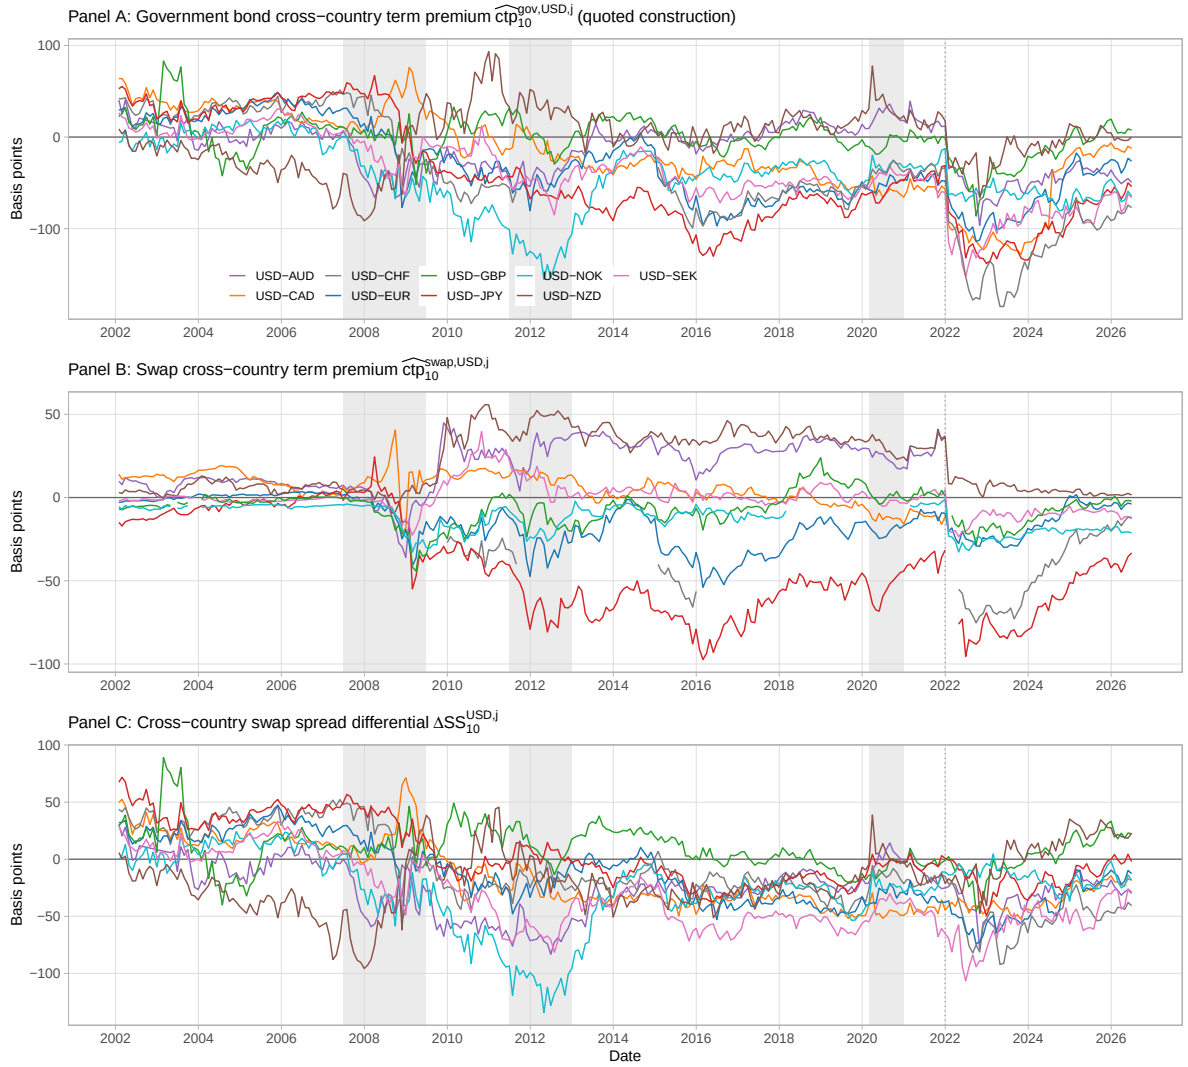


Figure 2: **Evolution of Ten-Year Cross-Country Term Premia and Swap Spread Differentials.** This figure shows the monthly time series of the three additive components of the bridge identity $\widehat{ctp}_{10}^{gov,j,USD} = \widehat{ctp}_{10}^{swap,j,USD} + \Delta SS_{10}^{USD,j}$ for the nine USD-based currency pairs, in basis points. On euro-cross quote-outage months, verified as zero-quote windows in the raw data, no closed triangle exists and the extracted swap term premium of Panel B is not reported, breaking at every unidentified month, while the quoted series of Panels A and C continue. Panel A plots the cross-country government bond term premium $\widehat{ctp}_{10}^{gov,j,USD}$, constructed from directly quoted series as $x + \Delta SS$ on the full quoted sample; it equals Panel B plus Panel C up to the pair-specific residual on identified months. Panel B plots the cross-country swap term premium $\widehat{ctp}_{10}^{swap,j,USD}$ recovered from the triangulation. Panel C plots the cross-country swap spread differential $\Delta SS_{10}^{USD,j} \equiv SS_{10}^{USD} - SS_{10}^{j}$, with the swap spread defined as $SS_{t,N}^{c} \equiv \text{IRS}_{t,N}^{c} - r_{t,N}^{gov,c}$ in each currency. Under my sign convention, negative values indicate that USD demands more term compensation than the comparison currency. Gray shaded areas mark episodes for visual orientation only, the Global Financial Crisis (August 2007 to December 2009), the European Debt Crisis (May 2011 to December 2012), and COVID-19 (March 2020 to December 2020); they are not the sub-period windows used in the tables, which are fixed across the descriptive and reversal exhibits as pre-GFC 2002 to 2007, GFC 2008 to 2012, and post-GFC 2013 to June 2026. Sample: January 2002 to June 2026. Internet Appendix Figures IA.6 and IA.7 report the analogous panel at the five- and two-year tenors. Series are drawn continuously across data gaps of at most one year; gaps exceeding one year appear as breaks, the only such gap being the Swiss franc's exclusion during the minimum exchange rate regime of September 2011 to January 2015. The Swiss pair appears for descriptive completeness and remains outside the estimation panels.

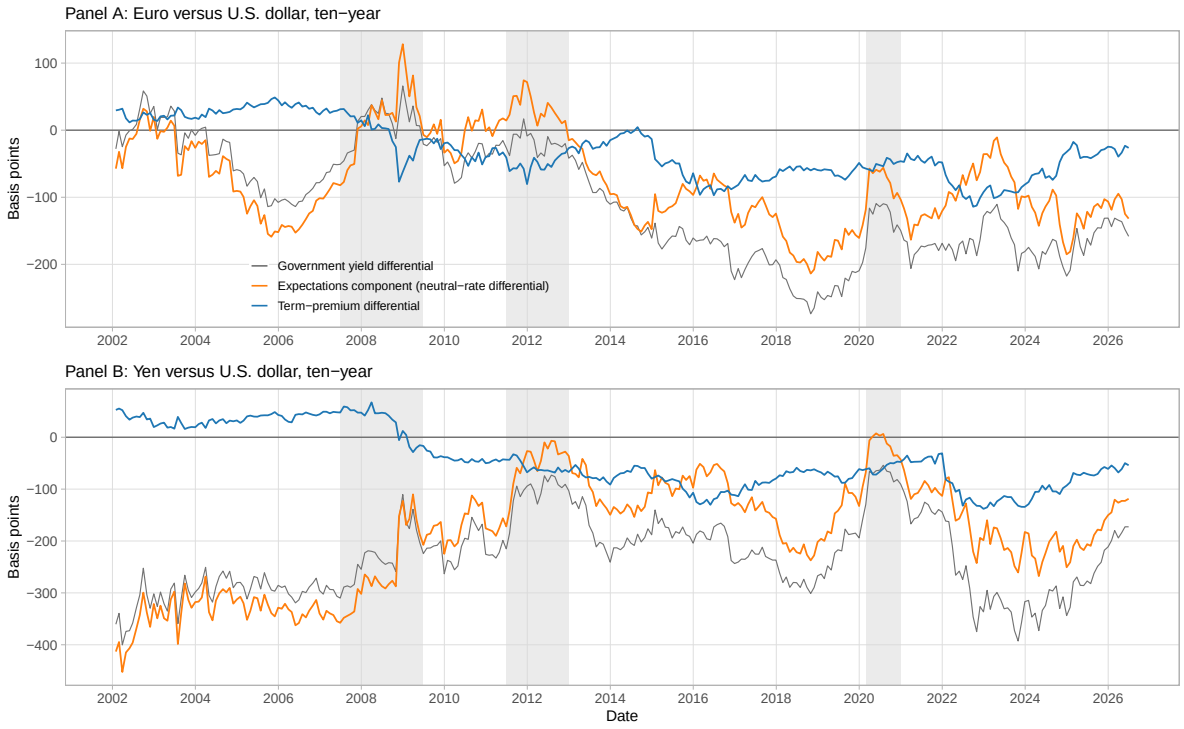


Figure 3: **The Ten-Year Government Yield Differential Split into Its Expectations and Term-Premium Components.** This figure plots, for the euro (Panel A) and yen (Panel B) pairs at the ten-year tenor, the observed benchmark-consistent government yield differential together with its exact split into the two series released with the paper: the recovered cross-country government term premium $\widehat{ctp}_{10}^{gov,j,USD}$ and the expectations component, the neutral-rate differential, defined as the yield differential less the recovered term-premium differential, so the two components sum to the yield differential in every month by construction and no estimation enters the figure. All series are monthly in basis points; under my sign convention negative values indicate that the U.S. quantity exceeds the partner quantity. The expectations component tracks the yield differential closely throughout, and the wedge between them, the term-premium differential, is positive before 2008 and negative from 2013 onward, the reversal of Table 3 seen from the identity's other side; the identification caveats of the recovered differential transfer to the expectations component with the opposite sign. Gray shaded areas mark the same episodes as Figure 2, for visual orientation only. Sample: January 2002 to June 2026, 294 months per pair. Both series, at all analysed tenors and countries, are posted as downloadable datasets on my website.

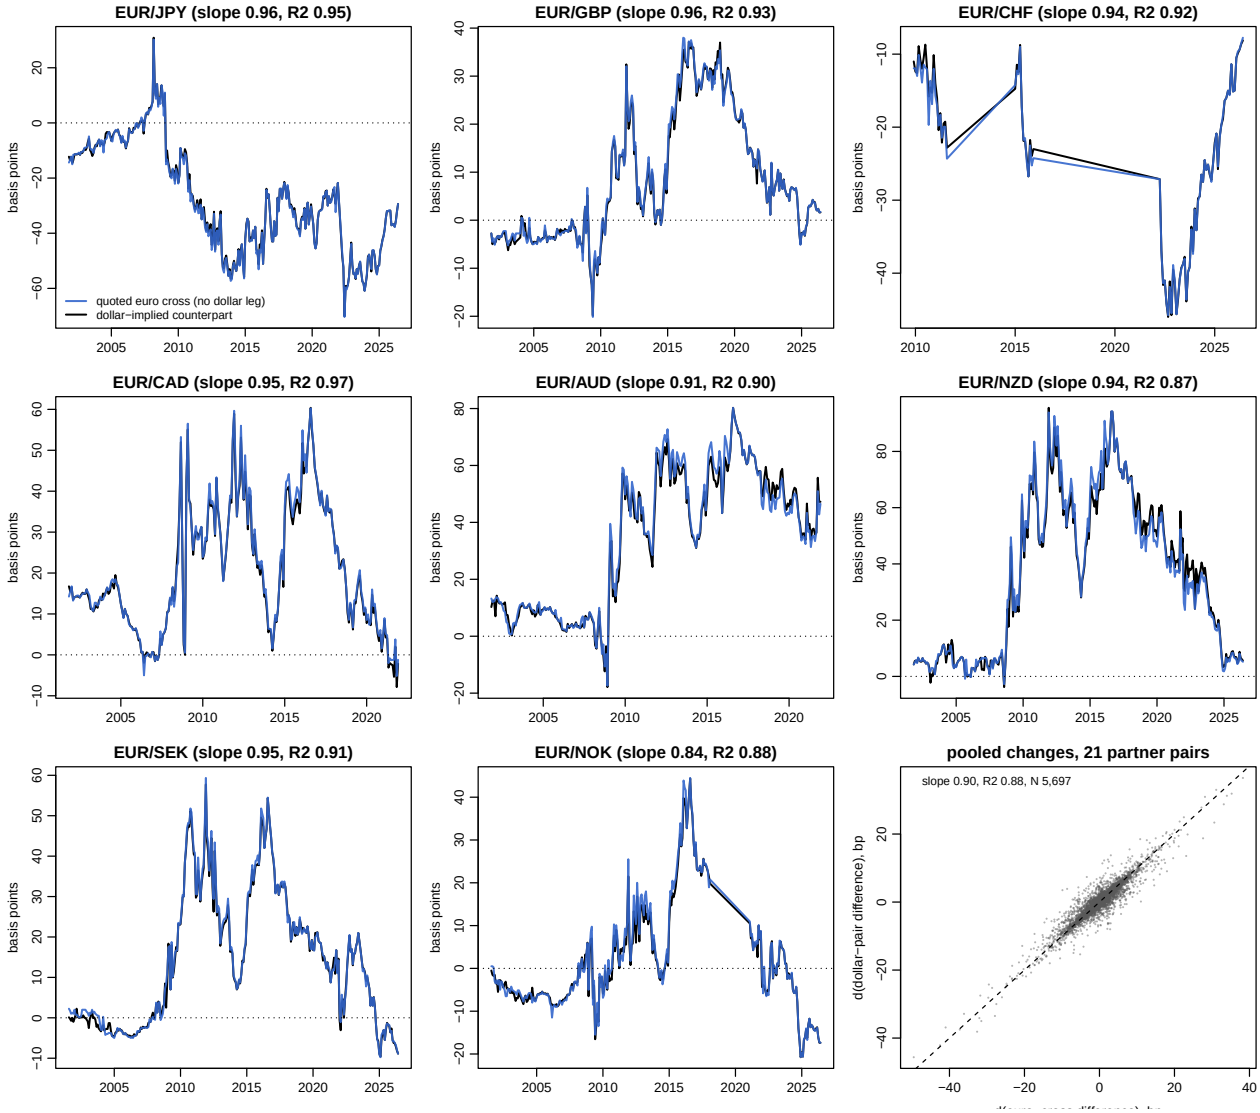


Figure 4: **The Dollar-Free Verification.** The first eight panels overlay, for each partner currency, the raw quoted ten-year euro-cross basis, a price containing no dollar leg, on its dollar-implied counterpart, the difference of the two raw dollar-pair bases; each title reports the slope and R^2 of the regression of monthly changes in the dollar-implied series on changes in the quoted euro cross. The two series coincide at level correlations of 0.993 to 0.999 with mean absolute differences of 0.5 to 2.4 basis points. The final panel pools the monthly changes across the twenty-one partner-pair differences against the forty-five-degree line. The euro-cross legs are the independently quoted series documented in Internet Appendix Section E.1; each panel covers the period over which its euro cross is quoted, the Swiss cross from December 2009 and the Canadian and Australian crosses to the 2021 benchmark transition. Monthly, to June 2026.

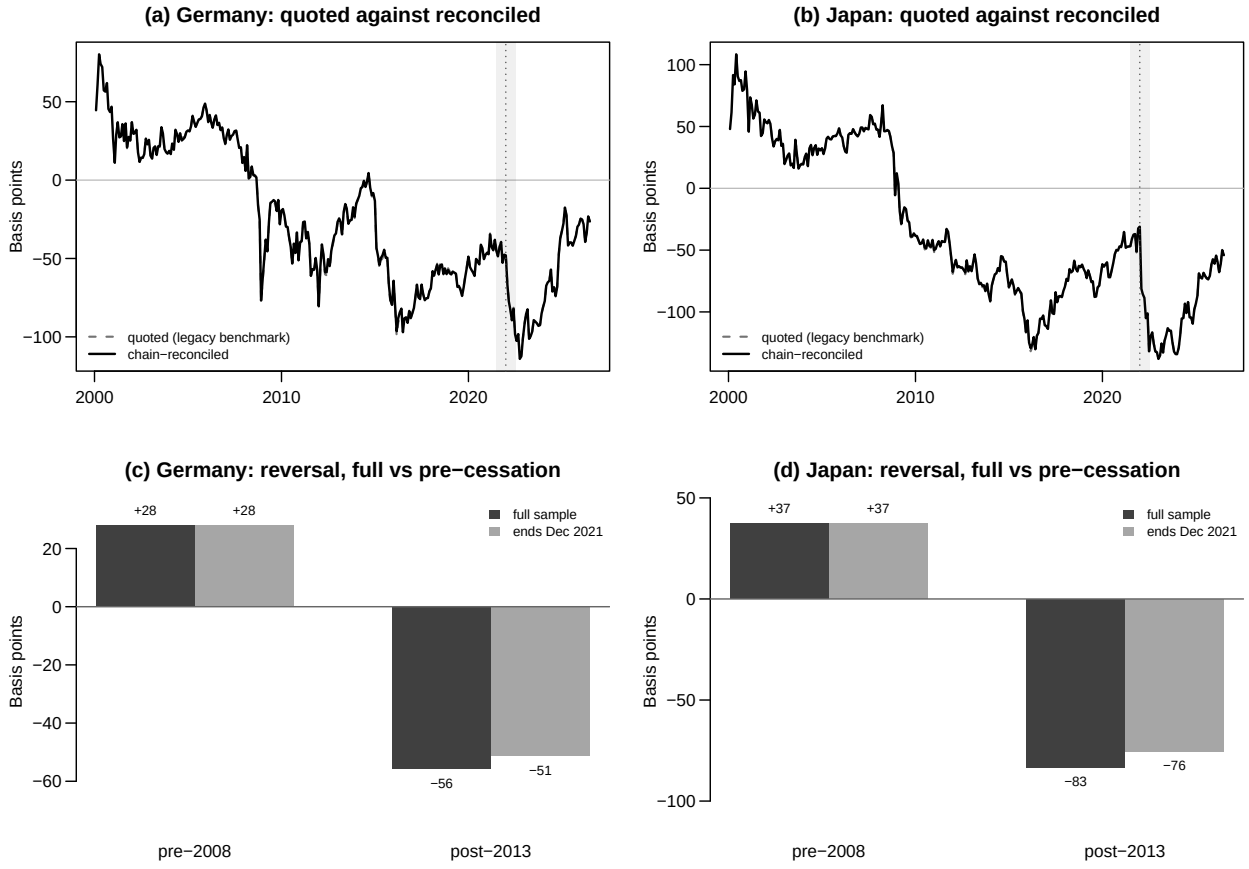


Figure 5: **The Benchmark Transition and the Reversal.** Panels (a) and (b) plot the ten-year cross-country government bond term-premium differential against the United States for Germany and Japan on two constructions: the series built from the quoted legacy benchmarks, dashed, and the chain-reconciled series used throughout the paper, solid. The shaded band and the vertical rule mark the end-December 2021 cessation of the sterling, yen, and Swiss franc legacy benchmarks, the earliest of the relevant dates. The two constructions are visually indistinguishable through the transition: averaged over the year on either side of cessation, the quoted series exceeds the reconciled one by 0.5 and 0.7 basis points for the euro pair and by 0.1 and 0.3 for the yen pair, so the step across the window is 0.2 basis points in both cases. Panels (c) and (d) report the reversal itself on the full sample and on a sample that ends in December 2021, before any of the cessation dates, in basis points. The pre-crisis window is January 2002 to December 2007 and the post-crisis window is January 2013 onward. The reversal is present in full on the pre-cessation sample: the euro differential moves from +27.9 to −51.3 basis points and the yen from +37.4 to −75.7, against +27.9 to −55.8 and +37.4 to −83.2 on the full sample. The benchmark transition therefore cannot be what produces the reversal, because the reversal is already there on data that predates it. The splice construction, the overlap-mean reconciliation, and the 18 to 50 basis point discontinuities that an untreated switch would introduce are in Internet Appendix Section D.2.

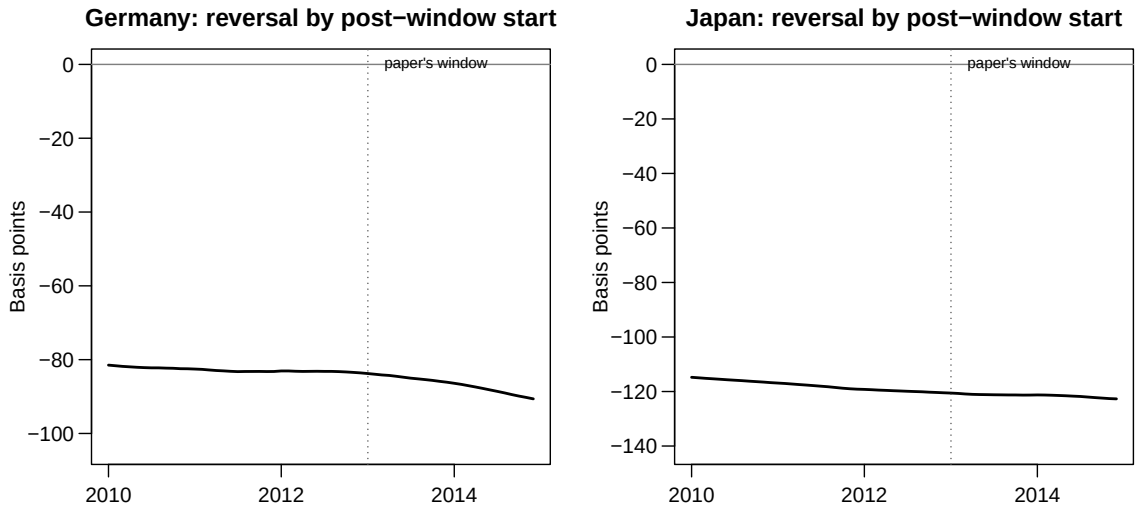


Figure 6: **The Reversal as the Post-Window Start Is Swept.** Each panel plots the difference between the mean ten-year government term-premium differential over a post-window beginning at the date on the horizontal axis and its mean over a fixed pre-window of January 2002 to December 2007, in basis points. The post-window start is swept month by month from January 2010 to December 2014, and the dotted rule marks the January 2013 window the paper's exhibits use. The exercise is not an attempt to date the break, which the single-break test cannot do reliably under this much persistence, but to establish that no reasonable partition removes the result. Over the whole sweep the German difference stays between -90.6 and -81.5 basis points and the Japanese between -122.7 and -114.8; across all eight partners every pair keeps its sign at every start, and a moving-block bootstrap of twenty-four months on the post-window observations keeps the sign at the adverse tail for every pair as well.

Table 1: **Cross-Country Government Bond Term Premium and Privately Held Sovereign Supply (Free-Float).** The dependent variable is the ten-year cross-country government bond term premium differential relative to the U.S. dollar, $\widehat{ctp}^{\text{gov},j,\text{USD}}_{t,10}$, in percentage points, constructed from directly quoted series as the bilateral cross-currency basis plus the cross-country swap spread differential, $x^{j/USD}_{10} + \Delta SS^{USD,j}_{10}$: by the containment result of Section 4.1, the pair-specific residual is bounded by the quoted bid-ask band, so, save the small π, the cross-currency basis is the cross-country swap term premium and the sum is the government bond term premium differential. No triangulation and no allocation restriction enter the construction. The supply variables are constructed on a privately held free-float basis: $\widetilde{d}_{k,t} \equiv d_{k,t} - c_{k,t}$ and $\widetilde{w}_{k,t} \equiv (d_{k,t}w_{k,t} - c_{k,t}w^{cb}_{k,t})/(d_{k,t} - c_{k,t})$. The sample pools the seven USD bilateral pairs, the six incumbent majors and the Norwegian krone, over January 2002 to June 2026 in columns (1) and (2); columns (3) through (5) end in May 2025, the final month of the published He et al. (2017) intermediary capital vintage. The Swedish krona, whose premium and debt-stock data are constructed identically, is excluded here because every column requires a free-float maturity series, which is not available for Sweden in my dataset; the Swedish pair enters every specification of the paper that does not require the maturity variables. All specifications include currency-pair fixed effects and are built sequentially: column (1) enters the four free-float supply and duration variables alone; column (2) adds the country-specific quantitative easing indicators QE^{US}_t and QE^{j}_t, constructed from each central bank's officially declared programme dates (Internet Appendix Section F.5); column (3) adds the standardised intermediary capital ratio HKM^{z}_t of He et al. (2017) and its interaction with U.S. quantitative easing; column (4) adds the general government fiscal balance as a share of GDP on both the U.S. and the comparison-country side, the flow counterpart of the free-float debt stock, together with the country-side macro factors (year-on-year real GDP growth rGDP, three-month interbank rate ir3m, term spread TS $\equiv$ ir10y $-$ ir3m where ir10y is the ten-year sovereign yield, and unemployment rate Unemp, each entered separately on the U.S. and partner-country sides and standardised on the pooled cross-country sample over January 2002 to June 2026 to zero pooled mean and unit pooled standard deviation, denoted by the z-superscript on the country identifier; sources documented in Internet Appendix Section F.5); column (5) adds the U.S. and partner-country year-on-year CPI inflation rates, standardised on the pooled cross-country sample in the same manner as the other country-side macro factors, together with the standardised broad dollar index, the standardised ICE BofA BBB credit spread, the standardised MOVE index of Treasury option-implied volatility, and the standardised VIX index of equity option-implied volatility. An Internet Appendix variant replaces the U.S. realised inflation control with ten-year expected inflation, leaving the preferred-habitat coefficients essentially unchanged (Internet Appendix Table IA.28). The decomposition of each $\widehat{\beta}^{\text{gov}}$ coefficient into its quoted-basis and cross-country swap spread differential components is reported in Internet Appendix Table IA.8 and summarised in Internet Appendix Table IA.17. Standard errors are Driscoll and Kraay (1998) kernel estimates with a Bartlett kernel and lag length set by the Schwert rule, $L = \lfloor 12\,(T/100)^{1/4} \rfloor$ on the estimation-window month count, which gives $L = 15$ at monthly frequency over the full sample. t-statistics are in parentheses.Notes: $^{*}\,p \leq .10$; $^{**}\,p \leq .05$; $^{***}\,p \leq .01$.

	(1)	(2)	(3)	(4)	(5)
$\log(\widetilde{d})_{\text{US}}$	-1.051^{***}	-1.150^{***}	-0.884^{***}	-1.279^{***}	-1.134^{***}
	(-10.72)	(-11.13)	(-9.75)	(-6.82)	(-4.57)
$\log(\widetilde{d})_{j}$	$+0.521^{***}$	$+0.479^{***}$	$+0.496^{***}$	$+0.175^{**}$	$+0.162^{**}$
	$(+6.87)$	$(+7.33)$	$(+8.04)$	$(+2.38)$	$(+2.14)$
$\widetilde{\text{WAM}}_{\text{US}}$	-0.060	$+0.004$	-0.134^{**}	-0.115^{***}	-0.062
	(-1.26)	$(+0.08)$	(-2.36)	(-2.59)	(-1.37)
$\widetilde{\text{WAM}}_{j}$	$+0.023$	$+0.039$	$+0.042$	$+0.029$	$+0.030$
	$(+0.99)$	$(+1.60)$	$(+1.56)$	$(+1.44)$	$(+1.50)$
QE^{US}_t		$+0.158^{***}$	$+0.263^{***}$	$+0.235^{***}$	$+0.255^{***}$
		$(+2.69)$	$(+2.87)$	$(+3.19)$	$(+3.73)$
QE^{j}_t		-0.198^{**}	-0.187^{**}	-0.288^{***}	-0.279^{***}
		(-2.35)	(-2.30)	(-4.73)	(-4.37)
HKM^{z}_t			$+0.146^{***}$	$+0.056$	$+0.050$

Continued on next page

Table 1 continued

	(1)	(2)	(3)	(4)	(5)
			(+3.69)	(+0.70)	(+0.85)
$\text{HKM}_t^z \times \text{QE}_t^{\text{US}}$			+0.107	+0.161*	+0.187**
			(+1.52)	(+1.78)	(+2.04)
$\text{rGDP}_t^{\text{US},z}$				+0.040	+0.030
				(+0.97)	(+0.96)
$\text{rGDP}_t^{j,z}$				−0.016	−0.016
				(−0.97)	(−0.95)
$\text{ir3m}_t^{\text{US},z}$				+0.116	+0.182***
				(+1.63)	(+3.17)
$\text{ir3m}_t^{j,z}$				−0.276***	−0.255***
				(−5.80)	(−4.23)
$\text{TS}_t^{\text{US},z}$				−0.047	−0.023
				(−1.06)	(−0.68)
$\text{TS}_t^{j,z}$				+0.033	+0.047
				(+0.88)	(+1.35)
$\text{Unemp}_t^{\text{US},z}$				−0.007	−0.037
				(−0.25)	(−1.09)
$\text{Unemp}_t^{j,z}$				+0.072	+0.065
				(+1.59)	(+1.60)
$\text{Fiscal bal}_t^{\text{US}}$				−0.006	−0.005
				(−0.68)	(−0.50)
Fiscal bal_t^{j}				−0.012	−0.014
				(−1.17)	(−1.34)
$\text{CPI}_t^{\text{US},z}$					−0.026
					(−0.81)
$\text{CPI}_t^{j,z}$					−0.033
					(−1.06)
Dollar_t^z					−0.045
					(−1.03)
BBB_t^z					−0.012
					(−0.27)
MOVE_t^z					−0.022
					(−1.24)
VIX_t^z					+0.052***
					(+2.61)
Pair fixed effects	Yes	Yes	Yes	Yes	Yes
Within-R^2	0.387	0.422	0.473	0.646	0.659
N	2,043	2,043	1,952	1,911	1,911

Table 2: **Descriptive Statistics Across the Term Structure.** This table reports the full-sample mean, time-series standard deviation, and number of monthly observations N by currency pair at the two-, five-, and ten-year tenors, for the three series studied in the paper. Panel A reports the cross-country government bond term premium $\widehat{ctp}_N^{gov,j,USD}$. Panel B reports the cross-country swap term premium $\widehat{ctp}_N^{swap,j,USD}$. Panel C reports the cross-country swap spread differential $\Delta SS_N^{USD,j} \equiv SS_N^{USD} - SS_N^{j}$, with the per-currency swap spread defined as $SS_{t,N}^{c} \equiv \mathrm{IRS}_{t,N}^{c} - r_{t,N}^{gov,c}$. Under the bridge identity, $\widehat{ctp}_N^{gov,j,USD} = \widehat{ctp}_N^{swap,j,USD} + \Delta SS_N^{USD,j}$. Panel B is formed as Panel A less Panel C, and all three panels are reported on the months where both legs are available, so the identity holds cell by cell; the common count for the Australian pair is 286 rather than 294. The six pairs shown are the six major pairs on which the validation families operate; the Swedish and Norwegian pairs enter the network recovery and appear in Table 3, and the Swiss franc is outside the estimation panels under the minimum-exchange-rate exclusion stated in Section 2.3. All values are in percentage points. The standard deviations reported are the time-series standard deviations of the level series for each pair at each tenor. Sample: January 2002 to June 2026, monthly frequency.

	2Y			5Y			10Y		
Pair	N	Mean	SD	N	Mean	SD	N	Mean	SD
Panel A: Government bond cross-country term premium $\widehat{ctp}_N^{gov,j,USD}$ *(pp)*									
EUR	294	−0.24	0.34	294	−0.27	0.40	294	−0.31	0.41
JPY	294	−0.22	0.38	294	−0.35	0.53	294	−0.42	0.57
GBP	294	−0.17	0.26	294	−0.10	0.28	294	+0.00	0.22
CAD	294	−0.09	0.46	294	−0.08	0.48	294	−0.18	0.45
AUD	286	+0.04	0.24	286	+0.01	0.23	286	−0.14	0.28
NZD	294	−0.04	0.37	294	+0.01	0.34	294	−0.01	0.32
Panel B: Swap cross-country term premium $\widehat{ctp}_N^{swap,j,USD}$ *(pp)*									
EUR	294	−0.18	0.18	294	−0.17	0.16	294	−0.15	0.14
JPY	294	−0.31	0.22	294	−0.41	0.29	294	−0.42	0.29
GBP	294	−0.06	0.11	294	−0.06	0.11	294	−0.07	0.10
CAD	294	−0.08	0.18	294	−0.02	0.16	294	+0.01	0.16
AUD	286	+0.11	0.08	286	+0.16	0.10	286	+0.18	0.14
NZD	294	+0.08	0.15	294	+0.15	0.15	294	+0.21	0.17
Panel C: Cross-country swap spread differential $\Delta SS_N^{USD,j}$ *(pp)*									
EUR	294	−0.06	0.21	294	−0.10	0.29	294	−0.16	0.31
JPY	294	+0.09	0.21	294	+0.06	0.28	294	−0.00	0.30
GBP	294	−0.10	0.23	294	−0.04	0.24	294	+0.07	0.20
CAD	294	−0.00	0.30	294	−0.06	0.34	294	−0.19	0.33
AUD	286	−0.06	0.20	286	−0.15	0.21	286	−0.31	0.25
NZD	294	−0.12	0.28	294	−0.15	0.26	294	−0.23	0.23

Table 3: **The Reversal, with Persistence-Robust Inference.** Sub-period means of the ten-year cross-country government bond term-premium differential $\widehat{ctp}_{10}^{gov,j,USD}$, in basis points, on the calendar windows fixed across the descriptive and reversal exhibits: pre-GFC 2002 to 2007, GFC 2008 to 2012, and post-GFC 2013 to June 2026. "Shift" is the post-GFC less pre-GFC difference in means, computed on unrounded means, so a printed shift can differ from printed-cell arithmetic in the terminal digit. The series are persistent, with first-order autocorrelations of 0.90 to 0.99, so three quantities replace a conventional standard error: the Newey-West t-statistic at the Schwert-rule lag on a post-period indicator; a moving-block bootstrap interval for the shift, blocks resampled within each sub-period; and the band in which the shift would fall under no break, blocks resampled from the pooled series. A shift outside its own no-break band is a rejection that survives the persistence. N_{eff} is the autocorrelation-implied number of independent months; with N_{eff} of 2.2 to 3.6 for the largest shifts, the Newey-West statistics are reported as descriptive summaries of magnitude relative to sampling variability, and the no-break band carries the inferential weight. The series are the benchmark-consistent construction, each leg carried on a single reference-rate path across the cessation dates, which is the construction used by the descriptive and executable-bands tables; the legacy quoted series for the Canadian and Australian pairs end with their reference quotes in December 2021, and the benchmark-consistent construction carries every pair to June 2026. Where both constructions exist they agree to within a basis point, on the overlapping window; the longer benchmark-consistent window is what moves the Canadian and Australian post-GFC means relative to the legacy quotes. The Australian post-GFC mean's sign is determined by the post-transition segment, 2022 to June 2026, whose start coincides with the December 2021 cessation-window break disclosed in Internet Appendix Table IA.29. For the sterling pair the three inference columns disagree, the shift lying marginally outside its no-break band and the bootstrap excluding zero while the Newey-West statistic does not reach conventional significance, and the Australian shift is negative but lies inside its no-break band; I read both as marginal and rest no headline claim on either. The Swiss franc is outside the estimation panels under the minimum-exchange-rate exclusion stated in Section 2.3 and is omitted under that standing exclusion.

	Sub-period means (bp)					Shift inference (95%)		
Pair	pre-GFC	GFC	post-GFC	Shift	NW t	Bootstrap CI	No-break band	N_{eff}
EUR	+27.9	−33.1	−55.8	−83.8	−10.83	[−93.9, −75.1]	[−30.3, +29.6]	3.6
JPY	+37.4	−28.2	−83.2	−120.6	−15.84	[−130.5, −111.0]	[−40.5, +38.1]	2.2
GBP	+11.5	+2.0	−5.2	−16.7	−1.74	[−29.8, −4.5]	[−14.6, +15.0]	15.3
CAD	+30.7	+9.5	−49.7	−80.4	−7.96	[−93.7, −69.9]	[−31.3, +30.4]	3.0
AUD	+4.0	−38.3	−11.5	−15.5	−1.76	[−26.2, −4.6]	[−16.9, +19.1]	9.2
NZD	−29.7	+18.4	+4.1	+33.7	+3.60	[+22.7, +44.8]	[−15.3, +16.1]	12.7
SEK	+9.9	−31.3	−61.7	−71.6	−9.08	[−82.0, −62.3]	[−27.8, +26.6]	4.8
NOK	+1.5	−85.4	−46.4	−47.9	−9.32	[−54.8, −40.6]	[−18.0, +17.2]	6.1

Table 4: **What the Cross-Currency Component Adds: a Component Horse Race.** The government-side measure is the sum of two observable wedges, the quoted long-tenor cross-currency basis x and the swap-spread differential ΔSS. Because the post-crisis behaviour of United States swap spreads is separately documented, this table asks what the basis adds by putting all four series through identical tests on identical samples. The fourth column is the long-maturity government deviation from covered interest parity, $\Delta SS - x$, the difference of the first two components where the measure of this paper is their sum, so the comparison between the last two columns shows which sign combination the economics selects. Each column is computed directly on its own series over that series' available months, so cross-column arithmetic holds only up to small differences in data availability. Panel A reports the change in each series from the January 2002 to December 2007 window to the January 2013 onward window, in basis points, and the mean across the eight partners. Panel B reports the United States free-float coefficient from the supply panel of Table 1, with pair fixed effects and Driscoll-Kraay standard errors, run with each series in turn as the dependent variable, and beneath it the ten-year response of each series to the maturity-composition news of the supply discussion, in basis points per standard deviation of DV01-weighted long-end issuance news at a held total, the two component cells being the additive decomposition of the composite cell by the bridge identity. Panel C reports the ten-year response of each series to the identified asset-purchase surprise factors, in basis points per standard deviation of surprise over the $[-1, +2]$ trading-day window, the Federal Reserve factor pooled across the eight pairs with clustering by announcement date and the European Central Bank factor on the euro pair, the target and guidance factors included as controls as in the policy analysis. Panel D reports how many of the six pairs with an independent affine-model estimate lie on the same side of zero as that model in the post-2009 window.

	Quoted basis x	Swap-spread differential ΔSS	Composite $x + \Delta SS$	CIP deviation $\Delta SS - x$
Panel A: the reversal, pre-2008 to post-2013 (bp)				
Germany	−21.5	−62.2	−83.8	−40.7
Japan	−56.6	−64.0	−120.6	−7.3
United Kingdom	−2.9	−13.8	−16.7	−10.9
Canada	−18.2	−62.2	−80.4	−43.9
Australia	+14.6	−29.6	−15.5	−43.8
New Zealand	+20.0	+13.7	+33.7	−6.3
Sweden	−1.6	−70.0	−71.6	−68.3
Norway	−7.6	−40.3	−47.9	−32.8
Mean across partners	−9.2	−41.1	−50.3	−31.8
Panel B: duration supply, U.S. free-float coefficient				
Coefficient	−0.257	−0.825	−1.088	−0.562
t-statistic	−3.88	−10.91	−8.37	−10.11
Composition news, ten-year response	−0.055 (−1.34)	−0.274 (−1.86)	−0.329 (−2.15)	−0.220 (−1.43)
Panel C: policy purchase surprises, ten-year response (bp per sd, t)				
Federal Reserve asset-purchase factor	+0.12 (1.77)	+0.56 (1.88)	+0.68 (2.12)	+0.47 (1.68)
ECB purchase factor, euro pair	−0.37 (−2.04)	−0.61 (−1.52)	−0.97 (−2.11)	−0.24 (−0.58)
Panel D: agreement with the independent affine model, post-2009				
Pairs on the same side of zero	4 of 6	5 of 6	5 of 6	4 of 6

Table 5: **The Sensitivity Frontier over the Currency-Node Null Space.** The currency-level component of the network is identified only up to one free level per currency, so the unidentified space has dimension $n-1$. This table reports worst cases for the headline results over that entire space rather than along the dollar direction alone; it is a calibrated sensitivity frontier over the stated family, not an identified set. Each currency c is allowed a long-horizon wedge $\lambda_{c,t} = \kappa_c\, n^{3M}_{c,t}$, where $n^{3M}_{c,t}$ is its node in the three-month network, solved from the three-month pair bases on the same seventeen pairs and the same incidence structure as the long-tenor system and available on 234 months. Each κ_c is free in $[-\bar{\kappa}, +\bar{\kappa}]$ and takes whichever sign is most adverse, independently across currencies, so the exercise spans all $n-1$ cross-sectional directions within this funding-indexed family of wedge paths. The short-end node is the natural scale because a funding wedge must appear at the horizon where funding is priced and a term premium is negligible, and κ_c is the pass-through into the ten-year node, measured at about 0.05. The grid is reported rather than chosen: $\bar{\kappa} = 1$ is full capitalisation, the ceiling a permanent funding state attains and which no stationary funding process exceeds, so a conclusion surviving there cannot be overturned by any capitalised funding wedge in any combination. The reversal is the change in the ten-year government differential from the January 2002 to December 2007 window to the January 2013 onward window, in basis points. The supply row is the U.S. free-float coefficient of the panel of Table 1, recomputed with the wedge removed from the dependent variable, the adversarial combination of signs across all currencies found by search over the sign space.

		Worst-case interval at $\bar{\kappa}$ equal to			
Reversal, ten-year	Observed	0.10	0.25	0.50	1.00
Germany	−83.8	[−88, −80]	[−94, −73]	[−105, −63]	[−126, −42]
Japan	−120.6	[−125, −117]	[−131, −111]	[−141, −100]	[−161, −80]
Canada	−80.4	[−82, −78]	[−85, −75]	[−90, −70]	[−100, −61]
Sweden	−71.6	[−73, −70]	[−75, −68]	[−79, −64]	[−86, −57]
Norway	−47.9	[−50, −46]	[−53, −43]	[−58, −38]	[−68, −28]
United Kingdom	−16.7	[−19, −15]	[−22, −12]	[−27, −7]	[−37, +3]
Australia	−15.5	[−18, −13]	[−22, −9]	[−28, −3]	[−40, +9]
New Zealand	+33.7	[+23, +44]	[+7, +60]	[−20, +87]	[−73, +141]
U.S. free-float supply coefficient, worst case over the same space					
Coefficient	−0.953	−0.910	−0.846	−0.740	−0.526
t-statistic	−5.38	−5.31	−5.03	−4.21	−2.28

Table 6: **External Validation of the Term-Premium Interpretation.** This table collects the evidence on the recovered differentials that comes from outside the cross-currency swap market entirely. Panel A asks whether the recovered differential improves on a benchmark of zero. The target is the expected average short-rate differential implied by professional forecasters, which carries no swap-market information of any kind; the recovered column subtracts the recovered term-premium differential from the observed government yield differential before comparing with the survey, and the benchmark column assumes the term-premium differential is zero. Lower root mean squared error and higher correlation favour the recovered measure. The survey legs are the European Central Bank Survey of Professional Forecasters and the Bank of England Survey of External Forecasters, each differenced against the corresponding projection in the U.S. Survey of Professional Forecasters. Panel B reports sub-period means of the recovered and the survey-implied differential on identical dates for the one cell whose pre-2008 window carries more than a handful of rounds, so that the reversal itself can be checked against expectations data that no part of the construction touches. Panel C compares post-2009 means of the recovered ten-year differential against a Gaussian affine term-structure model estimated country by country in the manner of Joslin et al. (2011) with the small-sample bias correction of Bauer et al. (2012), a method that shares no data construction and no identifying restriction with the measurement of this paper, and in the DNS column against a second model family outside the affine class, a dynamic Nelson-Siegel factor model in the manner of Diebold and Li (2006) with first-order vector-autoregressive factor dynamics and no no-arbitrage restriction, evaluated on the identical dates; the yen factor system's largest autoregressive root is 1.001, marginally nonstationary, and is used as estimated. The affine model reproduces the sign of the post-crisis configuration in five of the six pairs, sterling the miss, which is also the pair whose recovered post-2009 mean on the comparison dates lies within four basis points of zero; the factor model reproduces it in all six, sterling included; the same-side and correlation columns refer to the affine comparison. Across the six pairs the Spearman rank correlation between the recovered and the model post-2009 means is 0.94 for the factor model and 0.49 for the affine model. Monthly correlations with either model are low and are reported as such: the models corroborate the configuration and its cross-sectional ordering, not the high-frequency variation. Panels A and B are in basis points, and the full batteries are in Internet Appendix Section H.3 and Table IA.41.

Panel A: does the recovered differential beat a benchmark of zero against survey expectations?

			Correlation		RMSE (bp)		
Cell	N	Span	Recovered	Zero	Recovered	Zero	ΔRMSE
USD/EUR, 1 year	95	2002:10–2026:04	0.868	0.862	85.9	96.3	−10.4
USD/GBP, 1 year	59	2004:11–2020:01	0.974	0.945	24.9	36.3	−11.4
USD/GBP, 2 years	16	2005:02–2020:01	0.964	0.971	27.9	27.1	+0.8

Panel B: the reversal against survey-implied expectations, USD/EUR at one year

Window	N	Recovered	Survey-implied
Pre-2008	21	+9.7	+1.1
2008–2012	20	−46.3	−99.3
Post-2013	54	−23.8	−41.8
Change, pre-2008 to post-2013		−33.5	−42.9

Panel C: post-2009 ten-year means against an affine model and a dynamic factor model (bp)

Pair	N	Recovered	JSZ-BRW	JSZ-OLS	DNS	Corr.	Same side
EUR	299	−52.1	−30.0	−67.0	−64.3	−0.20	yes
JPY	293	−74.6	−16.3	−64.0	−86.4	−0.38	yes
GBP	291	−3.4	+19.5	−7.6	−5.3	−0.33	no
CAD	242	−26.2	−82.7	−60.1	−42.4	+0.26	yes
AUD	234	−5.7	−46.3	−64.2	−59.3	+0.38	yes
NZD	296	+11.3	+20.5	+6.7	+21.5	+0.13	yes

Table 7: **Breakeven Currency-Level Wedges for the Headline Quantities.** For each headline quantity, the table reports the smallest dollar-wide wedge path λ_t, the unidentified direction of Proposition 1, that would overturn the stated result, and its size relative to two verified anchors from the observed funding benchmark, defined throughout as the cross-pair average of the quoted one-year USD cross-currency basis over the eight dollar pairs (EUR, JPY, GBP, CAD, AUD, NZD, SEK, NOK): its sustained post-2009 mean of −11.6 basis points ("sustained") and the largest twelve-month moving average in absolute value, 26.3 basis points ("worst episode"). The supply row sets the required wedge loading on log U.S. free-float debt-to-GDP against the observed loading of the funding benchmark, the former taken from the column (1) estimate of Table 1 on the seven-pair panel and the latter estimated on the eight-pair funding panel, both with Driscoll and Kraay (1998) standard errors. The event row compares the benchmark differential's response of +12.5 basis points around the 20 December 2022 Bank of Japan announcement, the same quantity and window reported in Table 8, against the funding benchmark defined above, the cross-pair average one-year basis, which is the series a dollar-wide wedge moves and which shifted +1.2 basis points over the same $[-1, +2]$ window; the yen pair's own one-year basis moved +4.9 basis points, as Table 8 reports, which would place the ratio at 2.6 rather than 10.7.

Headline quantity	Point estimate	Overturning wedge	Multiple of benchmark
Germany, post-2013 mean	−56 bp	Constant dollar-wide wedge of 56 bp	4.8× sustained; 2.1× worst episode
Japan, post-2013 mean	−83 bp	Constant dollar-wide wedge of 83 bp	7.2× sustained; 3.2× worst episode
Germany, reversal	−84 bp change	Immune to constant wedges; requires a wedge *change* of 84 bp at the break	7.2× sustained level
Japan, reversal	−120 bp change	Immune to constant wedges; requires a wedge *change* of 120 bp at the break	10.3× sustained level
Canada, Sweden, Norway reversals	−48 to −80 bp changes	Wedge *changes* of 48 to 80 bp at the break, one per pair	4.1× to 6.9× sustained level
U.S. supply coefficient	−1.051	Wedge loading +1.05 on log U.S. debt, opposite in sign to the observed benchmark's −0.23	Wrong sign at 4.5× magnitude
BoJ event, 20 Dec. 2022	+12.5 bp (two-day)	Wedge jump of +12.5 bp inside the window	10.7× the observed move (+1.2 bp; the ratio is computed on the unrounded response and move)

Table 8: **The Pre-Specified Event Universe.** Two-day announcement-window responses ($[-1,+2]$ closes) of the daily quoted-construction ten-year benchmark differential of the named pair, or the cross-pair mean over the six core pairs for the Federal Reserve rows (both with all six signs positive); the direction is stated ex ante; the daily-standard-deviation column scales by the standard deviation of all same-length changes; the one-year basis column reports the movement of the named pair's own quoted one-year cross-currency basis over the identical $[-1,+2]$ window, the funding placebo, and is a pair-specific series rather than the cross-pair average funding benchmark of Table 7; the column beside it scales that movement by the standard deviation of all same-length changes of the same one-year series, showing the funding placebo itself elevated at about three daily standard deviations in the 2008 to 2009 crisis windows and small at the 2015 and 2022 to 2023 events; the final column is the percentile of the absolute response among all same-length windows of the same series, the reverse-timing placebo.

Event	Date	Pair	Pred.	Response (bp)	(daily sd)	1Y basis (bp)	(daily sd)	Placebo (pctile)
Fed QE1 announcement	2008-11-25	all six	+	+12.5	3.8	−4.4	2.7	99
BoE gilt QE launch	2009-03-05	GBP	−	−35.6	8.0	−8.6	3.0	100
Fed QE1 expansion	2009-03-18	all six	+	+8.3	2.5	+5.0	3.1	97
ECB expanded APP	2015-01-22	EUR	−	−8.1	2.0	+0.2	0.1	95
BoJ YCC band widening	2022-12-20	JPY	+	+12.5	3.2	+4.9	1.8	99
BoJ YCC reference change	2023-10-31	JPY	+	+9.7	2.5	+3.2	1.2	97

Table 9: The paper's epistemic structure: observations, identities, the assumption, and interpretations.

Observed evidence	Seventeen independently quoted long-tenor bases; dealer bid-ask spreads; swap spreads; near-transitivity of the network, with pair residuals of 0.2 to 0.9 basis points lying inside the quoted band.
Algebraic identities	Currency-level components cancel around triangles; the swap-to-government bridge through the swap-spread differential; the annuity convention, with the par-swap wedge bounded at 1.6 basis points.
The identifying assumption	Pair-specificity of frictions: currency-level variation, including any dollar-wide funding component, is allocated to the currency legs.
Results that do not depend on the allocation	The pair-level content and its containment within quoted spreads; the replication of the dollar-implied differentials on the euro crosses, where no dollar leg appears, and the appearance of the December 2022 Bank of Japan surprise in that dollar-free network; the capitalisation ceiling, which caps the observed ten-year component whatever mixture it contains. Reallocating between a term premium and a dollar-wide wedge moves none of these.
Interpretations, conditional on the assumption	The currency-level component measures cross-country term premia; the post-2008 reversal, the supply and reserve-demand pricing, and the exploratory dollar result follow; each is bounded against the assumption's failure, with breakeven wedges for every headline in Table 7, the reversal surviving unless a dollar-wide non-premium component moved by seven to ten times the sustained level of the observed funding benchmark at the break.

Internet Appendix for "Model-Free Cross-Country Term Premia"

This appendix is organised as follows. Section A carries the full methodological development, the extraction, the standard errors, the formal assumptions, and the identification-robustness batteries; Section B the stochastic discount factor representation and the covered-interest-parity route; Section C the sample construction; Section D the data sources, the benchmark-consistency construction, and the supply-data documentation; Section E additional composition, break, and attribution results; Section F the drivers panel column by column with its robustness; Section G additional supply results and inference; Section H the external validation and the policy events; Section I the dollar dimension; and Section J the interpretation of what the recovered premium bundles.

A Framework: Methodology and Theoretical Interpretation

This appendix provides the complete mathematical framework underlying the triangular extraction methodology of Section 2 of the main text, the stochastic discount factor interpretation of the recovered cross-country term-premium differentials, the formal assumptions and identifying conditions, the treatment of the LIBOR-to-SOFR benchmark transition, the empirical validation of the bilateral frictions, the time-series validation of the swap-to-government-bond bridge identity, the analytical and bootstrap standard errors for the recovered cross-country term premia and cross-country swap spread differential at every tenor, and the regularisation and leave-one-currency-out checks on the identification.

