BOLETIM
ECONOMIA EMPÍRICA
v. 6, n. 2 (2026) | ISSN: 2675-3391

Volatility Spillovers and Network Connectedness among Global Systemically Important Banks

Rodrigo Alves dos Santos[a,†]
[a]Instituto Brasileiro de Ensino, Desenvolvimento e Pesquisa, Brazil • [†]Autor correspondente: rodrigo-mc2@hotmail.com

INFORMAÇÕES DO ARTIGO
Palavras-chave:
Systemic risk
Financial contagion
Connectedness
G-SIBs
TVP-VAR
Realized volatility

ABSTRACT

We map the transmission of risk across the world's Global Systemically Important Banks (G-SIBs) by measuring volatility connectedness, the natural carrier of contagion, rather than return co-movement. Using Garman–Klass realized volatility built from daily ranges for 19 USD-traded G-SIBs over 2007–2026, we estimate the Diebold–Yilmaz connectedness network under three specifications – a generalized VAR, a time-varying-parameter VAR (TVP-VAR) and a LASSO/Elastic-Net VAR – and at two sampling frequencies, weekly and daily. Total connectedness is high and stable, at about 91% (weekly) and 87% (daily) of the typical bank's forecast-error variance. The network exhibits a robust hub structure: the net transmitters of risk are the Anglo-American, Swiss and Dutch banks (UBS, JPMorgan, Barclays, ING, Citigroup, Morgan Stanley, HSBC), while the net receivers are the three Japanese G-SIBs (Mizuho, Mitsubishi UFJ, Sumitomo Mitsui) together with Deutsche Bank, Santander and Toronto-Dominion. This structure is invariant to the estimator, the VAR lag, the forecast horizon and the sampling frequency, ruling out a non-synchronous-trading artefact. A currency-robustness check shows that the US-transmitter and Japanese-absorber poles survive removal of the dollar overlay, while the net positions of the cross-listed European banks are sensitive to denomination. We trace the robust structure to the banks' funding models, revisit and qualify the common "US-dominance" finding, and offer a market-based, contagion-oriented complement to the regulatory G-SIB framework.

1. Introduction

Large, internationally active banks are tightly interlinked through interbank funding, common exposures, derivatives and correlated investor expectations. When one institution is stressed, the disturbance can propagate to others – a spillover – and, in the limit, become systemic. Measuring which banks transmit and which absorb such shocks is therefore central to macroprudential policy, to the calibration of capital surcharges, and to portfolio and hedging decisions (Billio et al., 2012; Diebold and Yilmaz, 2014).

Two episodes in March 2023 make the question concrete. Silicon Valley Bank failed after rising policy rates impaired a long-duration bond portfolio and triggered a deposit run – the largest US bank failure since 2008. Within days, long-standing problems at Credit Suisse – losses from Archegos and Greensill and repeated governance failures – culminated in a state-brokered takeover by UBS. Credit Suisse was on the G-SIB list; the speed with which idiosyncratic distress became a system-wide confidence shock is exactly the transmission we quantify.

Our empirical question is: across the world's systemically important banks, how much of each bank's risk is imported from the others, and which banks sit at the centre of the network? We make three design choices that distinguish this paper from much of the applied connectedness literature on banks. First, we model volatility rather than returns: contagion is a second-moment phenomenon, and a range-based realized-volatility proxy follows the global-bank application of Demirer et al. (2018). Second, we confront the non-synchronous trading problem head-on. A daily panel that mixes Tokyo, London and New York closing times can make later-closing markets appear to lead earlier-closing ones purely as a timing artefact (Forbes and Rigobon, 2002); we therefore estimate the network at both the weekly frequency – which overlaps the trading windows, following Diebold and Yilmaz (2009) – and the daily frequency, and compare. Third, we estimate three complementary models (VAR, TVP-VAR, LASSO-VAR) so that the headline features can be checked against the choice of estimator.

Our contribution is fourfold. (i) We provide an updated, risk-based connectedness map of the G-SIBs through 2026, spanning the global financial crisis, the European debt crisis, COVID-19, the 2022–23 tightening cycle and the March-2023 bank stress. (ii) We document a robust hub structure – Anglo-American/Swiss/Dutch transmitters, Japanese and continental-European absorbers – that is stable across estimators, lags, horizons and sampling frequencies. (iii) We address the two measurement concerns most likely to bias a cross-border bank panel: the structure is not driven by non-synchronicity, and a currency check shows the US-transmitter and Japanese-absorber poles are robust to removing the dollar overlay, though the cross-listed European banks' net positions are denomination-sensitive. (iv) We interpret and qualify the "US-dominance" narrative: once risk is measured on volatility, the single-name dominance of JPMorgan gives way to a cluster of Anglo-European hubs whose common feature is a wholesale-funded, capital-markets business model, while the deposit-funded Japanese G-SIBs emerge as the system's principal risk absorbers.

The remainder proceeds as follows. Section 2 sets out the methodology, Section 3 the data, Section 4 the results, Section 5 robustness, and Section 6 concludes with policy implications.

2. Methodology

2.1. Measuring risk: range-based realized volatility

Let $P_t^O, P_t^H, P_t^L, P_t^C$ be the open, high, low and close of a bank's traded instrument on day t. We estimate the daily variance with the Garman–Klass (Garman and Klass, 1980) range estimator,

$$\sigma_t^2 = \tfrac{1}{2}\left(\ln\frac{P_t^H}{P_t^L}\right)^2 - (2\ln 2 - 1)\left(\ln\frac{P_t^C}{P_t^O}\right)^2, \tag{1}$$

several times more efficient than the squared close-to-close return because it exploits the intraday range; being a function of log price ratios it is also invariant to the currency unit, a property we exploit in Section 5.2. We aggregate to realized variance at the chosen frequency, $RV_w = \