A.1 Cash-Flow Construction and the Par-Swap Representation

This section records the cash-flow basis for treating the long-tenor cross-currency basis swap spread as a maturity-specific pricing wedge and quantifies the annuity-weighting approximation involved in the par-swap representation. A standard long-tenor cross-currency basis swap exchanges the two floating money-market legs with an exchange of notionals at initiation and at maturity, with the quoted basis paid as a spread on one floating leg, and both legs discounted on collateralised, overnight-indexed curves under standard credit support annexes. Because each floating leg telescopes, under its own collateralised discounting, to par

at each reset, the package prices the relative valuation of the two collateral currencies over the full tenor, and the quoted spread is the per-period annuitised compensation for that relative valuation: the spread $x_{t,N}$ satisfies an annuity identity in which $x_{t,N}$ multiplied by the domestic-leg annuity equals the present value of the cumulative wedge. The recovered cross-country term-premium differentials inherit this annuitised, per-period expression, which is the same convention in which swap spreads and yields are quoted, so the bridge identity combines like-for-like quantities.

The alternative representation expresses the same package as the difference between two fixed-rate par swaps, and the two conventions differ by an annuity-weighting term that depends on the level and slope of the discount curves.

Proposition 2 (Par-swap representation bound)**.** *For a ten-year annual-exchange structure with short rates at two percent, zero-curve slopes between* -100 *and* $+300$ *basis points, and basis levels within* ± 50 *basis points, the wedge between the quoted-spread convention and the fixed-for-fixed par difference is at most* 1.6 *basis points in absolute value.*

Proof. Direct computation over the stated grid, with the wedge below one basis point for the configurations that dominate the sample. □ The approximation is therefore second order relative to the term premia documented in the paper, and all empirical work uses the quoted-spread convention throughout, so no conversion enters the construction.

Table IA.1: **Free-Float Supply and the Benchmark Cross-Country Government Bond Term Premium.** This table re-estimates the five-column free-float specification of Table 1, with identical regressors, fixed effects, and Driscoll and Kraay (1998) standard errors, changing the dependent variable to the extracted premium: the ten-year cross-country government bond term premium constructed as the extracted cross-country swap term premium plus the cross-country swap spread differential, $\widehat{ctp}_{10}^{swap,j,USD} + \Delta SS_{10}^{USD,j}$, reported only on pair-months where the pair-specific residual is identified by a closed triangle; unidentified months, euro-cross quote outages including the post-2021 Australian and Canadian samples, are excluded rather than filled by any convention. The near coincidence of every cell with its counterpart in Table 1 shows the headline is invariant to replacing the quoted basis with the extracted swap term premium. The panel comprises the same seven USD bilateral pairs as Table 1, the six incumbent majors and the Norwegian krone; no free-float maturity series is available for the Swedish krona in my dataset, so it enters the specifications without maturity columns elsewhere in the paper and is excluded here. Sample: January 2002 to June 2026 in columns (1) and (2); columns (3) through (5) end in May 2025 with the published He et al. (2017) vintage. t-statistics in parentheses. Notes: $^{*}\,p \leq .10$; $^{**}\,p \leq .05$; $^{***}\,p \leq .01$.

	(1)	(2)	(3)	(4)	(5)
$\log(\widetilde{d})_{\text{US}}$	-1.066^{***}	-1.154^{***}	-0.952^{***}	-1.335^{***}	-1.216^{***}
	(-9.89)	(-9.55)	(-11.09)	(-8.87)	(-5.78)
$\log(\widetilde{d})_{j}$	$+0.581^{***}$	$+0.532^{***}$	$+0.535^{***}$	$+0.283^{***}$	$+0.267^{***}$
	$(+6.96)$	$(+7.02)$	$(+7.66)$	$(+2.94)$	$(+2.81)$
$\widetilde{\text{WAM}}_{\text{US}}$	-0.013	$+0.049$	-0.070^{*}	-0.073^{*}	-0.046
	(-0.38)	$(+1.04)$	(-1.78)	(-1.86)	(-1.21)
$\widetilde{\text{WAM}}_{j}$	$+0.018$	$+0.036$	$+0.041$	$+0.015$	$+0.017$
	$(+0.67)$	$(+1.36)$	$(+1.39)$	$(+0.66)$	$(+0.70)$
QE_t^{US}		$+0.157^{***}$	$+0.270^{***}$	$+0.262^{***}$	$+0.266^{***}$
		$(+2.60)$	$(+2.86)$	$(+4.49)$	$(+4.56)$
QE_t^{j}		-0.219^{***}	-0.211^{***}	-0.267^{***}	-0.262^{***}
		(-2.76)	(-2.69)	(-4.69)	(-4.43)
HKM_t^{z}			$+0.111^{***}$	$+0.002$	$+0.007$
			$(+3.71)$	$(+0.03)$	$(+0.13)$
$\text{HKM}_t^{z} \times \text{QE}_t^{\text{US}}$			$+0.121^{*}$	$+0.180^{**}$	$+0.188^{**}$
			$(+1.74)$	$(+2.40)$	$(+2.38)$
$\text{rGDP}_t^{\text{US},z}$				$+0.019$	$+0.016$
				$(+0.52)$	$(+0.53)$
$\text{rGDP}_t^{j,z}$				-0.010	-0.012
				(-0.57)	(-0.65)
$\text{ir3m}_t^{\text{US},z}$				$+0.163^{**}$	$+0.196^{***}$
				$(+2.42)$	$(+3.52)$
$\text{ir3m}_t^{j,z}$				-0.262^{***}	-0.241^{***}
				(-5.92)	(-4.10)
$\text{TS}_t^{\text{US},z}$				-0.026	-0.014

Table IA.1 continued

	(1)	(2)	(3)	(4)	(5)
				(−0.59)	(−0.36)
$TS_t^{j,z}$				+0.039	+0.051
				(+1.00)	(+1.35)
$Unemp_t^{US,z}$				−0.026	−0.040
				(−0.86)	(−1.20)
$Unemp_t^{j,z}$				+0.029	+0.032
				(+0.75)	(+0.84)
Fiscal bal_t^{US}				−0.007	−0.005
				(−0.84)	(−0.53)
Fiscal bal_t^{j}				−0.021**	−0.022**
				(−2.23)	(−2.38)
$CPI_t^{US,z}$					−0.002
					(−0.05)
$CPI_t^{j,z}$					−0.035
					(−1.24)
$Dollar_t^{z}$					−0.023
					(−0.59)
BBB_t^{z}					−0.003
					(−0.09)
$MOVE_t^{z}$					−0.012
					(−0.80)
VIX_t^{z}					+0.034**
					(+1.98)
Pair fixed effects	Yes	Yes	Yes	Yes	Yes
Within-R^2	0.400	0.445	0.482	0.659	0.666
N	1,888	1,888	1,823	1,782	1,782

A.2 The Core Decomposition

Relocated from the main text's Section 2.2, the risk-neutral representation of each fixed swap rate:

$$IRS_{t,N}^{c} = \frac{1}{N}\mathbb{E}_t^{\mathbb{Q}^c}\left[\sum_{k=1}^{N} r_{t+k-1}^{(1),c}\right] \quad \text{for } c \in \{i, j\}. \tag{IA.1}$$

Relocated from the main text's Section 2.2, the full substitution and regrouping display behind its Equation (5):

Substituting Equations (IA.1) and (4) into Equation (3) and rearranging,

$$\begin{aligned} x_{t,N}^{j/i} &= \frac{1}{N}\mathbb{E}_t^{\mathbb{Q}^j}\Big[\sum r^{(1),j}\Big] - \frac{1}{N}\mathbb{E}_t^{\mathbb{Q}^i}\Big[\sum r^{(1),i}\Big] - \frac{1}{N}\mathbb{E}_t^{\mathbb{P}}\Big[\sum r^{(1),j} - \sum r^{(1),i}\Big] + \pi_{t,N}^{j/i} \\ &= \underbrace{\frac{1}{N}\Big(\mathbb{E}_t^{\mathbb{Q}^j} - \mathbb{E}_t^{\mathbb{P}}\Big)\Big[\sum r^{(1),j}\Big]}_{=tp_{t,N}^{j}} - \underbrace{\frac{1}{N}\Big(\mathbb{E}_t^{\mathbb{Q}^i} - \mathbb{E}_t^{\mathbb{P}}\Big)\Big[\sum r^{(1),i}\Big]}_{=tp_{t,N}^{i}} + \pi_{t,N}^{j/i}, \end{aligned} \tag{IA.2}$$

where the second line regroups the risk-neutral expectations from the swap rate differentials with the physical-measure expectations from the FX forward premium, term by term, applying Definition 1 to each currency to identify the difference between the risk-neutral and physical-measure expected short-rate paths as the term premium of that currency. This is the additive decomposition of the cross-currency swap basis stated in the main text.

The observed cross-currency basis between currencies i and j for maturity N is decomposed as:

$$x_{t,N}^{j/i} = ctp_{t,N}^{j,i} + \pi_{t,N}^{j/i} \tag{IA.3}$$

where $ctp_{t,N}^{j,i} = tp_{t,N}^{j} - tp_{t,N}^{i}$ represents the cross-country term premium (the differential in swap term premia) and $\pi_{t,N}^{j/i}$ represents the pure parity deviation (pair-specific frictions). This decomposition separates currency-specific swap term premia conditions from bilateral arbitrage frictions.

A.3 Triangular Arbitrage Identification

Relocated from the main text's Section 2.1: This is the central identifying observation that the methodology rests on, and it differentiates the construction I develop here from triangular arbitrage in FX spot markets, which is studied as a measure of bilateral liquidity and market efficiency (Huang et al., 2024; Somogyi, 2025): in FX spot markets the cancelling quantity is the spot exchange rate and the residual measures bilateral spot mispricing, whereas in my setting the cancelling quantity is the term premium because the long-tenor basis swap embeds term-structure information over the life of the contract.

For any three currencies i, j, k, the triangular inconsistency $\Delta_{t,N}^{i,j,k} = x_{t,N}^{j/i} + x_{t,N}^{k/j} - x_{t,N}^{k/i}$ satisfies:

$$\begin{aligned}\Delta_{t,N}^{i,j,k} &= [(tp^j - tp^i) + \pi^{j/i}] + [(tp^k - tp^j) + \pi^{k/j}] - [(tp^k - tp^i) + \pi^{k/i}] \\ &= \underbrace{tp^j - tp^i + tp^k - tp^j - tp^k + tp^i}_{=0} + \pi^{j/i} + \pi^{k/j} - \pi^{k/i} \\ &= \pi^{j/i} + \pi^{k/j} - \pi^{k/i} \end{aligned} \tag{IA.4}$$

The currency-specific swap term premia cancel out exactly. This property is central to the identification: while neither component is directly observable, triangular inconsistencies reveal pure parity deviations alone.

A.4 Extraction: Full Development

I compute the unique solution that contains no currency-level shift via the Moore-Penrose pseudoinverse,

$$\hat{\boldsymbol{\pi}}_{t,N} \;=\; \mathbf{D}^{+}\,\boldsymbol{\Delta}_{t,N} \;=\; \mathbf{D}^{\top}\left(\mathbf{D}\mathbf{D}^{\top}\right)^{-1}\boldsymbol{\Delta}_{t,N}, \tag{IA.5}$$

which is the unique solution to the triangular system that is orthogonal to the null space of currency-specific shifts. The minimum-norm projection is therefore the algorithmic implementation of the swap-term-premium definition rather than a free convention: any currency-c-specific compensation that the swap rate embeds is by Equation (1) part of $tp_{t,N}^{c}$ and not part of $\pi_{t,N}^{j/i}$, and the pseudoinverse delivers the corresponding allocation in the data.[10] Internet Appendix A.7 works through the algebra in detail. The identified set can be stated formally.

[10] A global dollar-funding premium common to all USD bilaterals would, by Definition 1, be content of the U.S. swap term premium $tp_{t,N}^{swap,USD}$ rather than of $\pi^{j/i}$, since the corresponding currency-level shift vector lies in the null space of $\mathbf{D}$ and is therefore allocated to the term-premium component by the L2 projection. The empirical content of the recovery is in any case robust to the selection rule within the affine solution set: Internet Appendix Table IA.30 reports time-series correlations above 0.998 between the L2 baseline and three alternative rules at the ten-year tenor, including a USD-EUR zero-pinning rule that fixes the most-liquid pair-specific friction at zero by structural choice.

Proposition 3 (Observational equivalence and the role of the restriction)**.** *Let* $\mathbf{x}_{t,N} \in \mathbb{R}^m$ *stack the observed bases,* $\mathbf{A}$ *the* $m \times n$ *incidence matrix whose rows difference currency-level values, and write* $\mathbf{x}_{t,N} = \mathbf{A}\,\mathbf{tp}_{t,N} + \boldsymbol{\pi}_{t,N}$*. The set of decompositions consistent with the observed bases and with every triangular equation is exactly*

$$\left\{ (\mathbf{tp}_{t,N} + \boldsymbol{\delta},\ \boldsymbol{\pi}_{t,N} - \mathbf{A}\boldsymbol{\delta}) : \boldsymbol{\delta} \in \mathbb{R}^{n-1} \right\},$$

so a currency-level vector can be reallocated between the two components without changing any observable. The pair-specificity restriction, $\boldsymbol{\pi}_{t,N} \perp \mathrm{col}(\mathbf{A})$*, selects the unique element of this set, and the minimum-norm projection computes it.*

Proof. For any $\boldsymbol{\delta}$, $\mathbf{A}(\mathbf{tp} + \boldsymbol{\delta}) + (\boldsymbol{\pi} - \mathbf{A}\boldsymbol{\delta}) = \mathbf{A}\,\mathbf{tp} + \boldsymbol{\pi} = \mathbf{x}$, and since every triangular sum annihilates $\mathrm{col}(\mathbf{A})$, the triangle equations are unaffected; conversely any two consistent decompositions differ by an element of $\mathrm{col}(\mathbf{A})$. Uniqueness under the restriction follows from the orthogonal decomposition $\mathbb{R}^m = \mathrm{col}(\mathbf{A}) \oplus \mathrm{col}(\mathbf{A})^{\perp}$; Internet Appendix A.7 gives the computation. □

The proposition makes the epistemic division exact: everything before the restriction is algebra, the restriction resolves the $(n-1)$-dimensional reallocation freedom, nine dimensions in my ten-currency network, and Section 4.2 bounds the consequences of its failure in the one economically salient direction.

The cross-country term premium differentials are then recovered algebraically by going back to the additive decomposition of the basis and subtracting the identified pair-specific frictions from the observed bases,

$$\widehat{ctp}^{\,j,i}_{t,N} \;=\; x^{j/i}_{t,N} \;-\; \hat{\pi}^{j/i}_{t,N}. \tag{IA.6}$$

Because $\hat{\boldsymbol{\pi}}_{t,N}$ exactly satisfies $\mathbf{D}\hat{\boldsymbol{\pi}}_{t,N} = \boldsymbol{\Delta}_{t,N}$ and contains no currency-specific component by construction, the residuals $\{\widehat{ctp}^{\,j,i}_{t,N}\}$ can be expressed exactly as differences of currency-specific values, $\widehat{ctp}^{\,j,i}_{t,N} = tp^{j}_{t,N} - tp^{i}_{t,N}$, with no estimation residual. The recovery is algebraic, not statistical. For currency pairs that are not directly quoted in the basis swap market, the cross-country term premium differential is identified by transitivity from the recovered currency-specific values.

From swap to cross-country government bond term premia. The triangulation delivers the cross-country swap term premium $\widehat{ctp}^{\,swap,j,USD}_{t,N}$ for each currency pair and tenor. To obtain the corresponding cross-country government bond term premium $\widehat{ctp}^{\,gov,j,USD}_{t,N}$, I use a within-currency identity that converts the swap rate into the government bond yield

via the conventional swap spread. Define the swap spread in currency k at maturity N as $SS^k_{t,N} \equiv IRS^k_{t,N} - r^{gov,k}_{t,N}$, the gap between the N-period swap rate and the N-period government bond yield in currency k (Feldhütter and Lando, 2008; Duffie and Singleton, 1997; Jermann, 2020; Klingler and Sundaresan, 2019; Boyarchenko et al., 2018; Hanson et al., 2022). Substituting $IRS^k_{t,N} = r^{gov,k}_{t,N} + SS^k_{t,N}$ into the swap term premium of currency k from Definition 1 delivers the within-currency identity $tp^{swap,k}_{t,N} = tp^{gov,k}_{t,N} + SS^k_{t,N}$. Applying it to both currencies in the cross-country swap term premium yields the bridge identity from the swap to the cross-country government bond term premium,

$$ctp^{gov,j,USD}_{t,N} \;=\; ctp^{swap,j,USD}_{t,N} \;+\; \Delta SS^{USD,j}_{t,N}, \tag{IA.7}$$

where $\Delta SS^{USD,j}_{t,N} \equiv SS^{USD}_{t,N} - SS^{j}_{t,N}$ is the cross-country swap spread differential, computed directly from observable government bond yields and swap rates. Adding it to the cross-country swap term premium recovered from the triangulation delivers the cross-country government bond term premium for each currency pair and tenor.

The bridge does more than convert units, and displaying what it eliminates clarifies the construction. Writing the swap term premium of each currency as the swap rate less the physical-measure expectation and the swap spread as the swap rate less the government yield, the swap rates cancel term by term in the sum, leaving

$$\widehat{ctp}^{gov,j,USD}_{t,N} \;=\; \big(y^{gov,j}_{t,N} - \tfrac{1}{N}E^{\mathbb{P}}_t[\bar{r}^{\,j}]\big) \;-\; \big(y^{gov,USD}_{t,N} - \tfrac{1}{N}E^{\mathbb{P}}_t[\bar{r}^{\,US}]\big), \tag{IA.8}$$

a quantity in government yields and expectations alone. Consequently, any wedge that resides in the swap rates themselves, benchmark credit premia in the legacy IBOR regime, secured-benchmark distortions after the transition, or any dollar-specific pricing of the swap curve, enters the swap term premium with one sign and the swap spread differential with the opposite sign and cancels identically in the cross-country government bond term premium; this holds by algebra, not by estimation, and is formalised as Corollary 1. The residence of a candidate funding wedge therefore admits a complete classification with three cells: a wedge in the swap rates cancels exactly in the cross-country government bond term premium; a pair-level wedge in the FX forward leg is the residual π, which the containment result bounds within quoted transaction costs; and a currency-level wedge in the FX forward leg is the single surviving exposure, which is precisely where the excess-sensitivity bound, the tenor gradient of the funding loading, and the breakeven wedges of Table 7 operate. For the cross-country government bond term premium, the funding criticism thus reduces to one channel, and that channel is the one the paper tests and bounds.

Worked numerical example. Internet Appendix Section A.9 walks through the recovery on the USD-EUR-JPY triangle for a representative date (1 June 2014, at the ten-year horizon), using actual basis quotes from my sample, as a concrete numerical illustration of the identification procedure.

The methodology relies on two structural assumptions. The first is the additive decomposition of Equation (5), in which currency-specific term premia and pair-specific frictions enter the cross-currency basis additively rather than multiplicatively or with non-linear interactions. The second is that the pair-specific friction vector $\boldsymbol{\pi}_{t,N}$ is genuinely pair-specific, in the sense that it carries no currency-level component: $\boldsymbol{\pi}_{t,N}$ is orthogonal to the subspace of currency-level shifts spanned by vectors with components $\alpha^i_{t,N} - \alpha^j_{t,N}$ for any currency-specific function α. This orthogonality, formalised in Internet Appendix Section A.11 as Assumption 2, is the structural identifying restriction that the L2 minimum-norm projection of Equation (IA.5) enforces algorithmically. The institutional grounds for the restriction are direct: bilateral frictions arise on the specific dealer book that intermediates each currency pair through pair-specific balance-sheet usage, leverage-ratio constraints, G-SIB capital surcharges, central counterparty margining, and bilateral CSA arrangements, and these constraints are organised pair by pair rather than as currency-level factors.

The minimum-norm projection in Equation (IA.5) is the algorithmic implementation of this combined identifying restriction. The bilateral content of $\boldsymbol{\pi}$ that the data identify lies in the row space of $\mathbf{D}$ and has norm bounded by the magnitude of the triangular inconsistencies, $\|\boldsymbol{\pi}^{\text{bilat}}\| \leq \sigma_{\min}^{-1}\|\boldsymbol{\Delta}_{t,N}\|$, which is sub-basis-point in my sample. The level component, residing in the null space of $\mathbf{D}$, is fixed at zero by Assumption 2, the restriction that $\boldsymbol{\pi}$ carries no currency-level component; the swap-term-premium definition, which collects every currency-c-specific compensation embedded in the swap rate into $tp^c_{t,N}$, motivates the restriction but, absent independently observed physical-measure short-rate expectations, does not by itself enforce it.

The structural restriction $\boldsymbol{\pi}^{\text{level}} = \mathbf{0}$ is robust to $\pi^{j/i}_{t,N}$ being defined in Equation (4) as a residual in the FX forward premium rather than in the swap rate. FX forwards do not trade liquidly beyond the one-year horizon, while the methodology recovers cross-country term premium differentials at the two-, five-, and ten-year tenors. At these tenors the cross-currency basis swap is the only quoted instrument that prices the FX leg, and it does so jointly with the term-structure leg under a single coherent pricing measure, so any currency-level content of the long-horizon FX forward premium is co-priced with, and absorbed into, $tp^c_{t,N}$ through the swap-rate decomposition. The pair-specific residual $\pi^{j/i}_{t,N}$ retains its bilateral-pair interpretation. Section B sketches the SDF representation that delivers this joint pricing;

Internet Appendix Section B.2 develops it in full.

The post-2008 breakdown of covered interest parity and the wedge it has opened in the aggregate cross-currency basis are documented in Du et al. (2018b), Borio et al. (2016), and a body of subsequent work characterising the underlying mechanisms.[11] My additive decomposition is complementary to this literature in that it provides a route for separating the basis into a currency-specific term-structure component and a bilateral pair-specific component, leaving open which component carries the dealer-balance-sheet, regulatory, and credit-market channels that the CIP literature documents at any given tenor and date. In my long-tenor sample $|\pi|/|x|$ averages 2.3 percent across the twenty-four pair-tenor combinations, a pattern consistent with a small bilateral pair-specific component and a dominant currency-specific component.

The bilateral differentials I use in the empirical analysis are invariant to the choice of normalisation. Re-estimating with $tp_{t,N}^{EUR} = 0$ instead of $tp_{t,N}^{USD} = 0$ leaves each cross-country differential, each regression coefficient, each R^2, each principal component loading, and each driver estimate numerically identical to four decimal places. The economic content of the recovered swap term premia differentials resides in the differentials, not in the level normalisation.

Standard errors. The triangulation is a deterministic linear mapping from the observed bases quoted at each date to the recovered cross-country term premium differentials at each date, so the relevant standard errors capture the cross-sectional measurement uncertainty in the recovered point estimate at each date given input-quote uncertainty in the bases at that date. I compute these by exact linear variance propagation: the recovered cross-country term premium differential between any two currencies i and j at date t is $\widehat{ctp}_{t,N}^{j,i} = \mathbf{e}_{ij}^{\top}(\mathbf{I}-\mathbf{D}^{+}\mathbf{D})\,\mathbf{x}_{t,N}$, where the projection $(\mathbf{I}-\mathbf{D}^{+}\mathbf{D})$ acts on the observed basis vector $\mathbf{x}_{t,N}$ and the unit vector $\mathbf{e}_{ij}$ selects the corresponding pair (Internet Appendix Section A). The induced variance of $\widehat{ctp}_{t,N}^{j,i}$ at each date follows directly from this linear form given the input-quote covariance of $\mathbf{x}_{t,N}$. A complementary cross-sectional bootstrap check, which within each date resamples the quoted bilateral bases around their observed values to capture within-date input-quote measurement uncertainty and re-runs the triangulation on each replicate, yields similar magnitudes, with analytical and bootstrap standard errors averaging 0.12 and 0.09 basis points respectively for

[11] At money-market and short-tenor maturities, Andersen et al. (2019), Cenedese et al. (2021), Avdjiev et al. (2019), Rime et al. (2022), Correa et al. (2025), and Kloks et al. (2025) document that bilateral cross-currency bases are non-zero, persistent, and respond to dealer-balance-sheet usage, leverage-ratio constraints, regulatory cycles, and bank-level dollar-funding dynamics, with further descriptive analysis in Wong and Zhang (2018). At long tenors, Liao (2020) extends the dealer-balance-sheet channel through credit-market frictions.

the ten-year USD-EUR cross-country swap term premium. These cross-sectional measurement standard errors at each date are distinct from and an order of magnitude smaller than the time-series sampling standard errors of the subsample means reported in Table 2, the relevant magnitude for assessing the reserve-currency claims, with the corresponding persistence-robust inference reported in Table 3. Internet Appendix Section A provides the complete derivations and per-pair results.

Robustness of the recovered cross-country term-premium differentials. Three robustness exercises confirm that the recovered cross-country term premium differentials are stable features of the data rather than artefacts of methodological choices. First, varying the regularisation parameter of the Moore-Penrose pseudoinverse across six orders of magnitude, from 10^{-8} to 10^{-2}, changes the recovered ten-year cross-country differentials by less than 0.02 basis points per date for each pair, with a per-date mean deviation below 0.005 basis points (Internet Appendix Section A.16). The numerical insensitivity reflects the well-conditioned structure of the 8×8 Gram matrix $\mathbf{D}\mathbf{D}^\top$ that underlies the pseudoinverse: its eigenvalue 2 has multiplicity seven alongside a single eigenvalue of 10, so its condition number is 5, so the ridge term $\lambda\mathbf{I}$ is small relative to the spectral content of $\mathbf{D}\mathbf{D}^\top$ across the entire range of regularisation parameters I study, ridge regularisation converges numerically to the exact pseudoinverse, and the recovered differentials inherit the exact identification of the friction system under the structural pair-specificity restriction.

Second, replacing the L2 minimum-norm projection with three alternative selection rules including L1 minimum-norm via linear programming, weighted-L2 with weights in the magnitude of the bilateral basis, and zero-pinning of the most-liquid USD-EUR pair, leaves the recovered $\widehat{ctp}^{j,i}_{t,N}$ time series with cell-by-cell time-series correlations above 0.998 with the L2 baseline at the ten-year tenor and mean absolute deviations below one basis point at every pair-tenor cell across the sample, with every empirical claim of the paper preserved (Internet Appendix Section A.17). Beyond these mechanical rules, an economically calibrated transaction-cost metric, a feasible generalised least squares estimator whose friction weights are identified from the triangle covariances themselves, and hub-exclusion tests that drop every dollar pair or every euro pair all reproduce the baseline term premia with correlations above 0.999 and differences below 0.2 basis points; Internet Appendix Section E.1 reports each in detail. No rule within the solution set can address the currency-level null space, which the sensitivity-frontier analysis of Section 4.2 bounds instead. Third, the cross-sectional ordering of the recovered differentials is preserved across pre-crisis, crisis, and post-crisis subsamples, with magnitudes amplifying post-2008 in a manner consistent with the regime change docu-

mented in Table 2 and with the doubling of the U.S.-side variance share reported in Internet Appendix Section E.3.

The recovered cross-country government bond term premia are also unaffected by the LIBOR-to-SOFR benchmark transition: the short-rate input in Definition 1 is held fixed on a currency-level basis across the LIBOR-to-SOFR change, which is the convention I maintain throughout the paper. Under this convention, by construction the benchmark shift cancels in the bridge from the swap to the cross-country government bond term premium, since the level shift in $\widehat{ctp}_N^{swap,j,USD}$ at the cessation date is offset by an equal and opposite shift in $\Delta SS_N^{USD,j}$, leaving $\widehat{ctp}_N^{gov,j,USD}$ invariant; the same cancellation applies to the analogous SONIA, TONA, and AONIA transitions for the GBP, JPY, and AUD pairs.

A.5 System Formulation and Hub-and-Spoke Structure

With $n = 10$ currencies, there are $n_p = \binom{10}{2} = 45$ potential bilateral pairs and up to $\binom{10}{3} = 120$ possible triangular relationships. In practice, I observe seventeen independently quoted cross-currency basis swap spreads: nine USD-based pairs (USD against EUR, JPY, GBP, CHF, CAD, AUD, NZD, SEK, and NOK) and eight EUR-based pairs (EUR against JPY, GBP, CHF, CAD, AUD, NZD, SEK, and NOK), with quoting histories and outages varying by cross as documented in the main text. I exclude mechanically constructed cross-rates to ensure each basis observation provides independent information.

This hub-and-spoke structure, centred on the USD-EUR pair, generates eight triangular relationships of the form USD-EUR-X for each of the eight partner currencies X. Stacking these constraints yields:

$$\mathbf{D}\boldsymbol{\pi}_{t,N} = \boldsymbol{\Delta}_{t,N} \tag{IA.9}$$

where $\boldsymbol{\pi} = (\pi^{EUR/USD}, \pi^{JPY/USD}, \ldots, \pi^{NZD/EUR})^\top$ is the 17×1 vector of bilateral frictions for the independently quoted pairs, $\boldsymbol{\Delta} = (\Delta_{JPY}, \Delta_{GBP}, \ldots, \Delta_{NOK})^\top$ is the 8×1 vector of observed triangular inconsistencies, and $\mathbf{D}$ is the 8×17 incidence matrix.

A.6 Construction of the Incidence Matrix

Each row of $\mathbf{D}$ corresponds to one currency triad and contains exactly three non-zero entries. To illustrate, consider a simplified four-currency system with USD, EUR, JPY, GBP. The six currency pairs are (1) USD/EUR, (2) USD/JPY, (3) USD/GBP, (4) EUR/JPY, (5) EUR/GBP, (6) JPY/GBP. The four triangular relationships yield (note that the fourth row

is linearly dependent on the first three, so the effective rank is three):

$$\mathbf{D}_{\text{4-ccy}} = \begin{bmatrix} 1 & -1 & 0 & 1 & 0 & 0 \\ 1 & 0 & -1 & 0 & 1 & 0 \\ 0 & 1 & -1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & -1 & 1 \end{bmatrix} \tag{IA.10}$$

For instance, the first row encodes $\pi^{EUR/USD} - \pi^{JPY/USD} + \pi^{JPY/EUR} = \Delta^{USD,EUR,JPY}$. For my full ten-currency system, $\mathbf{D}$ is 8×17 with the same three-entry-per-row structure.

A.7 The Null Space and Minimum-Norm Solution

As a concrete illustration with four currencies, suppose the dollar, euro, yen, and sterling network shows bases of -20, -60, and -10 basis points on the three dollar pairs and exactly transitive euro crosses. Any constant δ added to all three dollar-pair term-premium differentials and simultaneously subtracted as a common dollar-leg friction, for instance $\delta = -9$ basis points representing a dollar funding premium, reproduces every quote and every triangle identically: the data cannot distinguish term premia of $(-20, -60, -10)$ with zero dollar friction from term premia of $(-11, -51, -1)$ with a nine-basis-point dollar funding wedge. This is the null-space reallocation of Proposition 3 in its simplest form, and it is exactly the direction the sensitivity-frontier analysis bounds and the excess-sensitivity restriction tests.

The matrix $\mathbf{D}$ has a non-trivial null space of dimension 9 (= 17 unknowns − 8 equations). Any component of the form $\pi^{j/i} = \lambda^i - \lambda^j$ for arbitrary currency-specific values $\lambda^1, \ldots, \lambda^{10}$, which are themselves pinned only up to a common constant, lies in the null space, because such vectors produce zero triangular sums by the same algebraic cancellation that eliminates cross-country term premia. The standard inverse $(\mathbf{D}'\mathbf{D})^{-1}$ therefore does not exist on the triangular equations alone.

To make the economic content of the null space concrete, consider the following example. Suppose a dealer operates with a balance-sheet constraint that imposes a uniform additional cost on every USD-against-j position the dealer intermediates, for every counterparty currency $j \in \{\text{EUR}, \text{JPY}, \text{GBP}, \text{CAD}, \text{AUD}, \text{NZD}\}$. This friction enters every USD-based bilateral basis with the same magnitude, regardless of the partner currency. Algebraically, the friction takes the form $\pi^{j/USD} = \lambda^{USD}$ for every j, which is exactly a currency-specific shift on USD. Such a vector lies in the null space of $\mathbf{D}$: substituting into the triangular sum $\pi^{j/i} + \pi^{k/j} - \pi^{k/i}$ for any triangle that involves USD, the λ^{USD} contributions cancel exactly, leaving zero on the right-hand side. The triangular equations therefore cannot detect such a friction. Under my

identifying restriction, this is the correct allocation: a friction that prices uniformly against every counterparty currency is, by the definition of currency-specificity I adopt, content of the USD-specific component of the swap rate, which the swap-term-premium definition of Equation (1) absorbs into $tp_{t,N}^{USD}$, rather than content of any bilateral pair-specific friction. The minimum-norm projection of Equation (IA.5) delivers this allocation algorithmically.

The triangular equations $\mathbf{D}\boldsymbol{\pi} = \boldsymbol{\Delta}$ leave the level of $\boldsymbol{\pi}$ underdetermined by these $n-1$ dimensions, the dimension of the space of currency-specific shifts $\pi^{j/i} = \lambda^i - \lambda^j$. Combining the triangular equations with the structural restriction that $\boldsymbol{\pi}$ contains no currency-level shift, motivated institutionally by the pair-specific organisation of cross-currency dealer-book frictions and made formally explicit through the swap-term-premium definition in Equation (1), which writes $tp_{t,N}^c$ as the wedge between the swap rate and the physical-measure expected short-rate path and therefore absorbs every currency-c-specific compensation embedded in the swap rate into $tp_{t,N}^c$, makes the level of $\boldsymbol{\pi}$ exactly identified: the triangular equations identify the bilateral content of $\boldsymbol{\pi}$, and the pair-specificity restriction identifies the level. The minimum-norm pseudoinverse,

$$\hat{\boldsymbol{\pi}} = \mathbf{D}^+\boldsymbol{\Delta} = \mathbf{D}^\top(\mathbf{D}\mathbf{D}^\top)^{-1}\boldsymbol{\Delta}, \tag{IA.11}$$

which can equivalently be computed via the singular value decomposition of $\mathbf{D}$, is the algorithmic implementation of this combined identifying restriction. The solution lies in the row space of $\mathbf{D}$, orthogonal to the null space of currency-specific shifts, and is the unique solution that satisfies both the triangular equations and Assumption 2's requirement that $\boldsymbol{\pi}$ contains no currency-level shift.

The pair-specific friction system is therefore exactly identified under the additive decomposition of Equation (5) combined with Assumption 2: the triangular equations identify the bilateral content of $\boldsymbol{\pi}$, and the structural restriction that any currency-level swap-rate compensation is part of $tp_{t,N}^c$, which the swap-term-premium definition of Equation (1) motivates, identifies the level. All pair-specific friction values are then uniquely determined, and the identification is exact under, and only under, that restriction.

A.8 Recovery of Currency-Specific Values

Having identified pair-specific frictions under the swap-term-premium definition, I recover cross-country term premia as $\hat{ctp}^{j,i} = x^{j/i} - \hat{\boldsymbol{\pi}}^{j/i}$. Any $\hat{\boldsymbol{\pi}}$ satisfying $\mathbf{D}\hat{\boldsymbol{\pi}} = \boldsymbol{\Delta}$ implies that the recovered cross-country term-premium vector has zero triangular sums and therefore can

be written as pairwise differences of currency-specific values, $\hat{ctp}^{j,i} = tp^j - tp^i$, for some set $\{tp^1, \ldots, tp^n\}$; this is the integrability property and it follows from the linear-algebraic structure of $\mathbf{D}$ alone. The no-currency-level-shift restriction implied by the swap-term-premium definition then pins down the unique set of $\{tp^c\}$ (and hence the unique selection within the $\hat{\boldsymbol{\pi}}$-solution set) by allocating any currency-c-specific compensation in the basis to tp^c rather than to $\boldsymbol{\pi}$. The system of n_p equations is consistent by construction and no estimation residual remains; the recovery is algebraic, not statistical.

With the normalisation $tp^{USD} = 0$, the individual currency-specific swap term premia are uniquely determined. For pairs not directly observed (e.g., JPY/GBP), cross-country term premia are inferred by transitivity: $\hat{ctp}^{JPY,GBP} = \hat{ctp}^{JPY,USD} + \hat{ctp}^{USD,GBP}$, which holds exactly.[12]

A.9 Worked Numerical Example

To make the identification concrete, this subsection walks through the recovery on the USD-EUR-JPY triangle on 1 June 2014 at the ten-year horizon, using actual basis quotes from my sample. The implied basis quotes are $x_{10}^{USD,EUR} = -1.25$ basis points, $x_{10}^{USD,JPY} = -55.00$ basis points, and $x_{10}^{EUR,JPY} = -53.70$ basis points. The triangular inconsistency is

$$\Delta_{10}^{USD,EUR,JPY} = -1.25 + (-53.70) - (-55.00) = +0.05 \text{ basis points},$$

well within the bid-ask spread of any constituent quote. The global minimum-norm solution of Equation (IA.5) of the main text, applied to the full system of pairs quoted on that date, delivers three pair-specific frictions on this triangle, $\hat{\pi}_{10}^{USD,EUR} = +0.17$, $\hat{\pi}_{10}^{USD,JPY} = +0.06$, and $\hat{\pi}_{10}^{EUR,JPY} = -0.06$ basis points, none of which inherits any currency-specific shift. Individual friction magnitudes can exceed the absolute value of any single triangular inconsistency because the global solution satisfies all triangular constraints simultaneously. The displayed values are rounded to two decimal places and the unrounded global solution satisfies the system exactly.

Subtracting these frictions from the observed bases yields recovered cross-country swap term premia of -1.42 basis points for the dollar-euro pair, -55.06 for the dollar-yen pair, and -53.64 for the euro-yen pair, which satisfy the triangular identity $-1.42 + (-53.64) - (-55.06) \approx 0$

[12] I do not report empirical results for the transitivity-implied non-USD non-EUR pairs in this paper, as my empirical analysis focuses on the nine USD bilateral pairs and the eight EUR bilateral pairs that are independently quoted in the cross-currency basis swap market. The full series for the transitivity-implied pairs are reported in the data release described after the conclusion.

to two-decimal-place rounding precision. Under the normalisation $tp_{10}^{USD} = 0$, I read directly $tp_{10}^{EUR} = -1.42$ and $tp_{10}^{JPY} = -55.06$ basis points on this date. The bridge to the cross-country government bond term premium adds the corresponding ΔSS value at that date through Equation (IA.7) of the main text. The decomposition is repeated independently at each month and at each tenor across the 2002 to 2026 sample.

A.10 Standard Errors

The triangulation extraction is a deterministic linear mapping from the vector of observed bilateral bases to the recovered cross-country term premium differentials at each date. I report standard errors that capture the cross-sectional measurement uncertainty in the recovered point estimate at each date given input-quote uncertainty in the bases at that date, distinct from the time-series sampling standard errors that govern inference on the subsample means reported in Table 2 of the main text.

Analytical standard errors. Let $\mathbf{\Sigma}_t$ denote the cross-sectional variance-covariance matrix of input-basis quote uncertainty at date t, approximated empirically. Under the additive decomposition (Assumption 1) combined with the structural pair-specificity restriction (Assumption 2) of Internet Appendix Section A.11, the recovered cross-country term premium differential for pair (i, j) is $\hat{ctp}_t^{j,i} = \mathbf{e}_{ij}^\top(\mathbf{I} - \mathbf{D}^+\mathbf{D})\mathbf{x}_t$, where $\mathbf{e}_{ij}$ is the selector vector for pair (i, j) and $(\mathbf{I} - \mathbf{D}^+\mathbf{D})$ is the projection onto the null space of $\mathbf{D}$ in which the recovered *ctp* values reside. Its variance is:

$$\mathrm{Var}(\hat{ctp}_t^{j,i}) = \mathbf{e}_{ij}^\top(\mathbf{I} - \mathbf{D}^+\mathbf{D})^\top \mathbf{\Sigma}_t(\mathbf{I} - \mathbf{D}^+\mathbf{D})\mathbf{e}_{ij}. \tag{IA.12}$$

Empirically the analytical standard error for the ten-year USD-EUR cross-country term premium averages 0.12 basis points across the sample.

Bootstrap standard errors. As a complementary check, I implement a cross-sectional bootstrap that, within each date, resamples the quoted bilateral bases around their observed values to capture within-date input-quote measurement uncertainty and re-runs the full triangulation extraction on each replicate. The standard deviation of the recovered cross-country swap term premium across replicates at each date provides the date-by-date bootstrap standard error, and the reported magnitude of 0.09 basis points is the time-series average of these date-by-date bootstrap standard errors at the ten-year USD-EUR pair, which is similar in order of magnitude to the analytical cross-sectional standard error of 0.12 basis points.

Inference on subsample means. The cross-sectional measurement standard errors above

are not the appropriate inferential quantity for the reserve-currency claim of Table 2, which compares pre-GFC and post-GFC subsample means of the recovered series. The time-series sampling uncertainty of each subsample mean is the relevant magnitude for the comparison, and is captured descriptively by the persistence-robust inference reported in Table 3, namely a Newey-West t-statistic at the Schwert-rule lag, a moving-block bootstrap interval for the shift, and the no-break band against which that shift is judged. The cross-sectional measurement standard error at each date is small relative to the magnitude of the pre-GFC versus post-GFC mean shift, so the reserve-currency reversal cannot be attributed to cross-sectional measurement error in the recovered series. I report the Welch statistics as descriptive markers of the economic magnitude of the pre-GFC to post-GFC mean shift rather than as formal inferential rejections, since the Welch test assumes independent observations and the monthly recovered series are persistent.

A.11 Formal Assumptions

The methodology relies on two formal assumptions, the first definitional and the second the single economically substantive identifying restriction referenced throughout the paper:

Assumption 1 (Additive decomposition). The cross-currency basis decomposes additively as $x^{j/i}_{t,N} = ctp^{j,i}_{t,N} + \pi^{j/i}_{t,N}$, where $ctp^{j,i}_{t,N}$ is the cross-country term premium differential between currencies i and j and $\pi^{j/i}_{t,N}$ is a pair-specific friction on the dealer book intermediating pair (i, j), with no non-linear interaction terms. This is standard in asset pricing, where systematic components and frictions are modelled additively, and is consistent with practitioner decompositions of spread products.

Remark (Measurability of triangular inconsistencies). Triangular no-arbitrage provides the relevant benchmark against which to measure market frictions. I do not impose that triangular inconsistencies are zero, which would assume perfect triangular arbitrage, but instead measure them empirically and attribute observed inconsistencies to pair-specific frictions.

Assumption 2 (Pair-specificity of bilateral basis frictions). The pair-specific friction vector $\boldsymbol{\pi}_{t,N} \in \mathbb{R}^m$, where m counts the pairs quoted at date t, lies in the row space of the triangular incidence matrix $\mathbf{D}$, equivalently, $\boldsymbol{\pi}_{t,N}$ is orthogonal to the $(n-1)$-dimensional null space $\mathcal{N}(\mathbf{D})$ of currency-level shifts across the n quoted currencies: for any currency-specific function λ generating the m-vector $\boldsymbol{\lambda}_{t,N}$ with components $\lambda^i_{t,N} - \lambda^j_{t,N}$ across bilateral pairs, $\langle \boldsymbol{\pi}_{t,N}, \boldsymbol{\lambda}_{t,N} \rangle = 0$. In particular, $\boldsymbol{\pi}_{t,N}$ does not admit a representation as a non-trivial currency-level shift, $\pi^{j/i}_{t,N} \neq \lambda^i_{t,N} - \lambda^j_{t,N}$ for any non-zero λ; the orthogonality condition is strictly stronger than this non-representation statement, since orthogonality requires that $\boldsymbol{\pi}_{t,N}$ contain no

currency-level component at all rather than merely failing to be a pure null-space element. The orthogonality is precisely the condition under which the triangular constraints, combined with this restriction, deliver $\boldsymbol{\pi}_{t,N}$ as the unique L2 minimum-norm solution of Equation (IA.5).

The institutional grounds for this restriction are the pair-specific organisation of dealer balance-sheet constraints, including G-SIB capital surcharges, leverage-ratio usage, central counterparty margining, and bilateral CSA arrangements, all of which apply pair by pair rather than as currency-level factors. The assumption is also self-enforcing through the swap-term-premium definition: even if a particular friction were to operate at the currency level rather than pair by pair, it would by construction be absorbed into the currency-specific swap term premium $tp^c_{t,N}$ as a term-premium attribute rather than misallocated to $\pi^{j/i}_{t,N}$, since the swap term premium is defined to collect every currency-c-specific compensation embedded in the swap rate. This is the structural identifying restriction that resolves the null space of the triangular system and identifies the level of $\boldsymbol{\pi}$. The minimum-norm projection of Equation (IA.5) is the algorithmic implementation: the recovered $\hat{\boldsymbol{\pi}}$ lies in the row space of $\mathbf{D}$, orthogonal to the null space of currency-specific shifts, and is the unique solution that satisfies both the triangular equations and Assumption 2.

Notably, the methodology does not require: (1) any specific model for how swap term premia evolve over time; (2) assumptions about the correlation structure between currencies; (3) assumptions about the specific sources of swap term premia; (4) complete markets or perfect arbitrage; (5) specific distributional assumptions for any variables; or (6) stationarity of the decomposition components. This list of non-requirements is important because it underscores the nonparametric nature of the extraction, model-free in the sense of requiring no parametric term-structure model though not assumption-free, and distinguishes it from parametric approaches that necessarily impose some or all of these restrictions. The recovered pair-specific frictions are also economically small, averaging 2.3 percent of the bilateral cross-currency basis across the twenty-four pair-tenor combinations in my sample, the only larger shares being two ten-year cells with near-zero mean bases, indicating that international arbitrage in the long-tenor cross-currency basis swap market enforces triangular consistency tightly and that the recovered cross-country term premium component is the first-order content of the observed bases.

A.12 Potential Limitations and How They Are Addressed

Measurement error in cross-currency bases. If basis spreads are measured with error, the minimum-norm solution provides robustness by distributing the impact of any single mea-

surement error across the entire connected network of triangular relationships, minimising the total adjustment vector length. I verify this by comparing decompositions using different currency subsets.

Non-linear effects. The additive decomposition may miss non-linear interactions between currency-specific and pair-specific components. Such interactions would require a specific mechanism where bilateral arbitrage costs depend on currency-specific funding levels. I find no evidence of such effects in robustness tests.

Currency selection. Results could depend on which currencies are included. I address this by: (1) using all major currencies with liquid cross-currency swap markets; (2) showing stability when dropping any single currency (correlations with full-sample estimates exceed 0.96, with average absolute differences below 2.5 basis points for ten-year term premia); and (3) demonstrating that the decomposition for existing pairs is insensitive to the addition or removal of peripheral currencies.

Hub-and-spoke structure. The reliance on USD-EUR as the hub pair means that every triangular equation on the network passes through at least one of these two currencies. The structural pair-specificity restriction of Section 2.3 restricts the recovered $\boldsymbol{\pi}$ to the row space of the incidence matrix and assigns currency-level shifts to the term-premium vector, so the hub-and-spoke topology cannot mechanically place currency-level information into the pair-specific frictions. The recovered bilateral frictions are also economically negligible in my sample (median absolute value 0.23 basis points, mean 0.44, across all pairs and tenors), and I verify that the empirical content of the recovery is robust to the selection rule across the family of L1, weighted-L2, and zero-pinned alternatives reported in Internet Appendix Table IA.30.

A.13 Empirical Validation: Bilateral Frictions

The recovered bilateral frictions provide a direct empirical check on the methodology. If the cross-sectional structure of observed bases is strongly consistent with currency-specific term premia and small bilateral frictions, the recovered $\boldsymbol{\pi}$ should be economically small relative to the observed bilateral basis. Across the twenty-four pair-tenor combinations covered by Table IA.12, the share $|\pi|/|x|$ averages 2.3 percent and lies below 2.1 percent for twenty-two of the twenty-four cells, with the two larger values at the ten-year Canadian and Swedish cells, where near-zero means of the bilateral basis inflate the share even though the absolute frictions themselves are modest in basis points (under five basis points at every month-end across all pairs and tenors). Under the additive decomposition combined with the swap-term-premium definition, the cross-section of observed bases is therefore strongly consistent

with a currency-specific structure, with bilateral frictions an empirically small residual. The small magnitude of frictions also has an independent economic implication: triangular arbitrage mechanisms function effectively in long-tenor cross-currency swap markets, even during periods when bilateral bases are wide due to funding constraints.

A.14 Swap-to-Government-Bond Bridge: Time-Series Validation

Two consequences of the swap-rate invariance corollary of the main text are developed here. First, it rationalises the forecast asymmetry documented in the main text: the forecastable funding content of the cross-country swap term premium, twelve basis points per hundred, is consistent with residing largely in the benchmark legs of the swap rates, which the bridge removes, while the cross-country government bond term premium, from which those legs have cancelled, carries none. Second, in observables the measured differential and the long-maturity government deviation from covered interest parity are the two signed combinations of the same components, the deviation carrying the quoted basis as a hedging charge where the measured differential carries it as premium content, so the corollary is also the algebraic bridge between this paper's term-premium reading and the convenience reading.

Before the validation, I display why the bridge is the right conversion, since the swap rates that define both differentials cancel in it identically. Within a currency the two term premia are measured against the same expected short-rate path, so they differ by the swap spread alone,

$$tp_{t,N}^{\text{swap},c} - tp_{t,N}^{\text{gov},c} \;=\; \left(s_{t,N}^{c} - E_{t,N}^{c}\right) - \left(y_{t,N}^{c} - E_{t,N}^{c}\right) \;=\; s_{t,N}^{c} - y_{t,N}^{c} \;=\; SS_{t,N}^{c}. \tag{IA.13}$$

Differencing Equation (IA.13) between partner j and the United States and collecting terms,

$$\begin{aligned} ctp_{t,N}^{\text{gov},j,\text{US}} &= \left(s_{t,N}^{j} - E_{t,N}^{j}\right) - \left(s_{t,N}^{\text{US}} - E_{t,N}^{\text{US}}\right) + \left(s_{t,N}^{\text{US}} - y_{t,N}^{\text{US}}\right) - \left(s_{t,N}^{j} - y_{t,N}^{j}\right) \\ &= \left(y_{t,N}^{j} - y_{t,N}^{\text{US}}\right) - \left(E_{t,N}^{j} - E_{t,N}^{\text{US}}\right) \;=\; ctp_{t,N}^{\text{swap},j,\text{US}} + \Delta SS_{t,N}^{\text{USD},j}, \end{aligned} \tag{IA.14}$$

where $s_{t,N}^{j}$ and $s_{t,N}^{\text{US}}$ each enter twice with opposite signs and vanish. The bridge therefore carries the swap-market quantity to the government-bond quantity without leaving any swap rate behind: what remains is the government yield differential net of the expected-path differential, which is the definition of the cross-country government bond term-premium differential. The same cancellation is what distinguishes the recovered differential from the long-maturity government deviation from covered interest parity, which combines the identical three quoted components with the basis entering as the cost of the hedge, so that the two differentials part

ways by exactly twice the quoted basis, as Section 1 states; that factor of two is a feature of the comparison between the two constructions and not of the bridge itself.

The main text's Figure 1 displays the bridge-identity decomposition of the cross-country government bond term premium into its swap and swap-spread-differential components, and the time-series co-movement between the two additive components is visible pair by pair in that figure. Table IA.12 reports the time-series correlation $\rho(\Delta SS, ctp^{swap})$ between the cross-country swap spread differential $\Delta SS_N^{USD,j}$ and the cross-country swap term premium $ctp_N^{swap,j,USD}$, the two additive components of the bridge identity, for each of the six currency pairs and the three reported tenors. The correlation is positive at the ten-year tenor for EUR ($\rho = +0.66$), JPY ($\rho = +0.85$), and CAD ($\rho = +0.70$), the values Table IA.12 prints. Because the two differentials are defined with opposite currency orderings, namely $ctp_N^{swap,j,USD} = tp_N^{swap,j} - tp_N^{swap,\mathrm{US}}$ on the partner-minus-U.S. side and $\Delta SS_N^{USD,j} = SS_N^{\mathrm{US}} - SS_N^{j}$ on the U.S.-minus-partner side, a positive correlation between the two cross-country differentials translates into the underlying U.S. variables $tp_N^{swap,\mathrm{US}}$ and SS_N^{US} moving in opposite algebraic directions.

This is consistent with the standard intermediary-based reading in which a tightening of dealer balance-sheet capacity simultaneously widens the U.S. cross-country swap term premium and compresses the U.S. swap spread, with the two opposite-signed responses on the U.S. side both reflecting the same underlying intermediary stress on the dealer balance sheet. The correlation is weak or negative for GBP, AUD, and NZD in the full-sample series, indicating that the ΔSS channel is orthogonal to or partly offsets the swap signal in these smaller-scale markets at the ten-year horizon.

A.15 Triangulation Standard Errors and Significance Across Maturities

This subsection summarises the standard errors and corresponding statistical significance for the cross-country term premia at the maturities reported in Table 2 of the main text (two, five, and ten years). For the recovered cross-country term-premium differentials $\widehat{ctp}_N^{swap,j,USD}$ and $\widehat{ctp}_N^{gov,j,USD}$, the standard errors are derived from exact linear variance propagation through the projection matrix $(\mathbf{I} - \mathbf{D}^{+}\mathbf{D})$ as described in Internet Appendix Section A, and measure the cross-sectional precision of each date's recovered point estimate; they are typically below 0.5 basis points across pairs and tenors and do not account for the time-series variation in the means, which is the proper quantity for the inference on subsample means reported in Table 2. The cross-country swap spread differential $\Delta SS_N^{USD,j}$ is directly observable from quoted swap rates and government bond yields and therefore does not carry an extraction standard error

of the triangulation type; I report it alongside the recovered cross-country term-premium differentials.

A.16 Regularisation Sensitivity

The recovery of the pair-specific frictions $\boldsymbol{\pi}_{t,N}$ in Equation (IA.5) of the main text uses the Moore-Penrose pseudoinverse of the incidence matrix $\mathbf{D}$, which is equivalent in my finite-dimensional setting to the limit of the Tikhonov-regularised least-squares solution

$$\hat{\boldsymbol{\pi}}_{t,N}(\lambda) \;=\; \mathbf{D}^{\top}\big(\mathbf{D}\mathbf{D}^{\top} + \lambda\mathbf{I}\big)^{-1}\boldsymbol{\Delta}_{t,N} \tag{IA.15}$$

as $\lambda \to 0$, where $\lambda > 0$ is a non-negative ridge parameter applied to the 8×8 Gram matrix $\mathbf{D}\mathbf{D}^{\top}$. The regularisation parameter does not enter the pair-specific frictions directly; it is an algorithmic stabiliser on the matrix inversion that controls how the recovery responds to small perturbations in the input bases. Setting $\lambda = 0$ delivers the exact pseudoinverse and the minimum-norm solution that the main text uses; setting $\lambda > 0$ shrinks the recovered $\boldsymbol{\pi}_{t,N}$ toward zero by a factor that depends on the eigenvalues of $\mathbf{D}\mathbf{D}^{\top}$ relative to λ.

The 8×8 Gram matrix $\mathbf{D}\mathbf{D}^{\top}$ is well-conditioned by construction in my ten-currency system: its eigenvalue 2 has multiplicity seven alongside a single eigenvalue of 10, and its condition number is 5, so the smallest eigenvalue is 2 and the ridge term $\lambda\mathbf{I}$ is small relative to the spectral content of $\mathbf{D}\mathbf{D}^{\top}$ across the entire range $\lambda \in [10^{-8}, 10^{-2}]$ that I study, with $\lambda/\sigma^2_{\min} \leq 5 \times 10^{-3}$ at the largest λ. Ridge regularisation across this range therefore converges numerically to the exact pseudoinverse, and the recovered cross-country term premia inherit the exact identification of the friction system under the structural pair-specificity restriction. Table IA.31 reports the sample-mean recovered ten-year $\widehat{ctp}_{10}^{swap,j,USD}$ at each of five values of λ and confirms the conditioning argument empirically: the sample-mean recovered values are indistinguishable to three decimal places across six orders of magnitude of λ, with the maximum per-date deviation between $\lambda = 10^{-8}$ and $\lambda = 10^{-2}$ below 0.02 basis points for every USD pair at every tenor and a per-date mean deviation below 0.005 basis points.

A.17 Alternative Selection Rules within the Affine Solution Set

The triangular system $\mathbf{D}\boldsymbol{\pi}_{t,N} = \boldsymbol{\Delta}_{t,N}$ has an affine solution set of the form $\hat{\boldsymbol{\pi}}^{\mathrm{L2}}_{t,N} + \mathcal{N}(\mathbf{D})$, where the L2 minimum-norm solution sits in the row space of $\mathbf{D}$ and $\mathcal{N}(\mathbf{D})$ is the $(n-1)$-dimensional null space spanned by currency-level shifts $\pi^{j/i} = \lambda^i - \lambda^j$. The row-space component of the recovery, $\boldsymbol{\pi}^{\mathrm{bilat}}_{t,N}$, is exactly identified by the triangular constraints and is therefore algebraically identical across any selection rule, with $\|\boldsymbol{\pi}^{\mathrm{bilat}}_{t,N}\|$ bounded by $\sigma^{-1}_{\min}\|\boldsymbol{\Delta}_{t,N}\|$ and sub-basis-point in my sample. Selection rules differ only in how they place the null-space component $\boldsymbol{\pi}^{\mathrm{level}}_{t,N}$ within $\mathcal{N}(\mathbf{D})$.

The L2 minimum-norm rule places it at the origin, which is the algorithmic implementation of Assumption 2 that the main text adopts. The recovered cross-country term premium differential $\widehat{ctp}^{j,i}_{t,N}$ is constructed from the full friction vector and therefore depends on both the row-space component (algebraically invariant across rules) and the null-space component (which differs by rule), so its invariance to the selection rule is an empirical rather than algebraic statement. This Internet Appendix subsection re-runs the recovery under three alternative rules within the affine solution set and reports the maximum absolute deviation in the recovered $\widehat{ctp}^{j,i}_{t,N}$ relative to the L2 baseline, pair by pair and tenor by tenor, to demonstrate that the empirical content of the recovery is robust to the selection rule across the family of rules a referee might propose.

The three alternatives are: (i) L1 minimum-norm via linear programming, minimising $\|\boldsymbol{\pi}\|_1$ subject to $\mathbf{D}\boldsymbol{\pi} = \boldsymbol{\Delta}_{t,N}$, which produces sparse solutions concentrating mass on a small number of pairs and is the closest analogue to a LASSO-style alternative within the deterministic affine solution set; (ii) weighted L2, $\hat{\boldsymbol{\pi}}^{\mathrm{wL2}}_{t,N} = \arg\min_{\boldsymbol{\pi}} \boldsymbol{\pi}^\top \mathbf{W}\boldsymbol{\pi}$ subject to $\mathbf{D}\boldsymbol{\pi} = \boldsymbol{\Delta}_{t,N}$, with the diagonal weight matrix $\mathbf{W}$ scaling each pair-specific friction by the inverse of the absolute level of the corresponding bilateral basis at that date, so that the rule pushes friction mass onto larger-basis pairs that are most plausibly intermediated through dealer-balance-sheet channels; and (iii) USD-EUR zero-pinning, in which the most-liquid pair-specific friction $\hat{\pi}^{EUR/USD}_{t,N}$ is fixed to zero on theoretical grounds and the remaining frictions are identified from the triangular system under L2 minimum norm, since the residual system on the date's quoted subnetwork retains a null space of dimension $n-2$ that the L2 selection criterion resolves uniquely.

Table IA.30 reports three complementary statistics on the dispersion of the recovered cross-country swap term premium $\widehat{ctp}^{swap,j,USD}_{t,N}$ across the three alternative rules versus the L2 baseline. Panel A reports the maximum absolute time-series deviation, pair by pair and tenor by tenor, in basis points. Panel B reports the time-series correlation between the recovered

$\widehat{ctp}_{t,N}^{swap,j,USD}$ under each alternative rule and the L2 baseline at the ten-year tenor. Panel C reports the maximum subsample-mean shift across the three alternative rules at the ten-year tenor, computed for the full sample, the pre-GFC subsample (2002 to 2007), and the post-GFC subsample (2013 onwards). The corresponding deviations and subsample-mean shifts in the cross-country government bond term premium $\widehat{ctp}_{t,N}^{gov,j,USD}$ are identical at every pair-tenor cell, since the bridge identity adds the same directly observable cross-country swap spread differential $\Delta SS_{t,N}^{USD,j}$ on both sides of the L2-versus-alternative comparison; the time-series correlations on the cross-country government bond term premium are not invariant under the same identity in general because the addition of $\Delta SS_{t,N}^{USD,j}$ rescales the variances of both series, but in the data they remain above 0.998 at every ten-year pair-tenor cell.

Three readings emerge from the three panels and they are mutually reinforcing. The time-series correlations in Panel B are above 0.998 for every USD pair under every alternative rule at the ten-year tenor, indicating that the within-pair time-series dynamics that drive the regime-shift, predictive-regression, and panel claims of the paper are essentially identical across the four solvers. The subsample-mean shifts in Panel C are below half a basis point for the four most economically central pairs (EUR, JPY, GBP, CAD) under every alternative rule, in every subsample; the larger USD-AUD and USD-NZD shifts of around one basis point post-GFC reflect smaller absolute basis levels at those pairs and remain small relative to the post-GFC subsample-mean magnitudes. The maximum point-by-point deviations in Panel A reach a few basis points at individual extreme observations and are concentrated at the September 2008 Lehman week and the March 2020 COVID stress, when the triangular inconsistencies $\boldsymbol{\Delta}_{t,N}$ are themselves elevated; at these extreme dates the L1 sparsity reallocates friction mass onto a small subset of pairs while the L2 baseline distributes it more diffusely, which mechanically opens the gap. Outside these extreme dates the four solvers track each other within a small fraction of a basis point per cell.

The four headline empirical claims of the paper are invariant to the selection rule across all three alternatives. The post-2008 sign reversal in $\widehat{ctp}_{10}^{gov,j,USD}$ for EUR, JPY, GBP, and CAD is preserved with post-GFC subsample means within at most half a basis point of the Table 2 baseline under all three rules, well within the magnitude of the pre-GFC-to-post-GFC shift that the reserve-currency claim documents. The four preferred-habitat coefficients in Table 1 change by at most a few percent of their baseline values across the three rules, with all signs preserved and statistical significance unaffected, since the panel regression with currency-pair fixed effects absorbs the constant component of any currency-level shift that an alternative rule introduces. The methodology's empirical content therefore does not depend on the L2 selection rule that the main text adopts as the algorithmic implementation of the swap-term-

premium definition.

A.18 Leave-One-Currency-Out and Random Feasible Allocations

Two further exercises probe the recovery from opposite directions: the first removes information from the network, the second adds feasible contamination to the solution.

Leave-one-currency-out. For each currency in turn, every pair containing that currency is removed from the monthly ten-year system, the reduced triangular system is re-solved under the minimum-norm rule, and the recovered dollar-pair differentials are compared with the full-network baseline on identical observations. Table IA.32 reports the results. Across all nine exclusions the pooled correlation with the baseline lies between 0.9996 and 1.0000 and the mean absolute deviation between 0.02 and 0.48 basis points, against baseline standard deviations of twenty to twenty-seven basis points. The euro exclusion is the instructive extreme: removing the euro removes every cross quote, the network degenerates to a pure dollar star with no cycles, and the recovery collapses to the raw bases with the price friction identically zero; that this benchmark still agrees with the full recovery at correlation 0.9996 restates, from the network side, the result documented in the main text that triangulation reallocates little of the quoted basis. The supply headline, re-estimated under each leave-one-out recovery against the baseline on the same reduced panel, matches the baseline to the third decimal in every case; the variation in the coefficient across exclusions, from -0.90 when the yen pairs are removed to -1.31 when the New Zealand pairs are removed, is present identically in the baseline column and therefore reflects panel composition rather than recovery instability.

Random feasible allocations. Every member of the affine solution set differs from the restricted recovery by a currency-level shift $\lambda_{c,t}$ with the dollar normalised to zero: adding any such shift to the recovered differentials leaves every observed basis and every triangular equation unchanged. I draw random feasible members by simulating per-currency AR(1) shifts with persistence 0.97, matching the term premia, and unconditional scale κ times the pooled within-currency standard deviation of the baseline differentials, for $\kappa \in \{0.25, 0.5, 1\}$, with two hundred draws at each scale, and re-run the two central validation tests on each draw. The supply coefficient survives every draw at every scale, with medians of -1.08, -1.05, and -1.06 against the baseline -1.066 and one hundred percent of draws significant with the correct sign: independent currency-level contamination adds noise to the dependent variable but cannot erase a common supply signal identified from nearly two thousand pair-month observations. The extracted-signal carry alpha of the replication materials, by contrast, degrades monotonically in the contamination scale, placing the restricted recovery at the top of the feasible distribution; its cell-by-cell record stands in the replication materials with the rest of the removed return battery. The correct reading of the two results together is stability rather than knife-edge uniqueness: the validation results are not a property of one fortunate point in the solution set, and the discrimination against systematic relabellings, which mean-zero random shifts cannot mimic by construction, is supplied by the sensitivity-frontier bounds of the main text and by the joint falsification standard, which any relabelling must reproduce in its entirety.

B Stochastic Discount Factor Interpretation

This appendix develops the stochastic discount factor representation underlying the additive decomposition of the cross-currency basis swap that the main-text triangulation exploits. The aim is to make explicit, in SDF language, why a single coherent measure is needed to price the cross-currency swap, what the basis prices as the joint effect of two distinct economic operations on that single measure, and how this representation maps onto the additive decomposition into a currency-specific and a pair-specific component. The exposition is self-contained and complements the restriction-based decomposition developed in Section 2 of the main text.

B.1 Cross-Currency Swap Pricing Requires a Single Measure

A fixed-for-fixed cross-currency swap between domestic currency i, taken to be USD throughout the paper, and foreign currency j specifies four streams of contractual cash flows. At inception $t = 0$, party A delivers the domestic notional N^i and receives the foreign notional $N^j = N^i/S_0$, with S_0 the spot exchange rate in USD per unit of currency j. At each coupon date t_k, for $k = 1, \ldots, n$, party A receives the fixed domestic coupon $c^i\, N^i\, \delta_k$ in currency i and pays the fixed foreign coupon $c^j\, N^j\, \delta_k$ in currency j, with δ_k the day-count fraction. At maturity $T = t_n$, the original notionals are re-exchanged at the same spot rate S_0. Two contractual features are central. First, the foreign notional is sized at S_0 at inception, so the two initial notional flows are economically equivalent at spot. Second, the maturity re-exchange occurs at S_0 rather than at any forward rate, so the contract embeds no forward exchange rate in its notional specification, and only the coupon stream and the terminal spot-rate notional re-exchange remain to reconcile the values of the two legs.

The two counterparties cannot agree on a zero-PV exchange of the two cash flow streams unless both streams are valued in a common numeraire and corresponding measure: a change of numeraire is pinned down by the relative numeraire and cannot be applied to one leg in isolation, so the comparison would not be well posed. Coherent bilateral clearing therefore requires that both legs be valued in a single common numeraire/measure before their present values are equated. The choice of which measure is itself a convention, not an identifying assumption, with the basis appearing as a spread on the foreign coupon under the domestic measure and, equivalently, as the same scalar with opposite sign on the domestic coupon under the foreign measure; this is an equivalent pricing representation rather than a re-attribution of the contractual cash flows to a different leg. I work under the domestic measure throughout and the next subsection makes explicit how foreign cash flows are translated under the domestic measure and what compensating measure-change adjustments enter the SDF representation. The triangular construction itself does not depend on the details of the measure-change algebra, only on the cross-currency basis being the equilibrium price of these joint operations and admits the additive decomposition of Equation (5).

B.2 Two Operations in the SDF Representation

The two-operation structure of the SDF representation has its institutional source in the fixed-for-fixed contractual mechanic established in Section B.1. The two notional principals are exchanged at the inception spot rate S_0 at $t = 0$ and re-exchanged at the same S_0 at maturity T, with no FX reset during the life of the contract. As time elapses between $t = 0$ and T, the

domestic-currency leg, the fixed domestic coupon stream plus the domestic notional returned at T, marks to market through the domestic yield curve, while the foreign-currency leg, the fixed foreign coupon stream plus the foreign notional returned at T, marks to market through the foreign yield curve. Because the terminal FX exchange is locked at the inception spot, no FX revaluation channel transmits movements in the foreign yield curve into the domestic-currency value of the contract through changes in the maturity spot rate; the foreign yield curve enters the contract's domestic-currency value through its independent term-structure dynamics rather than through co-moving spot FX path uncertainty.

The contract carries two independent term-structure exposures over its life, and the domestic SDF that prices the contract loads on both. The two formal operations decomposed in this section, the FX numeraire change on money market accounts and the SDF pricing of the long-tenor zero-coupon bond in each currency, are the formal counterpart of this institutional mechanic, and the cross-country differential of the embedded term premia is the cross-country term premium $ctp^{j,i}_{t,N}$ that the triangulation recovers.

Let $M^i_{0,T}$ denote the domestic stochastic discount factor between zero and T, in the sense of Hansen and Richard (1987) and Cochrane (2005). The fundamental pricing identity for a domestic cash flow of one unit deliverable at T is

$$\mathbb{E}_0\big[M^i_{0,T}\big] \;=\; P^i(0,T), \tag{IA.16}$$

where $P^i(0,T)$ is the domestic zero-coupon discount factor at horizon T. The pricing identity for a foreign cash flow of one unit translated into the domestic currency at the realised exchange rate S_T is

$$\mathbb{E}_0\big[M^i_{0,T}\, S_T\big] \;=\; F^{j/i}_{0,T}\, P^i(0,T), \tag{IA.17}$$

with $F^{j/i}_{0,T}$ the FX forward rate at horizon T.

First operation: change of FX numeraire on money market accounts. The CCS pricing identity makes use of a change of numeraire between the foreign and domestic money market accounts, B^j_t and B^i_t. The central point is that the money market accounts grow only at the local instantaneous risk-free rates and are by construction free of any term-structure compensation: $B^c_t = \exp\big(\int_0^t r^c_s\, ds\big)$. The numeraire-change algebra is therefore an FX-translation operation on a quantity that contains zero term-premium content, not an operation that itself prices term-horizon risk. In the terminology of Geman et al. (1995), the standard change of

numeraire takes the Radon-Nikodym derivative

$$\left.\frac{d\mathbb{Q}^j}{d\mathbb{Q}^i}\right|_{\mathcal{F}_t} = \frac{S_t}{S_0}\frac{B_0^i\,B_t^j}{B_0^j\,B_t^i}, \tag{IA.18}$$

which relates the foreign and domestic risk-neutral measures through the spot exchange rate path and the two money market accounts. In a complete market the forward premium priced by this Radon-Nikodym derivative is uniquely pinned down by tradable FX forwards, with the term structure of forward quotes determining the relevant characteristics of the derivative process. At long horizons, where liquid forward contracts do not trade, the forward premium priced by the derivative is identified jointly with the basis from the cross-currency swap quotes that do trade.

Second operation: term-premium content of the long-tenor zero-coupon bond. Even after the foreign cash flows have been translated into domestic units of account through the FX numeraire change of money market accounts, the SDF $M_{0,T}^i$ prices a domestic long-tenor cash flow and a domestic-translated foreign long-tenor cash flow differently because the two cash flows carry different term-horizon priced risks. The two long-tenor zero-coupon bond prices in each currency, $P^i(0,T) = \mathbb{E}_0[M_{0,T}^i]$ for the domestic currency under its own SDF and $P^j(0,T) = \mathbb{E}_0[M_{0,T}^j]$ for the foreign currency under its own SDF, are separately given by the within-currency pricing identities; the role of Equation (IA.16) in my argument is to map the priced foreign long-tenor cash flow into the domestic SDF representation through the FX numeraire change. By no-arbitrage and Definition 1, each long-tenor yield decomposes into the expected short-rate path of its currency under the physical measure $\mathbb{P}$ together with the priced compensation for term-structure risk, the latter being the term premium of that currency. The wedge between the two embedded term premia is the cross-country differential $tp_{0,T}^j - tp_{0,T}^i$, which is precisely the cross-country term premium $ctp_{t,N}^{j,i}$ recovered by the triangulation.

This separation of the two operations is the substantive economic content of the SDF representation. The first operation, FX numeraire change on money market accounts, translates units of account between currencies but contains no term-premium content; the second operation, the SDF pricing of the long-tenor zero-coupon bond in each currency, prices currency-specific term-horizon risk through the term-premium content embedded in $P^c(0,T)$. The cross-currency basis $x_{t,N}^{j/i}$ is the equilibrium price of the joint effect of these two operations on par swap rates, and it is the second operation that the additive decomposition isolates as $ctp_{t,N}^{j,i}$.

B.3 Mapping to the Additive Decomposition of the Basis

The SDF representation of the previous subsection maps onto the additive decomposition adopted in the main text,

$$x_{t,N}^{j/i} \;=\; \left[tp_{t,N}^{j} \;-\; tp_{t,N}^{i}\right] \;+\; \pi_{t,N}^{j/i} \;=\; ctp_{t,N}^{j,i} \;+\; \pi_{t,N}^{j/i}, \tag{IA.19}$$

in which the currency-specific component $tp_{t,N}^{c}$ collects the currency-c-specific term-horizon risk premium and the pair-specific residual $\pi_{t,N}^{j/i}$ carries no currency-level component, formalised as the orthogonality of $\boldsymbol{\pi}_{t,N}$ to the subspace of currency-level shifts in Assumption 2 of Internet Appendix Section A.11. Under the SDF representation, the two components carry distinct interpretations that reflect the two-operation structure of swap pricing.

The currency-specific component $tp_{t,N}^{c}$ corresponds to the projection of the domestic SDF onto currency-c-specific term-horizon risk factors that enter the term-premium content of the long-tenor zero-coupon bond price, including duration compensation against currency-c monetary policy shocks, inflation compensation against currency-c inflation, and priced risks that are specific to the swap market itself, including funding and collateralisation risks on the dealer book that intermediates the swap, the opportunity cost of posting eligible collateral under bilateral or central counterparty arrangements, and risks related to the choice and transition of the floating-rate reference benchmark (Georgievska, 2019; Duffie and Singleton, 1997; Jermann, 2020; Klingler and Sundaresan, 2019; Andersen et al., 2019). These swap-market frictions enter the swap term premium as currency-c-specific compensations because, for instance, the dealer franchise, the collateral and funding markets, and the regulatory environment under which the swap is intermediated are organised at the currency level rather than at the bilateral pair level.

The pair-specific component $\pi_{t,N}^{j/i}$ corresponds to the residual that carries no currency-level component, formalised in Assumption 2 as the orthogonality of $\boldsymbol{\pi}_{t,N}$ to the subspace of currency-level shifts, and that the SDF projection does not absorb. Bilateral arbitrage frictions that operate on the specific dealer book intermediating each currency pair naturally reside in this component. These frictions are pair-specific in incidence. For instance, balance-sheet usage, leverage-ratio constraints, G-SIB capital surcharges, central counterparty margining, and bilateral CSA arrangements may mostly operate on the specific currency pair the dealer is intermediating, and the dealer's balance-sheet cost of carrying a USD-EUR position is structurally distinct from the cost of carrying a USD-JPY position because the two positions occupy different lines on the dealer's risk-weighted assets and require different posted collateral profiles. The triangular cancellation in Section 2 of the main text identifies the com-

bination of pair-specific residuals on the right-hand side of the decomposition without using any structure on the SDF beyond the additive decomposition itself, and the structural pair-specificity restriction of Assumption 2 of Internet Appendix Section A.11 resolves the level of $\boldsymbol{\pi}$ exactly under the swap-term-premium definition that absorbs every currency-c-specific swap-rate compensation into $tp^c_{t,N}$.

B.4 From Covered Interest Parity to the Decomposition

The two-operation representation can be read directly off the covered-interest-parity relation, which makes concrete why the term-premium differential is what the long-tenor basis carries. Write the N-period par swap rate in currency c as the sum of an expectations component and a term premium,

$$s^c_{t,N} \;=\; E^c_{t,N} + tp^c_{t,N}, \qquad E^c_{t,N} \;\equiv\; \frac{1}{N}\sum_{k=0}^{N-1} \mathbb{E}_t\left[r^c_{t+k}\right], \tag{IA.20}$$

where $E^c_{t,N}$ is the average expected short rate over the life of the contract under the physical measure $\mathbb{P}$ and $tp^c_{t,N}$ is the compensation the marginal holder requires for the term-horizon risk of the fixed leg, in the sense of Definition 1, that is, the priced wedge between the physical and the risk-neutral expectation of the rate path.

The first operation of Section B.2, the FX numeraire change on the money market accounts, prices only the expectations dimension of the two legs: the money market accounts grow at the local instantaneous risk-free rates and carry no term-premium content, so the forward premium $\rho^{j/i}_{t,N}$ that this operation pins down is the price of the short-rate, or expectations, dimension alone. Covered parity is then the statement that the two expectations legs, once the forward premium is removed, coincide up to a residual,

$$\pi^{j/i}_{t,N} \;\equiv\; \left(E^j_{t,N} - E^i_{t,N}\right) - \rho^{j/i}_{t,N}, \tag{IA.21}$$

which is the genuine deviation from riskless parity and would vanish if the two legs were riskless in every sense. What the first operation does not price, and so cannot remove, is the second operation: the term premium is the compensation for term-horizon risk, which is orthogonal to the exchange-rate and short-rate-level risk that the covered contract hedges, so the term-premium differential is the residual dimension that the long-tenor basis is left carrying once each leg has been placed on the same risk footing.

Substituting Equation (IA.20) into the definition of the basis and using Equation (IA.21),

$$\begin{aligned} x^{j/i}_{t,N} &= s^{j}_{t,N} - s^{i}_{t,N} - \rho^{j/i}_{t,N} \\ &= \left(E^{j}_{t,N} + tp^{j}_{t,N}\right) - \left(E^{i}_{t,N} + tp^{i}_{t,N}\right) - \rho^{j/i}_{t,N} \\ &= \underbrace{\left(tp^{j}_{t,N} - tp^{i}_{t,N}\right)}_{ctp^{j,i}_{t,N}} + \underbrace{\left[\left(E^{j}_{t,N} - E^{i}_{t,N}\right) - \rho^{j/i}_{t,N}\right]}_{\pi^{j/i}_{t,N}} = ctp^{j,i}_{t,N} + \pi^{j/i}_{t,N}, \end{aligned} \tag{IA.22}$$

which is Equation (IA.19), now derived from the covered relation rather than posited. The cross-country term-premium differential $ctp^{j,i}_{t,N} = tp^{j}_{t,N} - tp^{i}_{t,N}$ is read from partner currency j to base currency i, and the forward premium enters the parity residual $\pi^{j/i}_{t,N}$ rather than the term-premium differential, so the term premia survive the risk adjustment rather than cancelling.

Equation (IA.22) offers a reading of the long-horizon covered-interest-parity basis that complements the one in the literature that measures it at shorter tenors (Du et al., 2018b; Borio et al., 2016). Where the basis is measured at short tenors, the term-premium differential $tp^{j}_{t,N} - tp^{i}_{t,N}$ is negligible, so $x^{j/i}_{t,N} \approx \pi^{j/i}_{t,N}$ and reading the entire basis as the parity deviation is innocuous. At long tenors the differential is no longer negligible: two long-dated swaps in different currencies carry different duration term premia and are therefore not the same risk, so exchanging them one for one at the spot rate is not a riskless, like-for-like trade, and the component of the basis equal to $tp^{j}_{t,N} - tp^{i}_{t,N}$ is the fair compensation that equalises the risk of the two legs rather than a departure from parity. The deviation that remains once the legs are placed on the same risk footing is the residual $\pi^{j/i}_{t,N}$, which the containment result of Section 4.1 bounds within quoted transaction costs, so at the long end the larger part of what the riskless-parity accounting records as a deviation is a term-premium adjustment and the residual parity deviation is an order of magnitude smaller.

C Sample Construction

This appendix section records the exact sample behind every headline exhibit, consolidated in Table IA.13: the sample window, frequency, cross-section, and sources of each headline table and figure, together with the construction sources of the free-float supply series in Table IA.16. The positioning of the paper relative to the term-structure, covered-interest-parity, preferred-habitat, intermediary, swap-spread, convenience-yield, and carry-trade literatures is developed in the main-text related-literature section and at the point of each result.

D Data and Implementation

The variable inventory in full. The main text carries a summary; the complete inventory follows.

The macro-financial dataset includes the intermediary capital ratio of He et al. (2017), the broker-dealer leverage factor of Adrian et al. (2014), the VIX, the ICE BofA BBB U.S. Corporate Index effective yield, the broad nominal dollar index, the MOVE index, and country-specific quantitative easing indicators constructed as binary variables that take the value one in months during which the country's central bank operated an officially announced asset-purchase programme. The supply-side dataset includes general government gross debt-to-GDP and the par-weighted weighted-average maturity of outstanding sovereign debt for each country, alongside the corresponding central bank holdings series used to construct the privately held free-float supply variables of Table 1. Internet Appendix Section D documents sources, primary disclosures, and construction notes for the central bank holdings and par-weighted maturity series, and Internet Appendix Section D.1 provides a complete table of variable definitions, frequencies, and sources. Internet Appendix Section C records the sample window, frequency, cross-section, and sources behind every headline exhibit, together with the construction sources of the free-float supply series. The reconciliation of the constructed swap-spread differential with its directly quoted counterpart, matching in both levels and changes across the reference-rate transition, is reported in Internet Appendix Section A.14.[13]

This appendix documents the construction of the variables used in the empirical analysis of the main text, including the cross-country term premium series, the privately held free-float supply panel, the par-weighted maturity panel, the central bank holdings panel, and the macro-financial control variables. Section 3 of the main text references these constructions throughout.

D.1 Variable Definitions and Data Sources

Table IA.10 provides definitions and sources for all variables used in the empirical analysis.

[13] On availability, and stating only what is not already given above: the Bloomberg basis, zero-coupon curve and bid-ask series are licensed and cannot be redistributed, so the replication materials carry a checksum manifest of the raw inputs sufficient to verify that a rebuild used the same sources, rather than the inputs themselves. General government gross debt is from the IMF International Financial Statistics, U.S. foreign official holdings from the Federal Reserve Financial Accounts, with the valuation-adjusted official flow series from the final public release of Bertaut and Judson (2014), and Treasury marketable-borrowing estimates from the quarterly refunding releases. The recovered cross-country term-premium series, the benchmark-consistent swap-spread and cross-currency basis panels, and the constructed free-float supply panel are released with the paper, together with the complete estimation and exhibit-generating code.

D.2 Benchmark Consistency: Full Construction and Diagnostics

Two facts relocated from the main text belong at the head of this section. Without the treatment developed here the recovered cross-country term-premium differentials jump by eighteen to fifty basis points at LIBOR cessation. And the benchmark-consistent differentials extend before the 2002 start of the headline sample where their inputs are quoted, with per-pair starts from January 2000; the pre-2002 extension is used only with the per-exhibit samples recorded in Internet Appendix Table IA.13.

The bridge from the cross-country swap term premium to its government bond counterpart adds the cross-country swap-spread differential $\Delta SS_N^{USD,j}$ to the recovered cross-country swap term premium. Both the cross-currency basis and the swap spread are quoted against a floating reference rate, and that reference changed over my sample as the major currencies moved from the interbank offered rate convention to the overnight risk-free rate convention during the benchmark transition of 2021 and 2022. I show here that the recovered cross-country government bond term premium is invariant to this change of reference rate, provided the cross-currency basis and the swap spread are constructed on the same convention at every date, and I describe the construction that enforces this property.

Let $\delta_{t,N}^c$ denote the tenor spread in currency c at tenor N, namely the difference at date t between the interbank-referenced rate and the overnight risk-free rate for the same currency and maturity. Moving the cross-currency basis for the pair USD/j from the interbank convention to the risk-free convention shifts its level by the differential of the two legs' tenor spreads, so that $x_N^{USD,j}(\text{RFR}) = x_N^{USD,j}(\text{IBOR}) + (\delta_{t,N}^{USD} - \delta_{t,N}^{j})$. The swap spread moves in the opposite direction, because the risk-free swap rate lies below the interbank swap rate by that currency's own tenor spread, so that $SS_{t,N}^c(\text{RFR}) = SS_{t,N}^c(\text{IBOR}) - \delta_{t,N}^c$ and therefore $\Delta SS_N^{USD,j}(\text{RFR}) = \Delta SS_N^{USD,j}(\text{IBOR}) - (\delta_{t,N}^{USD} - \delta_{t,N}^{j})$. Adding the two components, the tenor-spread differential cancels,

$$x_N^{USD,j}(\text{RFR}) + \Delta SS_N^{USD,j}(\text{RFR}) \;=\; x_N^{USD,j}(\text{IBOR}) + \Delta SS_N^{USD,j}(\text{IBOR}), \tag{IA.23}$$

so the recovered cross-country government bond term premium differential is the same economic quantity under either reference-rate convention. The cancellation is exact, and it does not require estimating or subtracting any $\delta_{t,N}^c$; it requires only that the basis and the swap spread be placed on a common convention at each date, since it is the same tenor-spread differential that enters the two components with opposite sign.

The cancellation also delivers a discipline for construction. If the two components switch convention on different dates, or if one is placed on the risk-free convention while the other

remains on the interbank convention, the tenor-spread differential no longer nets out and the recovered cross-country term-premium differential inherits a level step at the transition that is an artefact of the reference-rate change rather than a movement in the underlying term premium. I therefore construct both legs of every currency on a single reference-rate path. Each leg that transitions switches from its interbank quote to its overnight risk-free successor on the same currency-specific transition date, and the two segments are joined by a level reconciliation over their overlap window, so that the constructed series carries no step attributable to the change of convention. For the United States, Japan, the United Kingdom, Switzerland, and Canada the transition follows the respective interbank-rate cessation, the end of December 2021 for the sterling, yen, and Swiss franc markets and the end of June 2023 for the U.S. dollar, and for Canada the later cessation of the domestic offered rate in June 2024 is treated as a second segment joined by the same reconciliation; the euro, Australian dollar, New Zealand dollar, Swedish krona, and Norwegian krone term reference rates persist and require no switch, as the replication note below records. The interest rate swap legs additionally carry the standard three-month to six-month tenor adjustment where the local swap convention requires it.

Imposing this discipline removes the transition artefact. In the uncorrected series, in which the basis and the swap spread are joined without the common-convention reconciliation, the ten-year recovered differential steps by between eighteen and fifty basis points across pairs at the December 2021 cessation, a magnitude an order larger than a typical monthly movement in the series and of a single sign across pairs, which identifies it as a convention artefact rather than a repricing. After the reconciliation the step falls to within the range of ordinary monthly variation for the pairs whose reference legs are quoted continuously across the transition. The correction changes only the alignment of the two segments at the transition dates and leaves the within-regime variation of the series unchanged, so the recovered term premium differentials are continuous across the reference-rate change while retaining their economic content. On the months where both the legacy quoted and the benchmark-consistent ten-year government differentials exist, the per-pair mean absolute gap between the two constructions is 0.59 basis points for the euro pair, 0.46 for the yen, 0.23 for sterling, 0.35 for the Canadian dollar, 0.75 for the Australian dollar, 0.93 for the New Zealand dollar, 0.30 for the Swedish krona, and 0.32 for the Norwegian krone, with per-pair maxima of 1.7 to 4.6 basis points, so the two constructions differ materially only where the legacy quotes end, the Canadian and Australian pairs after December 2021.

D.3 Replication Note: Reference-Rate Consistency of the Swap-Spread Differential

The recovered cross-country government bond term premium differential is invariant to the floating reference-rate convention only when the cross-currency basis and the swap spread are constructed on the same convention at every date, as established in Section D.2. A researcher reconstructing the cross-country swap-spread differential $\Delta SS_N^{USD,j}$ should therefore observe three requirements, without which the reconstructed series inherits a level step at the reference-rate transition that is an artefact of the change of convention rather than a movement in the underlying term premium.

First, the interest rate swap leg used to form each swap spread must be quoted against the same floating reference as the cross-currency basis at the same date. Prior to each currency's interbank-rate cessation the swap and the basis are both on the interbank convention, and after it both are on the overnight risk-free convention. Second, the two legs must switch convention on the same currency-specific date, taken as the interbank-rate cessation date for that currency, namely the end of December 2021 for the sterling, yen, and Swiss franc markets, the end of June 2023 for the U.S. dollar, and the end of June 2024 for the Canadian dollar domestic offered rate, while the euro, Australian dollar, New Zealand dollar, Swedish krona, and Norwegian krone term reference rates persist and require no switch. Third, at each such date the post-transition segment must be level-reconciled to the pre-transition segment over their overlap window, and the interest rate swap leg must additionally carry the three-month to six-month tenor adjustment where the local convention is quoted on the six-month tenor, so that the constructed swap spread and the constructed basis carry the same reference-rate path and the tenor-spread differential cancels in the sum.

The accompanying rates database provides both constructions for transparency. It reports the benchmark-consistent swap spreads and cross-currency bases used to form the recovered term premium differentials, and, separately, the raw swap spreads switched at each currency's official cessation date without the overlap reconciliation, which display the convention step directly. Both are provided at daily and month-end frequency for ten currencies and five tenors, with a data dictionary that records for each series its definition, the source quotations, and the transition dates applied.

Quote quality. The long-tenor basis quotes are Bloomberg composite quotes, and I report two diagnostics. First, dealer bid and ask fields from the raw pull give mean quoted spreads on the ten-year bases of 2.9 basis points for the euro, 3.5 for the yen, 2.9 for sterling, 2.4 for

the Canadian dollar, 3.2 for the Australian dollar, and 6.0 for the New Zealand dollar, with medians of two to three basis points for the core pairs; spreads roughly double in the crisis windows and are an order of magnitude below the term premia documented in the main text. Second, the share of zero-change days in the daily mid series is 23 percent for the euro, 20 for the yen, 30 for sterling, 37 for the Canadian dollar, 34 for the Australian dollar, and 51 percent for the New Zealand dollar, so the New Zealand pair, which also carries the widest quoted spread, should be read with staleness in mind, consistent with its behaviour in the event-study exhibits; all headline analysis is at the monthly frequency, which mitigates staleness.

D.4 Gross Debt-to-GDP

The gross debt-to-GDP ratio $d_{k,t}$ for each of the nine countries in my panel is obtained from Bloomberg as the IMF International Financial Statistics general government gross debt-to-GDP series, accessed through the GDDI ticker family (GDDICCCG Index, where CCC is the three-digit IMF country code). Bloomberg aggregates these from the underlying IMF IFS submissions and publishes them as annual percentage ratios. For monthly panel use, I carry the most recent annual value forward within each calendar year. The IFS convention is source-consistent with the IMF World Economic Outlook annual nominal GDP series (dataflow `IMF.RES/WEO`, indicator `NGDP`), to which I cross-check the Bloomberg GDDI values directly through the IMF SDMX 3.0 API.

My use of the IMF IFS general-government convention at annual frequency differs from the Haver Analytics quarterly nominal GDP convention used in Du et al. (2018a) and Du et al. (2026). The currency-pair fixed effects in my panel absorb the time-invariant component of any pair-specific level wedge between the two conventions, leaving the residual difference, which is the difference between an annual step-function and a quarterly step-function within the year, in the within-pair time-series variation that identifies the preferred-habitat slope coefficients on $\log \widetilde{d}_{k,t}$. The residual intra-year stepping difference is small in magnitude relative to the within-pair time-series variation in the underlying free-float supply, and the four coefficients in column (1) of Table 1 carry the signs that the preferred-habitat framework predicts.

D.5 Par-Weighted Average Maturity

The par-weighted average maturity of total outstanding sovereign debt $w_{k,t}$ is constructed from the primary national debt-management disclosure for each country: the CRSP Treasury Master file for the United States, the Deutsche Finanzagentur debt stock file for Germany,

the Ministry of Finance Japan monthly outstanding Japanese government bond file for Japan, the UK Debt Management Office gilt portfolio file for the United Kingdom, the Bank of Canada Government of Canada securities outstanding file for Canada, the Australian Office of Financial Management Australian Government Securities portfolio file for Australia, and the New Zealand Debt Management Office New Zealand Government Bond portfolio file for New Zealand. In each case I compute face-value-weighted residual maturity across outstanding fixed-coupon nominal bonds, excluding inflation-linked issues and Treasury bills with original maturity below twelve months, to maintain comparability with the maturity definition used in Greenwood and Vayanos (2014) and Vayanos and Vila (2021).

D.6 Weighted Average Maturity: Sources, Extension, and Caveats

The weighted average maturity of the privately held free float is constructed as follows and extended to June 2026 by the same conventions as the original panel. For the United States, the series is the par-weighted remaining maturity of marketable Treasury debt from the Monthly Statement of the Public Debt, security by security, net of Federal Reserve System Open Market Account holdings, monthly throughout, including the extension. For the partner countries, the original panel transcribes annual average-term-to-maturity values from the OECD Sovereign Borrowing Outlook and national debt-office annual reports, carried to monthly frequency; the extension chains year-end anchors from the OECD Global Debt Report 2025 and 2026 editions onto the locked levels, with chart-transcribed values accurate to roughly two tenths of a year where exact figures are not stated in prose, linearly interpolated between anchors and carried flat into 2026, the same annual-to-monthly convention as the original panel.

Two caveats attach to the partner series: partner maturities carry genuine information at annual frequency only, in the extension as in the original, and for Japan and Canada, whose free-float maturities embed a central-bank portfolio adjustment, chaining gross-debt changes onto free-float levels is a conservative approximation that understates the free-float maturity rise during quantitative tightening. Seam continuity at the splice is within one hundredth of a year for every country. The raw anchor files, the method note, and the spliced series are included in the replication package.

D.7 Central Bank Holdings and Free-Float Construction

The free-float supply variables used in the panel regressions of Section 4.3.1 are constructed by netting out central bank holdings from the gross stock of sovereign debt and by recomputing

the par-weighted maturity on the privately held free float as a debt-stock-weighted average, with the explicit construction of the par-weighted free-float maturity given in Equation (IA.24) below. The construction follows the preferred-habitat logic of Greenwood and Vayanos (2014) under which the portion of the outstanding stock held by the central bank as a non-yield-sensitive holder at scale should be subtracted from the supply variable that enters the pricing equation, with the same logic applied to the maturity-weighted supply measure on the privately held free float. Central bank holdings cannot be drawn from a single supranational source and must be constructed from the primary national disclosure for each country in my panel. I summarise the data sources here.

For the United States, Treasury holdings are taken from the Federal Reserve System Open Market Account weekly Treasury holdings file and aggregated to month-end. For Japan, Japanese government bond holdings are taken from the Bank of Japan Flow of Funds Accounts (Bank of Japan code series), which is published quarterly; I use the quarterly disclosures and forward-fill to monthly frequency. For the United Kingdom, holdings are taken from the Bank of England Asset Purchase Facility quarterly portfolio disclosure and interpolated to monthly frequency. For Canada, holdings are taken from the Bank of Canada weekly balance sheet, aggregated to month-end. For Australia, holdings are taken from the Reserve Bank of Australia monthly disclosure of Australian Government Securities held under the Bond Purchase Program. For New Zealand, holdings are taken from the Reserve Bank of New Zealand monthly disclosure of holdings under the Large Scale Asset Purchase programme. For Germany, Eurosystem holdings of German federal government securities are taken from the European Central Bank's monthly Public Sector Purchase Programme and Pandemic Emergency Purchase Programme disclosures, supplemented by the Deutsche Bundesbank's monthly balance sheet, with the German share of Eurosystem holdings aggregated to month-end and matched to the corresponding outstanding stock of Bund debt.

The central-bank-adjusted debt-to-GDP ratio is computed as $\widetilde{d}_{k,t} \equiv d_{k,t} - c_{k,t}$, where $d_{k,t}$ is the gross debt-to-GDP ratio for country k at month t and $c_{k,t}$ is the corresponding central bank holdings ratio constructed from the sources above. The par-weighted average maturity of the privately held free float is computed as

$$\widetilde{w}_{k,t} \equiv \frac{d_{k,t}\, w_{k,t} - c_{k,t}\, w^{\mathrm{cb}}_{k,t}}{d_{k,t} - c_{k,t}}, \tag{IA.24}$$

where $w_{k,t}$ is the par-weighted maturity of total outstanding sovereign debt and $w^{\mathrm{cb}}_{k,t}$ is the par-weighted maturity of central bank holdings. To preserve internal consistency, the denominator of $c_{k,t}$ is the same annual IMF International Financial Statistics nominal GDP series that

enters the gross debt ratio $d_{k,t}$, so that $\widetilde{d}_{k,t}$ is a clean subtraction within a single denominator convention.

For the United Kingdom and New Zealand, I use the par-weighted residual maturity of central bank holdings $w^{\mathrm{cb}}_{k,t}$ as published directly by each central bank: the Bank of England Asset Purchase Facility quarterly portfolio reports and the Reserve Bank of New Zealand Large Scale Asset Purchase programme disclosures. For Canada, the par-weighted maturity series is reconstructed from the Bank of Canada's weekly balance sheet decomposed by maturity bucket: I compute a par-weighted average maturity from the bucket midpoints using the bucket-level holdings as weights, on the same monthly grid as the panel. I use the published or directly reconstructed series rather than rebuilding $w^{\mathrm{cb}}_{k,t}$ from the underlying purchase flow data because the resulting series moves in the same direction as the observable central bank portfolio dynamics over my sample, and the identification of the duration channel in the free-float panel depends on the maturity of the privately held free float moving correctly with the observable maturity of central bank holdings. For the remaining four countries (United States, Japan, Germany, Australia), $w^{\mathrm{cb}}_{k,t}$ is constructed from the primary disclosure files following the same convention as for $w_{k,t}$.

The demand-side and state variables introduced in the supply analysis of Section 4.3.3 are constructed as follows. Foreign official holdings of U.S. Treasury securities are taken from the Federal Reserve Financial Accounts (Z.1), which report the stock of Treasury securities held by foreign official institutions, namely foreign central banks, foreign governments, and international organisations, at quarterly frequency; I express the stock relative to U.S. nominal GDP and carry the quarterly observation forward within the quarter. Private foreign holdings are constructed as total rest-of-world Treasury holdings at market value less the foreign official component from the same source. The distinction matters for identification: foreign official institutions hold for reserve and policy purposes and are relatively insensitive to yields, whereas private foreign investors constitute the elastic margin, an ordering documented across investor groups by Fang et al. (2026), so the official component enters the supply regressions as an inelastic demand control while the private component serves as a placebo.

The fiscal-surge state variable is the trailing twelve-month sum of the U.S. federal budget deficit from the Monthly Treasury Statement, standardised to unit variance; the general government budget balance relative to GDP for each panel country enters as a fiscal control on both legs. The short-end dollar funding factor is the cross-sectional mean of the three-month cross-currency basis across the panel currencies, with the IBOR-era quotation chained onto its risk-free-rate successor at the benchmark transition by the same overlap reconciliation applied to the term legs, and is used in the funding-versus-term-premium discrimination of

Internet Appendix Table IA.18. All series and the exact source identifiers are listed in Internet Appendix Table IA.10.

E Composition and Descriptive Properties: Additional Results

The four claims of the main text, relocated here from the opening of its results section with their epistemic standing intact, are these. The paper's claims are four. The measurement claim is that long-tenor cross-currency prices are organised at the level of currencies rather than of pairs: the bilateral, pair-specific component is identified from prices alone and is small enough to sit inside dealers' quoted spreads, and the recovery is independent of how the network is solved. The currency-level component that remains is a residual of that identification rather than a quantity the prices allocate, and allocating it is the work of the restriction below. The interpretation claim is that under the pair-specificity restriction this component equals the cross-country swap term-premium differential, a reading Sections 4.3.4 and 4.2 test against the alternatives, most prominently expected dollar funding costs, and whose one unreachable direction, a funding wedge common to every dollar pair, it bounds. The reversal claim is that the United States' relative price for supplying long-duration sovereign exposure reversed: its long-horizon duration compensation sat below its peers' before the crisis and above them since, by 56 basis points against Germany and 83 against Japan at ten years on the post-2013 window. The economics claim is that the differentials behave as term premia should, priced by government debt duration supply with preferred-habitat signs and repriced within days of genuine purchase surprises.

What prices deliver, stated at its exact strength, is a long-horizon pricing component organised at the level of currencies rather than pairs. That component is a node potential, and no quantity of price data separates a node potential into term-premium compensation and a currency-attached wedge; the reading rests on the restriction, is stress-tested over the whole null space in Table 5, and is tested against evidence from outside the swap market. The epistemic standing of each of these layers is recorded, as the introduction noted, in Table 9.

Cross-sectional structure of the recovered differentials. The first principal component carries 68 percent of the unstandardised variation across the eight differentials on the identified-complete sample, and each core pair loads heavily on the average of the other seven, with leave-own-out R^2 of 0.85 for the Swedish krona, 0.71 for the euro, 0.63 for the yen, 0.37 for the Canadian dollar, and 0.31 for sterling, against 0.28, 0.07, and 0.02 for the Norwegian krone, the New Zealand dollar, and the Australian dollar. Identified-sample means range from -44 basis points for the Norwegian pair and -42 for the yen to 0 for sterling.

Sub-period behaviour of the swap-spread differential and the pairwise trajectories. Across the three sub-periods the cross-country swap spread differential ΔSS switches sign for five of the nine pairs (EUR, JPY, CAD, CHF, SEK) and, under the bridge identity, amplifies the swap-term-premium component for the six pairs on which the reversal lives (EUR, JPY, GBP, CAD, SEK, NOK) while for AUD and NZD it enters with the opposite sign and offsets rather than reinforces it. In the decomposition figures the post-2008 drift against EUR and CAD is driven primarily by the negative ΔSS with the swap-term-premium component more stable, for USD-JPY the two contribute roughly equally, and USD-GBP follows the same sign-reversal pattern with a more muted post-GFC trajectory, dropping sharply during the GFC subsample and stabilising at a smaller negative level rather than drifting further.

E.1 Identification Robustness: Metrics, Weights, and Hub Exclusion

Three features of the reconstructed short-tenor layer are recorded here for completeness. First, the currency-level components recovered at three months and at ten years have a cross-sectional correlation of 0.767 across the ten currencies when each is averaged over the post-2008 sample, so the ranking of currencies is similar at the two ends of the curve. The time-series association is much weaker, the monthly correlations running from 0.07 for sterling to 0.65 for the euro with a median near 0.30, and the pass-through of daily changes in the three-month basis into the ten-year basis has a cross-pair median of 0.022. The similar cross-section alongside the weak time-series association is what one expects if a common cross-sectional ordering of currencies, in safety and in intermediation cost alike, coexists with largely separate variation at the two ends.

Second, the reconstructed euro-New Zealand dollar forward carries erroneous prints before 2008, concentrated in 2000 to 2003 and in September 2007, with 23 of 6,938 days showing absolute deviations from the triangular-implied cross above 50 basis points; the New Zealand triangle is therefore not used before 2008, and from 2008 its cycle residual is 0.80 basis points, in line with the other pairs. Third, the three-month basis for the Canadian dollar correlates at -0.11 with the one-year quoted basis over the post-2008 sample, against positive correlations of 0.50 to 0.65 for the other currencies; I have not been able to account for this and record it rather than omit it.

The trivial member of the solution family, term premia zero and the entire price read as friction, fits every triangle by construction; its rejection by the evidence, on the required friction magnitudes and the validation results it would leave without a source, is stated in

the main text. A placebo simulation in which the data-generating process contains pure pair-specific funding and no currency-level component shows that the method does not manufacture one silently, the implied pair-residual share under that world being sixty-seven percent of quotes against the 2.7 percent observed on the identical mean-absolute convention, the mean across the twenty-four cells of the mean absolute residual over the mean absolute quote, so the fit diagnostics reject the pure-friction world by a factor of twenty-five; the ratio of signed means, the convention Table IA.12 prints and the main text reports, averages 2.3 percent across the same cells. Replacing the minimum-norm rule with least-absolute-deviations and Chebyshev selection leaves the recovered differentials essentially unchanged, at correlations of 1.000 and 0.9997 against the baseline and an extracted supply headline of -1.067 against -1.066, and a Bayesian rule that anchors each pair's residual to its quoted half-spread rather than minimising it delivers the same term premia at correlation 0.9999 with a headline of -1.064.

A fourth rule replaces the mechanical metrics with an economically calibrated one: a generalised least squares projection that weights each pair-specific friction by the inverse squared quoted half-spread of the corresponding basis, the natural metric if the frictions are quote noise bounded by transaction costs, as Section 4.1 documents. On the seventeen-pair network the liquidity-weighted term premia correlate above 0.999 with the baseline, with mean absolute differences below 0.15 basis points, so the recovery is invariant not only across mechanical selection rules but under the transaction-cost metric the data themselves justify.

A fifth estimator differs in principle rather than in norm: the cycle readings are pure combinations of the pair frictions, so their covariance matrix identifies friction second moments directly, with the euro-dollar leg, common to every triangle, identified exactly from the cross-triangle covariances at 0.58 basis points of monthly-change volatility, and each triangle's friction-variance sum identified at 0.67 to 1.91 basis points, ordered exactly as market liquidity and below the quoted half-spreads, a second independent confirmation of the noise reading. The feasible generalised least squares extraction built on these triangle-estimated variances, in which no norm is chosen by anyone and no external input enters, reproduces the baseline term premia with correlations above 0.999 and mean differences below 0.13 basis points.

A hub-exclusion test completes the invariance battery. Dropping every dollar pair leaves a euro-hub tree in which node differentials among the remaining currencies are read directly from the euro-cross quotes, with no projection anywhere; the full-network recovery reproduces these dollar-free direct readings with correlations of 0.9993 to 0.9996 and mean differences below 0.2 basis points, and the mirrored euro-exclusion gives the same result, so the recovered structure among non-dollar currencies is numerically identical whether or not any dollar quote exists.

What hub exclusion also does not, and need not, remove is the dollar from the funding of the market: dealers funding in dollars make markets in every cross, so dollar-centred intermediation may sit behind non-dollar quotes even where no dollar leg is quoted. The observation that disciplines this concern is that intermediation priced at the currency level cannot hide in the crosses alone: a balance-sheet cost of carrying currency-j duration prices the domestic currency-j swap curve through the same dealer, so any currency-level compensation it generates resides in tp^j by the term-premium definition, content of the premium under the global-marginal-investor reading of Section 4.3.4 rather than contamination of it, while whatever is specific to an individual cross is pair-level routing cost, which the triangular algebra identifies and bounds within quoted spreads. A mechanically synthetic account, in which non-dollar crosses were constructed from two dollar legs rather than independently priced, would imply exact coincidence between quoted euro crosses and their dollar-implied counterparts; the observed coincidence is tight but not mechanical, with deviations that sit inside dealers' quoted bands and vary independently across pairs.

Pair-specific configurations of the reserve sovereigns. The pairs are not homogeneous reserve assets: the German configuration reflects Bund scarcity under the Eurosystem's purchase programmes, euro-area collateral rules, and a redenomination tail with no Treasury counterpart, while the Japanese configuration reflects yield-curve control and a free float compressed by the Bank of Japan, with the yen exhibits essentially invariant to excluding the yield-curve-control window, as documented next. The free-float construction absorbs central-bank ownership directly and the pair-level quantitative-easing indicators enter the conditioning, so the headline is a statement about each bilateral pair against the Treasury market rather than about an undifferentiated reserve-asset class.

Yield-curve control and the yen exhibits. A structural split organises the yen's exposure to the administered regime: the swap-side recovery involves no Japanese government bond anywhere, since the administered curve enters the construction only through the swap-spread-differential leg of the bridge, so the swap-side exhibits are untouched by yield-curve control while the government-side exhibits inherit it through that one leg. The variants confirm that nothing headline rests on the administered window: the yen post-2013 government-side mean is -82.1 basis points excluding the September 2016 to March 2024 window against -83.2 including it, and the swap-side mean -61.2 against -61.4; the supply coefficient excluding the window is -0.80 with a Driscoll-Kraay t-statistic of -7.6 at the Schwert-rule lag, as reported in the main text; and the carry excluding the yen entirely earns $+2.87$ percent with a t-statistic of $+2.62$ in the leave-one-out battery. Consistent with the administered-market

principle stated in the framework, the regime is excluded from structural interpretation rather than modelled.

Table IA.27 reports the executable-band identified intervals referenced in the main text: the exact interval of the recovered ten-year differential when every quoted basis is perturbed within its quoted half-spread, computed date by date from the linearity of the recovery, with measured half-spreads for the dollar pairs and the month's maximum measured half-spread assigned conservatively to the euro crosses.

A directional, executable version of the containment result uses the dealer bid and ask quotes rather than mids: crossing every leg of a triangle at the adverse side of the quoted spread, the no-arbitrage band for each triangle reading is the sum of the three half-spreads, and between 98 and 100 percent of triangle-months lie inside that executable band, with the mean crossable profit in the remaining months below one basis point and a maximum of 1.1 basis points, so no executable triangle arbitrage survives transaction costs at any point in the sample. No rule within the solution set can address the currency-level null space, which the sensitivity-frontier analysis of Section 4.2 of the main text bounds instead.

E.2 Multiple-Break Analysis of the Reversal

This subsection reports the formal break analysis behind the main text's three-period split, on the cross-sectional average ten-year benchmark differential and pair by pair (Bai-Perron mean-shift estimation, minimum segment two years). Three results organise the evidence. First, the onset: the restricted two-break solution on the average series dates the first break at October 2007 with a ninety percent interval of September to November 2007, and the unrestricted solution at August 2007, so the shift begins with the funding-crisis onset, earlier than the calendar 2008 split, which strengthens the statement that the reversal has been in place since the crisis rather than emerging recently; the pre-break mean matches the calendar pre-crisis mean within a basis point.

Second, the segmentation within the negative regime: the information criterion selects a rich partition, seven breaks, whose second date, March 2011 with interval January to May 2011, lies near, though its interval does not contain, the single-break sup-F date of October 2011 reported in the main text, the two designs asking different questions, a single most-pronounced shift against a mean-shift partition, and both placing the consolidation in 2011, and whose later dates mark level shifts entirely within negative territory; every segment mean after the first break is negative under every segmentation considered, so the sign configuration does not depend on where the partition falls, and the headline German and Japanese sub-period means under the estimated onset differ from the calendar means by about a basis point.

Third, the caveat: the restricted two-break solution places its second date at December 2021 with an interval spanning July 2021 to February 2022, a window that coincides with the LIBOR cessation and the start of the hiking cycle; the benchmark transition is chain-reconciled leg by leg, and the coincidence is disclosed so that the date is not over-read. Pair by pair, a crisis-window break is estimated for five of the eight pairs, the euro at August 2008, the yen at January 2009, the krone at October 2008, the New Zealand dollar at November 2008, and the Canadian dollar at February 2010, with the Swedish krona at February 2011 beside the consolidation; sterling and the Australian dollar, whose differentials are small, place their estimated breaks at the 2021 transition window and after.

A benchmark-substitution exercise closes the question of whether the euro reversal is a property of Bunds rather than of the euro area. Under the shared euro short-rate expectation, replacing the German ten-year yield with another euro sovereign's shifts the recovered differential by exactly the alternative-minus-Bund yield spread, so the exercise needs only the two sovereign curves on a common convention, and both legs are taken from the OECD long-term-rate par composites on one source. The sub-period means move from $+27.9$, -33.1, and

−55.8 basis points on Bunds to +34.2, +23.0, and −5.9 against France and +33.4, +2.3, and −33.7 against the Netherlands, so the sign change from positive before 2008 to negative after 2013 survives both substitutions, at −40.0 and −67.1 basis points of pre-to-post movement against −83.8 on Bunds. The French cell is the conservative bound, because the French spread carries the euro-area sovereign-credit repricing, averaging +50 basis points post-2013 against +6 pre-crisis, which raises the measured partner premium mechanically and works against the reversal; the Dutch curve, whose credit distance from Germany is small, gives back seventeen basis points of the Bund reversal. The design and the verification gate on the locked Bund means are in the replication materials.

E.3 Currency-Level Attribution of the Differentials

The differentials are invariant by construction to the global term-premium factor, as stated in the main text; what remains to be allocated is the split of each dollar-pair differential between its two sides. Table IA.33 reports, for each dollar pair and subsample, the share of the monthly variance of the recovered differential attributable to each side under the zero-sum normalisation, together with the dollar-free verification on raw quotes: for each partner pair, the change in the difference of the raw dollar-pair bases regressed on the change in the corresponding raw euro-cross difference, which contains no dollar leg. The recovered differentials satisfy additivity by construction, so the raw quotes are the informative input for this test; the pooled slope of 0.90 with an R^2 of 0.88 establishes that the partner-side variation in dollar pairs is reproduced nearly one for one in quotes with no dollar leg, with the shortfall from one consistent with attenuation from quote-level measurement error.

E.4 Empirical Composition of the Cross-Currency Basis: Detailed Statistics

This subsection provides the complete pair-and-tenor statistics underlying the summary discussion in Section 4.1 of the main text. Table IA.12 reports, for each of the eight USD bilateral pairs at two-, five-, and ten-year tenors: the signed sample mean of the bilateral basis $x_{t,N}^{j/USD}$, the cross-country swap term premium $ctp_{t,N}^{j,USD}$, and the pair-specific friction $\pi_{t,N}^{j/USD}$ in basis points, with $x = ctp + \pi$ at every date by construction; the share $|\pi|/|x|$ in percent; and the time-series correlation $\rho(\Delta SS, ctp^{swap})$ between the two additive bridge-identity components of the cross-country government bond term premium. Internet Appendix Figure IA.2 plots the time-series scatter of $\pi^{j/USD}$ against $ctp^{j,USD}$ at the ten-year tenor for each of the eight

pairs, with the corresponding sample correlation reported in each panel title.

F Drivers: Additional Results

F.1 Drivers Panel: Column-by-Column Discussion

Columns (2) through (5) progressively enrich the specification by isolating one economically distinct conditioning channel at a time; the sign and significance pattern of the quantity and maturity coefficients across the columns is stated in the main text and not repeated here. Column (2) adds the country-specific quantitative easing indicators $\mathrm{QE}_t^{\mathrm{US}}$ and QE_t^j, constructed symmetrically across the nine countries from each central bank's officially declared quantitative easing programme dates: the indicator for country k takes the value one in months during which the country's central bank operated an officially announced asset-purchase programme, identified from the central bank's own programme announcement and termination disclosures, and zero otherwise. Because the privately held free-float construction $\widetilde{d}_{k,t}$ already nets out the central-bank-absorbed quantity of sovereign duration, the QE indicators in this specification do not capture the asset-purchase quantity channel itself, which is absorbed by the free-float supply variable; they instead capture the residual expectational content of an active asset-purchase regime, the perceived-risk and policy-stance signal that an officially declared programme conveys to marginal private holders over and above the duration absorbed at the date of measurement.

The two indicators enter with sharply contrasting signs, with $\mathrm{QE}_t^{\mathrm{US}}$ widening and QE_t^j compressing the bilateral, consistent with the expectational reading that an active U.S. programme signals additional perceived scarcity of U.S. duration in private hands beyond the contemporaneous quantity reduction, and the analogous signal on the partner-country side.

Column (3) adds the standardised He et al. (2017) intermediary capital ratio HKM_t^z and the interaction $\mathrm{HKM}_t^z \cdot \mathrm{QE}_t^{\mathrm{US}}$. The intermediary capital ratio enters as a conditioning state variable on the demand side rather than as an alternative driver of the bilateral cross-country term premium: in the preferred-habitat reading, the marginal yield-sensitive private holder is the dealer whose balance-sheet capacity governs the level of compensation it requires per unit of duration absorbed and the depth of its participation in the market for scarce U.S. duration. The HKM_t^z measure of intermediary equity capital adequacy scales this state, with higher intermediary capital corresponding to less constrained intermediaries with greater risk-bearing capacity. The positive level coefficient on HKM_t^z in column (3) maps onto this reading directly: when intermediaries are well-capitalised, they bid up the price of U.S. Treasuries relative to

peer-sovereign benchmarks at a given level of privately held free-float supply, widening the cross-country government bond term premium for U.S. Treasuries.

The interaction $\text{HKM}_t^z \cdot \text{QE}_t^{\text{US}}$ enters with a robustly positive coefficient at column (3) and remains so throughout columns (4) and (5), indicating that the strength of the intermediary-level effect is amplified during U.S. quantitative easing months. Section 4.3.2 below shows that the bridge decomposition allocates this interaction predominantly to the cross-country swap term premium rather than to the cross-country swap spread differential, consistent with a risk-premium-expectations reading: well-capitalised intermediaries operating during active U.S. asset-purchase regimes compress expected risk premia on U.S. duration through the cross-country swap term premium channel rather than through the within-country swap-spread channel that the negative-swap-spread literature (Klingler and Sundaresan, 2019; Jermann, 2020; Boyarchenko et al., 2018; Hanson et al., 2022) characterises for U.S. supply quantity effects. Internet Appendix Section F.4.5 reports a comparison of HKM_t^z with the broker-dealer leverage factor of Adrian et al. (2014) as alternative intermediary measures, and the capital-adequacy reading I adopt here is robust to the alternative.

Column (4) adds the country-side macro factors, comprising the year-on-year real GDP growth rate, the three-month interbank rate, the term spread between the ten-year and three-month yields, and the unemployment rate, each entered separately on the U.S. and partner-country sides and standardised on the pooled cross-country sample over the full January 2002 to June 2026 window to zero pooled mean and unit pooled standard deviation, so that a coefficient on $\text{ir3m}_t^{j,z}$ records the response of the bilateral cross-country government bond term premium in percentage points to a one-standard-deviation pooled movement in the partner-country three-month rate. The four free-float supply and duration coefficients retain their predicted signs and significance once the macro factors are included, and the within-pair R^2 rises from 0.473 at column (3) to 0.646 at column (4), indicating that the macro conditioning environment absorbs a substantial share of the residual time-series variation in the cross-country government bond term premium that is left unexplained by the supply, intermediary, and asset-purchase channels alone.

Column (5) adds the standardised broad dollar index, the standardised ICE BofA BBB credit spread, the standardised MOVE index of Treasury option-implied volatility, and the standardised VIX index as macro-financial conditioning factors. The four preferred-habitat coefficients retain their signs and significance, the intermediary capital interaction with QE_t^{US} remains statistically significant, and the broad dollar index loads negatively.

Section 4.3.3 develops the supply mechanism in full on the sample extended through June 2026, adding the reserve-demand controls and the state-dependence analysis, and absorbs the

supply discussion of this table; the conditioning channels of columns (2) through (5) and the bridge decomposition of Section 4.3.2 continue to operate on the five-column specification reported here.

F.2 The Elasticity Interpretation of the Quantity Coefficients

The empirical asymmetry between the U.S. and partner-country quantity coefficients in column (1), with the U.S. log free-float coefficient of -1.051 approximately twice the absolute magnitude of the partner-country coefficient of $+0.521$, is itself a signature of differential supply elasticity: in a symmetric setting in which both sovereigns optimised debt issuance against term-premium signals to the same degree, the two coefficients would be attenuated symmetrically, while the larger U.S. coefficient is consistent with U.S. supply being more inelastic to bond-market pricing than partner-country supply over my sample. Georgievska et al. (2025) provides direct evidence that U.S. Treasury issuance over the past two decades is approximately unresponsive to term-premium signals, in the sense that auction quantities and maturity composition do not adjust in the direction an optimising issuer would choose when term premia compress or widen; under such an exogenous-supply regime, the price response to changes in privately held U.S. Treasury quantities is the full unmediated preferred-habitat elasticity rather than the residual after supply accommodation, and the U.S. quantity coefficient is interpretable as a quantity-clearing price response rather than a partial-equilibrium response with Treasury offset.

I emphasise that reading the coefficient ratio as a relative demand elasticity requires the symmetric-issuance benchmark stated here; absent that benchmark the ratio is a reduced-form asymmetry, and none of the paper's conclusions depend on the structural reading. The current-inflation control on both legs enters the conditioned columns of Internet Appendix Table IA.1, where the fully conditioned supply loading is -1.216 ($t = -5.78$) and both inflation terms carry the signs an expectations channel implies, U.S. inflation raising and partner inflation lowering the differential, at marginal significance, with the series' coverage as recorded in Table IA.13; the expectations-based variant, using ten-year inflation swaps, is reported in Internet Appendix Section G.2, together with the coefficient-robustness frontier, the mean-group estimator, and the twelve-month first-difference specification, whose estimates are stated once there and not repeated here. Internet Appendix Section G.1 develops a three-currency preferred-habitat framework in which the supply loadings, their asymmetry as a function of relative issuance elasticity, the node structure of currency-level funding, and the transaction-cost containment of the pair frictions all arise in a single equilibrium. This interpretation

refines the Du et al. (2026) reading of the U.S. supply coefficient and aligns with the canonical Greenwood-Vayanos and Vayanos-Vila preferred-habitat framework: the compression of the U.S. Treasury term-premium advantage with rising U.S. supply reflects an environment in which the U.S. supply process is approximately exogenous to price signals, consistent with the standard inelastic-supply assumption maintained by Greenwood and Vayanos (2014) and Vayanos and Vila (2021).

The smaller positive coefficient on partner-country privately held debt is consistent with foreign sovereign issuers exhibiting a higher, although still limited, degree of responsiveness to domestic yield signals, which partially absorbs demand shocks through quantity adjustment. The reading is reinforced by three features of the panel design: the pair fixed effects absorb time-invariant cross-pair institutional differences and identify the supply coefficients from within-pair time-series variation only; the macro-financial factors of columns (4) and (5) absorb the broad demand-side shifts captured by the country-side macro environment, namely U.S. and partner-country real GDP growth, three-month interbank rates, term spreads, and unemployment, the broad dollar level, credit spreads, and volatility risk indicators, with the U.S. supply coefficient remaining large at -1.134 in column (5), of comparable magnitude to the column (1) baseline of -1.051 and confirming that the supply elasticity is robust to the inclusion of the macro-financial conditioning environment; and the intermediary capital ratio enters as a conditioning state variable on the demand-side without absorbing the supply coefficient. The supply coefficient also scales monotonically with tenor, at -0.68, -0.89, and -1.05 across the two-, five-, and ten-year horizons in the baseline specification and at -0.85, -1.02, and -1.13 in the fully conditioned one (Internet Appendix Section F.4.2), the pattern the duration channel of preferred-habitat models implies (Vayanos and Vila, 2021; Greenwood and Vayanos, 2014), in which a given supply shock commands more risk compensation per unit of duration at longer horizons.

F.3 Channel Decomposition: Discussion of Each Driver

A key methodological by-product of the decomposition is that the bridge identity reveals economically meaningful drivers of the cross-country term premium that the regression with $\widehat{ctp}_N^{gov,j,USD}$ as the dependent variable alone would obscure. A regressor whose net effect on the cross-country government bond term premium reads near zero may still be a substantive driver of one of the two bridge components, with a coefficient on the cross-country swap term premium that is offset by an opposite-signed coefficient on the cross-country swap spread differential, or with a contribution that lives entirely on one side of the bridge. Under the

bridge identity, the regression coefficient on $\widehat{ctp}_N^{gov,j,USD}$ for any regressor equals the sum of the corresponding coefficients on $\widehat{ctp}_N^{swap,j,USD}$ and $\Delta SS_N^{USD,j}$, so a small or insignificant net coefficient on the cross-country government bond term premium can mask economically meaningful and oppositely signed coefficients on the two bridge components. Several variables that read as statistically insignificant in the regression with $\widehat{ctp}_N^{gov,j,USD}$ as the dependent variable in Table 1 are statistically significant drivers of either $\widehat{ctp}_N^{swap,j,USD}$ or $\Delta SS_N^{USD,j}$ in Table IA.8, picking up the macro and risk-premium expectations content that operates through the cross-country swap term premium channel rather than through the cross-country swap spread differential. Without the decomposition into the two bridge components that my triangulation enables, this content would be invisible to the analyst.

The decomposition reveals three patterns on the supply and asset-purchase channels. First, the reduced-form covariance of the U.S. log free-float variable with the cross-country government bond term premium is concentrated in, though not confined to, the cross-country swap spread differential: a formal test of equality of the component loadings rejects, with the difference between the swap-spread and basis loadings at -0.66 with a t-statistic of -7.9: the column (1) coefficient of -1.051 in the regression with $\widehat{ctp}_{10}^{gov,j,USD}$ as the dependent variable decomposes into -0.194 in the corresponding regression with the quoted basis $x_{10}^{j/USD}$ as the dependent variable and -0.857 in the regression with $\Delta SS_{10}^{USD,j}$ as the dependent variable, an 81.5 percent ΔSS share, with the share remaining at 74.7 percent in the macro-augmented column (5). The partner-country log free-float coefficient carries the majority of its loading in the swap-spread component in the first three columns, with ΔSS shares of 67 to 71 percent, moderating to 43 percent once the macro and fiscal controls enter.

Second, the country-specific quantitative easing indicators load asymmetrically across the bridge channels: the partner-country quantitative easing indicator loads predominantly on the cross-country swap term premium, with 70 percent of its column (5) government-bond coefficient living in the quoted-basis regression, while the U.S. quantitative easing indicator loads mainly on the swap-spread component on the nine-currency sample, with 62 percent of its column (5) government-bond coefficient in the $\Delta SS_{10}^{USD,j}$ regression and 38 percent in the quoted-basis component, both components individually significant. The positive sign on $\mathrm{QE}^{\mathrm{US}}$ in Panel B of Table IA.8 remains consistent with relief of the prior post-2008 compression of SS_{US} during U.S. asset-purchase regimes, transmitting to the bilateral $\Delta SS^{USD,j} \equiv SS_{\mathrm{US}} - SS_j$, while non-USD asset-purchase regimes operate predominantly through the cross-country swap term premium.

Third, the intermediary-capital interaction with U.S. quantitative easing $\mathrm{HKM}_t^z \cdot \mathrm{QE}_t^{\mathrm{US}}$ divides its column (5) coefficient across the two components, 55.1 percent on the quoted-basis side,

with only that side individually significant ($t = 2.90$ against 1.23 on the swap-spread side), consistent with the risk-premium-expectations interpretation flagged in Section 4.3.1: the intermediary-balance-sheet amplification of the U.S. asset-purchase stance operates mainly through the cross-country swap term premium channel and the expectations of risk-premium compression it embeds, rather than through the U.S. swap spread.

The macro and macro-financial factors yield a complementary set of patterns that the decomposition makes visible. The partner-country three-month rate is the largest macro driver of the cross-country government bond term premium and concentrates 78.4 percent of its column (5) coefficient on the cross-country swap term premium, consistent with monetary-policy expectations on the partner-country side operating through expected-short-rate compensation embedded in the cross-country swap term premium rather than through the cross-country swap spread differential. The U.S. three-month rate enters significantly on the cross-country government bond term premium at column (5) and loads across both bridge channels, with shares of 59.3 percent on the cross-country swap term premium and 40.7 percent on $\Delta SS_{10}^{USD,j}$. The standardised broad dollar index loads on the cross-country swap term premium alone, at -0.078 ($t = -3.18$) against an insignificant opposite-signed $+0.033$ on the swap-spread component, a configuration in which a percentage attribution is uninformative, consistent with dollar-cycle conditioning of safe-asset risk premia operating through the cross-country swap term premium.

Four additional variables, the standardised intermediary capital ratio HKM_t^z entered as a level, U.S. and partner-country unemployment rates, and the standardised ICE BofA BBB credit spread, enter Table 1 with statistically insignificant coefficients on the cross-country government bond term premium in column (5) but enter the bridge decomposition with individually significant coefficients reported in Table IA.8, the first three loading on the cross-country swap term premium component (Panel A) and the credit-spread on the cross-country swap spread differential component (Panel B).

These variables capture macro and risk-premium expectations content that the decomposition allocates to one of the two bridge components and that the regression on $\widehat{ctp}_{10}^{gov,j,USD}$ alone would obscure. The cross-country swap term premium in particular emerges as the principal absorber of macro-financial expectations content, with the intermediary-capital interaction with U.S. asset purchases, both unemployment indicators, both quantitative easing indicators, the partner-country three-month rate, and the standardised broad dollar index all loading significantly on it; the cross-country swap spread differential is the principal locus of the U.S.-side preferred-habitat supply mechanism. The localisation of the U.S. supply mechanism and the U.S. asset-purchase quantity channel in $\Delta SS_N^{USD,j}$ is the cross-country

counterpart of the within-country negative-swap-spreads literature (Klingler and Sundaresan, 2019; Jermann, 2020; Boyarchenko et al., 2018; Hanson et al., 2022) and complements the dollar-prediction decomposition reported in Internet Appendix Section I.1, in which the dollar-prediction content of $\overline{ctp}^{gov,USD}$ also lives predominantly in the $\Delta SS_N^{USD,j}$ component of the bridge identity.

F.4 Drivers of Cross-Country Term Premia

The supply result is insensitive to the recovered pair-specific residuals. The quoted construction of the dependent variable is invariant to the residuals by definition, and rescaling every recovered residual by factors of two, five, and ten moves the extracted supply-only coefficient from -1.066 to -1.059, -1.037, and -1.000, with t-statistics of -9.8, -9.6, and -9.1, so even an order-of-magnitude understatement of the pair-specific frictions leaves every conclusion intact.

Holding foreign official reserve demand fixed, with the stock of Treasuries held by foreign official institutions from the Federal Reserve financial accounts expressed relative to GDP as a demand control, the supply-only coefficient is -0.930 ($t = -4.63$), so the relationship is not an artefact of omitted reserve demand; under the full conditioning set the supply coefficient attenuates to -0.438 ($t = -2.05$) while the official-demand control enters at -3.545 ($t = -3.02$) with the sign the reserve-currency channel implies, so part of the conditioned loading is shared with official demand. An effective float that nets out both price-insensitive holders, gross debt net of domestic central bank and foreign official holdings, gives -1.052 ($t = -10.71$), indistinguishable from the free-float baseline; the free-float specification with the demand control is retained as primary since it lets the data assign the demand coefficient rather than imposing a unit offset.

This subsection extends the five-column free-float supply specification of Table 1 of the main text along several dimensions. Table IA.21 reports the inference battery for the common U.S. regressor referenced in the main text: the point estimate is invariant across the rows because the estimator is unchanged and only the covariance treatment varies, and the two-way clustering treats the month dimension, in which the U.S. regressor is common, and the pair dimension jointly, Driscoll-Kraay standard errors at the Schwert-rule lag in the baseline. The full-pipeline bootstrap referenced in the main text, two hundred replicates with every quote perturbed within its quoted half-spread and the recovery, the dependent variable, and the regression rebuilt on each, gives the U.S. coefficient a bootstrap standard deviation of 0.0014 around an unbiased mean, one percent of the Driscoll-Kraay standard error at the

Schwert-rule lag.

This subsection also reports the gross-debt robustness, the term structure at the two- and five-year tenors, the pre-pandemic and post-pandemic stability, the comparison between the He-Kelly-Manela intermediary capital ratio and the Adrian-Etula-Muir broker-dealer leverage factor, and the official-demand flow margin on the valuation-adjusted series. The cross-country swap term premium versus cross-country swap spread differential channel decomposition, including the channel decomposition for the intermediary capital and macro-financial drivers, is presented in Section 4.3.2 of the main text rather than in this appendix subsection.

The three checks that close the main text's supply battery are recorded here. First, a reverse-timing placebo: with the current U.S. free-float stock and its twelve-month-ahead value entered together, the current stock retains its coefficient at -0.65 with a t-statistic of -3.0 while the future value is indistinguishable from zero, against the pattern reverse causality would produce. Second, a moving-block bootstrap of the supply regression places the ninety-five percent interval for the U.S. coefficient at $[-0.85, -0.08]$, excluding zero. Third, the bilateral construction carries information beyond the familiar within-U.S. swap-spread relation: partner-country swap spreads load on their own free floats weakly or not at all, while the triangulated differential loads significantly, so the differential is not a restatement of the U.S. leg alone.

Figure IA.8 provides visual evidence on the supply channel.

F.4.1 Robustness: Gross-Debt Specification

The supply measure that the preferred-habitat framework points to is the duration that the marginal yield-sensitive private holder must absorb, not the gross stock of outstanding sovereign debt. Central bank holdings sit on the books of an institution that is not yield-sensitive in the marginal-pricing sense and that is therefore not the marginal price-setter for sovereign duration; they enter the gross stock but do not enter the demand-supply equilibrium that determines the term premium. When the post-2008 era of large-scale asset purchases drove central bank holdings of sovereign debt to historically unprecedented levels in the United States, Germany, Japan, and the United Kingdom, the gross debt stock and the privately held free float decoupled materially, with central banks absorbing a substantial share of the marginal issuance flow. Constructing supply on the privately held free-float basis $\widetilde{d}_{k,t} \equiv d_{k,t} - c_{k,t}$ recovers the supply measure that the marginal private holder actually faces, and the same logic applies to the par-weighted maturity construction $\widetilde{w}_{k,t}$, which weights the privately held maturity profile rather than the total maturity profile.

Table IA.34 re-estimates the five-column specification of Table 1 with log gross sovereign debt-to-GDP on both sides in place of the privately held free-float measure, with no adjustment for central bank holdings. The maturity columns are omitted, since real gross average-maturity histories for the partner sovereigns are not part of my dataset and the free-float maturity is not a valid stand-in for a gross specification; because no maturity series is required, the Swedish and Norwegian pairs enter every column and the panel covers eight USD currency pairs with 2,120 pair-month observations on the main-supply specification.

The two quantity coefficients carry the theoretically predicted signs in all five columns. In the column (1) main-supply specification, $\log(d/\text{GDP})_{\text{US}} = -1.044$ ($t = -10.03$) and $\log(d/\text{GDP})_j = +0.370$ ($t = +5.22$) on the eight-pair panel of 2,120 observations, and the fully conditioned column (5) coefficient is -1.222 ($t = -6.66$). The U.S. log-debt coefficient under gross debt is close in absolute magnitude to the corresponding free-float coefficient of -1.051 in Table 1, and the partner-country log-debt coefficient is somewhat smaller in absolute magnitude than the corresponding free-float coefficient of $+0.521$. I do not interpret the direction of these magnitude shifts as a structural bias signature, since their joint algebraic sign is not pinned down by the substitution from $\widetilde{d}_{k,t}$ to $d_{k,t}$ alone and depends on the joint covariance structure of the omitted central-bank-holdings variables across pairs and time. I report the gross-debt specification here as a robustness exercise that demonstrates the qualitative pattern is unchanged under the cruder supply measure, and adopt the privately held free-float construction as my preferred specification on theoretical grounds, since the preferred-habitat mechanism operates on the supply that the marginal private holder must absorb rather than

on the gross stock that includes central-bank-absorbed supply.

The broader pattern across the five columns is preserved: the QE indicators enter with the same sharply contrasting signs across columns (2) through (5), with QE^{US} ranging between +0.170 and +0.285 (t-statistics between +2.79 and +5.97) and QE^{j} between −0.245 and −0.340 (t-statistics between −3.12 and −5.86); the intermediary capital interaction with QE^{US} enters significantly across columns (3) through (5), with the column (5) interaction coefficient at +0.196 ($t = +4.58$); and the macro factors of column (4) and the standardised volatility/credit/dollar factors of column (5) enter with the same signs and significance pattern as in the free-float specification. The column (5) U.S. log-debt coefficient under gross debt is −1.281 ($t = -6.88$), above the free-float column (5) coefficient of −1.134 in absolute magnitude once the macro factors and the standardised volatility/credit/dollar factors are jointly included.

F.4.2 Term Structure Robustness at 2 and 5 Years

Table IA.11 reports the five-column free-float panel specification of Table 1 estimated at the two-year (Panel A) and five-year (Panel B) tenors; the ten-year tenor is reported in the main-text table and is omitted here to avoid duplication. In column (1), the U.S. log free-float debt coefficient is −0.682 ($t = -7.22$) at the two-year tenor and −0.891 ($t = -8.38$) at the five-year tenor, scaling monotonically toward the ten-year value of −1.051. The partner-country log free-float debt coefficient is +0.238 ($t = +2.55$) at two years and +0.401 ($t = +4.63$) at five years, also scaling monotonically toward the ten-year value of +0.521. The absolute magnitudes of the log free-float debt coefficients scale monotonically with maturity, consistent with the preferred-habitat prediction that longer-duration assets are more sensitive to supply shocks because the risk compensation per unit of duration is larger at longer tenors and the marginal holder's substitution elasticity is lower further out the curve. The weighted-average maturity coefficients do not exhibit a comparable monotonic scaling across tenors and I do not interpret them through the same scaling lens.

The conditioning patterns in columns (2) through (5) parallel the ten-year main specification: QE indicators enter with opposing signs across the two countries, the intermediary capital channel contributes a small but generally significant level effect at column (3) and an interaction effect with U.S. quantitative easing that strengthens once macro and standardised controls are added in columns (4) and (5), and the macro factors of column (4) and the standardised volatility/credit/dollar factors of column (5) retain their predicted signs at both tenors. The column (5) U.S. log free-float coefficient at the two-year tenor is −0.848 ($t = -3.18$) and

at the five-year tenor is -1.022 ($t = -4.20$), scaling monotonically toward the column (5) ten-year value of -1.134, so the maturity scaling of the supply coefficient holds in the fully conditioned specification as well as in the baseline.

F.4.3 Pre-COVID Stability and the Post-Pandemic Structural Break

Table IA.35 reports the column (1) free-float preferred-habitat panel specification of Section 4.3.1 estimated separately on three nested cumulative windows that successively extend the sample end date through the COVID-19 pandemic (January 2002 to November 2018, January 2002 to February 2020, and January 2002 to December 2021) and on the post-pandemic window (March 2020 to June 2026), alongside the full-sample anchor for reference. The nested-window estimates are stable in sign, with the U.S. log free-float coefficient between -1.09 and -1.22 across the three windows. The post-pandemic window exhibits a structural break: the U.S. log free-float debt coefficient flips sign relative to the nested-window estimates, at $+1.60$ with a t-statistic of $+1.9$, marginally significant at the ten percent level on the window extended through the 2024 to 2026 quantitative tightening period, while the partner-country log free-float debt coefficient retains its predicted sign and statistical significance.

The reversal of the U.S. supply elasticity in the post-pandemic window is consistent with the conjunction of exceptional pandemic-era Treasury issuance, the post-2022 quantitative tightening regime, and ongoing reallocation of foreign official sector holdings away from U.S. Treasuries; I do not develop a structural account of the break here, which I leave to subsequent work, but note it as a salient empirical regularity; the standalone January 2022 to June 2026 window returns a strongly negative loading on its own, so the relation resumes after the pandemic window rather than ending with it, and the interpretive claims are conditioned on that window rather than restricted to the pre-pandemic sample.

F.4.4 Inflation Expectations as Controls

Column (5) of Table 1 conditions on realised year-on-year CPI inflation on both sides of each bilateral. Realised inflation is, however, an imperfect proxy for the inflation compensation embedded in nominal term premia, which responds to expected inflation and inflation risk over the tenor of the claim rather than to the trailing price index. Table IA.28 therefore re-estimates the column (5) specification with the U.S. realised inflation control replaced by the ten-year expected inflation series of the Federal Reserve Bank of Cleveland (Haubrich et al., 2012), which combines inflation swap prices, survey forecasts, and nominal yields in a term-structure model of expected inflation; the partner-country realised inflation control is retained, since comparable model-based long-horizon expectation series are not published for the six partner economies, and I note that a fully symmetric survey-based variant using long-run consensus expectations across the seven economies is a natural extension as those data are assembled. The table reports the specification for the cross-country government bond term premium and both bridge components on the identical sample, so the coefficients satisfy the bridge identity exactly. The U.S. free-float supply coefficient is -1.169 ($t = -5.02$), essentially unchanged from the realised-inflation specification, and the expected-inflation control itself enters negatively (-0.049, $t = -2.21$); the supply, asset-purchase, and rate-factor loadings are likewise stable across the two inflation treatments, indicating that the column (5) conclusions do not depend on whether inflation is measured in realised or expected terms.

F.4.5 Alternative Intermediary Constraint Measure: Broker-Dealer Leverage

Table IA.36 reports regressions comparing the He et al. (2017) intermediary capital ratio (HKM) against the Adrian et al. (2014) broker-dealer leverage factor (AEM) as drivers of the cross-country government bond term premium and its two components in the bridge from the swap to the cross-country government bond term premium. The HKM and AEM factors capture distinct dimensions of intermediary health: HKM measures the equity capital adequacy of primary dealers, while AEM measures flow-of-funds-based leverage shocks for securities broker-dealers that reflect risk appetite and funding conditions. The AEM factor is re-derived from the Federal Reserve Z.1 accounts through 2026-Q1, log broker-dealer leverage in levels, at quarterly frequency expanded to monthly via step function within quarter. Both factors are standardised on the full non-missing sample and entered alongside the country-side macro factors of column (4) and the standardised volatility, credit, and dollar factors of column (5) of Table 1, on a sample-matched panel with currency-pair fixed effects and Driscoll and Kraay (1998) kernel standard errors.

In the joint specification neither factor is distinguishable from zero in the pooled-panel $\widehat{ctp}^{gov}$ regression, at +0.04 ($t = +0.6$) for the capital ratio and +0.32 ($t = +1.3$) for the leverage factor, under the full macro, fiscal, and financial conditioning on a common complete-case sample that the Scandinavian coverage restricts to 125 pair-months; the discrimination between the two factors lives in the swap component, where the capital ratio loads at +0.065 ($t = +6.3$) while the leverage factor is indistinguishable from zero (+0.036, $t = +1.3$), and entered alone the leverage factor attracts the swap loading the joint specification attributes to the capital ratio (+0.170, $t = +3.2$). This table enters HKM as a main effect alongside AEM rather than the main effect plus quantitative-easing interaction of the headline panel, so the coefficients are reduced-form averages of the headline state-dependent total effect rather than its main-effect counterpart. I therefore retain HKM as the primary intermediary constraint measure throughout the paper, with the measured interpretation that the intermediary channel on the cross-country term premium is better captured by capital adequacy than by leverage appetite, the two factors correlated and the discrimination resting on the swap component.

F.4.6 Official-Demand Flows: The Valuation-Adjusted Margin

This subsubsection reports the design and full results behind the main text's official-demand flow sentence. The design was pre-registered as descriptive support with a positive ex-ante sign, official absorption of United States duration lowering the United States term premium and raising the differential, and was estimated only after the inputs had been retrieved, under the dual-computation lock. The data are the valuation-adjusted foreign official Treasury positions and flows of Bertaut and Judson (2014), chained across the authors' two release files at June 2011, where the two agree to 0.002 percent; the July 2021 release is the authors' final update and the series ends in December 2020, so the panel here runs to that date. The specification is the baseline free-float supply panel of Table 1 augmented with the twelve-month adjusted net official Treasury flow relative to GDP, six dollar pairs, pair fixed effects, and Driscoll and Kraay (1998) standard errors at the Schwert-rule lag. Table IA.37 reports the results.

The flow margin enters at +3.309 ($t = 2.30$), the registered sign at the five percent level, with the United States free-float coefficient at −1.054 ($t = -10.95$); the corresponding stock margin enters wrong-signed (−1.865, $t = -2.50$), and the same stock failure appears on the Financial Accounts official-holdings ratio over the full sample (−0.025, $t = -0.03$), the two stock series correlating 0.994 on their common span, so the stock result is a feature of stocks rather than of either source. On the Financial Accounts flow measure, the pre-stated fallback that extends past the 2020 endpoint of the valuation-adjusted series, the twelve-month change

carries the registered sign short of the ten percent level (+1.727, $t = 1.39$). I read the margin as flows pricing the differential where slow-moving stocks do not, the same flow-frequency pattern the announcement evidence carries, and the result remains descriptive: the flow series is not an instrument, and no causal claim rests on it.

F.5 Macro-Augmented Specification: Factors and Data Sources

The macro-augmented specifications reported in columns (4) and (5) of Table 1 extend the column (3) intermediary specification along two complementary dimensions. Throughout the panel, each variable in the conditioning set is standardised on the cross-section that is naturally available for that variable: country-specific macro series, which exist separately for each of the nine countries, are pooled across countries to deliver coefficients that are comparable across the U.S. and partner-country sides; single time-series variables, which take a common value across pairs at each date, are standardised on their own time series.

Column (4) adds the country-side macro factors, namely the year-on-year real GDP growth rate $\text{rGDP}_{k,t}$, the three-month interbank rate $\text{ir3m}_{k,t}$, the term spread $\text{TS}_{k,t} \equiv \text{ir10y}_{k,t} - \text{ir3m}_{k,t}$, where $\text{ir10y}_{k,t}$ is the ten-year sovereign yield, and the unemployment rate $\text{Unemp}_{k,t}$, each entered separately on the U.S. ($k = \text{US}$) and partner-country ($k = j$) sides. The four series are sourced from Bloomberg and the Federal Reserve Economic Database (FRED) at quarterly frequency for real GDP and unemployment and at monthly frequency for the rate variables, and forward-filled to monthly frequency for the panel. I standardise each macro variable on the pooled cross-country sample, computing a single mean and a single standard deviation from the union of all nine countries' observations over January 2002 to June 2026 and applying these pooled moments to each country's series, so that a unit movement in the standardised regressor corresponds to a one-pooled-standard-deviation movement in the underlying macro variable on a scale that is comparable across the two sides of each bilateral and across the four macro dimensions. The standardised series are denoted $\text{rGDP}_t^{k,z}$, $\text{ir3m}_t^{k,z}$, $\text{TS}_t^{k,z}$, and $\text{Unemp}_t^{k,z}$, with the country indicator preserved in the superscript and the z-suffix indicating pooled standardisation.

Column (5) adds the U.S. and partner-country year-on-year CPI inflation rates, sourced from Bloomberg at monthly frequency with the euro-area HICP series proxying the euro side and standardised on the pooled cross-country sample in the same manner as the other country-side macro factors, together with the standardised broad dollar index, the standardised ICE BofA BBB credit spread, the standardised MOVE index of Treasury option-implied volatility, and the standardised VIX index of equity option-implied volatility, the latter four sourced

from Bloomberg and FRED and standardised to zero mean and unit variance over the full January 2002 to June 2026 sample window. Internet Appendix Section F.4.4 reports the column (5) specification with the U.S. realised inflation control replaced by ten-year expected inflation.

The standardised He et al. (2017) intermediary capital ratio HKM_t^z that enters columns (3) through (5) uses the published series, accessed from the He-Kelly-Manela data webpage in its 27 June 2025 vintage ending May 2025, which sets the sample end for those columns, standardised to zero mean and unit variance over its available window. The country-specific quantitative easing indicators QE_t^{US} and QE_t^j that enter columns (2) through (5) are constructed symmetrically across the nine countries from each central bank's officially declared quantitative easing programme dates: the indicator for country k takes the value one in months during which the country's central bank operated an officially announced asset-purchase programme, identified from the central bank's own programme announcement and termination disclosures, and zero otherwise. The full schedule of programme date windows used to construct each country's indicator is given in Table IA.38; the central bank holdings series $c_{k,t}$ used in the privately held free-float construction is itself constructed from each country's primary central bank disclosure, as detailed in Section D.7.

G Supply: Additional Results

G.1 A Three-Currency Preferred-Habitat Framework

This section develops a compact equilibrium framework that organises three of the paper's empirical designs, namely the supply regressions and their coefficient asymmetry, the sensitivity-frontier bounds, and the transaction-cost containment of the pair-specific component. The framework combines preferred-habitat bond markets in the tradition of Vayanos and Vila (2021) and Greenwood and Vayanos (2014) with an intermediated cross-currency swap network.

There are three currencies $c \in \{u, e, f\}$. In each currency, long-duration government bonds carry interest-rate risk with variance σ^2 per unit of duration. Preferred-habitat investors in market c demand duration

$$q_c \;=\; a_c + b_c\, tp_c, \qquad b_c \geq 0, \tag{IA.25}$$

increasing in the term premium tp_c, and the government issues duration $s_c = \bar{s}_c - e_c\, tp_c$ with issuance elasticity $e_c \geq 0$, so that $\bar{s}_c$ is the exogenous component of government debt duration supply. Global arbitrageurs with mean-variance preferences and risk-bearing capacity $1/\gamma$ absorb the residual $z_c = s_c - q_c$ in each market, which requires compensation $tp_c = \gamma\sigma^2 z_c$. Market clearing gives

$$tp_c \;=\; \frac{\gamma\sigma^2\,(\bar{s}_c - a_c)}{1 + \gamma\sigma^2\,(b_c + e_c)}, \qquad \frac{\partial\, tp_c}{\partial\, \bar{s}_c} \;=\; \frac{\gamma\sigma^2}{1 + \gamma\sigma^2\,(b_c + e_c)} \;>\; 0. \tag{IA.26}$$

Equation (IA.26) delivers the supply predictions. The cross-country term premium $ctp^{u,f} = tp_f - tp_u$ loads negatively on exogenous U.S. government debt duration supply and positively on foreign supply, and the magnitude of each loading is decreasing in the sum of the local demand and issuance elasticities. Under symmetric habitat demand, $b_u = b_f$, the ratio of the two reduced-form coefficients identifies the relative issuance elasticity: a U.S. coefficient exceeding the partner coefficient in absolute value corresponds to $e_u < e_f$, that is, to relatively inelastic U.S. issuance. This is precisely the symmetric-issuance benchmark under which the coefficient asymmetry of Section 4.3.3 admits a structural reading, and the framework makes explicit that the reading is conditional on that benchmark.

The same arbitrageurs intermediate the cross-currency swap network. A bilateral position in pair (i, j) incurs a pair-specific balance-sheet cost κ_{ij} per unit, reflecting the credit support annex, netting set, and settlement infrastructure of that pair, and gross funding in currency c incurs a currency-level cost φ_c per unit, reflecting funding scarcity in that currency. Com-

petitive intermediation then prices the long-tenor basis as

$$x^{j/i} \;=\; \big(tp_j - tp_i\big) \;+\; \big(\varphi_j - \varphi_i\big) \;+\; \kappa_{ij}\,\omega_{ij}, \tag{IA.27}$$

where ω_{ij} is the equilibrium bilateral position. Equation (IA.27) reproduces the identification structure of Proposition 3 inside the model: the funding component $\varphi_j - \varphi_i$ is a node potential that cancels around every triangle, so triangular information identifies only the pair-specific terms $\kappa_{ij}\omega_{ij}$, and the allocation of φ between the recovered premium and the frictions is exactly the reallocation freedom the proposition describes. The framework also rationalises the bounding strategy of Section 4.2: at short tenors the term-premium differential vanishes while the funding component does not, so the short-end basis measures $\varphi_j - \varphi_i$ and provides the observable counterpart with which the sensitivity frontier is traced.

Finally, competition among dealers bounds the pair-specific component by transaction costs. If $|\kappa_{ij}\,\omega_{ij}|$ exceeded the quoted half-spread t_{ij} of the triangle-completing quotes, a rival intermediary could improve on the posted price of pair (i, j) through the two remaining legs and earn the difference net of costs, so in equilibrium

$$\big|\kappa_{ij}\,\omega_{ij}\big| \;\le\; t_{ij}, \tag{IA.28}$$

which is the containment property documented in Section 4.1, where between 80 and 99 percent of monthly pair residuals lie inside the quoted half-spread. Within the framework, the sufficient conditions for the currency-level component to equal the difference of conventional local-currency term premia are explicit: habitat demand and arbitrage pricing determine tp_c in each bond market through equation (IA.26), the basis prices the differential plus the funding and pair terms through equation (IA.27), and the currency-level component equals the term-premium differential exactly when the currency-level funding costs φ_c are zero, and up to the sensitivity-frontier bounds when they are limited by the observable short-end benchmark. The framework is deliberately reduced form in its demand and issuance schedules; its role is to show that a single coherent equilibrium delivers the supply loadings and their asymmetry, the node structure of the funding component together with the short-end bounding strategy, and the transaction-cost containment of the pair frictions, which are the three places where the paper's empirical designs and its identification argument meet.

G.2 Supply Inference Additions: Frontier, Mean Group, and Differences

Persistence-valid inference and the state-dependence race

Four exercises answer the inference objections directly, each on the six-pair free-float panel and its collapsed cross-pair mean, and each computed independently in R and in Python before being reported. First, the collapsed single series, which is the specification carrying the honest effective sample, delivers the negative supply loading under Newey-West inference but does *not* survive an Ibragimov-Muller group-t on eight non-overlapping blocks, where the statistic falls below unity against a critical value of 2.365; I report this plainly, and it is the reason the supply evidence is presented as descriptive throughout.

Second, the stochastic-trend structure is treated explicitly rather than left to the standard-error convention. Augmented Dickey-Fuller statistics with a constant and twelve fixed augmentation lags do not reject a unit root for the free-float supply series on either leg, the statistics lying between $+0.28$ and -1.96, nor for five of the six ten-year differentials, sterling the exception at -3.76, so the levels panel is read as a regression among integrated series and the question becomes whether the residual is stationary.

Per-pair Engle-Granger residual statistics from the pair-level levels relation, read against the MacKinnon three-variable critical values of -3.74 at five percent and -3.45 at ten, reject the no-cointegration null for the yen at -4.20 and Australia at -3.77, and for sterling at -4.72, though sterling's rejection restates its stationary differential rather than establishing a cointegrating relation; the euro is marginal at -3.70, and Canada at -3.29 and New Zealand at -2.33 do not reject, the group mean being -3.67. An error-correction specification on the pooled panel returns an adjustment coefficient on the lagged levels residual of -0.07 with a statistic of -2.98, marginal at ten percent against the corresponding non-standard critical values. Cointegration between the differential and duration supply is therefore supported for half the panel and not established for the rest, and I rest no cointegration claim on it: the diagnostics are the formal counterpart of the descriptive labelling of the supply panel used throughout.

Third, the state dependence is dated rather than terminal: the loading is strongly negative over 2002 to 2019, indistinguishable from zero over the pandemic window, and strongly negative again over January 2022 to June 2026 on its own. Fourth, the balance-sheet reading of the suspension does not survive a race: an indicator for the April 2020 to March 2021 leverage-relief window and an indicator for the pandemic window each carry the suspension on their

own, but entered together the pandemic window retains the interaction while the leverage-relief window contributes nothing, so the suspension is not the leverage constraint being lifted.

A fifth exercise confronts the long-horizon intermediation and hedging-demand alternative on observable proxies, with the two proxies chosen on data availability and the design and its prediction fixed before estimation. The first proxy is swap-market depth and receive-fixed activity, the BIS semiannual notional outstanding of single-currency interest-rate swaps by currency, which the survey breaks out for the euro, yen, sterling, and Canadian pairs; the second is the issuance-driven hedging stock, the BIS quarterly international debt securities outstanding by currency of denomination, available for all six supply pairs, with the euro-area aggregate for the euro. Each enters as the log level in the partner currency less the log level in dollars, carried forward within the reporting interval, so the test addresses the persistent level margin on which the alternative operates.

In the supply panel with pair fixed effects and Driscoll-Kraay errors, the proxies alone load only through the issuance stock, at $+0.574$ with a t-statistic of 6.25; adding the issuance stock to the supply specification leaves the United States free-float loading at -0.930 with a t-statistic of -8.97 on the six-pair panel, and adding both proxies leaves it at -1.041 ($t = -9.65$) against a four-pair baseline of -1.325 ($t = -9.19$), while the depth proxy carries no loading in any specification (t-statistics of 0.57 and 0.83). The pre-specified prediction was that the free-float loading survives the proxies with its sign and order of magnitude, and it does; the issuance stock's own positive association, larger relative international issuance in a currency accompanying a higher differential, carried no ex-ante sign and is reported as returned. The exercise is the direct form of the discrimination the alternative-readings table describes: the observable margins of long-horizon intermediation demand do not absorb the duration-supply pricing.

The intermediary-state interaction

The interaction referenced in the main text's channel discussion is recorded here. Interacting the U.S. free-float supply variable with the standardised intermediary capital ratio of He et al. (2017) on the six-pair free-float panel, over the 1,581 pair-months with capital-ratio coverage, gives an interaction coefficient of $+0.28$ with a t-statistic of $+4.2$, the supply loading attenuated when intermediary capital is abundant and steepest when it is scarce, with the main supply loading retaining its sign and significance in the interacted specification; the direction is the one a preferred-habitat account predicts when intermediaries are the marginal absorbers of duration, and the interaction does not displace the supply loading.

Frontier, mean group, and differences

Four additional inference checks complete the supply analysis. First, a coefficient-robustness frontier addresses an omitted currency-level component correlated with supply: subtracting λ times the cross-sectional mean one-year dollar basis from the dependent variable, which allocates a funding component of that size against the U.S. leg, and re-estimating the panel gives U.S. supply coefficients of -1.02, -0.93, -0.84, -0.75, -0.67, and -0.49 at λ of 0, 0.5, 1, 1.5, 2, and 3, with t-statistics from -9.3 to -5.1, so the association survives an omitted component three times the observable benchmark. The frontier runs on the six-pair quoted panel, whose $\lambda = 0$ baseline of -1.02 is the six-pair analogue of the seven-pair -1.051; the main text's sensitivity-frontier reallocation, -1.051 to -0.818 over $\lambda \in [0, 1]$, is the same allocation design on the seven-pair panel, so the two exercises differ in panel and not in construction.

Second, a mean-group estimator, averaging pair-by-pair time-series coefficients, gives -0.85 with a standard error of 0.34, five of six pair coefficients negative, confirming the panel estimate without pooling. Third, a twelve-month first-difference specification yields a coefficient near zero, an informative null: the supply relation is a low-frequency association between levels of supply and levels of the premium, as a preferred-habitat stock mechanism implies, and differencing removes the identifying variation, which is why the collapsed time-series and block-bootstrap designs of the main text are the appropriate persistence-robust checks. Fourth, an inflation-expectations control: with ten-year inflation swap rates for both legs, available for four pairs from 2004, the U.S. supply coefficient is -0.63 with a t-statistic of -3.6; the joint maturity-and-inflation specification is underpowered by the intersection of the two coverages and is not informative either way. On the extension pull of August 2026, with data through June 2026, the panel widens to five pairs, the yen covered from 2007 and Canada added: the U.S. supply coefficient is -1.38 with a Driscoll-Kraay t-statistic of -5.8 on 1,044 pair-months over 235 months, and both inflation-swap controls carry the expectations-channel signs ($+0.07$, $t = 0.8$, on the U.S. leg; -0.09, $t = -1.65$, on the partner leg).

The paper-wide multiplicity treatment is the reported-cells Holm-Bonferroni family stated in the main text, every t-statistic reported in the paper, in a table or in prose, counted once: thirty-five tests, of which twenty-six reach five percent uncorrected and twenty survive the correction, the supply battery in its entirety, the twenty-year European Central Bank purchase cell, and six cells of the survey-revisions orthogonality exercise. Earlier treatments, a twelve-test Benjamini-Hochberg family, the forty-two-test family that included the since-removed returns battery, and the twenty-three-test family of the intermediate build, are superseded by this family and stand only in the replication materials.

The extracted-premium record

Two features of the extracted-premium re-estimation of Table IA.1, referenced in the main text, are recorded here. The comparison-country supply term enters with the opposite sign throughout, the pattern one expects when the two sovereigns serve the same safe-asset demand, and the general government fiscal balances, which enter the macro column of both tables as the flow counterpart of the debt stock, are insignificant on both term premia, so it is the stock of the free float rather than the flow of the deficit that carries the pricing. The free-float maturity series enter every column; their construction, with the maturity panel extended through June 2026, is described in Section D.6.

G.3 The Sovereign Credit Differential Controlled Directly

The main text's statement that a sovereign credit reading is controlled directly is recorded here, under a design fixed before estimation in the replication materials. Ten-year sovereign credit-default-swap differentials, partner minus the United States, are built from senior contracts, the United States contract denominated in euros as its market convention requires, so the differential embeds a quanto component that is stated rather than adjusted; the United States contract quotes from September 2009, so every exercise runs on the post-2009 common sample of 202 months per pair and says so, and the pre-2008 side of the reversal is not testable against measured protection and is not claimed to be. Table IA.39 reports the panel.

Three findings dispose of the credit reading. First, the differentials are weakly related to the government differentials where it matters: in monthly changes the correlations run from -0.02 to $+0.17$ across the six pairs with quoted protection, and the level correlations, less informative under the series' persistence, from -0.07 to $+0.44$. Second, the supply loading is unchanged in point estimate by the control, -0.77 percentage points per log unit without it and -0.79 with it, and the CDS differential itself enters at -0.10 with a Driscoll-Kraay t-statistic of -0.9, indistinguishable from zero under persistence-robust inference; the control regression runs on the four pairs carrying both quoted protection and a constructed free-float series. Third, subtracting the full CDS differential from the government differential leaves the post-2013 configuration intact for every pair with quoted protection, and enlarges it for the yen, sterling, and the Australian dollar, whose protection costs more than United States protection. The subtraction is an upper bound on the credit component of a yield: both legs are foreign-currency contracts, the United States contract in euros and the partners' in dollars, each carrying a quanto premium relative to own-currency protection, and the CDS-bond basis and recovery assumptions inflate protection costs relative to the credit component priced into

a yield, so the correction is conservative for the differential.

TABLE IA.39 ABOUT HERE

G.4 Leverage Regulation: The Full Signed Event Universe

The main text rests the leverage-regulation check on the single announcement that is both plausibly unanticipated and free of confounding policy events, and this section records the full universe behind that choice. Ten dated U.S. leverage-regulation actions from July 2013 to June 2025 were assembled with their direction fixed ex ante, tightening or relaxation, and with the ex-ante prediction that a tightening moves the swap-spread differential down and a relaxation moves it up. Each event is read over the same $[-1, +2]$ window and against the same local placebo convention as the rest of the paper, the enumerated distribution of same-length windows within eighteen months either side. Two of the ten clear the registered informativeness screen, the July 2013 enhanced supplementary leverage ratio proposal and the April 2020 exclusion of Treasuries and reserves, and both are first introductions of their measure.

The July 2013 event carries the predicted sign, -6.1 basis points, with the basis quiet at seven percent of the composite's response. The April 2020 event is excluded from the main-text check because its window sits inside the pandemic interventions and cannot isolate leverage regulation; for the record it moved the swap-spread differential -2.5 basis points, against the relaxation prediction, with the basis quiet at three percent of the composite, so its allocation matches the identity's prediction while its sign does not. Across all ten actions five of ten predicted signs are realised, which is chance, and I rest nothing on the sign count; the surviving content of the exercise is the allocation margin, and on the one clean event both the allocation and the direction hold.

H External Validation and Identification: Additional Results

H.1 The Model-Based Benchmark

The model-based benchmark in full. The comparison whose headline appears in Table 6 is reproduced here in full.

A canonical model-based benchmark corroborates the configuration. An external check, pre-registered with its criterion in the replication materials, asks whether a canonical yields-only benchmark places the post-crisis configuration on the same side: the Joslin et al. (2011) Gaussian term structure model, estimated currency by currency on the identical five-node zero-coupon curves with an identical procedure and no surveys, with the Bauer et al. (2012) bias correction of the physical dynamics addressing the small-sample persistence understatement that distorts separately estimated term premia. The resulting ten-year term-premium differentials lie on the same side of zero as the recovered differentials in post-2009 mean for five of the six pairs, including the euro, at -30 model-based against -52 recovered basis points, and the yen, at -16 against -75; the one miss is sterling, whose recovered post-2009 mean is within four basis points of zero, and the uncorrected variant agrees on all six. The comparison is one of levels, so it is not invariant to how the currency-level component is allocated between a term premium and a dollar-wide wedge, and that is its purpose rather than a limitation: it brings evidence on that allocation from outside the price system, and it should not be read alongside the results that hold whatever the allocation. The corroboration is at the level of the signs of the post-crisis means, the pre-specified standard; monthly correlations are low, and negative for three of the six pairs, the euro, the yen, and sterling, at -0.20, -0.38, and -0.33, so no high-frequency tracking is present and none is claimed. A country-by-country implementation of the regression-based estimator of Adrian et al. (2013) is not undertaken; published estimates exist for the United States alone. Internet Appendix Section H.7 reports the design, the criterion, and the full comparison.

H.2 Transaction Evidence from the Swap Data Repository

The quote-quality discussion of the main text rests on institutional documentation and the independence harness; this subsection adds direct transaction evidence that the long-tenor contracts underlying my quotes trade actively. Under the CFTC's real-time public reporting regime, the DTCC swap data repository disseminates every reported interest-rate derivative transaction with a U.S. reporting party in daily public files. I collect the full rolling public window, July 2025 to July 2026, some 367 daily files, and identify cross-currency swaps as interest-rate swaps whose two legs carry different G10 notional currencies, in all three contract forms: float-for-float, the quoted basis instrument; fixed-for-float; and fixed-for-fixed, the form typical of issuance hedging. I retain new trades only, dissemination action type NEWT, deduplicated by dissemination identifier, so lifecycle events, amendments, corrections, and terminations of swaps executed earlier, are excluded exactly rather than by approximation. This yields 76,175 new transactions over the eleven complete months from August 2025 to June

2026, of which 70,204 are float-for-float. Tenor is computed from the effective and expiration dates of each contract.

Table IA.14 reports average monthly trade counts by pair and tenor bucket over the eleven complete months of the window. The tenor bracket containing my ten-year quotes, eight to twelve years, records 250 new float-for-float trades per month in euro-dollar, 204 in yen-dollar, and 164 in sterling-dollar, roughly forty per business day across the network's dollar pairs, and every dollar pair, including the thinnest antipodean and Scandinavian legs, records trades at that tenor every month. The float-for-float form, the instrument my quotes price, accounts for eighty-nine percent of bracket activity, with fixed-for-fixed issuance-hedging swaps making up most of the remainder. The euro crosses barely appear in these files because transactions without a U.S. reporting party fall under the European reporting regime rather than the CFTC's; the independence of those quotes is established separately, by the dealer documentation and the harness of Section E.1.

A natural concern is that a trailing-year window samples a comparatively tranquil market, and that trading might evaporate precisely in the stress episodes where the recovered term premia move most. The official weekly aggregates published by the CFTC from the same repositories, available from 2013, answer this directly, and in the opposite direction: cross-currency swap ticket counts intensify in stress. Against 848 tickets in a representative calm week of June 2019 and 1,213 in June 2021, the four weeks spanning the March 2020 dollar-funding crisis record 1,730, 1,930, 2,310, and 1,675 tickets, and the gilt-crisis week of October 2022 records 1,961; the invasion week of March 2022 records 1,511. Already in mid-2013 the reported cross-currency book stood above sixteen trillion dollars of notional outstanding.

Coverage caveats are disclosed for completeness: the weekly report carried cross-currency as its own asset class until October 2015, folded it invisibly into interest rates until the foreign-exchange section was introduced in October 2018, and folded it again with exotics in a 2023 redesign, so the ticket series above spans the 2019 to 2022 releases that separately identify the product, which happen to contain the two deepest stress episodes of the sample. The pattern accords with the mechanics of the episodes themselves: dollar-funding stress is transacted through precisely these instruments.

The transaction evidence thus anchors precisely the legs that carry the network's dollar structure: the long-tenor quotes I use are not indicative marks on dormant contracts but prices in an actively traded market, with new ten-year-bracket trades arriving many times every business day.

Traded levels against quoted levels, 2013 to 2020. The counts above establish that the contracts trade; they do not establish at what price. A historical collection answers that second question over a window the rolling public file no longer reaches. The same repository disseminated daily public slices from 2010, and a day-by-day collection of the interest-rate asset class, concatenated by year, yields usable files for 2013 through 2020. From 15.3 million disseminated rows I retain the float-for-float cross-currency form, the instrument my quotes price, keeping new price-forming trades whose two notional currencies are the dollar and one of the nine partner currencies and whose traded tenor, the interval between the effective and expiration dates, lies within forty-five days of two, five, or ten years. That leaves 71,220 transactions, and 65,611 after discarding the 7.9 percent of rows carrying an empty, zero, or implausible price. Monthly medians are formed by pair and tenor and a cell is retained only where at least twenty trades support it, which leaves 1,005 monthly cells between January 2013 and November 2020. The design, including the units convention, the sign convention, and the retention floor, was fixed before the data were read, and the one deviation from it, the raising of the floor from five trades to twenty, is recorded with its reason in the replication materials.

Two features of the comparison bear on how the numbers should be read. The first is that the sign convention was pinned by the data: the repository does not record which leg carries the spread, and I fixed in advance that the panel would be oriented to make the post-2015 euro and yen medians negative, as the post-crisis literature and my own convention require. No reorientation proved necessary, the medians arriving at -26.9 and -56.6 basis points as reported, so the agreement below is not an artefact of a convention fitted to it. The second is that reporting is not neutral across pairs. The repository captures transactions with a U.S. reporting party, so a pair that is thin here is thin in that reporting rather than thin in the market: the euro pair contributes 23,570 trades and a median of 56 a month, the yen 16,983 and 46, the Australian dollar 8,678 and 22, sterling 7,678 and 23, and the Canadian dollar 4,354 and 13, while the Swiss franc, the New Zealand dollar, the Swedish krona, and the Norwegian krone return medians of five a month or fewer and fall below the retention floor. The exercise is accordingly silent on those four pairs, and that silence carries no information about their markets or about the series this paper recovers for them.

Table IA.2 reports the comparison. Across the 1,005 cells the traded median differs from the quoted series by 1.2 basis points at the median, against a mean quoted half-spread of 1.65 basis points on the same cells, and lies inside that half-spread in 60 percent of them. Correlations between the traded and quoted monthly series run from 0.823 to 0.988 over the fourteen pair-tenor combinations the floor admits, and the median absolute difference ranges from 0.8 to 1.8 basis points, so no combination departs by as much as two basis points at the median. The correspondence is closest where reporting is deepest, which is the pattern a reporting-coverage account predicts. Read beside the executable band of 3.6 to 5.1 basis points that Section 4.1 builds by perturbing quotes within their own spreads, the comparison says that the prices at which these contracts actually changed hands over eight years sat about a basis point from the quoted series the measurement uses. Two limits bound the reading: public dissemination caps large notionals, so nothing here should be read as a volume estimate, and the collection ends in 2020, so the exercise speaks to a window rather than to the whole sample.

H.3 Survey Expectations

The main text reports that subtracting the recovered term-premium differential brings implied expectations closer to professional forecasters' own projections, and that the euro cell shows the reversal in the same direction as the recovered series. This section gives the construction of each leg, every cell, and the diagnostics the comparison requires.

What qualifies as a source. Only surveys of forecasters are used. The U.S. leg throughout is the Survey of Professional Forecasters, which reports the expected average three-month Treasury bill rate over the coming one, two, five and ten years, the horizons above one year being collected once a year in the first-quarter round. The euro-area leg is the published survey expectation of the policy rate twelve months ahead. The sterling leg is the Bank of England's Survey of External Forecasters, whose results are reported in a table captioned "averages of other forecasters' central projections" in each quarterly report, giving the expected policy rate at roughly one, two and three years ahead; the row is labelled the repo rate before 2006 and Bank Rate afterwards, and the number of horizon columns varies across editions, so each value is matched to its own column heading rather than to a fixed horizon. Expectations for monetary policy left that survey after 2021, so the sterling series ends in January 2020 on the extraction used here.

Constructing an expected average from a path. A term premium at horizon N is the yield less the expected *average* short rate over N years, so a survey that reports the expected policy rate *at* a horizon cannot be used directly: the difference between the two is the slope of the expected

path, which averages 29.1 basis points across the sterling sample and is systematically positive whenever rates are expected to rise. For sterling I therefore reconstruct the path, anchoring it at the observed three-month Treasury bill rate on the report date, running it piecewise-linearly through the survey's own points, and integrating over the horizon. No row required extrapolation beyond the survey's furthest point. The anchor is a Treasury bill rather than an overnight-index rate for the reason given in the main text: a bill is a government security, an overnight-index rate is a swap price, and the point of the exercise is a comparison that carries no swap-market information.

The incremental comparison. Table IA.3 reports it. The benchmark is the observed yield differential, that is, the implied expectations a reader obtains by assuming the term-premium differential is zero, and the recovered measure must beat it on both correlation and root mean squared error to count as an improvement. It does so in the euro cell and in the sterling one-year cell, and marginally fails in the sterling two-year cell.

The reversal. Table IA.4 reports sub-period means of the recovered and survey-implied term-premium differentials, taken on exactly the months where both exist. The windows are the paper's standard ones. Counts are printed beside every cell because they differ sharply: the euro cell holds twenty-one pre-2008 rounds, the sterling one-year cell twelve, and the sterling two-year cell three, and no reading should rest on the last of these.

Two qualifications, stated with the result. First, the comparison is a joint test of the surveys being unbiased and of the recovered quantity being a term premium, so a departure could reflect forecast error rather than mismeasurement. Scored against what short rates subsequently did, the sterling and U.S. legs give errors of 7.8 and 12.9 basis points at the one- and two-year horizons, small relative to the differences under discussion, and decomposing the sterling reversal window by window shows that the survey's movement is not the instrument's own error moving: net of the change in forecast bias the survey's fall is larger, not smaller. Second, adjacent rounds of an N-year expectation overlap, so any t-statistic formed on these averages would overstate significance and none is reported.

The one public long-horizon source: the panel of Wright (2011). The surveys above reach one and two years. The one public source of long-horizon forecasts across these countries is the Consensus Economics semiannual panel underlying Wright (2011), April and October rounds from 1990 to April 2009, whose replication archive carries each country's long-horizon growth and inflation forecasts and an imputed long-run nominal short rate. From it I construct a survey-anchored ten-year term premium with no model at all: the expected average short rate over ten years is a linear convergence from the current one-year zero to the imputed long-run rate over five years, flat at that rate thereafter, which integrates to one quarter of the one-

year zero plus three quarters of the long-run rate; the term premium is the ten-year zero less this expectation, both legs on this paper's own curves, and the differential is partner minus U.S. Because the panel ends in April 2009, the exercise can speak to the pre-2008 side of the reversal and not to the post-2009 configuration. Table IA.5 reports it on the fifteen rounds from 2002. The survey-anchored differential sits on the same side as the recovered measure in the pre-2008 window in five of eight pairs, and the headline cell is clean, the euro at +41.1 basis points against the recovered +27.8. The misses are three. The Japanese cell fails for a reason visible in the panel itself: the imputed long-run Japanese rate averages 2.93 percent across these rounds, between 1.85 and 3.44, while the ten-year JGB zero averages 1.47 percent, so any survey-anchored Japanese term premium is pushed deeply negative by construction; that long-horizon Japanese forecasts sat far above realised rates for two decades is precisely the stickiness that motivates a measurement estimating no expectations model in any country. The sterling disagreement sets −3.8 against +10.0, a near-zero cell on both sides. The New Zealand cell disagrees outright, +87.0 against −31.5, and the disagreement is recorded rather than set aside. The crisis fall from the pre-2008 window into 2008 to 2009 is matched in only three of eight pairs, a semiannual survey being too slow to track it. The comparison is a joint test in the sense of the qualification above, and no test statistic is reported on fifteen overlapping rounds.

The revisions margin: survey revisions and the two components. The exercises above compare levels; a more direct use of the same independent forecasts is the revisions margin, with the design and its prediction fixed before estimation. If the recovered differential carried residual expectations content, revisions of the survey expected-rate differentials would have to co-move with it; if the construction allocates expectations where they belong, those revisions load on the expectations component, the neutral-rate differential of the contribution statement, and not on the premium. For each survey leg, the euro one-year and the sterling one- and two-year legs against the matched-horizon U.S. Survey of Professional Forecasters projection, carried forward within rounds with a six-month staleness cap, the h-month revision of the survey differential is the regressor and the h-month changes of the two components at the matched tenor and at ten years are the dependent variables, at h of three, six, and twelve months with Newey-West standard errors at lag $h + 6$. Table IA.6 reports the twelve-month cells. The expectations component absorbs the survey revisions with loadings that rise toward one in the revision horizon, from 0.23 to 0.40 on the euro leg and from 0.56 to 0.87 on the sterling one-year leg across the three horizons, while the recovered differential is orthogonal to the same revisions, the euro loadings never as much as half a standard error from zero at any horizon or tenor and the sterling ten-year loadings flat throughout. Two deviations are disclosed rather than set aside: the sterling two-year leg is sparse and erratic at the three-month

horizon, sixty-three observations, and normalises by twelve months; and the sterling one-year premium loading is small and negative, -0.16 with a t-statistic of -4.6, the OPPOSITE sign to expectations contamination, which would have to mirror the expectations component's positive loading, and read accordingly as the premium moving against expectations news rather than carrying it. The exercise is necessary rather than sufficient, since a non-expectations wedge is not excluded and is bounded by the sensitivity frontier; what it establishes, from data that enter no part of the construction, is that independent expectation revisions move the expectations component nearly one for one and do not move the recovered premium. The correlation counterpart of the same margin, reported in the main text, has the revisions of the expectations component tracking the survey revisions at 0.36, 0.45, and 0.60 on the euro leg and 0.63, 0.81, and 0.93 on the sterling one-year leg at the three-, six-, and twelve-month horizons; the sterling two-year leg is sparse at the short horizons and reaches 0.82 at twelve months.

H.4 The Identification Battery: Full Detail

The construction read at the short end. Before the bound itself, the same combination of a basis and a swap-spread differential can be formed at three months, where by its own logic it cannot be measuring term risk, and the main text carries only the summary of what it contains there. At three months a term premium is negligible, so the identical combination of a basis and a swap-spread differential recovers whatever else the two markets price at that horizon. Reconstructing the three-month basis as above and taking the government leg from the fitted sovereign zero at the three-month node, the recovered differential averages -32.4 basis points for the euro pair and -48.5 for the yen pair since 2013, splitting into a basis component of -20.7 and -43.2 and a swap-spread differential of -11.6 and -5.1. The second component is the one that contains no basis at all, being each government's richness against its own overnight curve differenced across the two countries, and it is the short-horizon counterpart of the convenience yield the money-market literature measures: it stood at $+51.7$ basis points for the euro pair and $+23.9$ for the yen pair before 2008 and at -11.6 and -5.1 after 2013, so the richness of the United States bill against its own curve fell relative to both partners across the crisis, which complements the decline in the short-horizon Treasury premium that Du et al. (2018b) document.

Two disclosures belong with the row. A reconstructed basis inherits the money-market benchmark it is built on, and this row is on the overnight-index route; on the IBOR route and matched days the same quantity averages -33.1 basis points for the euro pair against -34.2

and -1.4 for the yen pair against -41.4, the yen gap being the two currencies' IBOR-to-overnight spreads entering the combination twice. And the three-month forward half-spreads, annualised, run from 4.4 basis points for the euro to 53.1 for the krone against long-tenor executable bands of 3.6 to 5.1, so the short-end row is measured far less precisely than the quotes the rest of the paper uses, cleanly for the euro, the yen, and sterling and weakly for the Scandinavian pairs; the euro overnight-index series begins in January 2007, so its pre-2008 average rests on that year alone, while the yen window runs from March 2002.

The excess-sensitivity bound invoked in the main text is the following.

Proposition 4 (Excess-sensitivity bound on the funding hypothesis)**.** *Suppose the currency-level component of the basis at tenor τ equals the expected average of a stationary funding process φ_t with monthly autocorrelation ρ, plus a constant risk premium. Then the five-year-five-year forward component implied between tenors, $2\,x_{10Y} - x_{5Y}$, satisfies two restrictions: its loading on the current funding state is at most $\frac{1}{60}\sum_{k=60}^{119}\rho^k$, and its unconditional volatility is strictly below that of the funding state, for any $\rho < 1$.*

A final battery dispenses with the allocation altogether. The economically salient alternatives to the term-premium reading, funding scarcity, collateral demand, and balance-sheet costs, are concentrated at short tenors or flat across them, because funding rolls over at short horizons; a term premium, by contrast, accumulates with duration. Any currency-level component that is flat across tenors therefore differences out of maturity-differenced constructions entirely, whatever its origin and however it would have been allocated.

Three results follow, none of which involves the projection or the restriction. First, the tenor slope of the raw benchmark construction, the ten-year minus the one-year value of $x + \Delta SS$ computed directly from quotes, exhibits the central reversal: the euro slope moves from $+21$ basis points pre-GFC to -19 post, and the yen from $+27$ to -40. Second, supply prices the slope: in the same panel design the U.S. free-float coefficient on the slope construction is -0.75 with a t-statistic of -6.6. Third, the supply loading on the recovered cross-country term-premium differential rises monotonically with tenor, from -0.67 at two years to -0.98 at five, -1.02 at ten, and -1.30 at twenty, which is the signature of priced duration risk and the opposite of a short-concentrated funding channel. The same panel design applied to the five-year-five-year forward component implied between tenors returns a U.S. free-float coefficient of -0.76 with a t-statistic of -6.4 on the eight-pair panel, the far-forward loading the main text reports. A further check points the same way: the long-tenor basis displays no quarter-end seasonality, with a window dummy of -0.1 basis points, unlike the very short tenors where regulatory funding windows operate.

Three policy-based and forecast-based tests complete the discrimination. First, the Bank of Japan's yield-curve-control band changes are identified shocks to the long end of one curve with no funding content, and the daily yen benchmark series responds exactly as a released term premium should: the two genuine surprises, the December 2022 widening and the October 2023 reference change, move it by +12.5 and +9.7 basis points over two days, or 3.2 and 2.5 daily standard deviations, while the anticipated July 2023 and fully telegraphed March 2024 announcements move it not at all, as announcement logic implies. The prediction generalises across central banks: domestic asset-purchase announcements should compress the announcing country's premium and move the differential in a signed direction, and five of seven major announcement windows across the Federal Reserve, the European Central Bank, and the Bank of England respond with the predicted sign, including the March 2009 gilt purchase launch at −36 basis points, eight daily standard deviations, the January 2015 ECB announcement at −8 basis points, and both 2008–2009 Federal Reserve announcements positive across six of six pairs, with the two exceptions, the pandemic-window PEPP and the partially anticipated 2013 taper testimony, disclosed as uninformative.

Second, the 2011–2012 Maturity Extension Program windows, duration shocks with no funding counterpart, show no clean response, but those windows coincide with the peak of the euro funding crisis and I read them as uninformative rather than negative, consistent with the limited announcement-frequency power disclosed throughout; I note that the formal break location estimate follows the program's start. Third, a forecast horse race separates the two readings in outcomes: controlling for the current short-end basis, the recovered cross-country term-premium differential carries no incremental forecast content for future funding realisations, with a t-statistic of 0.9 against the basis forecasting itself at 16.5; a disguised funding component would forecast funding, and this differential does not.

Decomposing each of these results across the two bridge components is instructive and consistent with the framework of Internet Appendix Section G.1. The maturity-differenced evidence is carried predominantly by the swap-spread side, where the supply shock has its direct incidence, since the preferred habitat is in the government market: the slope reversal and the slope-supply loading, at −0.79 with a t-statistic of −6.7, live in the swap-spread differential. The extracted swap side is not a passenger, however: U.S. supply prices the cross-country swap term premium alone at every tenor, at −0.22 to −0.35 with t-statistics of −3.8 to −5.1 and the largest loading at twenty years, and the swap side carried the majority of the best-identified policy shock, +8.2 of the +12.5 basis-point December 2022 response, with the October 2023 response arriving mainly through the spread. The pattern, direct supply incidence in the bond market transmitted to the spread, with the swap curve priced through arbitrage and moved

directly by policy shocks to it, is precisely what the equilibrium framework delivers.

For the cross-country swap term premium specifically, the funding question can be answered quantitatively rather than categorically, and I do so. In the largest funding shock on record, from late February to late March 2020, the one-year dollar bases widened by 19 and 41 basis points for the euro and yen pairs at their troughs while the ten-year bases moved by 1 and 13, ratios of 0.04 and 0.32, so the long-tenor differential responds to a pure funding event at a small fraction of the short end. The supply pricing of the cross-country swap term premium survives a direct control for intermediary capital, at -0.22 with a t-statistic of -3.4, closing the dealer-balance-sheet variant of the funding channel. And the cross-country swap term premium does carry modest incremental forecast content for future funding beyond the current short basis, twelve basis points per hundred with a t-statistic of 3.0, which is the signature of a bounded expected-funding share in the currency-level component; the cross-country government bond term premium carries none. Three independent metrics, the forecast pass-through of roughly twelve percent, the crisis response ratios of four to thirty-two percent, and the sensitivity-frontier reallocation, thus agree in placing the funding share of the swap-side currency-level component at a bounded minority, with every economic conclusion intact at the extremes of that range.

Direct short-tenor evidence completes the funding discrimination, reported in Table IA.22 under a design stated there. The pass-through of three-month funding-point changes decays monotonically along the curve for every pair, and the profile measures a different quantity from the two levels profiles reported in the main text: the levels loadings of the tenor-level bases on the funding state rise with tenor because both series carry persistent premium content in levels, while the change pass-through isolates the high-frequency funding transmission, which is the quantity the decay prediction disciplines; the discrimination in the main text rests on the exceedance-versus-bound comparison, to which both profiles are inputs rather than verdicts.

TABLE IA.22 ABOUT HERE

One node-space alternative remains to be ruled out, a forward dollar exchange-rate premium residing in the basis itself, and three pieces of evidence bound and test it, with the dollar-free euro-cross evidence of Section 4.1 adding a fourth: ninety percent of the partner-relative variation in dollar pairs appears, nearly one for one, in quotes that contain no dollar leg. Structurally, the instrument is covered and collateralised: both floating legs telescope to par in their own collateral currency, so exchange-rate risk enters only through the settlement of the spread leg, an exposure of the order of the spread multiplied by exchange-rate volatility, below one basis point (Internet Appendix Section A.1). Empirically, the defining property of

a currency risk premium is that it forecasts currency returns, and the swap-side currency-level component has essentially no dollar forecast power, with an R^2 of 0.001 in the encompassing comparison of Internet Appendix Section I.1.

And the best-identified policy event appears undiminished away from the dollar in two forms: the network-implied yen-euro benchmark differential, in which the dollar node cancels, responds by +15.0 basis points against +12.5 for the dollar pair, the euro differential having moved the opposite way inside the window, and the directly quoted euro-yen basis, a swap-side quote involving no dollar leg and no swap-spread leg, moves by +8.5 basis points against the +8.2 swap-side component of the dollar-pair response, its like-for-like comparator. The currency-level component of the basis, net of the bounded pair frictions, is therefore a term-premium differential and not a forward dollar premium. Table IA.23 collects every December 2022 quantity with its construction; the two distinct quantities that both equal +8.5 basis points, the euro-yen quote and the yen five-year-five-year forward component, are disambiguated there.

TABLE IA.23 ABOUT HERE

H.5 Leg Durations of the Basis Package: Construction and Statistics

The instrument's own construction supplies a further argument. By the fixed-for-fixed representation underlying Proposition 2, the basis package is equivalent to a position long the foreign par instrument and short the dollar par instrument, and the durations of the two legs differ mechanically whenever rate levels differ, because low-rate curves have longer annuities. I compute the ten-year par-bond Macaulay duration of every leg, month by month, from two curve families: the government zero curves, which are the relevant family for the cross-country government bond term premium, and the interest-rate swap curves themselves, the chain-reconciled par swap rates from which the swap spreads of the rates database are constructed, which are the relevant family for the cross-country swap term premium; Internet Appendix Section H.5 details every step of the construction, Table IA.40 reports the statistics, and Figure IA.9 plots the leg-duration mismatch against the swap term-premium differentials over time.

The mismatch structure is the same in both families and ordered exactly as rate levels: post-crisis, the yen leg is longer than the dollar leg by 0.84 years on the government curves and 0.79 on the swap curves, the euro leg by 0.56 and 0.44, sterling and Canada near zero, and the

antipodean legs shorter. The currency-level compensation in the basis is therefore compensation for holding a duration position, by the instrument's construction, and both term premia reprice per unit of duration across the crisis: dividing each premium by the mean duration of its own curve family's legs, the euro government premium moves from $+3.4$ to -5.7 basis points per year of duration and the yen from $+4.2$ to -8.2, while the euro cross-country swap term premium moves from $+0.1$ to -2.2 and the yen from -0.7 to -6.3, so the central reversal is a per-duration-unit repricing in both families, not an artefact of the leg mismatch.

The duration exposure is priced, with the incidence split exactly as the framework predicts: the cross-country government bond term premium, where the bond-market habitat lives, scales with the leg-duration mismatch, at 28.5 basis points per year of mismatch with a t-statistic of 4.5 under pair fixed effects and an average-duration control, while the cross-country swap term premium scales with the average duration of the package, at 14.0 basis points per year with a t-statistic of 7.4. Both are duration results; a funding component scales with neither, since funding costs do not vary with the bond durations of a package they do not reference. The most direct discrimination uses the time-series structure of funding itself and requires neither the projection nor any allocation: it is an inequality with measured inputs.

Table IA.2: **Relation Between Traded and Quoted Cross-Currency Basis Levels.** Monthly medians of reported cross-currency basis swap transactions compared with the quoted basis series used throughout the paper, by currency pair and tenor, January 2013 to November 2020. Transactions are the float-for-float cross-currency form disseminated by the DTCC swap data repository, restricted to new price-forming trades whose notional currencies are the dollar and one partner currency and whose traded tenor lies within forty-five days of the stated tenor; a pair-tenor-month cell is retained where at least twenty trades support it. Cells reports retained months and Median trades the typical number of transactions behind a monthly median. Correlation is between the traded and quoted monthly series, Median difference is the median absolute gap between them in basis points, and Within half-spread is the percentage of months in which that gap does not exceed the pair's own quoted half-spread. The Swiss franc, New Zealand dollar, Swedish krona, and Norwegian krone do not reach the retention floor, reporting being confined to transactions with a U.S. reporting party, and are therefore absent. Source: DTCC public price dissemination, collected as described in Internet Appendix Section H.2.

Pair	Tenor	Cells	Median trades per cell	Correlation	Median difference (basis points)	Within half-spread (%)
EUR/USD	2Y	92	61	0.979	1.3	61
EUR/USD	5Y	95	79	0.931	1.4	63
EUR/USD	10Y	92	78	0.975	1.4	65
JPY/USD	2Y	93	83	0.988	1.2	61
JPY/USD	5Y	91	55	0.983	1.5	59
JPY/USD	10Y	82	44	0.973	1.8	56
GBP/USD	2Y	49	30	0.937	1.2	63
GBP/USD	5Y	59	34	0.972	0.8	71
GBP/USD	10Y	63	33	0.977	1.1	73
AUD/USD	2Y	36	27	0.955	1.4	42
AUD/USD	5Y	76	34	0.823	0.9	57
AUD/USD	10Y	83	40	0.972	0.8	67
CAD/USD	2Y	39	25	0.831	1.8	26
CAD/USD	5Y	34	26	0.957	1.0	44

Table IA.3: **Do the Recovered Differentials Improve on a Benchmark of Zero?** The target is the survey-implied expected average short-rate differential, partner minus U.S. The recovered measure is the observed government yield differential less the recovered term-premium differential; the benchmark is the observed yield differential itself. A positive change in correlation and a negative change in root mean squared error, in basis points, are improvements.

			Correlation with survey			Root mean squared error		
Cell	N	Span	Recovered	Benchmark	Change	Recovered	Benchmark	Change
USD/EUR, 1 year	95	2002:10–2026:04	0.868	0.862	+0.006	85.9	96.3	−10.4
USD/GBP, 1 year	59	2004:11–2020:01	0.974	0.945	+0.028	24.9	36.3	−11.4
USD/GBP, 2 years	16	2005:02–2020:01	0.964	0.971	−0.008	27.9	27.1	+0.8

Table IA.4: **The Reversal Against Survey Expectations.** Sub-period means of the cross-country government term-premium differential, in basis points, partner minus U.S. The recovered columns are the paper's measure; the survey columns replace its expectations content with professional forecasters' projections. Both are computed on exactly the months where both legs exist. Windows are January 2002 to December 2007, January 2008 to December 2012, and January 2013 onward.

Cell	Window	N	Recovered	Survey
USD/EUR, 1 year	pre-2008	21	+9.7	+1.1
USD/EUR, 1 year	2008–2012	20	−46.3	−99.3
USD/EUR, 1 year	post-2013	54	−23.8	−41.8
USD/GBP, 1 year	pre-2008	12	−18.2	−12.1
USD/GBP, 1 year	2008–2012	20	−36.8	−20.7
USD/GBP, 1 year	post-2013	27	−4.8	−17.5
USD/GBP, 2 years	pre-2008	3	−1.8	−6.2
USD/GBP, 2 years	2008–2012	5	−22.5	+3.9
USD/GBP, 2 years	post-2013	8	−12.9	−21.1

Table IA.5: **The Recovered Differentials Against the Long-Horizon Survey Panel of Wright (2011).** Survey-anchored ten-year term-premium differentials, partner minus U.S., built from the Consensus Economics semiannual long-horizon panel in the replication archive of Wright (2011): the expected average short rate over ten years converges linearly from the current one-year zero to the panel's imputed long-run rate over five years and is flat thereafter, the term premium is the ten-year zero less this expectation on the same curves as the recovered series, and the recovered column is the benchmark-consistent differential on the identical survey months. The panel ends in April 2009; rounds are April and October from 2002. Basis points. The comparison is a joint test of survey accuracy and of the recovered quantity, and no test statistic is reported on overlapping semiannual rounds.

		Pre-2008 mean			2008–2009 mean		
Pair	Rounds	Recovered	Survey	Same side	Recovered	Survey	Corr.
EUR	15	+27.8	+41.1	yes	−16.8	+33.7	−0.03
JPY	15	+36.3	−114.2	no	+18.1	−104.8	−0.59
GBP	15	+10.0	−3.8	no	+0.2	+39.6	−0.08
CAD	15	+29.7	+40.9	yes	+37.9	+2.9	+0.45
AUD	14	+3.4	+39.8	yes	−59.0	+57.8	−0.32
NZD	15	−31.5	+87.0	no	−14.8	+119.4	+0.57
SEK	15	+8.1	+21.2	yes	−27.2	−18.9	+0.17
NOK	15	+1.1	+19.7	yes	−50.1	−23.1	+0.08

Table IA.6: **The Revisions Margin: Survey Expectation Revisions and the Two Components.** For each survey leg, the twelve-month revision of the survey expected-rate differential (partner minus US, the ECB Survey of Professional Forecasters euro leg and the Bank of England Survey of External Forecasters sterling legs, each less the matched-horizon U.S. SPF projection, carried forward within rounds with a six-month staleness cap) is the regressor, and the dependent variables are the twelve-month changes of the expectations component (the neutral-rate differential) and of the recovered term-premium differential, at the survey leg's matched tenor and at ten years, in percentage points on both sides. Newey-West t-statistics at lag eighteen in parentheses. Under the ex-ante prediction the survey revisions load on the expectations component and not on the premium; the sterling one-year premium cell carries the opposite sign to expectations contamination and is read as the premium moving against expectations news.

		Expectations component		Recovered differential	
Survey leg	N	Matched tenor	Ten-year	Matched tenor	Ten-year
Euro, one-year	273	+0.40	+0.12	−0.02	−0.00
		(3.06)	(1.22)	(−0.43)	(−0.03)
Sterling, one-year	173	+0.87	+0.32	−0.16	−0.03
		(22.15)	(3.62)	(−4.63)	(−0.57)
Sterling, two-year	103	+0.85	+0.43	+0.01	−0.02
		(13.76)	(3.72)	(0.33)	(−0.21)

This section documents, step by step, the construction of the leg durations used in the identification battery. For each currency c and month t, the inputs are the fitted government zero-coupon yields at the one-, two-, five-, and ten-year tenors, in percent, from the Bloomberg fitted sovereign curves described in Section E, and the quoted par swap spreads at the same tenors from the rates database, with the convention $SS = s - y$ so that the swap rate is the government yield plus the spread. Two curve families are built. The government family uses the zero yields directly. The swap family uses the interest-rate swap curves themselves: the rates database defines the swap spread at each tenor as one hundred times the difference between the chain-reconciled par swap rate and the fitted sovereign zero, so the swap rate at each node is recovered identically as the zero plus the spread divided by one hundred, and the resulting nodes are exactly the benchmark-consistent swap rates constructed in the database, carrying the same benchmark-transition treatment described there; the full reconciled swap-rate panel is included in the replication package. Both families deliver the same mismatch structure, documented in Table IA.40.

For each family, the four nodes are interpolated to integer tenors one through ten with a natural cubic spline, discount factors are formed as $\exp(-y_k k)$ for $k = 1, \ldots, 10$, the ten-year par coupon is $c_0 = (1 - d_{10})/\sum_k d_k$, and the Macaulay duration is $D = \sum_k k\, cf_k d_k / \sum_k cf_k d_k$ with cash flows $cf_k = c_0$ for $k < 10$ and $1 + c_0$ at maturity. The calculation is monthly, currency by currency, and the full panels for both families are included in the replication package.

H.6 Identified Policy-Surprise Factors: Design and Full Results

This subsection reports the design and full results behind the main text's policy-surprise-factor paragraph, pre-specified with ex-ante signs in the replication materials before any factor data was obtained, and computed under the dual-computation lock. Outcomes are $[-1, +2]$ trading-day changes of the daily quoted-construction benchmark differentials around each announcement date; factors are standardised to unit variance over the estimation sample; inference clusters by announcement date, since the factors are common across pairs. On the Federal Reserve side the factors are taken directly from the Swanson (2021) extended file, with the asset-purchase factor's orientation fixed empirically before any outcome regression, positive meaning a purchase easing, anchored on the March 2009 purchase expansion and the June 2013 taper testimony.

On the European Central Bank side the event-study database ships raw intraday swap-curve changes rather than factors, so the four factors are constructed by the authors' own method:

the press-release-window target factor is the first principal component of the release-window changes, and the press-conference window yields three principal components rotated so that only the first, timing, loads on the one-month swap and the third, the purchase factor, explains the minimal share of variance over the early, pre-asset-purchase sample, their identifying restriction, the second being forward guidance; the purchase factor's orientation is anchored on the 22 January 2015 asset-purchase announcement, their own largest purchase surprise, positive meaning a purchase easing. Table IA.26 reports the estimates at all four tenors for both legs, with the funding placebos beneath each panel.

Read against the pre-specification: the core predictions hold on both legs, the asset-purchase factor positive on the dollar side and negative on the euro side with a monotone tenor gradient there; the strict prediction that the ten-year asset-purchase response exceeds the target response per standard deviation fails on Federal Reserve point estimates (0.68 against 0.83, the target marginal at $t = 1.7$), while the qualitative separation is realised, the target response concentrating at the two-year tenor and fading beyond it; the Federal Reserve-side one-year-basis placebo is significant at roughly half the ten-year response, whereas the euro placebo is indistinguishable from zero; the European target factor enters positively and flat across tenors, a level response of euro yields to rate surprises, stated as returned under its two-sided registration; and timing and forward guidance are insignificant throughout, as expectations content should be in a premium differential. The single-surprise alternative of the San Francisco Fed's maintained dataset is retained in the replication materials as a logged robustness route.

H.7 A Canonical Model-Based Benchmark: Joslin-Singleton-Zhu with Bias Correction

The comparison referenced in the main text was pre-registered with its criterion in the replication materials before the first estimation run. The benchmark is the canonical Gaussian affine term structure model of Joslin et al. (2011), three factors, estimated currency by currency with an identical procedure on the identical constructed five-node zero-coupon government curves at one, two, five, ten, and twenty years, monthly from January 2000 (August 2000 for Australia) to June 2026, using yield curves alone and no surveys. The factors are the first three principal components of the five yields and are priced exactly; the Q-parameters, the long-run intercept and three real, distinct, ordered eigenvalues, are estimated by minimising the root-mean-squared fitting error of the five yields, which runs from 2.9 to 5.1 basis points across the seven currencies. The physical dynamics are a first-order vector autoregression on the factors,

bias-corrected as in Bauer et al. (2012) by residual-bootstrap mean-bias correction with one thousand replications on draws shared between the two independent implementations, with the stationarity adjustment of scaling the correction back until the corrected system is stable; for the yen the adjustment scales the correction fully back, so its corrected variant differs from the uncorrected only in anchoring the unconditional mean at the sample mean.

The ten-year term premium is the fitted ten-year yield less the average expected short-rate path over one hundred and twenty months, and differentials are expressed in the recovered series' convention. The criterion, fixed ex ante, is that the post-2009 mean of the corrected differential lies on the same side of zero as the recovered differential for at least four of the six pairs, including both the euro and the yen; it is met at five of six. The bias-corrected variant was fixed as the reported series in the registration before the first estimation run. Table IA.41 reports the comparison; the level correlations are reported for completeness and are low, negative for three of the six pairs, the euro, the yen, and sterling, at -0.20, -0.38, and -0.33, so the corroboration is claimed at the level of the signs of the post-crisis means and no further. The benchmark deliberately does not extend to a country-by-country implementation of the Adrian et al. (2013) regression-based estimator; published estimates exist for the United States alone.

Table IA.41 about here

H.8 The Empirical Sensitivity Frontier

This subsection reports the estimated counterpart of the calibrated sensitivity frontier, pre-specified with its criterion in the replication materials before estimation and computed under the dual-computation lock. The calibrated set imposes the reallocation at a stated λ; this exercise instead fits the currency-level wedge from a measurable proxy, removes the fitted part from the recovered differential, and recomputes the reversal on the residual, re-centred on its own mean so that only the pre-to-post change, and not the level, is compared. The design is deliberately generous to the alternative reading: any variation the proxy can explain is assigned to the wedge rather than to the term premium, so what survives is a lower bound on the term-premium content. The proxy of record is the currency's own three-month network node, which is the boundary funding wedge measured today and is constructed entirely outside the long-tenor prices whose interpretation is at issue; the pre-specified criterion was that the reversal survive in sign and retain at least half of the locked post-2013 gap.

On that specification the criterion is met for both headline pairs. The residual differential moves from $+14.2$ to -56.3 basis points for the euro pair and from $+18.1$ to -77.1 for the yen pair between the pre-2008 and post-2013 sub-periods, retaining 87 and 79 per cent of the raw gap, and the three-month node explains 0.149 and 0.153 of the ten-year differential's variance, a minority as the pre-registered third prediction anticipated and consistent with the measured pass-through of 0.022 from the boundary basis into the ten-year basis.

A second specification adds the intermediary capital ratio, the broad dollar index, and the MOVE index to the proxy set. There the fit rises to 0.610 and 0.510 and the retained share falls to 36 and 44 per cent, below the pre-specified half. I report that as returned and read it as over-control rather than as a failure of the criterion: a term premium loads on global risk states too, so assigning the whole of that common variation to a funding wedge removes term-premium content along with it, and the specification therefore bounds the term-premium content from below by more than the design intends. The option-adjusted credit spread was dropped from the control set before estimation because it carries thirty-six monthly observations from July 2023 over the three-month sample and its inclusion collapses the panel to twenty-three months; its absence is recorded rather than worked around.

I The Dollar Dimension: Full Analysis

I.1 Predicting the Dollar: Cross-Country Term Premia and Cross-Country Exchange Rate Dynamics

The theoretical framework of Jiang et al. (2021) links the U.S. convenience yield to exchange rate dynamics: the relative demand for dollar-denominated safe assets translates into dollar movements through a no-arbitrage mapping between convenience premia and currency returns. Their analysis is developed at the spot and short-maturity end of the term structure, where the Treasury basis measures the wedge between the U.S. Treasury yield and a currency-hedged foreign sovereign yield at a one-year or shorter maturity. My cross-country term premium is a complementary measure at the long end of the term structure: it is a long-horizon differential of own-currency term premia rather than a level wedge across two yields denominated in a common currency through an FX hedge, it is identified at the maturity where cross-country substitution is most active, and it captures the relative compensation investors demand for bearing long-horizon term risk across currencies. Whether and how the Jiang et al. (2021) convenience-to-exchange-rate mapping extends from the spot dimension to the long-horizon term-premium dimension is an empirical question that I answer directly in this section.

To make my comparison as directly comparable to Jiang et al. (2021) as possible, I follow their construction and use as the dependent variable the equally-weighted dollar basket against my six bilateral currencies, $\overline{s}_t^{eq} \equiv \frac{1}{6}\sum_j s_t^{USD,j}$, with each bilateral log exchange rate sign-aligned so that an increase corresponds to U.S. dollar appreciation against currency j. The cross-sectional average ten-year cross-country government bond term premium, $\overline{ctp}_t^{gov,USD} = \frac{1}{6}\sum_j \widehat{ctp}_{10,t}^{gov,j,USD}$, is the analogue of the cross-sectional average Treasury basis on which Jiang et al. (2021) build their main results. I control for the cross-sectional average ten-year interest rate differential $\overline{IRD}_{10,t}^{USD} = \frac{1}{6}\sum_j (r_{10,t}^{USD} - r_{10,t}^{j})$ to rule out a confounding interest-parity-based explanation. The corresponding specification with the cross-sectional average cross-country swap term premium as the predictor is reported in Table IA.7.

Table IA.7: **Reconciliation with Jiang et al. (2021): Equally-Weighted Dollar Basket Across Horizons.** This table reports the equally-weighted dollar-basket regressions that mirror the construction in Jiang et al. (2021), $\overline{s}_t^{eq} \equiv \frac{1}{6}\sum_j s_t^{USD,j}$ with each bilateral log exchange rate sign-aligned so that an increase corresponds to U.S. dollar appreciation. Panel A reports the contemporaneous monthly change-on-change specification paralleling Table III of Jiang et al. (2021): Δs_t^{eq} regressed on $\Delta\overline{ctp}_t^{\cdot,USD}$ and $\Delta\overline{\text{IRD}}_{10,t}^{USD}$, both unstandardised. Panel B reports the forward-predictive specification paralleling the Section IV.F discussion of Jiang et al. (2021): $\Delta s_{t,t+h}^{eq}$ regressed on the level of $\overline{ctp}_t^{\cdot,USD}$ and $\overline{\text{IRD}}_{10,t}^{USD}$, both standardised, at the spot ($h = 1$ month), three-month, and twelve-month horizons, with the full inference battery retained in the replication materials. The sign mapping with Jiang et al. (2021) is as follows: under my convention $\widehat{ctp}^{j,USD} \uparrow$ corresponds to Jiang's basis $x^{Treas} \downarrow$ (both indicate higher U.S. convenience), so a contemporaneous coefficient that is positive in my convention corresponds to Jiang's negative basis coefficient, and a forward-predictive coefficient that is negative in my convention corresponds to the positive basis predictive coefficient implied by their Proposition 1. The cross-country government bond term premium $\overline{ctp}_t^{gov,USD}$ delivers the predicted sign in the twelve-month forward specification; the contemporaneous monthly change-on-change coefficient is statistically indistinguishable from zero for the cross-country government bond term premium and enters with the opposite sign for the cross-country swap term premium, a pattern consistent with the funding-channel comovement of bases and the dollar dominating the convenience-revaluation channel at the monthly frequency. Bilateral counterparts to these aggregate regressions, paralleling Table VII of Jiang et al. (2021), are reported in the bilateral battery retained in the replication materials. Sample: predictor dates from January 2002 through June 2026 less the horizon. t-statistics in parentheses. Notes: $^{*}p \leq .10$; $^{**}p \leq .05$; $^{***}p \leq .01$. The cross-country government bond term premium is the quoted construction $x + \Delta SS$ averaged over the six major pairs; the swap column uses the quoted-basis average, representing the cross-country swap term premium up to the contained residual. Newey-West standard errors with six lags in Panel A and twelve in Panel B.

	$\overline{ctp}^{gov,USD}$	$\overline{ctp}^{swap,USD}$
Panel A: Contemporaneous changes		
$\Delta\overline{ctp}_t^{\cdot,USD}$	-0.022	-0.388^{***}
	(-0.77)	(-5.69)
$\Delta\overline{\text{IRD}}_{10,t}^{USD}$	$+0.046^{***}$	$+0.045^{***}$
	$(+4.18)$	$(+4.92)$
R^2	0.076	0.200
N	301	316
Panel B: Twelve-month forward change on the level of the premium		
$\overline{ctp}_t^{\cdot,USD}$	-0.060^{**}	$+0.110$

Continued on next page

Table IA.7 continued

	$\overline{ctp}^{gov,USD}$	$\overline{ctp}^{swap,USD}$
	(−2.01)	(+0.60)
$\overline{\text{IRD}}_{10,t}^{USD}$	−0.016	−0.003
	(−0.74)	(−0.12)
R^2	0.062	0.006
N	291	305

The negative predictive sign maps directly to the positive sign Jiang et al. (2021) obtain on the level Treasury basis at one year, after sign-convention translation. In the Jiang et al. (2021) convention, a more negative Treasury basis x^{Treas} corresponds to higher U.S. convenience, while in my convention a more positive $\widehat{ctp}^{j,USD}$ corresponds to higher U.S. convenience, so the two indicators of U.S. convenience move in opposite arithmetic directions and Jiang et al. (2021)'s positive forward-predictive coefficient maps directly to my negative one. The economic content is the same in both conventions: a more convenient U.S. dollar today raises the spot exchange rate today and induces expected depreciation in the future, which is what the negative coefficient on the level cross-country term premium in my forward-change regression captures.

On the bridge identity $\widehat{ctp}_N^{gov,j,USD} = \widehat{ctp}_N^{swap,j,USD} + \Delta SS_N^{USD,j}$, the content of the sign relation lives in the directly observable cross-country swap spread differential rather than in the triangulated component, which aligns with the Jiang et al. (2021) interpretation that the safe-asset content priced into the U.S. dollar is the wedge between the U.S. Treasury yield and the currency-hedged peer-country yield.

I close the section by reconciling the sign relation with the preferred-habitat duration channel of Section 4.3.1 and the quantity-based dollar-pricing mechanism of Jiang et al. (2021). The dollar-prediction content concentrates in the same cross-country swap spread differential component that absorbs the U.S.-side preferred-habitat duration channel and the U.S. asset-purchase quantity channel (Section 4.3.2); this differential is the equilibrium price that records the wedge in compensation per unit of U.S. Treasury duration that preferred-habitat demand requires relative to peer-sovereign duration. Because U.S. supply is approximately exogenous to bond-market pricing signals (Section 4.3.1 and Georgievska et al. (2025)), the market clears entirely through the price adjustment of preferred-habitat demand against the inelastic supply: changes in the privately held free-float of U.S. duration, whether from Treasury issuance

or central bank absorption, shift the quantity that yield-sensitive private holders must absorb at given prices, and $\Delta SS_N^{USD,j}$ adjusts to clear the new quantity.

The Jiang et al. (2021) mechanism supplies the complementary demand-schedule perturbation, in which a shift in foreign safe-asset demand for U.S. Treasuries changes the demand schedule itself at given supply, and $\Delta SS_N^{USD,j}$ adjusts accordingly. Both supply-quantity and demand-schedule perturbations clear through the same equilibrium price, with the price adjustment of preferred-habitat demand doing the clearing work in either case, which is why the dollar-prediction signal concentrates there. The cross-country panel identifies the supply-quantity perturbations directly; the demand-schedule perturbations, which would require flow-of-funds measurement of U.S. Treasury holdings by foreign sectors, are left to future work.

I.2 Predicting the Dollar: Reconciliation with Jiang et al. (2021)

The sign relation itself I state once, with its caveats, and rest nothing on it: at the twelve-month horizon the cross-sectional average cross-country government bond term premium predicts dollar depreciation with a negative coefficient that is significant under Newey-West and essentially unchanged under the Amihud and Hurvich (2004) Stambaugh-bias correction, but the Hodrick (1992) reverse-regression statistic is insignificant and the out-of-sample content is confined to the major-currency basket, so I treat the predictive dimension as an exploratory consistency check with Jiang et al. (2021) and Lustig et al. (2019) rather than a result of this paper; the complete inference battery is retained in the replication materials.

This subsection reconciles the sign relation stated above with the bilateral specification of Jiang et al. (2021), addressing the sign convention mapping, the aggregation conventions, and the aggregate-level evidence.

I.2.1 Reconciliation with Jiang et al. (2021): Sign Convention, Aggregation, and Specification

This subsection reconciles my cross-country convenience-to-exchange-rate evidence with the spot-convenience evidence in Jiang et al. (2021). I address two questions in turn. First, are my weighted-signal specifications constructed in a manner comparable to theirs? Second, how does the predictive content map across the spot, three-month and twelve-month horizons, and across pairs in the bilateral panel?

Sign convention mapping. Jiang et al. (2021) define the Treasury basis as $x_t^{Treas} = y_t^{\$} + (f_t^1 - s_t) - y_t^*$, where s_t is the log of foreign currency per dollar and f_t^1 is the corresponding one-period log forward rate. Their Lemma 5 shows that a more negative basis corresponds to a higher U.S.-Treasury convenience yield earned by foreign investors. My cross-country term premium is $\widehat{ctp}_{t,N}^{j,USD} \equiv tp_{t,N}^{j} - tp_{t,N}^{USD}$, so a more positive $\widehat{ctp}^{j,USD}$ corresponds to a relatively lower U.S. term premium and therefore to a more convenient U.S. Treasury at long maturities. The two term premia therefore carry the opposite arithmetic sign for the same economic content: a more positive $\widehat{ctp}$ in my convention corresponds to a more negative x^{Treas} in the Jiang et al. (2021) convention, with both indicating a more convenient U.S. Treasury. Under this mapping, a contemporaneous regression of Δs on Δx^{Treas} that delivers a negative coefficient in the Jiang et al. (2021) convention should deliver a *positive* coefficient in my convention; a forward-predictive regression of $\Delta s_{t,t+k}$ on the level of x_t^{Treas} that delivers a positive coefficient in their convention (which Proposition 1 of Jiang et al. (2021) predicts) should deliver a *negative* coefficient in mine.

Aggregation conventions. Jiang et al. (2021) aggregate by first computing country-level changes and then averaging across the cross-section, using an equally-weighted basket of log bilateral exchange rates against the dollar; both the average basis and the average exchange rate enter the regression at the same equal weights. I follow their construction in Table IA.7 of this Internet Appendix: the dependent variable is the h-month forward log change in the equally-weighted basket $\overline{s}_t^{eq} \equiv \frac{1}{6}\sum_j s_t^{USD,j}$, sign-aligned bilateral logs, and the predictor is the cross-sectional average ten-year cross-country term premium across the same six major pairs, the basket that matches their construction. Because the cross-sectional weighting on both sides of the regression is identical to Jiang et al. (2021)'s construction at one year, the aggregate-level forward-predictive results in my retained battery in the replication materials parallel the long-horizon predictive content discussed in Section IV.F of Jiang et al. (2021); the contemporaneous-regression analogue to their Table III, which uses concurrent monthly changes rather than forward-predictive changes, is reported in Panel A of Table IA.7 of this Internet Appendix for direct side-by-side comparison.

Aggregate reconciliation. Table IA.7 presents the aggregate-level reconciliation in two complementary forms on the equally-weighted dollar basket: Panel A reports the contemporaneous monthly change-on-change specification paralleling Table III of Jiang et al. (2021), and Panel B reports the forward-predictive specification at the spot, three-month, and twelve-month horizons paralleling the Section IV.F discussion. Both panels use the cross-country government bond term premium as the primary predictor and report the cross-country swap

term premium for completeness.

The aggregate-level evidence in Table IA.7 is consistent with the Jiang et al. (2021) framework at the long end of the term structure on the forward dimension and informative on the attribution. In the twelve-month forward specification, the level of $\overline{ctp}_t^{gov,USD}$ relates to $\Delta\overline{s}_{t,t+12}^{eq}$ with a coefficient of -0.060 at $t = -2.01$ and R^2 of 6.2 percent, which translates to the positive basis sign predicted by Proposition 1 of Jiang et al. (2021); the contemporaneous relation on the cross-country government bond term premium is statistically indistinguishable from zero. The quoted-basis average, standing in for the cross-country swap term premium, carries no forward relation at all, so the content of the sign relation lives in the cross-country swap spread differential component of the bridge identity rather than in the swap term premium itself, consistent with the attribution stated above. The equally-weighted-basket aggregation on the cross-country government bond term premium therefore extends the Jiang et al. (2021) sign structure to the long end of the term structure, at horizons and with the cross-sectional aggregation that match their main analysis, and I read it as a reconciliation rather than a predictive claim, per the inference caveats stated at the head of this section.

Internet Appendix Figures and Tables

Table IA.8: **Decomposition: Quoted Cross-Currency Basis and Cross-Country Swap Spread Differential Components.** For every column of Table 1, this table reports the identical regression estimated with the two directly quoted additive components of the dependent variable: the bilateral cross-currency basis $x_{10}^{j/USD}$ (Panel A), which by the containment result represents the cross-country swap term premium up to the small pair-specific residual, and the cross-country swap spread differential $\Delta SS_{10}^{USD,j}$ (Panel B). The estimation sample is identical to Table 1 row for row, so the decomposition is not sample dependent. Since the dependent variable of Table 1 is $x + \Delta SS$ by construction, the coefficients satisfy $\widehat{\beta}^{\text{gov}} = \widehat{\beta}^{x} + \widehat{\beta}^{\Delta SS}$ exactly in every cell on the unrounded coefficients, so printed-cell arithmetic can differ in the terminal digit, with no extracted quantity anywhere in the decomposition. The Panel A loadings are numerically close to the corresponding cross-country swap term premium loadings in Table IA.8, as the two differ only by the pair-specific residual. Specifications, fixed effects, and Driscoll and Kraay (1998) standard errors are identical to Table 1. Panel C reports the symmetric variance share of the composite carried by the basis-side component in each sub-period (the swap-spread differential carries the complement to one hundred), documenting that both components matter and that the division between them is time-varying and pair-varying. t-statistics in parentheses. Notes: $^{*}\, p \leq .10$; $^{**}\, p \leq .05$; $^{***}\, p \leq .01$.

	(1)	(2)	(3)	(4)	(5)
Panel A: Dependent variable: quoted cross-currency basis $x_{10}^{j/USD}$					
$\log(\widetilde{d})_{\text{US}}$	-0.194^{***}	-0.197^{***}	-0.093^{*}	-0.356^{***}	-0.287^{***}
	(-3.38)	(-3.58)	(-1.75)	(-3.85)	(-2.67)
$\log(\widetilde{d})_{j}$	$+0.170^{***}$	$+0.139^{***}$	$+0.160^{***}$	$+0.100^{**}$	$+0.092^{*}$
	$(+2.83)$	$(+2.67)$	$(+3.63)$	$(+2.01)$	$(+1.88)$
$\widetilde{\text{WAM}}_{\text{US}}$	-0.023	-0.002	-0.066^{**}	-0.060^{***}	$+0.021$
	(-0.80)	(-0.08)	(-2.15)	(-2.67)	$(+0.79)$
$\widetilde{\text{WAM}}_{j}$	-0.012	-0.001	$+0.001$	-0.013	-0.012
	(-1.18)	(-0.06)	$(+0.12)$	(-1.36)	(-1.25)
QE_t^{US}		$+0.057^{*}$	$+0.166^{***}$	$+0.144^{***}$	$+0.096^{***}$
		$(+1.81)$	$(+3.66)$	$(+3.42)$	$(+3.89)$
QE_t^{j}		-0.152^{***}	-0.157^{***}	-0.207^{***}	-0.194^{***}
		(-3.29)	(-3.45)	(-5.61)	(-4.81)
HKM_t^{z}			$+0.050^{**}$	$+0.008$	$+0.021$
			$(+2.49)$	$(+0.22)$	$(+0.79)$
$\text{HKM}_t^{z} \times \text{QE}_t^{\text{US}}$			$+0.135^{***}$	$+0.147^{***}$	$+0.103^{***}$
			$(+3.08)$	$(+3.05)$	$(+2.90)$
$\text{rGDP}_t^{\text{US},z}$				$+0.049^{**}$	$+0.026^{*}$
				$(+2.31)$	$(+1.70)$
$\text{rGDP}_t^{j,z}$				-0.014	-0.015
				(-1.41)	(-1.49)
$\text{ir3m}_t^{\text{US},z}$				$+0.081^{**}$	$+0.108^{***}$
				$(+1.97)$	$(+3.29)$

Continued on next page

Table IA.8 continued

	(1)	(2)	(3)	(4)	(5)
$\text{ir3m}_t^{j,z}$				-0.199^{***}	-0.200^{***}
				(-6.24)	(-5.35)
$\text{TS}_t^{\text{US},z}$				$+0.010$	$+0.025$
				$(+0.45)$	$(+1.29)$
$\text{TS}_t^{j,z}$				-0.048^{**}	-0.047^{**}
				(-2.33)	(-2.19)
$\text{Unemp}_t^{\text{US},z}$				$+0.042^{**}$	$+0.013$
				$(+2.09)$	$(+1.10)$
$\text{Unemp}_t^{j,z}$				-0.041^{**}	-0.043^{***}
				(-2.28)	(-2.69)
$\text{Fiscal bal}_t^{\text{US}}$				-0.002	-0.006
				(-0.43)	(-1.60)
Fiscal bal_t^{j}				-0.000	-0.001
				(-0.00)	(-0.24)
$\text{CPI}_t^{\text{US},z}$					$+0.001$
					$(+0.05)$
$\text{CPI}_t^{j,z}$					-0.022
					(-1.39)
Dollar_t^{z}					-0.078^{***}
					(-3.18)
BBB_t^{z}					$+0.007$
					$(+0.35)$
MOVE_t^{z}					-0.013
					(-1.51)
VIX_t^{z}					$+0.021^{**}$
					$(+2.36)$
Pair fixed effects	Yes	Yes	Yes	Yes	Yes
Within-R^2	0.131	0.212	0.289	0.491	0.525
N	2,043	2,043	1,952	1,911	1,911
Panel B: Dependent variable: $\Delta SS_{10}^{\text{USD},j}$ *(cross-country swap spread differential)*					
$\log(\widetilde{d})_{\text{US}}$	-0.857^{***}	-0.953^{***}	-0.791^{***}	-0.923^{***}	-0.847^{***}

Continued on next page

Table IA.8 continued

	(1)	(2)	(3)	(4)	(5)
	(−12.05)	(−14.53)	(−14.94)	(−8.79)	(−5.70)
$\log(\widetilde{d})_j$	+0.351***	+0.339***	+0.336***	+0.075	+0.070
	(+7.81)	(+7.29)	(+6.57)	(+1.61)	(+1.48)
$\widetilde{\text{WAM}}_{\text{US}}$	−0.037	+0.007	−0.068**	−0.055*	−0.083**
	(−1.24)	(+0.22)	(−2.03)	(−1.91)	(−2.52)
$\widetilde{\text{WAM}}_j$	+0.035*	+0.040**	+0.041*	+0.042***	+0.042***
	(+1.83)	(+1.98)	(+1.89)	(+2.61)	(+2.62)
QE_t^{US}		+0.102***	+0.098	+0.092**	+0.159***
		(+3.11)	(+1.61)	(+1.98)	(+2.89)
QE_t^j		−0.046	−0.030	−0.081***	−0.085***
		(−1.06)	(−0.73)	(−2.78)	(−2.92)
HKM_t^z			+0.096***	+0.048	+0.029
			(+3.57)	(+1.06)	(+0.71)
$\text{HKM}_t^z \times \text{QE}_t^{\text{US}}$			−0.028	+0.015	+0.085
			(−0.52)	(+0.25)	(+1.23)
$\text{rGDP}_t^{\text{US},z}$				−0.010	+0.004
				(−0.37)	(+0.19)
$\text{rGDP}_t^{j,z}$				−0.003	−0.001
				(−0.22)	(−0.09)
$\text{ir3m}_t^{\text{US},z}$				+0.034	+0.075**
				(+0.90)	(+2.09)
$\text{ir3m}_t^{j,z}$				−0.076***	−0.055*
				(−2.93)	(−1.68)
$\text{TS}_t^{\text{US},z}$				−0.057**	−0.048**
				(−1.98)	(−1.99)
$\text{TS}_t^{j,z}$				+0.081***	+0.094***
				(+3.05)	(+4.29)
$\text{Unemp}_t^{\text{US},z}$				−0.049*	−0.050*
				(−1.95)	(−1.85)
$\text{Unemp}_t^{j,z}$				+0.112***	+0.108***
				(+3.09)	(+3.14)
$\text{Fiscal bal}_t^{\text{US}}$				−0.005	+0.000

Continued on next page

Table IA.8 continued

	(1)	(2)	(3)	(4)	(5)
				(−0.72)	(+0.02)
Fiscal bal_t^j				−0.012	−0.013*
				(−1.53)	(−1.69)
$\text{CPI}_t^{\text{US},z}$					−0.027
					(−1.39)
$\text{CPI}_t^{j,z}$					−0.011
					(−0.51)
Dollar_t^z					+0.033
					(+1.08)
BBB_t^z					−0.019
					(−0.57)
MOVE_t^z					−0.009
					(−0.73)
VIX_t^z					+0.031**
					(+2.25)
Pair fixed effects	Yes	Yes	Yes	Yes	Yes
Within-R^2	0.440	0.457	0.491	0.636	0.646
N	2,043	2,043	1,952	1,911	1,911

Table IA.9: **Robustness Decomposition on the Extracted Premia.** For every column of the extracted-premium specification of Internet Appendix Table IA.1, this table reports the identical regression estimated with the two additive components of the extracted dependent variable: the extracted cross-country swap term premium $\widehat{ctp}_{10}^{swap,j,USD}$ (Panel A) and the cross-country swap spread differential $\Delta SS_{10}^{USD,j}$ (Panel B), on the identified sample of Internet Appendix Table IA.1 row for row, with the bridge identity holding exactly in every cell on the unrounded coefficients, so printed-cell arithmetic can differ in the terminal digit. The Panel A loadings are numerically close to the quoted-basis loadings of Table IA.8, the two differing only by the pair-specific residual and the identified-sample restriction. Specifications, fixed effects, and Driscoll and Kraay (1998) standard errors are identical to Internet Appendix Table IA.1. t-statistics in parentheses. Notes: * $p \leq .10$; ** $p \leq .05$; *** $p \leq .01$.

	(1)	(2)	(3)	(4)	(5)
Panel A: Dependent variable: extracted $\widehat{ctp}_{10}^{swap,j,USD}$					
$\log(\widetilde{d})_{\text{US}}$	-0.187^{***}	-0.184^{***}	-0.107^{**}	-0.366^{***}	-0.310^{***}
	(-3.13)	(-3.07)	(-1.98)	(-4.43)	(-3.44)
$\log(\widetilde{d})_{j}$	$+0.190^{***}$	$+0.157^{***}$	$+0.169^{***}$	$+0.132^{**}$	$+0.118^{**}$
	$(+3.06)$	$(+2.84)$	$(+3.57)$	$(+2.34)$	$(+2.11)$
$\widetilde{\text{WAM}}_{\text{US}}$	-0.000	$+0.020$	-0.036	-0.040^{*}	$+0.032$
	(-0.00)	$(+0.71)$	(-1.64)	(-1.93)	$(+1.40)$
$\widetilde{\text{WAM}}_{j}$	-0.018	-0.006	-0.003	-0.020^{*}	-0.020
	(-1.51)	(-0.52)	(-0.23)	(-1.78)	(-1.64)
QE_t^{US}		$+0.056^{*}$	$+0.169^{***}$	$+0.156^{***}$	$+0.099^{***}$
		$(+1.71)$	$(+3.85)$	$(+4.32)$	$(+4.56)$
QE_t^{j}		-0.162^{***}	-0.168^{***}	-0.200^{***}	-0.188^{***}
		(-3.68)	(-3.82)	(-5.48)	(-4.72)
HKM_t^{z}			$+0.037^{**}$	-0.008	$+0.007$
			$(+2.52)$	(-0.23)	$(+0.31)$
$\text{HKM}_t^{z} \times \text{QE}_t^{\text{US}}$			$+0.140^{***}$	$+0.153^{***}$	$+0.101^{***}$
			$(+3.62)$	$(+3.58)$	$(+3.27)$
$\text{rGDP}_t^{\text{US},z}$				$+0.037^{*}$	$+0.018$
				$(+1.77)$	$(+1.16)$
$\text{rGDP}_t^{j,z}$				-0.010	-0.013
				(-0.91)	(-1.17)
$\text{ir3m}_t^{\text{US},z}$				$+0.098^{**}$	$+0.107^{***}$
				$(+2.35)$	$(+3.10)$
$\text{ir3m}_t^{j,z}$				-0.192^{***}	-0.190^{***}
				(-6.09)	(-5.00)
$\text{TS}_t^{\text{US},z}$				$+0.018$	$+0.027$
				$(+0.79)$	$(+1.27)$

Table IA.9 continued

	(1)	(2)	(3)	(4)	(5)
$TS_t^{j,z}$				-0.050^{**}	-0.047^{**}
				(-2.34)	(-2.09)
$Unemp_t^{US,z}$				$+0.036^{**}$	$+0.011$
				$(+2.22)$	$(+0.93)$
$Unemp_t^{j,z}$				-0.055^{***}	-0.054^{***}
				(-3.05)	(-3.09)
Fiscal bal_t^{US}				-0.002	-0.006^{*}
				(-0.44)	(-1.67)
Fiscal bal_t^{j}				-0.004	-0.005
				(-1.33)	(-1.61)
$CPI_t^{US,z}$					$+0.011$
					$(+0.76)$
$CPI_t^{j,z}$					-0.025^{*}
					(-1.81)
$Dollar_t^{z}$					-0.071^{***}
					(-3.21)
BBB_t^{z}					$+0.013$
					$(+0.77)$
$MOVE_t^{z}$					-0.011
					(-1.31)
VIX_t^{z}					$+0.014^{**}$
					$(+2.00)$
Pair fixed effects	Yes	Yes	Yes	Yes	Yes
Within-R^2	0.129	0.229	0.299	0.486	0.514
N	1,888	1,888	1,823	1,782	1,782
Panel B: Dependent variable: $\Delta SS_{10}^{USD,j}$					
$\log(\widetilde{d})_{US}$	-0.879^{***}	-0.969^{***}	-0.846^{***}	-0.969^{***}	-0.905^{***}
	(-11.40)	(-12.83)	(-16.61)	(-11.75)	(-6.86)
$\log(\widetilde{d})_{j}$	$+0.390^{***}$	$+0.375^{***}$	$+0.366^{***}$	$+0.152^{**}$	$+0.149^{**}$
	$(+7.74)$	$(+7.31)$	$(+6.99)$	$(+2.38)$	$(+2.46)$
$\widetilde{WAM}_{US}$	-0.013	$+0.029$	-0.033	-0.033	-0.078^{**}
	(-0.48)	$(+1.02)$	(-1.05)	(-1.21)	(-2.37)

Table IA.9 continued

	(1)	(2)	(3)	(4)	(5)
$\widetilde{\text{WAM}}_j$	$+0.036^{*}$	$+0.042^{**}$	$+0.044^{**}$	$+0.036^{*}$	$+0.036^{**}$
	$(+1.78)$	$(+2.04)$	$(+1.98)$	$(+1.96)$	$(+1.97)$
QE_t^{US}		$+0.101^{***}$	$+0.101$	$+0.106^{**}$	$+0.167^{***}$
		$(+3.01)$	$(+1.54)$	$(+2.57)$	$(+3.29)$
QE_t^{j}		-0.057	-0.043	-0.067^{**}	-0.073^{***}
		(-1.38)	(-1.06)	(-2.55)	(-2.78)
HKM_t^{z}			$+0.074^{***}$	$+0.010$	-0.000
			$(+2.92)$	$(+0.28)$	(-0.01)
$\text{HKM}_t^{z} \times \text{QE}_t^{\text{US}}$			-0.019	$+0.027$	$+0.087$
			(-0.31)	$(+0.51)$	$(+1.38)$
$\text{rGDP}_t^{\text{US},z}$				-0.018	-0.002
				(-0.79)	(-0.11)
$\text{rGDP}_t^{j,z}$				-0.001	$+0.001$
				(-0.05)	$(+0.07)$
$\text{ir3m}_t^{\text{US},z}$				$+0.065^{*}$	$+0.089^{**}$
				$(+1.79)$	$(+2.56)$
$\text{ir3m}_t^{j,z}$				-0.069^{***}	-0.051
				(-2.70)	(-1.61)
$\text{TS}_t^{\text{US},z}$				-0.044	-0.041
				(-1.53)	(-1.56)
$\text{TS}_t^{j,z}$				$+0.088^{***}$	$+0.098^{***}$
				$(+3.31)$	$(+4.16)$
$\text{Unemp}_t^{\text{US},z}$				-0.062^{**}	-0.052^{*}
				(-2.17)	(-1.95)
$\text{Unemp}_t^{j,z}$				$+0.084^{***}$	$+0.086^{***}$
				$(+2.64)$	$(+2.69)$
$\text{Fiscal bal}_t^{\text{US}}$				-0.005	$+0.000$
				(-0.91)	$(+0.01)$
Fiscal bal_t^{j}				-0.017^{**}	-0.017^{**}
				(-2.14)	(-2.27)
$\text{CPI}_t^{\text{US},z}$					-0.012
					(-0.67)
$\text{CPI}_t^{j,z}$					-0.010

Table IA.9 continued

	(1)	(2)	(3)	(4)	(5)
					(−0.47)
Dollar_t^z					+0.048
					(+1.57)
BBB_t^z					−0.016
					(−0.56)
MOVE_t^z					−0.001
					(−0.10)
VIX_t^z					+0.020
					(+1.47)
Pair fixed effects	Yes	Yes	Yes	Yes	Yes
Within-R^2	0.451	0.470	0.491	0.645	0.652
N	1,888	1,888	1,823	1,782	1,782

Table IA.10: **Variable Definitions and Data Sources.** This table lists the variables used in the empirical analysis of the main text and the Internet Appendix, with construction details, frequency, and data sources. Cross-country term premia and bilateral frictions are extracted via the triangular methodology described in Section 2 of the main text. Government bond cross-country term premia $\widehat{ctp}_N^{gov,j,USD}$ are constructed as $\widehat{ctp}_N^{swap,j,USD} + \Delta SS_N^{USD,j}$ under the bridge identity, with the cross-country swap spread differential $\Delta SS_N^{USD,j} \equiv SS_N^{USD} - SS_N^j$ and the per-country swap spread $SS_{t,N}^c \equiv \text{IRS}_{t,N}^c - r_{t,N}^{gov,c}$. The privately held free-float supply variables on which the preferred-habitat regressions of Section 4.3.1 are constructed are net of central bank holdings of own-currency sovereign debt; the construction details for the central-bank-holdings adjustment, including the reconstruction of the par-weighted maturity of central bank holdings, are documented in Internet Appendix Sections D.7 and D.5.

Variable	Definition	Frequency	Source
Panel A: Cross-Country Term Premia and Frictions			
$\widehat{ctp}_N^{swap,j,USD}$	Cross-country swap term premium extracted from the triangulation, in pp	Monthly	Extracted
$\widehat{ctp}_N^{gov,j,USD}$	Cross-country government bond term premium, $\widehat{ctp}_N^{swap,j,USD} + \Delta SS_N^{USD,j}$, in pp	Monthly	Extracted
$\pi_N^{j/USD}$	Bilateral friction (pure parity deviation) extracted from the triangular system solution, in pp	Monthly	Extracted
$x_N^{j/USD}$	Bilateral cross-currency basis at tenor N, $x = ctp + \pi$, in pp	Monthly	Bloomberg cross-currency swap quotes
Panel B: Yields and Swap Rates (ten currencies: US, EA, JP, GB, CH, CA, AU, NZ, SE, NO)			
r_N^c	Zero-coupon sovereign yield at tenor $N \in \{2, 5, 10\}$ years, in pp	Monthly	Bloomberg country-specific zero curves
IRS_N^c	Single-currency interest-rate-swap rate at tenor N, in pp; for transitioning countries post-IBOR cessation I use the Reference Rate (RFR)-IRS rate	Monthly	Bloomberg
$SS_{t,N}^c$	Per-country swap spread at tenor N, $SS_{t,N}^c \equiv \text{IRS}_{t,N}^c - r_{t,N}^{gov,c}$, in pp	Monthly	Constructed
$\Delta SS_N^{USD,j}$	Cross-country swap spread differential, $SS_N^{USD} - SS_N^j$, in pp	Monthly	Constructed

Continued on next page

Table IA.10 continued

Variable	Definition	Frequency	Source
Panel C: Privately Held Free-Float Supply (Six Countries)			
$d_{c,t}$	Gross general-government sovereign debt as a share of GDP, in percent	Quarterly	IMF WEO; FRED (US)
$c_{c,t}$	Central bank holdings of own-currency sovereign debt as a share of GDP, in percent	Monthly	National central bank disclosures (Internet Appendix Section D.7)
$\widetilde{d}_{c,t}$	Privately held free-float debt-to-GDP, defined as $d_{c,t} - c_{c,t}$, in percent	Monthly	Constructed
$w_{c,t}$	Par-weighted average maturity of total outstanding sovereign debt, in years	Quarterly	National debt-management offices
$w^{cb}_{c,t}$	Par-weighted average maturity of central bank holdings, in years	Monthly	Reconstructed (Internet Appendix Section D.5)
$\widetilde{w}_{c,t}$	Par-weighted average maturity of the privately held free-float, $\widetilde{w}_{c,t} \equiv (d_{c,t}w_{c,t} - c_{c,t}w^{cb}_{c,t})/(d_{c,t} - c_{c,t})$, in years	Monthly	Constructed
Panel D: Foreign Official Reserve Demand, Fiscal and Funding State Variables			
$FO_{US,t}$	U.S. Treasury securities held by foreign official institutions (foreign central banks, governments, and international organisations), expressed as a share of U.S. nominal GDP. Enters the supply regressions as the inelastic reserve-demand control.	Quarterly	Federal Reserve Financial Accounts (Z.1), series FL263061130, via FRED (BOGZ1FL263061130Q); GDP from FRED

Continued on next page

Table IA.10 continued

Variable	Definition	Frequency	Source
$FP_{US,t}$	U.S. Treasury securities held by private foreign investors, constructed as total rest-of-world Treasury holdings at market value less foreign official holdings, as a share of GDP.	Quarterly	Federal Reserve Z.1, series LM263061105 less FL263061130, via FRED
$FOF_{US,t}$	Twelve-month adjusted net foreign official Treasury flow relative to GDP, the valuation-adjusted flow margin of the official-demand analysis (Internet Appendix Section F.4.6). Chained across the two release files at June 2011; final release, series ends December 2020.	Monthly	Bertaut and Judson (2014) public data release, Federal Reserve Board; GDP from FRED
$\widetilde{d}^{e}_{US,t}$	Effective float: gross debt-to-GDP net of both domestic central bank holdings and foreign official holdings, $d_{US,t} - c_{US,t} - FO_{US,t}$. Robustness variant that nets out both price-insensitive holders.	Monthly	Constructed; components as above
$b_{c,t}$	General government budget balance as a share of GDP, entering the supply regressions as a fiscal control on both the U.S. and comparison-country side.	Monthly	Bloomberg
$\text{Deficit}^{z}_{US,t}$	Trailing twelve-month sum of the U.S. federal budget deficit from the Monthly Treasury Statement, standardised to unit variance over the sample; the fiscal-surge state variable in the state-dependence regressions.	Monthly	U.S. Treasury Monthly Treasury Statement, via FRED (MTSDS133 FMS)

Continued on next page

Table IA.10 continued

Variable	Definition	Frequency	Source
f_t^{3M}	Short-end dollar funding factor: cross-sectional mean of the three-month cross-currency basis across the panel currencies, with the IBOR-era quotation chained onto its risk-free-rate successor at the benchmark transition by the overlap reconciliation of Section D.2.	Monthly	Bloomberg; construction in the public rates database
$\text{Slope}_t^{USD,j}$	Cross-country term-structure slope differential: the ten-year less two-year fitted sovereign zero-coupon slope of the comparison country less that of the United States.	Monthly	Bloomberg fitted sovereign curves (BLC2)
Panel E: Country-Specific Quantitative Easing Indicators			
QE_t^c	Indicator equal to one in months during which country c's central bank operated an officially announced quantitative easing programme, identified from the central bank's own programme announcement and termination disclosures, and zero otherwise	Monthly	Constructed (Internet Appendix Section F.5, Table IA.38)
Panel F: Intermediary and Macro-Financial Variables			
HKM_t^z	Standardised intermediary capital ratio of He et al. (2017), monthly, the published vintage ending May 2025 and standardised over the full sample	Monthly	He et al. (2017); FRB H.4.1
AEM_t^z	Standardised broker-dealer leverage factor of Adrian et al. (2014) (used in the Internet Appendix comparison)	Quarterly	Adrian et al. (2014); FRED
Dollar_t^z	Standardised nominal broad dollar index	Monthly	FRED (DTWEXBGS)
BBB_t^z	Standardised ICE BofA BBB U.S. corporate effective yield spread to Treasuries	Monthly	ICE; FRED

Continued on next page

Table IA.10 continued

Variable	Definition	Frequency	Source
MOVE_t^z	Standardised Merrill Lynch Treasury option-implied volatility index	Monthly	Bloomberg
VIX_t^z	Standardised CBOE volatility index	Monthly	Bloomberg
Panel G: Dollar-Prediction Variables			
$s_t^{USD,j}$	Log spot exchange rate, sign-aligned so that an increase corresponds to USD appreciation against j	Monthly	FRED FX series
$\overline{s}_t^{eq}$	Equally-weighted dollar basket, $\overline{s}_t^{eq} \equiv \frac{1}{6}\sum_j s_t^{USD,j}$	Monthly	Constructed
$\overline{IRD}_{N,t}^{USD}$	Cross-sectional average N-year interest rate differential, $\frac{1}{6}\sum_j (r_{N,t}^{USD} - r_{N,t}^{j})$	Monthly	Constructed

Table IA.15 records the margins of the ten-year differential by pair and regime, the raw government yield differential, the swap spread differential, and the recovered swap and government bond differentials, on the windows stated in its caption.

Table IA.16 lists, country by country, the sources used in the free-float construction.

Table IA.11: **Cross-Country Government Bond Term Premium and Privately Held Sovereign Supply Across Tenors (2-Year and 5-Year).** The five-column free-float panel specification of Table 1 estimated separately at the two-year (Panel A) and five-year (Panel B) tenors. The dependent variable is the cross-country government bond term premium differential at the relevant tenor, in percentage points. The column structure and sample mirror Table 1. All specifications include currency-pair fixed effects; standard errors are Driscoll and Kraay (1998) kernel estimates with Bartlett kernel and Schwert-rule lag. t-statistics in parentheses. Notes: $^{*}\,p \leq .10$; $^{**}\,p \leq .05$; $^{***}\,p \leq .01$. Sample: January 2002 to June 2026 in columns (1) and (2); columns (3) through (5) end in May 2025 with the published He et al. (2017) vintage. Throughout, the quoted bilateral cross-currency basis x is used to represent the cross-country swap term premium, which it equals up to the pair-specific residual bounded within quoted half-spreads.

	(1)	(2)	(3)	(4)	(5)
Panel A: Two-year tenor					
$\log(\widetilde{d})_{\text{US}}$	-0.682^{***}	-0.723^{***}	-0.571^{***}	-0.936^{***}	-0.848^{***}
	(-7.22)	(-6.89)	(-3.36)	(-3.94)	(-3.18)
$\log(\widetilde{d})_{j}$	$+0.238^{**}$	$+0.220^{**}$	$+0.239^{***}$	$+0.119$	$+0.094$
	$(+2.55)$	$(+2.49)$	$(+2.61)$	$(+1.17)$	$(+0.85)$
$\widetilde{\text{WAM}}_{\text{US}}$	-0.166^{***}	-0.139^{**}	-0.220^{***}	-0.249^{***}	-0.120^{**}
	(-2.73)	(-2.15)	(-3.25)	(-3.75)	(-2.45)
$\widetilde{\text{WAM}}_{j}$	$+0.047^{**}$	$+0.053^{**}$	$+0.055^{**}$	$+0.021$	$+0.017$
	$(+2.24)$	$(+2.41)$	$(+2.20)$	$(+0.93)$	$(+0.72)$
QE_t^{US}		$+0.067$	$+0.048$	$+0.099$	$+0.069$
		$(+1.47)$	$(+0.52)$	$(+1.06)$	$(+1.09)$
QE_t^{j}		-0.085	-0.055	-0.111^{**}	-0.088
		(-1.23)	(-0.84)	(-2.11)	(-1.64)
HKM_t^{z}			$+0.103$	-0.001	$+0.126^{**}$
			$(+1.36)$	(-0.01)	$(+2.14)$
$\text{HKM}_t^{z} \times \text{QE}_t^{\text{US}}$			-0.048	$+0.046$	$+0.074$
			(-0.54)	$(+0.38)$	$(+0.86)$
$\text{rGDP}_t^{\text{US},z}$				-0.015	$+0.007$
				(-0.35)	$(+0.28)$
$\text{rGDP}_t^{j,z}$				$+0.011$	$+0.008$
				$(+0.54)$	$(+0.45)$
$\text{ir3m}_t^{\text{US},z}$				$+0.125$	$+0.021$
				$(+1.54)$	$(+0.29)$
$\text{ir3m}_t^{j,z}$				-0.302^{***}	-0.293^{***}
				(-5.35)	(-3.87)
$\text{TS}_t^{\text{US},z}$				$+0.022$	-0.039

Continued on next page

Table IA.11 continued

	(1)	(2)	(3)	(4)	(5)
				(+0.34)	(−0.64)
$TS_t^{j,z}$				−0.040	−0.028
				(−0.70)	(−0.43)
$Unemp_t^{US,z}$				−0.047	−0.017
				(−1.00)	(−0.50)
$Unemp_t^{j,z}$				−0.031	−0.032
				(−0.65)	(−0.87)
Fiscal bal_t^{US}				−0.002	−0.006
				(−0.20)	(−0.50)
Fiscal bal_t^{j}				−0.002	−0.005
				(−0.16)	(−0.48)
$CPI_t^{US,z}$					+0.001
					(+0.01)
$CPI_t^{j,z}$					−0.063*
					(−1.73)
$Dollar_t^{z}$					−0.154***
					(−2.96)
BBB_t^{z}					+0.131**
					(+2.28)
$MOVE_t^{z}$					+0.008
					(+0.41)
VIX_t^{z}					+0.063***
					(+3.06)
Pair fixed effects	Yes	Yes	Yes	Yes	Yes
Within-R^2	0.267	0.273	0.299	0.395	0.431
N	2,043	2,043	1,952	1,911	1,911
Panel B: Five-year tenor					
$\log(\widetilde{d})_{US}$	−0.891***	−0.980***	−0.811***	−1.202***	−1.022***
	(−8.38)	(−9.40)	(−6.80)	(−5.48)	(−4.20)
$\log(\widetilde{d})_{j}$	+0.401***	+0.368***	+0.388***	+0.134	+0.113
	(+4.63)	(+4.65)	(+4.91)	(+1.63)	(+1.38)

Continued on next page

Table IA.11 continued

	(1)	(2)	(3)	(4)	(5)
$\widetilde{\text{WAM}}_{\text{US}}$	-0.126^{**}	-0.072	-0.168^{***}	-0.162^{***}	-0.061
	(-2.49)	(-1.38)	(-2.60)	(-2.78)	(-1.36)
$\widetilde{\text{WAM}}_{j}$	$+0.010$	$+0.022$	$+0.025$	$+0.008$	$+0.006$
	$(+0.45)$	$(+0.95)$	$(+0.97)$	$(+0.41)$	$(+0.28)$
QE_t^{US}		$+0.134^{**}$	$+0.161$	$+0.115$	$+0.131^{*}$
		$(+2.55)$	$(+1.55)$	$(+1.37)$	$(+1.88)$
QE_t^{j}		-0.153^{*}	-0.128	-0.213^{***}	-0.193^{***}
		(-1.80)	(-1.56)	(-3.44)	(-3.00)
HKM_t^{z}			$+0.109^{**}$	$+0.019$	$+0.087$
			$(+2.24)$	$(+0.19)$	$(+1.56)$
$\text{HKM}_t^{z} \times \text{QE}_t^{\text{US}}$			$+0.009$	$+0.046$	$+0.092$
			$(+0.12)$	$(+0.45)$	$(+0.98)$
$\text{rGDP}_t^{\text{US},z}$				$+0.039$	$+0.042$
				$(+0.88)$	$(+1.41)$
$\text{rGDP}_t^{j,z}$				-0.006	-0.006
				(-0.34)	(-0.42)
$\text{ir3m}_t^{\text{US},z}$				$+0.136^{*}$	$+0.137^{*}$
				$(+1.66)$	$(+1.75)$
$\text{ir3m}_t^{j,z}$				-0.286^{***}	-0.273^{***}
				(-4.84)	(-3.54)
$\text{TS}_t^{\text{US},z}$				-0.022	-0.041
				(-0.40)	(-0.87)
$\text{TS}_t^{j,z}$				$+0.015$	$+0.029$
				$(+0.31)$	$(+0.65)$
$\text{Unemp}_t^{\text{US},z}$				$+0.005$	$+0.001$
				$(+0.13)$	$(+0.02)$
$\text{Unemp}_t^{j,z}$				$+0.021$	$+0.014$
				$(+0.47)$	$(+0.39)$
$\text{Fiscal bal}_t^{\text{US}}$				-0.007	-0.008
				(-0.79)	(-0.77)
Fiscal bal_t^{j}				-0.011	-0.014
				(-1.00)	(-1.34)

Continued on next page

Table IA.11 continued

	(1)	(2)	(3)	(4)	(5)
$\text{CPI}_t^{\text{US},z}$					-0.031
					(-0.79)
$\text{CPI}_t^{j,z}$					-0.051
					(-1.46)
Dollar_t^z					-0.112^{**}
					(-2.23)
BBB_t^z					$+0.049$
					$(+0.97)$
MOVE_t^z					$+0.000$
					$(+0.00)$
VIX_t^z					$+0.070^{***}$
					$(+3.25)$
Pair fixed effects	Yes	Yes	Yes	Yes	Yes
Within-R^2	0.394	0.417	0.451	0.599	0.630
N	2,043	2,043	1,952	1,911	1,911

J Interpretation: What the Recovered Premium Bundles

What prices can and cannot decide. Separating a physical-measure term premium from a broader currency-level price factor clarifies which questions prices can decide. A natural conjecture is that realised returns could perform the separation, on the reasoning that a premium must be earned as pecuniary return while a convenience wedge is consumed as services. The conjecture fails, instructively: whatever the cause of a low price, risk compensation demanded by the marginal holder or convenience forgone by inframarginal ones, the price mechanically implies the same expected pecuniary excess return for whoever holds the instrument, so realised returns confirm the price-implied premium under either attribution. Realised returns separate premium content from expectations content, which the construction here already accomplishes without an expectations model; they cannot attribute the premium to its sources. The general statement is that under the expected-excess-return definition, decomposing a term premium into risk compensation, convenience, and intermediary shadow costs is an attribution exercise no price-based method performs, and none of the existing estimates performs it either: the model-based term premia of Adrian et al. (2013) and Kim and Wright (2005) for the United States bundle the safety and liquidity convenience of Treasuries exactly as Krishnamurthy and Vissing-Jorgensen (2012) document, and the literature has never required that bundle to be resolved before calling those estimates term premia. Held to a symmetric standard, then, three questions can be asked of the recovered differential: whether it is premium rather than expectations, decided without an expectations model by construction; whether it is currency-level rather than pair-level, decided by the triangular containment within quoted spreads; and what mixture the premium bundles, which is undecidable from prices for this measurement and for every other. What distinguishes the treatment here is that the bundle is stated openly, its pair-level part is bounded, every testable implication is tested, and the consequences of the untestable direction are priced by the sensitivity frontier.

The taxonomy of the currency-level component then closes completely, with each element addressed by evidence of the appropriate kind: (1) expected rolled funding costs, the one component that is a cash-flow cost rather than compensation, generate testable restrictions and are rejected by an order of magnitude by the excess-sensitivity bound below; (2) transitory balance-sheet and regulatory-window effects are absent at long tenors and controlled directly through intermediary capital; (3) any component residing in the swap rates themselves, including benchmark credit and collateral remuneration, cancels exactly in the government bridge by Corollary 1; and (4) persistent priced compensation for currency-attached duration risk,

whatever its microfoundation in hedging demand, reserve demand, or habitat, is a term premium under any definition in which the marginal investor is global, which is precisely the definition the recovered cross-country term-premium differential's own behaviour supports: it is priced by government debt duration supply measured on the globally available free float, and it transmits policy shocks to the announcing country's curve within days. What remains disputable after the evidence is nomenclature for component (4), not economics.

The convergence of the two vocabularies at the long end can be made precise. Define the horizon-N currency premium against the forward as the expected depreciation less the forward premium. At long horizons, where the expected real exchange rate is stationary so that expected depreciation converges to the expected short-rate differential (Chinn and Meredith, 2004), Equation (4) implies this premium converges to the pair residual $\pi^{j/i}_{t,N}$, which the containment result bounds within quoted transaction costs. Persistent long-horizon compensation for currency-attached risk therefore cannot reside in the forward wedge; it relocates, by the definition of the term premium, into the term premia of the two currencies, so that at the long end the currency premium and the cross-country term premium differential are not rival quantities but the same quantity under two names, and the quantity measured here is the long-horizon form of the currency risk premium itself.

Two implications for term-structure modelling follow. Single-country models are estimated as if this international margin were absent, yet the partner-side supply loadings, the common global component, and the clustering of model-based estimators around their shared expectations machinery all indicate that the margin is priced; accordingly, the measured differentials constitute over-identifying moment conditions that any country-level system should be required to match, and the natural design for the next generation of estimated models is joint multi-country estimation in which the cross-country restrictions hold by construction, a direction I leave to future work.

One alternative reading of the currency-level component survives every test: a persistent, priced compensation for holding dollar-denominated duration whose microfoundation one might prefer to call safety value, collateral value, or the cost of committing funds for the term rather than duration risk. It survives because it shares every measurable property of a term premium, being persistent, earned in realised returns, and priced by the supply of the asset that carries it, and no regression can separate quantities that agree on every observable.

The surviving dispute therefore concerns the granularity of a label, not the existence, size, sign history, pricing, or behaviour of the measured differential, and I find the inclusive reading the natural one: the term premium is compensation for bearing the risks of the term, which comprise both the duration risk of the long rate and the risk of locking funds to the term,

so persistent priced currency-attached compensation of either kind is term-premium content under the agnostic definition of Section 1. Under finer labels nothing is lost, since every result can be read as a statement about the priced currency-level component of long-tenor international fixed-income pricing, and every conclusion about the reversal and the supply channel stands unchanged. The attribution of Section 4.1 bounds any such reading further: a dollar-side relabelling claims at most the roughly one third of the differentials' variation that the dollar's own movement contributes, and the remaining partner-side variation is reproduced nearly one for one in quotes with no dollar leg.

FIGURE IA.3 ABOUT HERE

TABLE IA.13 ABOUT HERE

TABLE IA.30 ABOUT HERE

TABLE IA.31 ABOUT HERE

TABLE IA.32 ABOUT HERE

TABLE IA.18 ABOUT HERE

TABLE IA.29 ABOUT HERE

TABLE IA.27 ABOUT HERE

The independence of the euro-cross quotes also holds on their dynamics, tested in an error-correction design fixed before estimation in the replication materials. For each cross with a stored daily history, the quote and its dollar-implied counterpart are modelled as a bivariate error-correction system in daily changes, the cointegrating vector imposed at one-to-minus-one as the levels containment justifies, two lags of each series in each equation, the lagged level gap entering both equations, Newey-West standard errors at ten lags, and the zero-change fraction of each series beside every row, since a stale quote attenuates measured adjustment mechanically. Table IA.24 reports the ten-year battery.

The prediction that separates an independently discovered price from a repackaging of the dollar legs is the implied series' adjustment coefficient: a repackaged quote corrects toward the implied series while the implied series never moves toward the quote. The data show two-sided adjustment. The quote's adjustment coefficient is negative in all nine rows and the implied series' coefficient is positive in all nine, and the implied series' correction toward the quote, the coefficient a repackaged quote cannot generate, is individually significant in six: the legacy sterling cross (-0.157, $t = -5.8$, against $+0.119$, $t = +5.1$), the pre-transition New Zealand composite (-0.027, $t = -3.2$, against $+0.015$, $t = +2.3$), the pre-transition Norwegian cross ($+0.024$, $t = +2.3$), the Norwegian native successor ($+0.428$, $t = +4.8$), the post-transition yen cross ($+0.508$, $t = +2.2$), and the Swiss cross ($+0.187$, $t = +3.3$), with Gonzalo-Granger component shares of the quote between 0.37 and 0.76 on those rows, so

the quoted crosses carry permanent-component weight of their own rather than being derived from the dollar legs.

Three rows are reported as returned rather than claimed. The Swedish native's discovery share sits mostly on the implied side (0.15), its quote correcting strongly toward the implied series ($t = -3.0$) while the reverse coefficient is imprecise; the sterling successor's two adjustment coefficients are large, near-offsetting, and individually imprecise on its short history, so its independence rests on the documented quoting mechanism and the admission diagnostics rather than on this battery, its levels containment at 0.1 basis points being equally consistent with genuine quotes kept tight by arbitrage; and the legacy Norwegian quote's post-transition segment is uninformative, both series stale on more than a tenth of days, with the native successor row superseding it. The Norwegian battery is split at April 2022, where the legacy cross's reference convention changes; the split removes the convention splice that depressed the pooled levels fit previously reported, the pre-transition segment conforming at a slope of +1.00 and an R^2 of 0.96. The Swiss ten-year history enters from December 2009, and the Swiss five-year series before 2010 quotes on a different convention, with early prints reaching hundreds of basis points, and stays excluded. The levels fit in the table is daily, and therefore sits below the monthly correlations of Figure 4, whose 0.993 to 0.999 range is computed on month-end observations where daily staleness nets out.

TABLE IA.24 ABOUT HERE

FIGURE IA.2 ABOUT HERE

FIGURE IA.6 ABOUT HERE

FIGURE IA.7 ABOUT HERE

FIGURE IA.4 ABOUT HERE

FIGURE IA.5 ABOUT HERE

FIGURE IA.10 ABOUT HERE

FIGURE IA.11 ABOUT HERE

TABLE IA.33 ABOUT HERE

TABLE IA.38 ABOUT HERE

TABLE IA.21 ABOUT HERE

FIGURE IA.8 ABOUT HERE

TABLE IA.35 ABOUT HERE

TABLE IA.28 ABOUT HERE

TABLE IA.36 ABOUT HERE

TABLE IA.37 ABOUT HERE

TABLE IA.14 ABOUT HERE

TABLE IA.26 ABOUT HERE

TABLE IA.40 ABOUT HERE

FIGURE IA.9 ABOUT HERE

TABLE IA.17 ABOUT HERE

TABLE IA.25 ABOUT HERE

TABLE IA.20 ABOUT HERE

TABLE IA.12 ABOUT HERE

TABLE IA.15 ABOUT HERE

TABLE IA.16 ABOUT HERE

TABLE IA.34 ABOUT HERE

FIGURE IA.1 ABOUT HERE

TABLE IA.19 ABOUT HERE

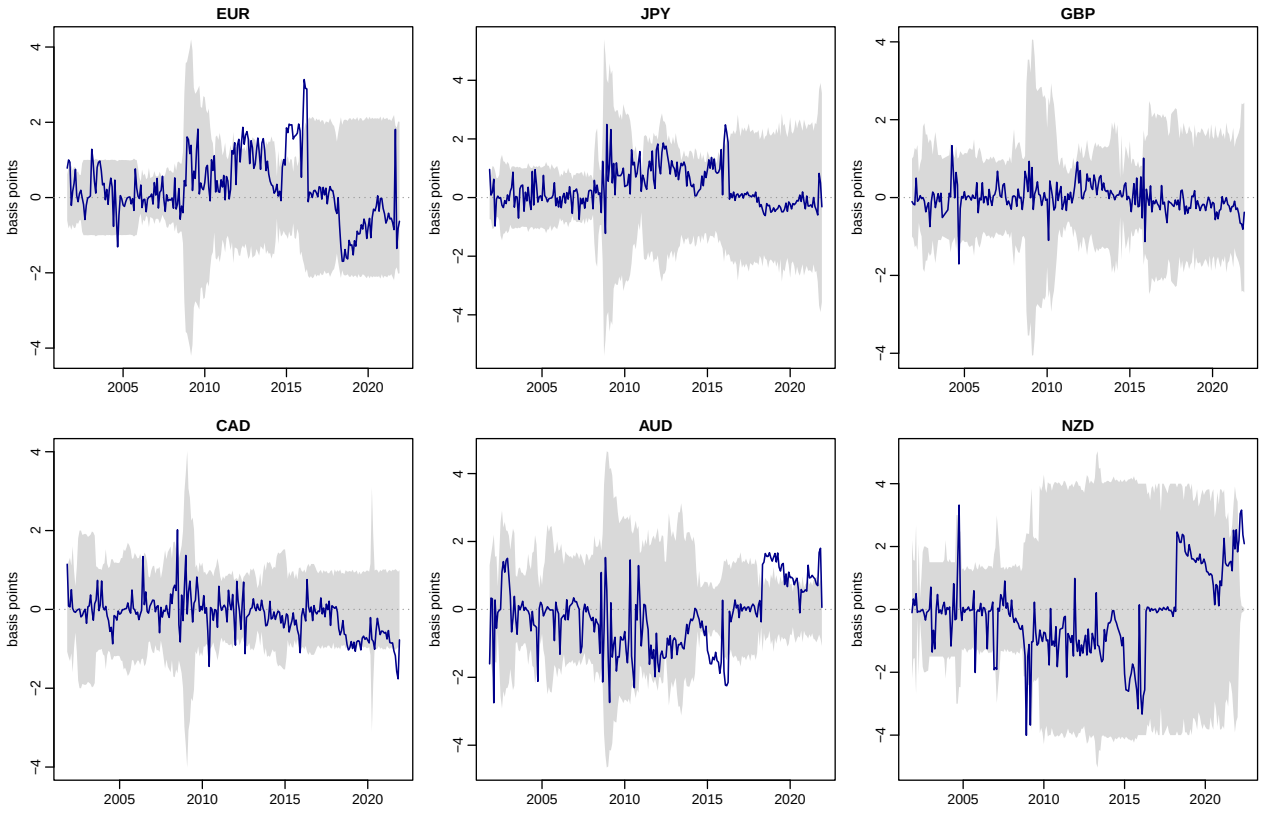


Figure IA.1: *Pair-specific residuals and the quoted bid-ask band. Each panel plots the signed ten-year pair-specific residual $\hat{\pi}_{j,t}$ in basis points (solid line) against the band formed by plus and minus the mean quoted half-spread on the same ten-year basis in each month (shaded), constructed from the dealer bid and ask fields of the underlying Bloomberg quotes; the sample is monthly.*

What hub exclusion cannot test is the level of the excluded hub's own node, which is exactly where the sensitivity-frontier bounds and the excess-sensitivity restriction of Section 4.2 operate. Two features of the residual series deserve explicit note. The sterling, Canadian, and Australian euro crosses terminate cleanly at the benchmark transition, with no stale quotes carried into any triangle, so the network thins to two surviving cycles after 2021; for this reason the selection-rule and regularisation exhibits, Tables IA.30 and IA.31, are estimated to December 2021 and characterise the recovery on the full-depth network, with the thinner network after the transition governed by the degeneracy discussion here. The surviving cycles are the yen cross's risk-free-rate successor and the New Zealand composite; with fewer cycles, and with the New Zealand composite mixing benchmark conventions before its declassification, the residual allocation concentrated visibly in the euro and New Zealand legs, producing occasional spikes in the euro residual of up to thirteen basis points. The successor quotations and the convention declassification resolve this, and the search for them doubles as direct institutional evidence on quote independence. Dealer documentation and the stored price

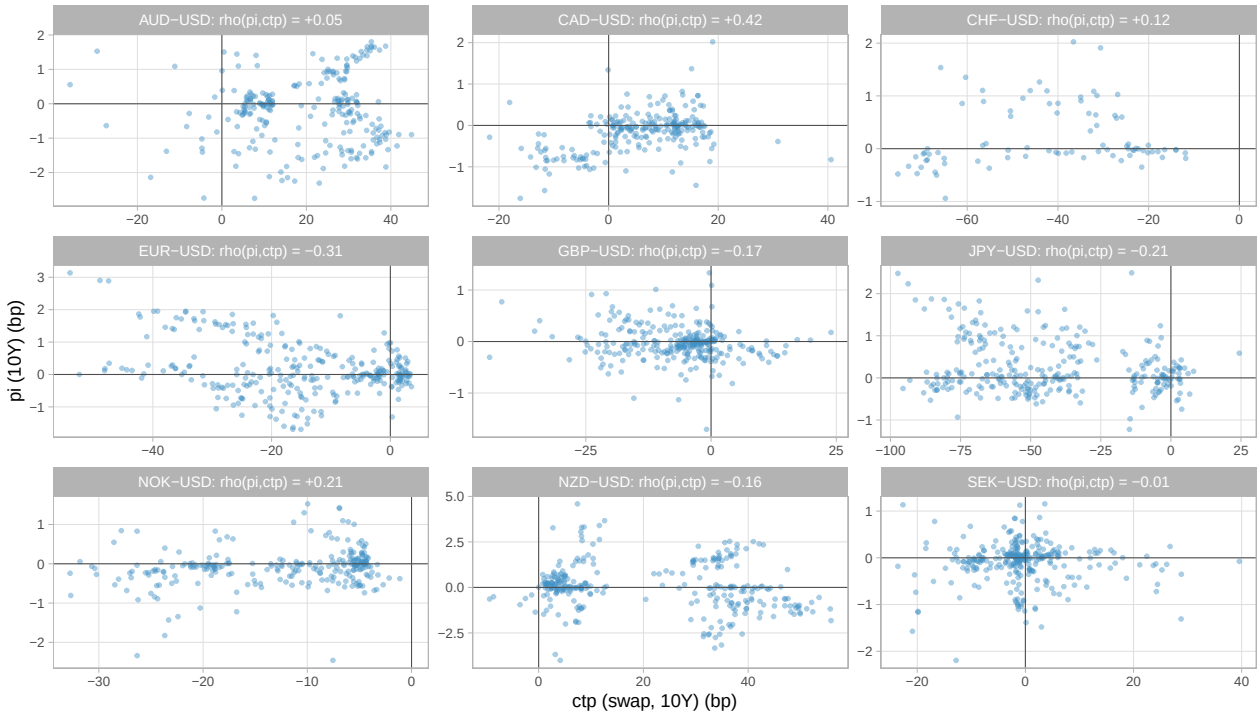


Figure IA.2: **Pair-Specific Component versus Cross-Country Term Premium at the Ten-Year Tenor.** Each panel plots the monthly time series of the ten-year pair-specific component $\pi_{t,10}^{j/USD}$ against the corresponding cross-country term premium $ctp_{t,10}^{j,USD}$ in basis points, for one of the nine USD bilateral pairs. The panel title reports the sample time-series correlation $\rho(\pi, ctp)$. Sample: January 2002 to June 2026. Internet Appendix Figures IA.10 and IA.11 report the analogous scatter at the five- and two-year tenors.

record together distinguish three classes of euro cross after the transition. The first is the sterling successor, a euro short-rate-versus-SONIA basis quoted continuously since February 2022: it carries a stored daily price history and fails every synthetic criterion in a formal test, regressing its quotes on the dollar-implied euro-sterling differential gives a slope of 0.86, an R^2 of 0.93, and one basis point of genuine residual variation, establishing it as independently discovered price information that restores a second cycle to the post-transition network. The same stored-history criterion admits the native risk-free-rate euro crosses for the Swedish and Norwegian currencies, quoted daily since April 2022 across all five tenors; regressing their quotes on the dollar-implied differentials gives slopes of 0.93 to 0.99 with residual standard deviations of 1.4 to 2.1 basis points, and they replace the convention-adjusted legacy quotes as the post-transition legs of the Swedish and Norwegian triangles. Transaction records from the swap data repository, collected and documented in Internet Appendix Section H.2, complete the picture: the ten-year tenor bracket records hundreds of new trades per month across the major dollar pairs, and every dollar leg of the network, including the thinnest, records

trades at that tenor every month, so the long-tenor quotes are prices in an actively traded market rather than indicative marks, and official weekly aggregates from the same repositories show that this activity intensifies rather than evaporates in stress, with ticket counts in the worst weeks of March 2020 running at more than twice their calm-period pace. The second is the family of successor tickers for the Canadian, Australian, and New Zealand euro crosses documented as constructed through dollar overnight-rate legs: these carry no stored price history at all, the vendor computing them on demand, so they are identifiable as synthetic both by documentation and by the absence of any retrievable series, and they are excluded. The third is the incumbent New Zealand euro cross itself, a multi-contributor composite that continues quoting daily through the transition with a full stored history. Its post-transition prints mix two benchmark conventions: contributors alternate between a Euribor-referenced and an ESTR-referenced euro leg, and the two modes sit exactly one observed Euribor-to-ESTR fixed-leg wedge apart, the composite's one-day level shift of -19 basis points in March 2022 matching the contemporaneous ten-year wedge of 20 to 22 basis points and the roughly twelve-basis-point alternation of its five-year prints in early 2025 matching the thirteen-basis-point five-year wedge. The mixture is itself a fingerprint of genuine multi-contributor quoting, since a series synthesised from the dollar legs would be internally convention-consistent, and I retain the cross as an independent equation after assigning each print to the ESTR convention by proximity to the running level of the declassified series, before the thirty-one-day quote carry, so that no mixed-convention print enters any triangle. On the declassified quotes the independence test retains clear residual variation at five years, with a residual standard deviation of 2.4 basis points, while at ten years the corrected quote tracks the dollar-implied differential within 1.6 basis points, a borderline verdict I disclose; the declassification, not the exclusion, is what the evidence supports, and excluding the cross instead changes no result in the paper because the pair-level series then coincides with the parity-free benchmark. With the restored sterling cycle and the declassified New Zealand cross the post-transition network is re-solved throughout: the euro residual's maximum falls from thirteen to under three basis points and the New Zealand residual's from nine to under five, every quoted basis, recovered differential, and currency-node series extends through June 2026, and the pre-transition solution is unchanged at every date. After the transition the Canadian and Australian dollar nodes are connected to the network by their dollar legs alone; on such a spoke the triangular allocation is degenerate, the residual is zero under the minimum-norm convention, and the recovered pair differential coincides with the parity-free benchmark that the paper carries as its transparent working construction, so the series continue with the identification regime flagged in the data and stated in the exhibit captions, while the residual itself is reported only where the pair sits in a closed triangle.

Table IA.12: **Decomposition of the Cross-Currency Basis: Cross-Country Term Premium and Pair-Specific Friction.** For each USD bilateral pair and tenor, the table reports the sample (signed) mean of the bilateral cross-currency basis $x_N^{j/USD}$, the cross-country swap term premium $ctp_N^{j,USD}$, and the pair-specific friction $\pi_N^{j/USD}$, all in basis points, with $x = ctp + \pi$ at every date by construction. The column $|\pi|/|x|$ reports the absolute mean of the pair-specific friction as a share of the absolute mean of the bilateral basis, in percent. The complementary share $|ctp|/|x|$ is not reported in the table and need not equal $1 - |\pi|/|x|$, since the two components ctp and π can carry opposite signs at the same date and therefore $|x| = |ctp+\pi|$ is in general distinct from $|ctp|+|\pi|$. The column $\rho(\Delta SS, ctp^{swap})$ reports the time-series correlation between the cross-country swap spread differential and the cross-country swap term premium, the two additive components of the bridge identity from the swap to the cross-country government bond term premium. Sample: January 2002 to June 2026, monthly frequency; the friction statistics and the reported N use triangle-identified months only, since on a post-Libor spoke the residual is zero by the minimum-norm convention rather than an estimate, while the basis and term premium means use all quoted months.

Tenor	Pair	Mean x (bp)	Mean ctp (bp)	Mean π (bp)	$\lvert\pi\rvert/\lvert x\rvert$ (%)	$\rho(\Delta SS, ctp^{swap})$	N
Panel A: 2-year tenor							
2Y	USD-EUR	−17.64	−17.79	+0.15	0.84	+0.511	294
2Y	USD-JPY	−31.34	−31.44	+0.10	0.33	+0.590	290
2Y	USD-GBP	−6.41	−6.43	+0.02	0.28	+0.028	293
2Y	USD-CAD	−8.42	−8.38	−0.05	0.56	+0.841	240
2Y	USD-AUD	+10.47	+10.62	−0.19	1.80	+0.379	240
2Y	USD-NZD	+7.92	+8.05	−0.15	1.91	+0.421	244
2Y	USD-SEK	−14.17	−14.21	+0.04	0.26	+0.321	290
2Y	USD-NOK	−17.14	−17.11	−0.03	0.18	+0.495	289
Panel B: 5-year tenor							
5Y	USD-EUR	−16.95	−17.15	+0.20	1.16	+0.539	294
5Y	USD-JPY	−40.74	−40.92	+0.19	0.46	+0.754	290
5Y	USD-GBP	−6.04	−6.00	−0.04	0.62	+0.153	293
5Y	USD-CAD	−2.16	−2.14	−0.02	0.97	+0.831	240
5Y	USD-AUD	+15.56	+15.71	−0.18	1.17	+0.020	240
5Y	USD-NZD	+15.10	+15.13	−0.03	0.21	+0.321	294
5Y	USD-SEK	−7.07	−6.97	−0.10	1.35	+0.482	293
5Y	USD-NOK	−12.81	−12.71	−0.10	0.79	+0.523	289
Panel C: 10-year tenor							
10Y	USD-EUR	−15.05	−15.21	+0.16	1.07	+0.663	294
10Y	USD-JPY	−42.43	−42.68	+0.25	0.58	+0.845	291
10Y	USD-GBP	−6.53	−6.49	−0.03	0.52	−0.022	289
10Y	USD-CAD	+0.83	+0.96	−0.15	18.30	+0.695	240
10Y	USD-AUD	+17.46	+17.63	−0.20	1.15	−0.094	240
10Y	USD-NZD	+21.33	+21.29	+0.03	0.16	+0.216	294
10Y	USD-SEK	−0.51	−0.41	−0.10	19.67	+0.148	293
10Y	USD-NOK	−11.80	−11.72	−0.09	0.80	+0.487	255

Table IA.13: Sample construction for the headline exhibits. Frequencies are monthly unless stated; the recovered cross-country term-premium differentials are benchmark-consistent throughout.

Exhibit	Sample	Principal inputs
Recovered term premia (levels, composition)	Jan. 2002 – Jun. 2026	Bloomberg bases, swap spreads
Extended benchmark-consistent panel	Aug. 2001 – Jun. 2026	Rates database (chain-reconciled)
Supply and reserve-demand panel	Jan. 2002 – Jun. 2026 (294 months; inputs from Aug. 2001)	Free float, Z.1 official holdings
Official-demand flow margin (IA)	Aug. 2001 – Dec. 2020	Bertaut-Judson valuation-adjusted flows
Intermediary-augmented supply columns (3)–(5)	to May 2025	Published He et al. (2017) vintage; WAM to Jun. 2026
Dollar reconciliation (Jiang et al.)	Aug. 2001 – Jun. 2026 less horizon	Equally weighted spot basket
Refunding-surprise events	2000 – 2026 (99 events)	Treasury releases, daily series
Supply, current-inflation control (col. 4)	Dec. 2001 – Jun. 2026	CPI coverage; JP to Jun. 2021; AU/NZ quarterly

Table IA.14: **Reported Cross-Currency Basis Swap Transactions by Pair and Tenor.** Average monthly counts of cross-currency basis swap transactions reported to the DTCC swap data repository under CFTC real-time public dissemination, August 2025 to June 2026, by currency pair and tenor bucket. Transactions are interest-rate swaps with different G10 notional currencies on the two legs; new trades only, action type NEWT deduplicated by dissemination identifier, so amendments, corrections, and terminations are excluded exactly. The first four columns cover the float-for-float form, the quoted basis instrument; the final column reports all contract forms at the ten-year bracket. Tenor is the difference between expiration and effective dates. Euro crosses without a U.S. reporting party fall under the European reporting regime and are not captured here. Source: DTCC public price dissemination, collected as described in Internet Appendix Section H.2.

	Float-for-float (quoted instrument)				All forms
Pair	0–3Y	3–8Y	8–12Y	12Y+	8–12Y
EUR/USD	636	551	250	139	297
JPY/USD	969	642	204	64	219
GBP/USD	303	330	164	181	176
AUD/USD	231	142	75	34	80
CHF/USD	189	122	63	31	72
CAD/USD	425	164	51	21	56
NOK/USD	91	33	6	1	6
NZD/USD	52	36	6	0	6
SEK/USD	114	43	5	0	6

Table IA.15: Margins of the ten-year differential by pair and regime. Entries are means in basis points over the pre-GFC (January 2002 to June 2007), GFC (July 2007 to June 2009), and post-GFC (July 2009 to June 2026) windows, which are the crisis-dated windows of this table alone and are deliberately narrower than the calendar windows used in the main text, so the sub-period means here differ from those reported there for the same series: the raw ten-year government yield differential $y_j - y_{US}$, the cross-country swap spread differential $SS_{US} - SS_j$, and the recovered swap and government bond cross-country term premia.

Pair	Regime	Yield diff.	ΔSS	Swap premium	Gov. premium
EUR	pre	−41	28	1	28
	GFC	12	3	−11	−8
	post	−139	−32	−21	−52
JPY	pre	−301	41	−6	36
	GFC	−223	36	−8	28
	post	−209	−18	−59	−76
GBP	pre	17	14	−2	12
	GFC	63	12	−15	−3
	post	−26	4	−7	−3
CAD	pre	18	22	11	33
	GFC	−2	22	10	32
	post	−38	−37	−3	−40
AUD	pre	106	−2	8	7
	GFC	178	−34	−5	−40
	post	60	−40	23	−16
NZD	pre	151	−32	6	−26
	GFC	211	−41	2	−39
	post	94	−18	29	11

Table IA.16: Free-float supply construction by country. The free float is gross marketable central government debt net of the domestic central bank's holdings of own-government securities; all series are monthly, interpolated from the source frequency where necessary.

Country	Gross marketable debt	Central bank holdings
United States	Treasury MSPD	Federal Reserve SOMA
Germany	Finanzagentur / Bundesbank	ECB PSPP and PEPP, German allocation
Japan	Ministry of Finance JGB statistics	Bank of Japan accounts
United Kingdom	Debt Management Office	Bank of England Asset Purchase Facility
Canada	Bank of Canada / Statistics Canada	Bank of Canada purchase programmes
Australia	Australian Office of Financial Management	Reserve Bank of Australia
New Zealand	New Zealand Debt Management	Reserve Bank of New Zealand LSAP
Sweden	Gross general government debt series	Riksbank holdings (Statistics Sweden)
Norway	Norges Bank debt key figures	None; state own holdings netted

Table IA.17: **Channel Allocation Summary: How Each Driver Splits Across the Two Bridge-Identity Components.** For every regressor in column (5) of the free-float specification, this table reports the column-(5) coefficient on the cross-country government bond term premium ($\widehat{\beta}^{\text{gov}}$), on the quoted bilateral cross-currency basis, representing the cross-country swap term premium ($\widehat{\beta}^{x}$), and on the cross-country swap spread differential ($\widehat{\beta}^{\Delta SS}$), drawn from Tables 1 and IA.8. The coefficients satisfy $\widehat{\beta}^{\text{gov}} = \widehat{\beta}^{x} + \widehat{\beta}^{\Delta SS}$ exactly. Coefficients are reported with their statistical significance asterisks, omitting t-statistics for compactness. The final column summarises which bridge component carries the variable's significant loading: "$\widehat{\beta}^{\Delta SS}$" or "$\widehat{\beta}^{x}$" indicates that the listed component carries the larger absolute coefficient and is statistically significant; "$\widehat{\beta}^{x}$ only" or "$\widehat{\beta}^{\Delta SS}$ only" flags variables whose $\widehat{\beta}^{\text{gov}}$ coefficient is statistically insignificant but whose $\widehat{\beta}^{x}$ or $\widehat{\beta}^{\Delta SS}$ component is individually significant, identifying drivers that the bridge decomposition reveals but that the regression on the cross-country government bond term premium alone would obscure; "balanced" flags variables whose two components are both individually significant with absolute magnitudes within a factor of 1.25 of each other; and "—" marks variables with no individually significant coefficient, for which no carrier is assigned. The identity holds on unrounded coefficients, so a printed row can differ from printed-cell arithmetic in the terminal digit. Sample: monthly, January 2002 to June 2026, as in Table 1. Notes: $^{*}\,p \le .10$; $^{**}\,p \le .05$; $^{***}\,p \le .01$.

Driver	$\widehat{\beta}^{\text{gov}}$	$\widehat{\beta}^{x}$	$\widehat{\beta}^{\Delta SS}$	Carrier
$\log(\widetilde{d})_{\text{US}}$	-1.134^{***}	-0.287^{***}	-0.847^{***}	$\widehat{\beta}^{\Delta SS}$
$\log(\widetilde{d})_{j}$	$+0.162^{**}$	$+0.092^{*}$	$+0.070$	$\widehat{\beta}^{x}$
$\widetilde{\text{WAM}}_{\text{US}}$	-0.062	$+0.021$	-0.083^{**}	$\widehat{\beta}^{\Delta SS}$ only
$\widetilde{\text{WAM}}_{j}$	$+0.030$	-0.012	$+0.042^{***}$	$\widehat{\beta}^{\Delta SS}$ only
$\text{QE}_{t}^{\text{US}}$	$+0.255^{***}$	$+0.096^{***}$	$+0.159^{***}$	$\widehat{\beta}^{\Delta SS}$
QE_{t}^{j}	-0.279^{***}	-0.194^{***}	-0.085^{***}	$\widehat{\beta}^{x}$
HKM_{t}^{z}	$+0.050$	$+0.021$	$+0.029$	—
$\text{HKM}_{t}^{z} \times \text{QE}_{t}^{\text{US}}$	$+0.187^{**}$	$+0.103^{***}$	$+0.085$	$\widehat{\beta}^{x}$
$\text{rGDP}_{t}^{\text{US},z}$	$+0.030$	$+0.026^{*}$	$+0.004$	$\widehat{\beta}^{x}$ only
$\text{rGDP}_{t}^{j,z}$	-0.016	-0.015	-0.001	—
$\text{ir3m}_{t}^{\text{US},z}$	$+0.182^{***}$	$+0.108^{***}$	$+0.075^{**}$	$\widehat{\beta}^{x}$
$\text{ir3m}_{t}^{j,z}$	-0.255^{***}	-0.200^{***}	-0.055^{*}	$\widehat{\beta}^{x}$
$\text{TS}_{t}^{\text{US},z}$	-0.023	$+0.025$	-0.048^{**}	$\widehat{\beta}^{\Delta SS}$ only
$\text{TS}_{t}^{j,z}$	$+0.047$	-0.047^{**}	$+0.094^{***}$	$\widehat{\beta}^{\Delta SS}$
$\text{Unemp}_{t}^{\text{US},z}$	-0.037	$+0.013$	-0.050^{*}	$\widehat{\beta}^{\Delta SS}$ only
$\text{Unemp}_{t}^{j,z}$	$+0.065$	-0.043^{***}	$+0.108^{***}$	$\widehat{\beta}^{\Delta SS}$
$\text{Fiscal bal}_{t}^{\text{US}}$	-0.005	-0.006	$+0.000$	—
$\text{Fiscal bal}_{t}^{j}$	-0.014	-0.001	-0.013^{*}	$\widehat{\beta}^{\Delta SS}$ only
$\text{CPI}_{t}^{\text{US},z}$	-0.026	$+0.001$	-0.027	—
$\text{CPI}_{t}^{j,z}$	-0.033	-0.022	-0.011	—
Dollar_{t}^{z}	-0.045	-0.078^{***}	$+0.033$	$\widehat{\beta}^{x}$ only
BBB_{t}^{z}	-0.012	$+0.007$	-0.019	—
MOVE_{t}^{z}	-0.022	-0.013	-0.009	—
VIX_{t}^{z}	$+0.052^{***}$	$+0.021^{**}$	$+0.031^{**}$	$\widehat{\beta}^{\Delta SS}$

Panel C: basis-side share of the composite's monthly variance, by sub-period (percent)

Pair	2002–2007	2008–2012	2013–2026
USD-EUR	12	34	33
USD-JPY	20	60	52
USD-GBP	−2	36	28
USD-CAD	24	0	49

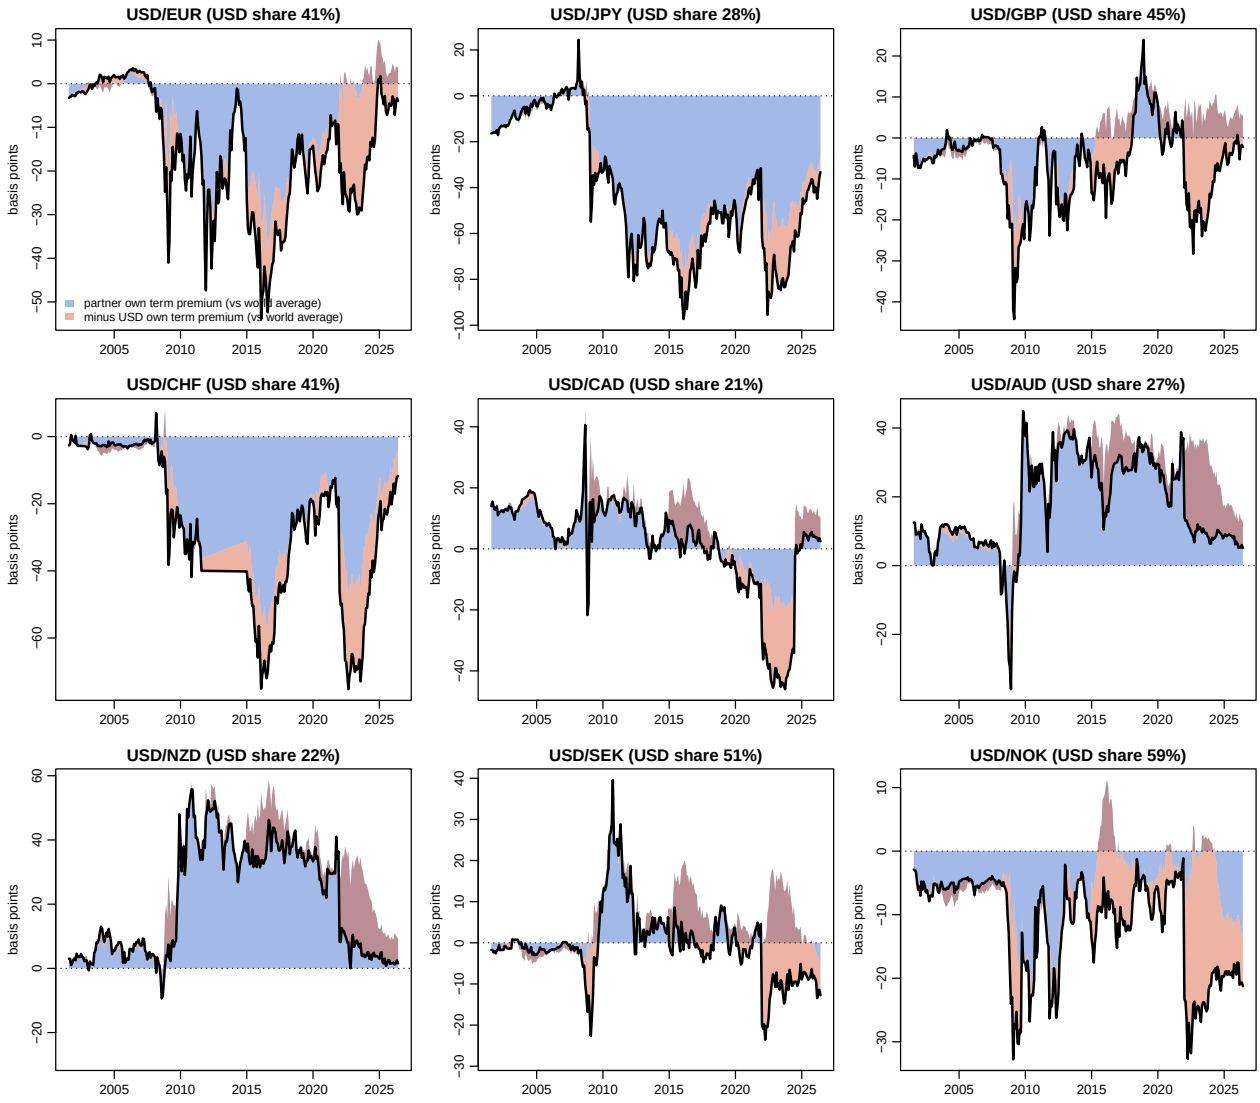


Figure IA.3: **Currency-Level Decomposition of the Cross-Country Term Premium Differentials.** Each panel decomposes the recovered ten-year cross-country swap term premium differential (solid line) into the partner country's own term premium and the negative of the U.S. own term premium (shaded areas summing to the line), where each currency's own term premium is measured relative to the cross-currency average at each date, the component invariant to the global term-premium factor that the differentials strip by construction. The currency-level series are recovered from the full quoted network by date-by-date least squares under the zero-sum normalisation. The nine panels cover every dollar pair, and each title reports the share of the monthly variance of the differential attributable to the dollar side, highest for the Scandinavian pairs at 51 and 59 percent and lower for the reserve majors and commodity currencies at 21 to 45 percent. Monthly, January 2002 to June 2026.

Table IA.18: **Robustness of the Supply Relationship and the Funding-versus-Term-Premium Discrimination.** Panel A reports robustness of the U.S. free-float supply coefficient. Every Panel A row runs the two-variable specification, U.S. and comparison-country supply with currency-pair fixed effects, on the eight-pair benchmark panel, the six incumbent majors plus the Swedish and Norwegian pairs, over 295 months, which is why the base coefficient of -1.052 on $N = 2{,}352$ differs from the seven-pair free-float column of Table 1 (-1.051, $N = 2{,}043$) and from the identified-sample column of Internet Appendix Table IA.1 (-1.066, $N = 1{,}888$); the leave-one-out rows drop one pair from the eight-pair frame, and the pre-2020 row truncates it. "Orthogonalised vs global supply" replaces $\log D^{ff}_{US}$ with the residual from projecting it on the cross-sectional mean comparison-country supply, the direct test of whether a common global supply factor drives the coefficient; the relation is essentially unchanged. The leave-one-out rows drop each currency pair in turn. Panel B addresses whether the recovered cross-country term-premium differential is a repackaging of short-end dollar funding: it regresses the cross-sectional average recovered cross-country swap term premium differential at each tenor on the short-end dollar funding factor, defined as the cross-sectional mean three-month cross-currency basis. A loading that declines toward zero as the tenor lengthens is the signature of a term premium rather than a funding deviation. Standard errors are Driscoll-Kraay with a Newey-West kernel (Panel A) and Newey-West (Panel B), each with a Schwert-rule lag; t-statistics are in parentheses. $^{*}\,p \le .10$; $^{**}\,p \le .05$; $^{***}\,p \le .01$. Panel A sample: monthly, January 2002 to June 2026. Panel B sample: monthly, 2008 to 2022, over which the three-month basis is quoted on a consistent convention. Throughout, the quoted bilateral cross-currency basis x is used to represent the cross-country swap term premium, which it equals up to the pair-specific residual bounded within quoted half-spreads.

Panel A: U.S. free-float supply coefficient			
	$\log D^{ff}_{US}$	t	N
Base (pair FE)	-1.052^{***}	(-9.08)	2,352
Orthogonalised vs global supply	-1.051^{***}	(-4.65)	2,352
Excluding EUR	-1.057^{***}	(-8.90)	2,057
Excluding JPY	-0.899^{***}	(-8.71)	2,057
Excluding GBP	-1.108^{***}	(-9.10)	2,057
Excluding CAD	-0.972^{***}	(-8.48)	2,057
Excluding AUD	-1.036^{***}	(-9.08)	2,065
Excluding NZD	-1.210^{***}	(-9.94)	2,057
Excluding SEK	-1.090^{***}	(-8.85)	2,057
Excluding NOK	-1.050^{***}	(-8.53)	2,057
Pre-2020 subsample	-1.091^{***}	(-9.61)	1,728
Panel B: loading on the short-end dollar funding factor			
	Loading	t	R^2
Two-year tenor	0.122^{***}	(3.34)	0.140
Five-year tenor	0.064^{*}	(1.88)	0.049
Ten-year tenor	0.063^{**}	(2.15)	0.053

Table IA.19: **State Dependence of the Supply Relationship: The Fiscal-Surge Channel.** This table reports panel regressions of the ten-year cross-country government bond term premium differential $\widehat{ctp}_{10}^{gov,j,USD}$ on U.S. log free-float supply $\log D_{US}^{ff}$ interacted with three candidate state variables. Deficit_z is the standardised trailing twelve-month U.S. federal deficit; COVID is an indicator for March 2020 to December 2021; QE is an indicator for active U.S. large-scale asset purchases. Each interacted state variable is also included in levels (coefficients suppressed for brevity), and all specifications include currency-pair fixed effects and the comparison-country supply term $\log D_j^{ff}$. Column (2) shows that the relationship is suspended during the COVID window and recovers outside it. Column (1) enters the deficit-pace interaction, which is positive but not significant on the full benchmark-consistent sample; column (3) races the two, and the COVID window retains significance while the deficit measure does not, indicating that on the full sample the state dependence is dated by the pandemic window. Column (4) shows that quantitative easing does not account for the shift once the COVID window is controlled for. Standard errors are Driscoll-Kraay with a Newey-West kernel and a Schwert-rule lag; t-statistics are in parentheses. $^{*}p \leq .10$; $^{**}p \leq .05$; $^{***}p \leq .01$. Sample: monthly, January 2002 to June 2026.

	(1) Deficit	(2) COVID	(3) Race	(4) QE
$\log D_{US}^{ff}$	-1.126^{***}	-1.134^{***}	-1.084^{***}	-1.155^{***}
	(-8.92)	(-9.80)	(-8.46)	(-9.63)
$\log D_{US}^{ff} \times$ Deficit_z	0.107		-0.095	
	(1.18)		(-0.67)	
$\log D_{US}^{ff} \times$ COVID		0.921^{***}	1.315^{**}	0.771^{***}
		(7.98)	(2.20)	(3.34)
$\log D_{US}^{ff} \times$ QE				0.172
				(0.79)
$\log D_j^{ff}$	0.519^{***}	0.522^{***}	0.514^{***}	0.517^{***}
	(8.51)	(8.53)	(8.75)	(8.65)
State variable in levels	Yes	Yes	Yes	Yes
Currency-pair FE	Yes	Yes	Yes	Yes
N	2,352	2,352	2,352	2,352
Within R^2	0.427	0.449	0.453	0.455

Table IA.20: **Channel Allocation Summary on the Extracted Premia.** For every regressor in column (5) of the extracted-premium specification, this table reports the column-(5) coefficient on the extracted cross-country government bond term premium ($\widehat{\beta}^{\text{gov}}$), on the triangulation-extracted cross-country swap term premium ($\widehat{\beta}^{\text{swap}}$), and on the cross-country swap spread differential ($\widehat{\beta}^{\Delta SS}$), drawn from Tables IA.1 and IA.9. The coefficients satisfy $\widehat{\beta}^{\text{gov}} = \widehat{\beta}^{\text{swap}} + \widehat{\beta}^{\Delta SS}$ exactly on the identified sample. The final column applies the classification rule of Table IA.17, with the extracted swap term premium as the first component. Notes: $^{*}\, p \le .10$; $^{**}\, p \le .05$; $^{***}\, p \le .01$.

Driver	$\widehat{\beta}^{\text{gov}}$	$\widehat{\beta}^{\text{swap}}$	$\widehat{\beta}^{\Delta SS}$	Carrier
$\log(\widetilde{d})_{\text{US}}$	-1.216^{***}	-0.310^{***}	-0.905^{***}	$\widehat{\beta}^{\Delta SS}$
$\log(\widetilde{d})_{j}$	$+0.267^{***}$	$+0.118^{**}$	$+0.149^{**}$	$\widehat{\beta}^{\Delta SS}$
$\widetilde{\text{WAM}}_{\text{US}}$	-0.046	$+0.032$	-0.078^{**}	$\widehat{\beta}^{\Delta SS}$ only
$\widetilde{\text{WAM}}_{j}$	$+0.017$	-0.020	$+0.036^{**}$	$\widehat{\beta}^{\Delta SS}$ only
QE_t^{US}	$+0.266^{***}$	$+0.099^{***}$	$+0.167^{***}$	$\widehat{\beta}^{\Delta SS}$
QE_t^{j}	-0.262^{***}	-0.188^{***}	-0.073^{***}	$\widehat{\beta}^{\text{swap}}$
HKM_t^{z}	$+0.007$	$+0.007$	-0.000	—
$\text{HKM}_t^{z} \times \text{QE}_t^{\text{US}}$	$+0.188^{**}$	$+0.101^{***}$	$+0.087$	balanced
$\text{rGDP}_t^{\text{US},z}$	$+0.016$	$+0.018$	-0.002	—
$\text{rGDP}_t^{j,z}$	-0.012	-0.013	$+0.001$	—
$\text{ir3m}_t^{\text{US},z}$	$+0.196^{***}$	$+0.107^{***}$	$+0.089^{**}$	balanced
$\text{ir3m}_t^{j,z}$	-0.241^{***}	-0.190^{***}	-0.051	$\widehat{\beta}^{\text{swap}}$
$\text{TS}_t^{\text{US},z}$	-0.014	$+0.027$	-0.041	—
$\text{TS}_t^{j,z}$	$+0.051$	-0.047^{**}	$+0.098^{***}$	$\widehat{\beta}^{\Delta SS}$
$\text{Unemp}_t^{\text{US},z}$	-0.040	$+0.011$	-0.052^{*}	$\widehat{\beta}^{\Delta SS}$ only
$\text{Unemp}_t^{j,z}$	$+0.032$	-0.054^{***}	$+0.086^{***}$	$\widehat{\beta}^{\Delta SS}$
$\text{Fiscal bal}_t^{\text{US}}$	-0.005	-0.006^{*}	$+0.000$	$\widehat{\beta}^{\text{swap}}$ only
Fiscal bal_t^{j}	-0.022^{**}	-0.005	-0.017^{**}	$\widehat{\beta}^{\Delta SS}$
$\text{CPI}_t^{\text{US},z}$	-0.002	$+0.011$	-0.012	—
$\text{CPI}_t^{j,z}$	-0.035	-0.025^{*}	-0.010	$\widehat{\beta}^{\text{swap}}$ only
Dollar_t^{z}	-0.023	-0.071^{***}	$+0.048$	$\widehat{\beta}^{\text{swap}}$ only
BBB_t^{z}	-0.003	$+0.013$	-0.016	—
MOVE_t^{z}	-0.012	-0.011	-0.001	—
VIX_t^{z}	$+0.034^{**}$	$+0.014^{**}$	$+0.020$	—

Table IA.21: **Inference on the Common U.S. Supply Regressor.** The U.S. log free-float coefficient from the column (1) specification of Table 1 (dependent variable the benchmark ten-year cross-country government bond term premium, pair fixed effects, $N = 2{,}043$ pair-months, seven pairs, 294 months) under alternative covariance treatments. Sample: monthly, January 2002 to June 2026.

	Coefficient	t-statistic
Baseline (Driscoll-Kraay, Schwert lag 15)	-1.051	(-10.72)
Two-way clustered (month and pair)	-1.051	(-4.39)

Table IA.22: **The Funding-Decay Profile: Pass-Through of Three-Month Funding-Point Changes Along the Curve.** Per pair, monthly changes of the month-end six-month, one-year, five-year, and ten-year dollar-pair bases regressed on monthly changes of the month-end three-month basis, one regression per cell, Newey-West standard errors at six lags, t-statistics in parentheses. The three- and six-month quotes are from the August 2026 extension pull, with data through June 2026; each pair's sample starts where its three- and six-month quotes begin, later than the 2008 start of the cross-sectional mean three-month funding factor of Table IA.18 Panel B, which is built on a consistent convention over 2008 to 2022 and measures a levels loading rather than this change pass-through. Stale is the fraction of zero-change months of the three-month point; the one-month tenor is unquoted, and the Swedish and Norwegian three-month legs are not in the extension pull.

Pair	N	Sample	6M	1Y	5Y	10Y	Stale
EUR	153	2013-10 to 2026-06	+0.36	+0.22	+0.06	+0.03	36%
			(+4.6)	(+5.7)	(+1.7)	(+1.2)	
JPY	153	2013-10 to 2026-06	+0.40	+0.21	+0.10	+0.07	36%
			(+4.6)	(+2.7)	(+1.9)	(+1.5)	
GBP	149	2014-02 to 2026-06	+0.43	+0.16	+0.06	−0.01	66%
			(+9.3)	(+6.2)	(+3.0)	(−0.7)	
CAD	137	2015-02 to 2026-06	+0.52	+0.30	+0.16	+0.16	38%
			(+7.1)	(+7.1)	(+3.8)	(+3.8)	
AUD	153	2013-10 to 2026-06	+0.61	+0.32	+0.07	+0.01	36%
			(+14.0)	(+6.0)	(+1.3)	(+0.2)	
NZD	121	2016-06 to 2026-06	+0.71	+0.47	+0.16	+0.10	43%
			(+13.5)	(+10.3)	(+4.8)	(+3.0)	
Median			+0.48	+0.26	+0.08	+0.05	

Table IA.23: **The December 2022 Event: Every Reported Quantity and Its Construction.** All responses are over the identical $[-1, +2]$ trading-day window around the 20 December 2022 Bank of Japan announcement, in basis points. Two distinct quantities equal +8.5: the directly quoted euro-yen basis, a cross quote with no dollar leg, and the yen five-year-five-year forward component of the dollar-pair basis; every occurrence of 8.5 in the text refers to one of these two, named at each use.

Quantity	Construction	Response
Yen benchmark differential	Daily quoted USD-JPY basis plus the swap-spread differential	+12.5
of which the swap side	Quoted USD-JPY basis alone	+8.2
Network-implied yen-euro benchmark differential	USD-JPY less USD-EUR benchmark series; the dollar node cancels	+15.0
Directly quoted euro-yen basis	Euro-cross quote; swap side only, no dollar leg, no swap-spread leg	+8.5
Yen five-year-five-year forward component	$2\,x_{10Y} - x_{5Y}$ on the quoted USD-JPY basis	+8.5
Own-pair funding placebo	Quoted one-year USD-JPY basis	+4.9
Dollar-wide funding benchmark	Cross-pair average one-year basis over the eight dollar pairs	+1.2

Table IA.24: **Error-Correction Independence of the Euro-Cross Quotes, Ten-Year Tenor.** Bivariate error-correction system in daily changes for each quoted euro cross q and its dollar-implied counterpart m, the cointegrating vector imposed at one-to-minus-one: each equation carries the lagged level gap $q_{t-1} - m_{t-1}$, two own lags, and two cross lags, with Newey-West standard errors at ten lags. α_q and α_m are the gap coefficients in the quote and implied equations; a quote repackaged from the dollar legs implies $\alpha_m = 0$, so the t-statistic on α_m is the asymmetry test. GG is the Gonzalo-Granger component share of the quote, a permanent-component weight rather than an information share, $\alpha_m/(\alpha_m - \alpha_q)$. Own lags reports the sum of the two own-lag coefficients in each equation, quote and implied. Levels columns report the regression of the quote on the implied counterpart over the row's sample; the levels fit is daily and sits below the monthly correlations of Figure 4 because daily staleness nets out at month-end. Stale is the fraction of zero-change days, quote and implied. N is the levels count; the error-correction equations lose three initial observations to lags. The Norwegian battery is split at April 2022, where the cross's reference convention changes. The successor rows, the sterling successor, the Swedish and Norwegian natives (all quoted from early 2022), and the New Zealand composite, are from the August 2026 extension pull; the New Zealand row runs on the composite's pre-transition segment, where the quote sits on a single reference convention, since its post-transition prints alternate conventions and enter the network only after the declassification assignment described in the text, which is a network step rather than a raw quote. t-statistics in parentheses.

			Levels			Stale	Adjustment			Own lags
Cross	N	Sample	Slope	R^2	sd	q/m	α_q	α_m	GG	q/m
Sterling (legacy)	4,708	2001-11 to 2021-12	+1.00	1.00	0.9	10/7	−0.157	+0.119	0.43	−0.35/ − 0.28
							(−5.8)	(+5.1)		
Sterling (successor)	1,140	2022-02 to 2026-06	+1.00	1.00	0.1	15/15	−0.468	+0.429	0.48	+0.98/ − 0.73
							(−1.4)	(+1.4)		
New Zealand, pre-transition	5,257	2001-11 to 2021-12	+1.02	0.99	3.1	26/14	−0.027	+0.015	0.37	−0.30/ − 0.33
							(−3.2)	(+2.3)		
Norway, pre-transition	3,725	2001-08 to 2022-03	+1.00	0.96	2.8	20/9	−0.018	+0.024	0.57	−0.40/ − 0.49
							(−1.4)	(+2.3)		
Norway, native successor	1,093	2022-04 to 2026-06	+1.00	1.00	0.3	16/14	−0.139	+0.428	0.76	+0.06/ + 0.07
							(−1.5)	(+4.8)		
Norway, legacy post-transition	1,100	2022-04 to 2026-06	+1.31	0.78	5.6	12/14	−0.044	+0.005	0.09	−0.44/ + 0.12
							(−1.1)	(+1.5)		
Sweden, native successor	1,093	2022-04 to 2026-06	+1.00	1.00	0.6	15/15	−0.689	+0.121	0.15	−0.18/ + 0.21
							(−3.0)	(+1.1)		
Yen, post-transition	1,081	2022-04 to 2026-06	+1.00	1.00	0.3	5/6	−0.291	+0.508	0.64	+0.05/ − 0.07
							(−1.2)	(+2.2)		
Switzerland	1,717	2009-12 to 2026-06	+0.98	1.00	0.8	8/7	−0.148	+0.187	0.56	−0.25/ − 0.18
							(−2.3)	(+3.3)		

Table IA.25: Disposition of alternative readings of the currency-level component. Each row names a candidate alternative to the term-premium interpretation and the evidence that addresses it; the nature of each piece of evidence is stated explicitly, and each disposition opens with its category in the four-way classification: ruled out theoretically, rejected empirically, bounded quantitatively, or observationally equivalent.

Alternative	Disposition and evidence
Expected rolled funding costs (synthetic dollar funding)	REJECTED EMPIRICALLY. Rejected by inequality: the observed ten-year components exceed in absolute size the three-month funding cost they would have to capitalise, which no funding process can deliver at any persistence (Proposition 4); crisis response ratios of 0.04 to 0.32 in March 2020.
Dealer balance-sheet funding	REJECTED EMPIRICALLY. Controlled directly: supply pricing survives the intermediary-capital control; the pair residuals do not load on intermediary capital or the VIX.
Corporate issuance-hedging flows at long tenors	BOUNDED QUANTITATIVELY; RESIDUALLY OBSERVATIONALLY EQUIVALENT. The credit-migration reading of Liao (2020) ties long-tenor bases to issuance-driven hedging demand: transitory issuance pressure is bounded by the far-forward volatility restriction and the event placebos; a persistent, priced issuance-demand component is currency-attached long-horizon compensation, absorbed into the swap term premium under the global-marginal-investor reading (Section 4.3.4) rather than a contaminant outside it.
Benchmark credit and any swap-rate-resident component	RULED OUT THEORETICALLY. Cancelled algebraically in the cross-country government bond term premium (Corollary 1); consistent with the funding-forecast asymmetry between the swap and government term premia.
Sovereign credit differential in the government leg	REJECTED EMPIRICALLY. Controlled directly on measured protection: CDS differentials small and weakly correlated on the post-2009 common sample, the supply loading unchanged by their inclusion, and the post-2013 configuration intact, the yen's larger, under subtraction of the full CDS differential (Internet Appendix Section G.3).
Forward dollar exchange-rate premium	RULED OUT THEORETICALLY AND REJECTED EMPIRICALLY. Excluded structurally (covered, collateralised legs; sub-basis-point settlement exposure) and empirically (no dollar forecast power; dollar-neutral carry; the December 2022 event undiminished in dollar-free differentials).
Pair-specific frictions and quote noise	BOUNDED QUANTITATIVELY. Bounded by the market's own band: 80 to 99 percent of residuals inside the quoted half-spread; second moments identified from triangle covariances agree; invariant across five selection rules and hub exclusion.
Leg-duration mismatch artefact	RULED OUT THEORETICALLY. Excluded by construction and scaling: the reversal holds per unit of duration in both curve families; compensation scales with the duration exposure, which funding does not reference.
Currency-level intermediation and demand mechanisms (collateral demand, regulatory netting asymmetries, benchmark-specific hedging pressure, persistent receive-fixed demand, swap-line expectations, dealer inventory	REJECTED EMPIRICALLY; BOUNDED QUANTITATIVELY; RESIDUALLY OBSERVATIONALLY EQUIVALENT. Transitory, short-concentrated forms are rejected by the far-forward loading and volatility bounds; currency-level compensation these mechanisms generate prices the domestic curve through the same dealer balance sheet and is content of the swap term premium under the global-marginal-investor reading (Section 4.3.4); persistent long-horizon forms are adjudicated

Table IA.26: **Announcement-Window Responses to Identified Policy-Surprise Factors.** This table reports regressions of $[-1,+2]$ trading-day changes of the daily quoted-construction benchmark government term-premium differentials, in basis points, on identified announcement-surprise factors standardised to unit variance, with standard errors clustered by announcement date. Panel A: 163 Federal Reserve announcements inside the sample, factors of Swanson (2021) (target, forward guidance, asset purchases; purchase factor oriented so positive is a purchase easing, anchored empirically as described in Section H.6), pooled across the eight dollar pairs. Panel B: 128 European Central Bank announcements, the four factors of Altavilla et al. (2019) constructed from their event-study database by the authors' method as described in Section H.6, euro pair. The placebo rows report the same regressions on the named pair's own quoted one-year cross-currency basis. Pre-specified ex-ante signs: asset purchases positive in Panel A and negative in Panel B; all outcomes reported as returned. Significance: *, **, *** for $p \leq .10, .05, .01$.

	Panel A: Federal Reserve, pooled eight pairs (bp per sd)			
	2y	5y	10y	20y
Target	$+0.98^{**}$ (2.2)	-0.06 (-0.1)	$+0.83^{*}$ (1.7)	$+0.64$ (1.2)
Forward guidance	$+0.81^{*}$ (1.8)	$+0.85^{*}$ (1.9)	$+0.27$ (0.8)	$+0.39$ (0.9)
Asset purchases	$+0.22$ (1.0)	$+1.24^{**}$ (2.1)	$+0.68^{**}$ (2.1)	$+0.65^{***}$ (2.7)
Placebo, one-year basis: asset purchases		$+0.31^{**}$ (2.1)		
	Panel B: European Central Bank, euro pair (bp per sd)			
	2y	5y	10y	20y
Target	$+0.98^{**}$ (2.1)	$+1.43^{***}$ (2.9)	$+1.30^{***}$ (2.9)	$+1.26^{***}$ (3.3)
Timing	$+0.04$ (0.1)	-0.57 (-1.1)	-0.52 (-1.2)	$+0.01$ (0.0)
Forward guidance	$+0.69$ (1.1)	$+0.20$ (0.4)	-0.28 (-0.5)	-0.71 (-1.5)
Asset purchases	-0.53 (-1.1)	-0.67^{*} (-1.7)	-0.97^{**} (-2.1)	-1.39^{***} (-3.1)
Placebo, one-year basis: asset purchases		-0.10 (-0.3)		

Table IA.27: **Executable-Band Identified Intervals for the Reversal.** Sub-period means of the ten-year cross-country government bond term premium with the mean executable band in parentheses, the exact half-width of the interval of recoveries attainable when every quoted basis is perturbed within its quoted half-spread; the final column reports the share of post-2013 months in which the monthly interval excludes zero. Pre-crisis is 2002 to 2007 and post-crisis 2013 to June 2026. The full-sample mean band, averaged over all months rather than within the sub-periods printed here, runs from 3.6 to 5.1 basis points across the dollar pairs, the figure the main text quotes. The underlying prices are Bloomberg dealer-contributed composite quotes, indicative rather than firm, so the bands measure quoted spreads rather than guaranteed execution.

Pair	Pre-crisis mean (bp, ± band)	Post-crisis mean (bp, ± band)	Interval excludes zero (share of post months)
USD-EUR	$+27.9$ (± 2.8)	-55.8 (± 5.0)	96%
USD-JPY	$+37.4$ (± 2.6)	-83.2 (± 4.8)	100%
USD-GBP	$+11.5$ (± 2.7)	-5.2 (± 4.6)	76%
USD-CAD	$+30.7$ (± 2.8)	-49.7 (± 3.5)	100%
USD-AUD	$+4.0$ (± 2.9)	-11.5 (± 3.7)	80%
USD-NZD	-29.7 (± 2.9)	$+4.1$ (± 5.6)	73%
USD-SEK	$+9.9$ (± 3.1)	-61.7 (± 5.6)	100%
USD-NOK	$+1.5$ (± 3.0)	-46.4 (± 5.2)	100%



Figure IA.4: **Bridge-Identity Decomposition of the Cross-Country Government Bond Term Premium at the Five-Year Tenor.** Each panel plots, for one of the nine USD bilateral pairs, the monthly time series of the five-year cross-country government bond term premium $\widehat{ctp}_5^{gov,j,USD}$ (black line) constructed as $x + \Delta SS$ from directly quoted series, with its quoted components drawn as stacked bands, the bilateral cross-currency basis (representing the swap term premium) from zero (blue) and the cross-country swap spread differential stacked on it (green), the outer envelope coinciding with the government line exactly, all in basis points. Sample: January 2002 to June 2026.



Figure IA.5: **Bridge-Identity Decomposition of the Cross-Country Government Bond Term Premium at the Two-Year Tenor.** Each panel plots, for one of the nine USD bilateral pairs, the monthly time series of the two-year cross-country government bond term premium $\widehat{ctp}_2^{gov,j,USD}$ (black line) constructed as $x + \Delta SS$ from directly quoted series, with its quoted components drawn as stacked bands, the bilateral cross-currency basis (representing the swap term premium) from zero (blue) and the cross-country swap spread differential stacked on it (green), the outer envelope coinciding with the government line exactly, all in basis points. Sample: January 2002 to June 2026.

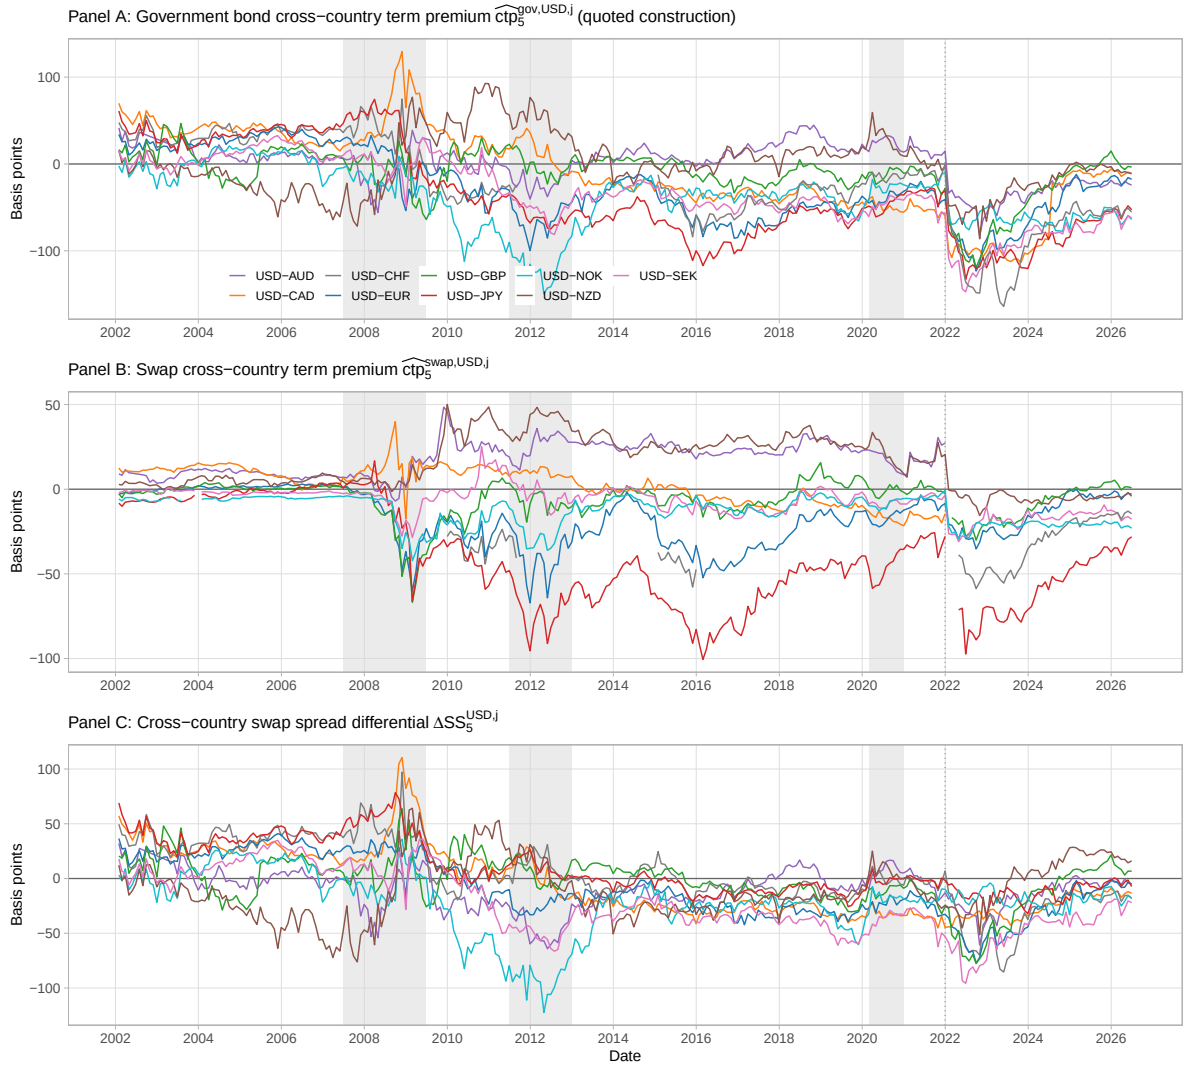


Figure IA.6: **Five-Year Cross-Country Term Premia and Swap Spread Differentials.** Three-panel time-series view paralleling Figure 2 of the main text at the five-year tenor. On euro-cross quote-outage months, verified as zero-quote windows in the raw data, no closed triangle exists and the extracted swap term premium of Panel B is not reported, breaking at every unidentified month, while the quoted series of Panels A and C continue. Panel A plots the cross-country government bond term premium $\widehat{ctp}_5^{gov,j,USD}$, Panel B the cross-country swap term premium $\widehat{ctp}_5^{swap,j,USD}$, and Panel C the cross-country swap spread differential $\Delta SS_5^{USD,j}$, in basis points. Sample: January 2002 to June 2026.

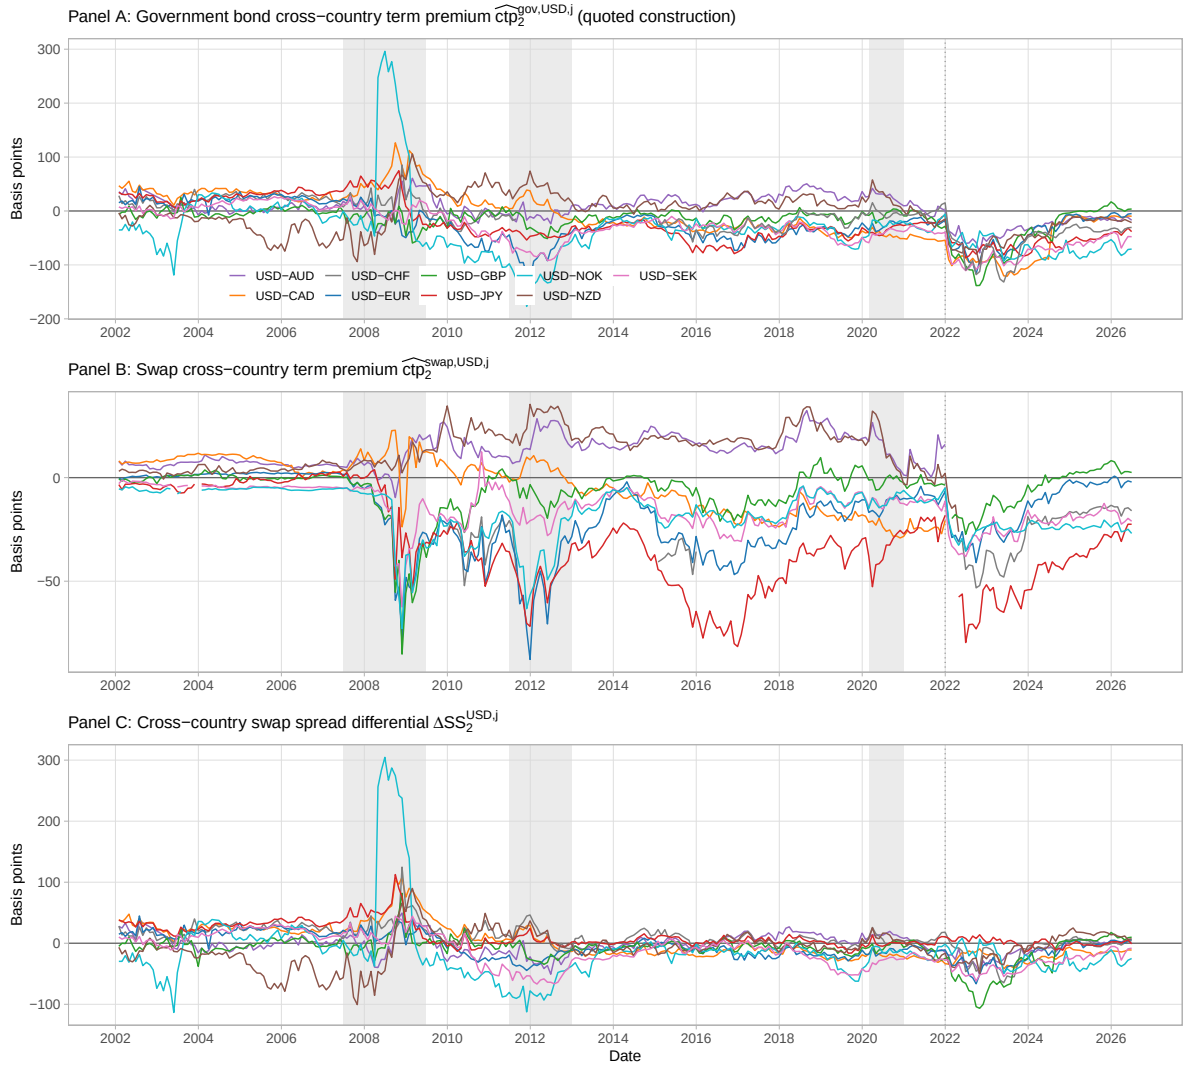


Figure IA.7: **Two-Year Cross-Country Term Premia and Swap Spread Differentials.** Three-panel time-series view paralleling Figure 2 of the main text at the two-year tenor. On euro-cross quote-outage months, verified as zero-quote windows in the raw data, no closed triangle exists and the extracted swap term premium of Panel B is not reported, breaking at every unidentified month, while the quoted series of Panels A and C continue. Panel A plots the cross-country government bond term premium $\widehat{ctp}_2^{gov,j,USD}$, Panel B the cross-country swap term premium $\widehat{ctp}_2^{swap,j,USD}$, and Panel C the cross-country swap spread differential $\Delta SS_2^{USD,j}$, in basis points. Sample: January 2002 to June 2026.

Table IA.28: **Inflation Expectations as Controls: The Column (5) Specification with U.S. Expected Inflation.** The column (5) free-float specification of Table 1 with the U.S. realised year-on-year CPI control replaced by the standardised ten-year expected inflation series of the Federal Reserve Bank of Cleveland (Haubrich et al., 2012), denoted $\mathrm{E}\pi_t^{\mathrm{US},z}$; the partner-country realised inflation control is retained. The dependent variable is, in turn, the ten-year cross-country government bond term premium differential, the cross-country swap term premium, and the cross-country swap spread differential, on the identical sample, so the coefficients satisfy the bridge identity $\beta^{\mathrm{gov}} = \beta^{\mathrm{swap}} + \beta^{\Delta SS}$ exactly. All specifications include currency-pair fixed effects; standard errors are Driscoll and Kraay (1998) kernel estimates with Bartlett kernel and Schwert-rule lag. t-statistics in parentheses. Notes: $^{*}\,p \le .10$; $^{**}\,p \le .05$; $^{***}\,p \le .01$. Sample: January 2002 to May 2025. Throughout, the quoted bilateral cross-currency basis x is used to represent the cross-country swap term premium, which it equals up to the pair-specific residual bounded within quoted half-spreads.

	$\widehat{ctp}_{10}^{\,gov,j,USD}$	$\widehat{ctp}_{10}^{\,swap,j,USD}$	$\Delta SS_{10}^{USD,j}$
$\log(\widetilde{d})_{\mathrm{US}}$	-1.169^{***}	-0.286^{***}	-0.883^{***}
	(-5.02)	(-2.71)	(-6.47)
$\log(\widetilde{d})_j$	$+0.157^{**}$	$+0.092^{*}$	$+0.065$
	$(+2.00)$	$(+1.85)$	$(+1.37)$
$\widetilde{\mathrm{WAM}}_{\mathrm{US}}$	-0.056	$+0.020$	-0.076^{**}
	(-1.29)	$(+0.80)$	(-2.37)
$\widetilde{\mathrm{WAM}}_j$	$+0.030$	-0.012	$+0.042^{***}$
	$(+1.52)$	(-1.26)	$(+2.66)$
$\mathrm{QE}_t^{\mathrm{US}}$	$+0.235^{***}$	$+0.097^{***}$	$+0.138^{***}$
	$(+4.22)$	$(+3.59)$	$(+3.45)$
QE_t^{j}	-0.279^{***}	-0.194^{***}	-0.085^{***}
	(-4.32)	(-4.82)	(-2.84)
HKM_t^{z}	$+0.073$	$+0.021$	$+0.052$
	$(+1.17)$	$(+0.74)$	$(+1.32)$
$\mathrm{HKM}_t^{z} \times \mathrm{QE}_t^{\mathrm{US}}$	$+0.184^{**}$	$+0.104^{***}$	$+0.080$
	$(+2.33)$	$(+2.81)$	$(+1.45)$
$\mathrm{rGDP}_t^{\mathrm{US},z}$	$+0.024$	$+0.026^{*}$	-0.001
	$(+0.82)$	$(+1.69)$	(-0.06)
$\mathrm{rGDP}_t^{j,z}$	-0.016	-0.015	-0.001
	(-0.90)	(-1.49)	(-0.06)
$\mathrm{ir3m}_t^{\mathrm{US},z}$	$+0.220^{***}$	$+0.110^{***}$	$+0.109^{***}$
	$(+3.76)$	$(+3.43)$	$(+2.99)$
$\mathrm{ir3m}_t^{j,z}$	-0.251^{***}	-0.200^{***}	-0.051
	(-4.16)	(-5.34)	(-1.56)
$\mathrm{TS}_t^{\mathrm{US},z}$	$+0.005$	$+0.026$	-0.021
	$(+0.14)$	$(+1.43)$	(-1.05)
$\mathrm{TS}_t^{j,z}$	$+0.050$	-0.047^{**}	$+0.097^{***}$
	$(+1.42)$	(-2.18)	$(+4.25)$
$\mathrm{Unemp}_t^{\mathrm{US},z}$	-0.035	$+0.013$	-0.048^{**}
	(-1.12)	$(+0.99)$	(-2.05)
$\mathrm{Unemp}_t^{j,z}$	$+0.069$	-0.043^{**}	$+0.113^{***}$
	$(+1.58)$	(-2.58)	$(+3.14)$
$\mathrm{Fiscal\ bal}_t^{\mathrm{US}}$	-0.006	-0.006	-0.000
	(-0.53)	(-1.55)	(-0.02)
$\mathrm{Fiscal\ bal}_t^{j}$	-0.014	-0.001	-0.013^{*}
	(-1.37)	(-0.24)	(-1.73)
$\mathrm{E}\pi_t^{\mathrm{US},z}$	-0.049^{**}	-0.001	-0.048^{***}
	(-2.21)	(-0.16)	(-2.68)
$\mathrm{CPI}_t^{j,z}$	-0.041	-0.021^{*}	-0.020
	(-1.57)	(-1.80)	(-1.05)
Dollar_t^{z}	-0.054	-0.078^{***}	$+0.024$
	(-1.27)	(-3.22)	$(+0.85)$
BBB_t^{z}	$+0.010$	$+0.006$	$+0.004$
	$(+0.25)$	$(+0.33)$	$(+0.14)$
MOVE_t^{z}	-0.029	-0.013	-0.016
	(-1.51)	(-1.48)	(-1.18)
VIX_t^{z}	$+0.054^{***}$	$+0.021^{**}$	$+0.033^{**}$
	$(+2.93)$	$(+2.40)$	$(+2.55)$
Within-R^2	0.659	0.525	0.647
N	1,911	1,911	1,911

Table IA.29: **Estimated Break Dates.** Bai-Perron mean-shift break dates with ninety percent confidence intervals, minimum segment two years, monthly 2002 to June 2026; the average series is the cross-sectional mean of the eight ten-year benchmark differentials.

Series	Estimated break dates [90% interval]
Average, two breaks	2007-10 [2007-09, 2007-11]; 2021-12 [2021-07, 2022-02]
Average, selected (7)	first two: 2007-08 [2007-07, 2007-09]; 2011-03 [2011-01, 2011-05]
USD-EUR	2008-08 [2008-06, 2008-09]; 2015-01 [2014-07, 2015-05]
USD-JPY	2009-01 [2008-12, 2009-02]; 2013-03 [2012-09, 2013-05]
USD-GBP	2021-12 [2021-10, 2022-02]; 2024-01 [2023-12, 2024-03]
USD-CAD	2010-02 [2010-01, 2010-03]; 2018-10 [2017-11, 2018-11]
USD-AUD	2013-05 [2013-03, 2014-06]; 2021-12 [2021-11, 2022-01]
USD-NZD	2008-11 [2008-09, 2009-02]; 2021-12 [2021-08, 2022-05]
USD-SEK	2011-02 [2011-01, 2011-04]; 2021-12 [2021-09, 2022-01]
USD-NOK	2008-10 [2008-08, 2008-11]; 2013-05 [2013-03, 2013-10]

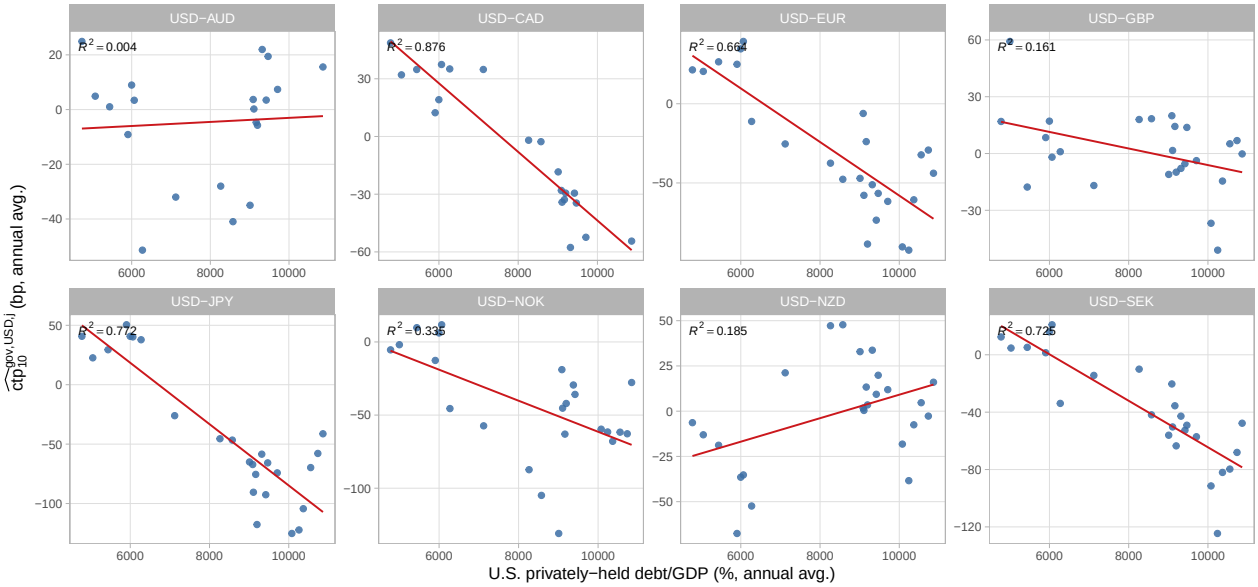


Figure IA.8: **Cross-Country Government Bond Term Premium and U.S. Privately-Held Debt-to-GDP, by Pair (10Y).** Each panel reports a scatter plot of the ten-year cross-country government bond term premium $\widehat{ctp}_{10}^{gov,j,USD}$ (vertical axis, basis points) against the U.S. privately held debt-to-GDP ratio (horizontal axis, percent) for one of the eight USD bilateral pairs, computed as annual averages over the full sample. The privately held debt-to-GDP measure is gross U.S. federal debt-to-GDP net of Federal Reserve sovereign holdings, constructed on the same basis as the supply variables in Table 1. The red line is the bivariate ordinary-least-squares fit and the panel title reports the corresponding R^2. Sample: January 2002 to June 2026, annual averages.

Table IA.30: **Sensitivity to Alternative Selection Rules within the Affine Solution Set.** This table reports three complementary statistics on the dispersion of the recovered cross-country swap term premium $\widehat{ctp}_{t,N}^{swap,j,USD}$ under three alternative selection rules within the affine solution set of the triangular system, relative to the L2 minimum-norm baseline used in the main text. Panel A reports the maximum absolute time-series deviation across the sample, in basis points. Panel B reports the time-series correlation between the recovered $\widehat{ctp}_{t,N}^{swap,j,USD}$ under each alternative rule and the L2 baseline at the ten-year tenor. Panel C reports the maximum subsample-mean shift across the three alternative rules, computed as $\max_{\text{alt}} |\overline{ctp}^{\text{alt}} - \overline{ctp}^{\text{L2}}|$ for each pair, where $\overline{ctp}$ is the subsample mean of the recovered ten-year cross-country swap term premium. The three alternatives are: *L1*, the L1 minimum-norm solution, $\arg\min \|\boldsymbol{\pi}\|_1$ subject to $\mathbf{D}\boldsymbol{\pi} = \boldsymbol{\Delta}$, computed via linear programming; *wL2*, the weighted-L2 solution with the diagonal weight matrix scaling each pair-specific friction by the inverse of the absolute level of the corresponding bilateral basis at that date; and *Zero-pin*, the unique solution that fixes $\hat{\pi}_{t,N}^{EUR/USD} = 0$ at every date and identifies the remaining frictions from the triangular system under L2 minimum norm. The corresponding deviations in the cross-country government bond term premium $\widehat{ctp}_{t,N}^{gov,j,USD}$ are identical at every pair-tenor cell, since the bridge identity adds the same directly observable cross-country swap spread differential on both sides of the comparison. Sample: January 2002 to December 2021.

Panel A. Maximum absolute deviation $|\hat{\pi}^{\text{alt}} - \hat{\pi}^{\text{L2}}|$ in basis points

	2Y			5Y			10Y		
Pair	L1	wL2	Zero-pin	L1	wL2	Zero-pin	L1	wL2	Zero-pin
USD-EUR	1.14	0.47	1.93	1.18	0.69	1.87	2.05	1.53	3.13
USD-JPY	1.94	1.17	0.96	2.81	0.99	0.93	2.49	1.47	1.56
USD-GBP	1.91	0.99	0.96	2.39	1.85	0.93	1.18	1.10	1.56
USD-CAD	2.33	1.35	0.96	1.49	1.44	0.93	1.88	0.91	1.56
USD-AUD	2.54	1.63	0.96	2.63	1.52	0.93	2.65	1.87	1.56
USD-NZD	2.09	1.36	0.96	1.90	1.29	0.93	3.73	3.65	1.56

Panel B. Time-series correlation of $\widehat{ctp}_{10}^{swap,j,USD}$ under each alternative rule with the L2 baseline

Pair	L1	wL2	Zero-pin
USD-EUR	0.999	1.000	0.998
USD-JPY	1.000	1.000	1.000
USD-GBP	1.000	1.000	0.999
USD-CAD	0.999	1.000	0.999
USD-AUD	0.999	1.000	0.999
USD-NZD	0.998	1.000	1.000

Panel C. Maximum subsample-mean shift $\max_{\text{alt}} |\overline{ctp}_{10}^{\text{alt}} - \overline{ctp}_{10}^{\text{L2}}|$ in basis points

Pair	Full sample	Pre-GFC	Post-GFC
USD-EUR	0.30	0.17	0.22
USD-JPY	0.34	0.09	0.36
USD-GBP	0.15	0.09	0.22
USD-CAD	0.28	0.09	0.43
USD-AUD	0.88	0.20	0.99
USD-NZD	1.04	0.28	1.33

Table IA.31: **Sensitivity to Regularisation Parameter.** This table reports the sample-mean recovered ten-year cross-country swap term premium $\widehat{ctp}_{10}^{swap,j,USD}$ for each USD bilateral pair under Tikhonov regularisation $\hat{\boldsymbol{\pi}}_{t,N}^{(\lambda)} = \mathbf{D}^\top(\mathbf{D}\mathbf{D}^\top + \lambda\mathbf{I})^{-1}\boldsymbol{\Delta}_{t,N}$ at five values of the regularisation parameter λ, spanning six orders of magnitude. The $\lambda \to 0$ limit recovers the L2 minimum-norm baseline used in the main text. Recovered values are essentially indistinguishable across the entire range: the maximum per-date deviation between $\lambda = 10^{-8}$ and $\lambda = 10^{-2}$ is below 0.02 basis points for every USD pair at every tenor, with the per-date mean deviation below 0.005 basis points. The numerical insensitivity reflects the well-conditioned structure of the 8×8 Gram matrix $\mathbf{D}\mathbf{D}^\top$, whose eigenvalue 2 has multiplicity seven alongside a single eigenvalue of 10, so that its condition number is 5, so the ridge term $\lambda\mathbf{I}$ is small relative to the smallest eigenvalue across the entire range of λ studied. Sample: January 2002 to December 2021. The end date here predates the headline panel by approximately two years; the L2 baseline values reported in this table therefore differ slightly from the headline subsample-mean values in Table 2 and Table IA.12, which use the full January 2002 to June 2026 sample.

	Regularisation Parameter (λ)				
Pair	10^{-8}	10^{-6}	10^{-4}	10^{-3}	10^{-2}
USD-EUR	−18.205	−18.205	−18.205	−18.205	−18.205
USD-JPY	−44.604	−44.604	−44.604	−44.604	−44.602
USD-GBP	−5.886	−5.886	−5.886	−5.886	−5.886
USD-CAD	+4.919	+4.919	+4.919	+4.919	+4.918
USD-AUD	+22.146	+22.146	+22.146	+22.146	+22.145
USD-NZD	+28.797	+28.797	+28.797	+28.797	+28.796

Table IA.32: **Leave-One-Currency-Out of the Triangular Network.** For each excluded currency, every pair containing that currency is removed from the monthly ten-year system, the reduced triangular system is re-solved under the L2 minimum-norm rule, and the recovered dollar-pair swap term premium differentials are compared with the full-network baseline on identical observations. Corr and MAD report the pooled correlation and mean absolute deviation (basis points) across the remaining dollar pairs; sd reports the baseline standard deviation for scale. The final two columns report the extracted-design preferred-habitat headline coefficient on log U.S. free-float supply (Driscoll-Kraay t-statistics at the Schwert-rule lag in parentheses) under the leave-one-out recovery and under the baseline recovery restricted to the same reduced panel. Excluding the euro removes every cross quote, so the network degenerates to a dollar star with no cycles and the recovery collapses to the raw bases; the leave-one-out panel is then slightly larger because the raw basis is available in months where the full extraction is not.

Excluded	Corr	MAD (bp)	sd (bp)	Supply, LOO	Supply, baseline
EUR	0.9996	0.48	26.5	−1.055 (−10.4)	−1.066 (−9.4)
JPY	1.0000	0.07	20.0	−0.896 (−10.5)	−0.895 (−10.5)
GBP	1.0000	0.03	27.0	−1.086 (−9.6)	−1.086 (−9.7)
CHF	1.0000	0.02	24.7	—	—
CAD	1.0000	0.04	26.4	−1.008 (−7.2)	−1.008 (−7.2)
AUD	1.0000	0.09	24.9	−1.055 (−8.5)	−1.056 (−8.5)
NZD	1.0000	0.14	24.1	−1.306 (−10.9)	−1.306 (−10.8)
SEK	1.0000	0.05	27.0	−1.066 (−9.9)	−1.066 (−9.9)
NOK	1.0000	0.04	26.8	−1.037 (−10.2)	−1.038 (−10.2)

Table IA.33: **Currency-Level Variance Shares and the Dollar-Free Verification.** Panel A: share of the monthly variance of each recovered ten-year dollar-pair differential attributable to the U.S. side, $[\mathrm{var}(\Delta tp^{USD}) - \mathrm{cov}(\Delta tp^{USD}, \Delta tp^{j})]/\mathrm{var}(\Delta ctp^{j,USD})$, under the zero-sum normalisation; the partner share is the complement. The full-sample and post-2008 columns coincide at the printed integer precision in every row, a genuine coincidence rather than a duplication: the post-2008 window contributes most of the pooled variance of monthly changes, so the full-sample shares (mean 37.1 percent at one decimal) sit within a percentage point of the post-2008 shares (mean 37.5). Panel B: pooled and range statistics for the dollar-free regressions of changes in raw dollar-pair base differences on changes in the corresponding raw euro-cross differences. Monthly, January 2002 to June 2026.

Panel A: U.S.-side variance share			
Pair	Full sample	Pre-2008	Post-2008
USD/EUR	41%	27%	41%
USD/JPY	28%	11%	28%
USD/GBP	45%	11%	45%
USD/CHF	41%	12%	41%
USD/CAD	21%	9%	21%
USD/AUD	27%	15%	27%
USD/NZD	22%	13%	22%
USD/SEK	51%	30%	51%
USD/NOK	59%	25%	59%
Mean	37%	17%	37%
Panel B: dollar-free verification (raw quotes)			
Pooled slope (21 pairs, $N = 5{,}697$)	0.90 ($R^2 = 0.88$)		
Per-pair slope range	0.80 to 0.96		
Per-pair R^2 range	0.74 to 0.95		

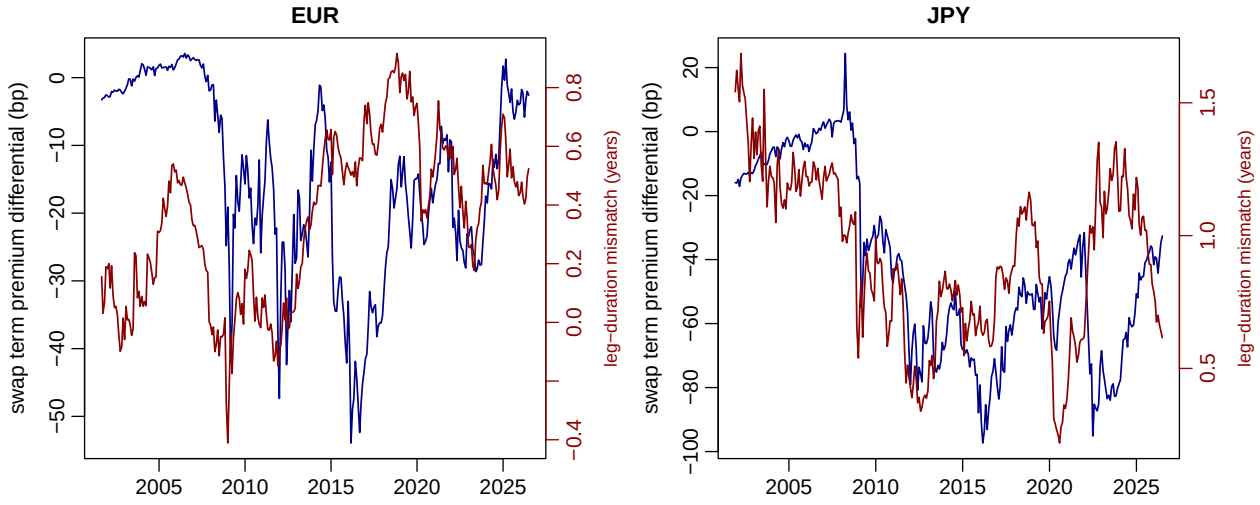


Figure IA.9: *Leg-duration mismatch and the swap term-premium differential. Each panel plots the recovered ten-year swap term-premium differential in basis points (dark blue, left axis) and the swap-curve leg-duration mismatch $D_j - D_{US}$ in years (dark red, right axis), monthly. Durations are Macaulay durations of the ten-year par bond on each currency's swap curve, constructed as described in this section.*

Table IA.34: **Robustness: Gross-Debt Specification of the Preferred-Habitat Panel.** This table re-estimates the five-column specification of Table 1 with log gross sovereign debt-to-GDP on both sides in place of the privately held free-float measure; the maturity columns are omitted, since real gross average-maturity histories for the partner sovereigns are not part of my dataset and the free-float maturity is not a valid stand-in for a gross specification. Because no maturity series is required, the Swedish and Norwegian pairs enter every column, so the panel comprises eight pairs. The dependent variable, controls, fiscal balances, fixed effects, sample bounds, and Driscoll and Kraay (1998) standard errors are otherwise identical to Table 1. Throughout, the quoted bilateral cross-currency basis x is used to represent the cross-country swap term premium, which it equals up to the pair-specific residual bounded within quoted half-spreads. t-statistics in parentheses. Notes: $^{*}\,p \leq .10$; $^{**}\,p \leq .05$; $^{***}\,p \leq .01$.

	(1)	(2)	(3)	(4)	(5)
$\log(d/\text{GDP})_{\text{US}}$	-1.044^{***}	-1.016^{***}	-0.955^{***}	-1.379^{***}	-1.222^{***}
	(-10.03)	(-9.05)	(-8.27)	(-8.96)	(-6.66)
$\log(d/\text{GDP})_{j}$	$+0.370^{***}$	$+0.384^{***}$	$+0.347^{***}$	$+0.146^{**}$	$+0.145^{**}$
	$(+5.22)$	$(+6.06)$	$(+5.34)$	$(+2.32)$	$(+2.09)$
QE_t^{US}		$+0.158^{**}$	$+0.317^{***}$	$+0.331^{***}$	$+0.345^{***}$
		$(+2.40)$	$(+4.16)$	$(+3.97)$	$(+5.08)$
QE_t^{j}		-0.183^{**}	-0.220^{***}	-0.286^{***}	-0.282^{***}
		(-2.39)	(-2.89)	(-4.68)	(-4.51)
HKM_t^{z}			$+0.129^{***}$	$+0.058$	$+0.044$
			$(+4.39)$	$(+0.82)$	$(+0.89)$
$\text{HKM}_t^{z} \times \text{QE}_t^{\text{US}}$			$+0.087$	$+0.201^{**}$	$+0.251^{***}$
			$(+1.33)$	$(+2.12)$	$(+2.74)$
$\text{rGDP}_t^{\text{US},z}$				$+0.037$	$+0.029$
				$(+1.03)$	$(+1.08)$
$\text{rGDP}_t^{j,z}$				-0.021	-0.020
				(-1.53)	(-1.52)
$\text{ir3m}_t^{\text{US},z}$				$+0.089$	$+0.140^{***}$
				$(+1.24)$	$(+2.84)$
$\text{ir3m}_t^{j,z}$				-0.304^{***}	-0.269^{***}
				(-6.00)	(-4.72)
$\text{TS}_t^{\text{US},z}$				-0.070	-0.045
				(-1.46)	(-1.28)
$\text{TS}_t^{j,z}$				$+0.016$	$+0.037$
				$(+0.42)$	$(+0.97)$
$\text{Unemp}_t^{\text{US},z}$				$+0.026$	-0.020
				$(+0.80)$	(-0.56)
$\text{Unemp}_t^{j,z}$				$+0.079^{**}$	$+0.069^{**}$
				$(+2.34)$	$(+2.03)$
$\text{Fiscal bal}_t^{\text{US}}$				$+0.004$	$+0.001$
				$(+0.42)$	$(+0.06)$
Fiscal bal_t^{j}				-0.015	-0.015
				(-1.37)	(-1.50)
$\text{CPI}_t^{\text{US},z}$					-0.035
					(-1.43)
$\text{CPI}_t^{j,z}$					-0.027
					(-1.19)
Dollar_t^{z}					-0.035
					(-1.13)
BBB_t^{z}					$+0.001$
					$(+0.03)$
MOVE_t^{z}					-0.029^{*}
					(-1.79)
VIX_t^{z}					$+0.046^{***}$
					$(+2.87)$
Pair fixed effects	Yes	Yes	Yes	Yes	Yes
Within-R^2	0.384	0.428	0.472	0.676	0.691
N	2,120	2,120	2,016	1,975	1,975

Table IA.35: **Subsample Stability of the Preferred-Habitat Panel: Pre-Pandemic vs. Post-Pandemic.** The dependent variable is the ten-year cross-country government bond term premium differential $\widehat{ctp}_{t,10}^{\text{gov},j,\text{USD}}$ in percentage points. Each column reports the column (1) Main supply specification of Table 1, with the four privately held free-float supply variables $\log(\widetilde{d})_{\text{US}}$, $\log(\widetilde{d})_j$, $\widetilde{\text{WAM}}_{\text{US}}$, and $\widetilde{\text{WAM}}_j$, estimated on a different sample window. Column 1 reports the full-sample estimate through June 2026; columns 2–4 report nested pre-pandemic windows that successively extend the sample end date through the COVID-19 pandemic and the post-pandemic monetary tightening; column 5 isolates the post-pandemic window from March 2020 through June 2026. All specifications include currency-pair fixed effects; standard errors are Driscoll and Kraay (1998) kernel estimates with Bartlett kernel and Schwert-rule lag. t-statistics in parentheses. Notes: $^{*}\,p \leq .10$; $^{**}\,p \leq .05$; $^{***}\,p \leq .01$. Throughout, the quoted bilateral cross-currency basis x is used to represent the cross-country swap term premium, which it equals up to the pair-specific residual bounded within quoted half-spreads.

	Full sample	Pre-pandemic windows			Post-pandemic
	2002-01 to 2026-06	2002-01 to 2018-11	2002-01 to 2020-02	2002-01 to 2021-12	2020-03 to 2026-06
$\log(\widetilde{d})_{\text{US}}$	-1.051^{***}	-1.217^{***}	-1.188^{***}	-1.086^{***}	$+1.599^{*}$
	(-10.72)	(-8.84)	(-9.34)	(-8.89)	$(+1.93)$
$\log(\widetilde{d})_j$	$+0.521^{***}$	$+0.800^{***}$	$+0.728^{***}$	$+0.594^{***}$	$+0.510^{***}$
	$(+6.87)$	$(+7.51)$	$(+7.68)$	$(+6.10)$	$(+5.35)$
$\widetilde{\text{WAM}}_{\text{US}}$	-0.060	-0.060	-0.032	$+0.007$	-0.545^{***}
	(-1.26)	(-1.18)	(-0.82)	$(+0.20)$	(-3.74)
$\widetilde{\text{WAM}}_j$	$+0.023$	-0.008	$+0.005$	$+0.024$	$+0.038^{**}$
	$(+0.99)$	(-0.31)	$(+0.21)$	$(+1.04)$	$(+2.10)$
N	2,043	1,410	1,513	1,665	530
Within-R^2	0.387	0.398	0.401	0.356	0.293

Table IA.36: **HKM Capital Ratio versus AEM Leverage Factor.** Panel regressions of the ten-year cross-country government bond term premium and its two bridge-identity components on the standardised intermediary capital ratio of He et al. (2017) and the standardised broker-dealer leverage factor in the spirit of Adrian et al. (2014), on the eight-pair panel. The leverage factor is re-derived from the Federal Reserve Z.1 financial accounts: total financial assets over net worth of security brokers and dealers (series FL664090005 and FL664190005), log leverage in levels, standardised over the estimation window and step-carried from quarterly to monthly frequency, vintage through 2026-Q1. Both specifications include the country-side macro factors and the general government fiscal balances of column (4) of Table 1 together with the standardised broad dollar, BBB spread, MOVE, and VIX controls, and currency-pair fixed effects; standard errors are Driscoll and Kraay (1998) kernel estimates with Bartlett kernel and Schwert-rule lag. Panel A reports the joint specification, bounded at May 2025 by the published He et al. (2017) vintage; Panel B reports the leverage factor alone, through June 2026. N counts the pair-months of the complete-case panel, which the Scandinavian macro and fiscal coverage and, after December 2021, the quoted Canadian and Australian legacy series restrict well below the full frame. t-statistics in parentheses. Notes: $^{*}\,p \leq .10$; $^{**}\,p \leq .05$; $^{***}\,p \leq .01$. Throughout, the quoted bilateral cross-currency basis x is used to represent the cross-country swap term premium, which it equals up to the pair-specific residual bounded within quoted half-spreads.

	Dependent variable		
	$\widehat{ctp}_{10}^{\,gov,j,USD}$	$\widehat{ctp}_{10}^{\,swap,j,USD}$	$\Delta SS_{10}^{USD,j}$
Panel A: Joint specification (HKM and AEM)			
HKM_t^z	$+0.038$	$+0.065^{***}$	-0.027
	$(+0.56)$	$(+6.34)$	(-0.44)
AEM_t^z	$+0.323$	$+0.036$	$+0.286$
	$(+1.26)$	$(+1.27)$	$(+1.16)$
R_w^2	0.589	0.791	0.414
N (pair-month)	125	125	125
Panel B: AEM only specification			
AEM_t^z	$+0.066$	$+0.170^{***}$	-0.104
	$(+0.38)$	$(+3.24)$	(-0.79)
Pair fixed effects	Yes	Yes	Yes
Macro-financial controls	Yes	Yes	Yes
N (pair-month)	201	201	201

Table IA.37: **Official-Demand Flows on the Valuation-Adjusted Series.** Panel regressions of the ten-year recovered government term-premium differential on the official-demand margins constructed from the valuation-adjusted foreign official Treasury data of Bertaut and Judson (2014), on the six-pair supply panel with the United States and partner free-float controls retained, pair fixed effects, and Driscoll and Kraay (1998) kernel standard errors at the Schwert-rule lag. The flow margin is the twelve-month adjusted net official Treasury flow relative to GDP; the stock margin is the official Treasury position relative to GDP. The valuation-adjusted series is chained across the authors' two release files at June 2011 and ends with their final release in December 2020, which bounds the sample. The design was pre-registered as descriptive support with a positive ex-ante sign on official demand; both margins are reported as returned. t-statistics in parentheses. Notes: $^{*}\,p \leq .10$; $^{**}\,p \leq .05$; $^{***}\,p \leq .01$.

	Coefficient	N (pair-month)
Twelve-month official flow / GDP	$+3.309^{**}$	1,356
	$(+2.30)$	
Official position / GDP	-1.865^{**}	1,356
	(-2.50)	
Free-float controls	Yes	
Pair fixed effects	Yes	
Effective months	229	
Sample end	Dec. 2020	

Table IA.38: **Officially Declared Quantitative Easing Programme Date Windows.** Programme date windows used to construct the country-specific quantitative easing indicators QE_t^c for $c \in \{\mathrm{US}, \mathrm{DE}, \mathrm{JP}, \mathrm{GB}, \mathrm{CA}, \mathrm{AU}, \mathrm{NZ}\}$. The indicator equals one in any month falling within at least one of the country's listed programme windows and zero otherwise. Date windows are taken directly from each central bank's own programme announcement and termination disclosures. The September 2011 to December 2012 Maturity Extension Program is not coded as a purchase window in this indicator, its announcement windows being treated in the event analysis, and the November 2019 restart of net Asset Purchase Programme purchases is not coded as a separate window; the windows listed here are the indicator as estimated.

Central bank	Programme	Start	End
Federal Reserve (US)	QE1	Nov 2008	Mar 2010
	QE2	Nov 2010	Jun 2011
	QE3	Sep 2012	Oct 2014
	Pandemic asset purchases	Mar 2020	Mar 2022
ECB / Eurosystem (DE)	APP / PSPP	Mar 2015	Dec 2018
	PEPP	Mar 2020	Jun 2022
Bank of Japan (JP)	QQE (and YCC continuation)	Apr 2013	Mar 2024
Bank of England (GB)	Asset Purchase Facility (APF)	Mar 2009	Oct 2012
	Brexit-vote programme	Aug 2016	Feb 2017
	Pandemic programme	Mar 2020	Dec 2021
Bank of Canada (CA)	Government Bond Purchase Program (GBPP)	Apr 2020	Oct 2021
RBA (AU)	Bond Purchase Program (BPP)	Nov 2020	Feb 2022
RBNZ (NZ)	Large Scale Asset Purchase (LSAP)	Mar 2020	Jul 2021

Table IA.39: **The Sovereign Credit Differential Against the Government Term-Premium Differential.** Month-end ten-year sovereign CDS differentials, partner minus the United States, in basis points, on the post-2009 common sample (202 months per pair; the United States contract quotes from September 2009). Post-2013 columns are means over January 2013 to June 2026. The final column subtracts the full CDS differential from the government differential month by month, an upper bound on what credit can absorb; differences are computed on unrounded inputs, so a printed cell can differ from printed-column arithmetic in the terminal digit. The Canadian and New Zealand contracts returned no stored history from the vendor source used and are absent.

Pair	N	CDS diff. post-2013	Gov. diff. post-2013	Corr.	Gov. less CDS, post-2013
EUR	202	−5.0	−55.8	+0.06	−50.8
JPY	202	+48.0	−83.2	−0.07	−131.1
GBP	202	+4.6	−5.2	+0.36	−9.9
AUD	202	+38.5	−11.5	+0.31	−50.0
SEK	202	−8.3	−61.7	+0.21	−53.4
NOK	202	−9.9	−46.4	+0.44	−36.5

Table IA.40: Ten-year leg durations of the basis package by curve family. Full-sample mean Macaulay durations of the ten-year par bond on the government and swap curves, and the post-2009 mean leg-duration mismatch against the dollar leg in each family, in years.

	Mean duration (years)		Post-2009 mismatch vs. USD	
Currency	Government	Swap	Government	Swap
USD	8.66	8.60	—	—
EUR	9.06	8.95	+0.56	+0.44
JPY	9.60	9.57	+0.84	+0.79
GBP	8.71	8.69	+0.11	+0.14
CAD	8.73	8.67	+0.14	+0.02
AUD	8.42	8.29	−0.22	−0.33
NZD	8.26	8.16	−0.34	−0.38

Table IA.41: **Recovered versus Model-Based Ten-Year Term-Premium Differentials.** Post-2009 means, in basis points, of the ten-year cross-country government term-premium differential (foreign minus USD): the recovered series against the Joslin et al. (2011) benchmark with the Bauer et al. (2012) bias correction (JSZ-BRW) and without it (JSZ-OLS), estimated per currency on the identical five-node zero-coupon curves as described in the text. "Same side" applies the pre-registered criterion, agreement of the sign of the post-2009 mean between the recovered and the bias-corrected series. "Corr." is the level correlation over the common monthly sample of N months. Sample: monthly, differentials from August 2001 (Canada 2006, Australia 2007) to June 2026, so N runs to 299 months where coverage is complete.

Pair	N	Recovered	JSZ-BRW	JSZ-OLS	Corr.	Same side
EUR	299	−52.1	−30.0	−67.0	−0.20	yes
JPY	293	−74.6	−16.3	−64.0	−0.38	yes
GBP	291	−3.4	+19.5	−7.6	−0.33	no
CAD	242	−26.2	−82.7	−60.1	+0.26	yes
AUD	234	−5.7	−46.3	−64.2	+0.38	yes
NZD	296	+11.3	+20.5	+6.7	+0.13	yes

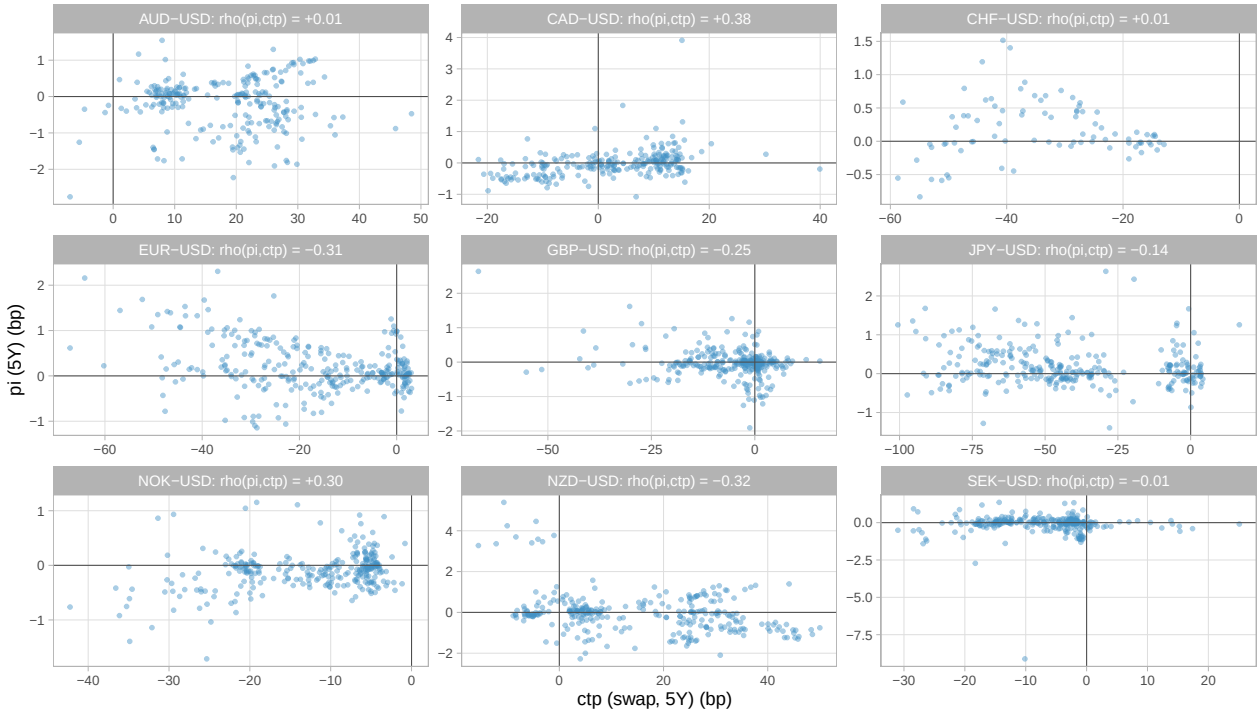


Figure IA.10: **Pair-Specific Component versus Cross-Country Term Premium at the Five-Year Tenor.** Each panel plots the monthly pair-specific component $\pi_{t,5}^{j/USD}$ against the corresponding cross-country swap term premium $ctp_{t,5}^{j,USD}$ in basis points, for one of the nine USD bilateral pairs. The panel title reports the sample time-series correlation $\rho(\pi, ctp)$. Sample: January 2002 to June 2026.

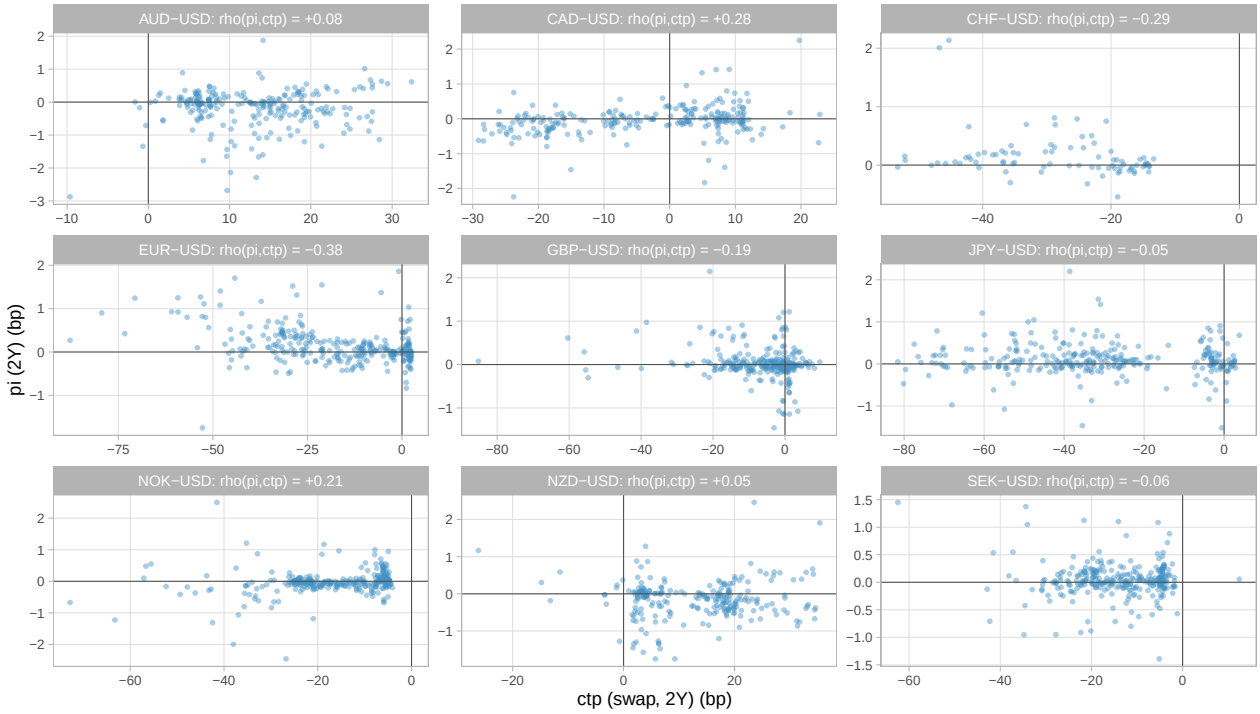


Figure IA.11: **Pair-Specific Component versus Cross-Country Term Premium at the Two-Year Tenor.** Each panel plots the monthly pair-specific component $\pi_{t,2}^{j/USD}$ against the corresponding cross-country swap term premium $ctp_{t,2}^{j,USD}$ in basis points, for one of the nine USD bilateral pairs. The panel title reports the sample time-series correlation $\rho(\pi, ctp)$. Sample: January 2002 to June 2026.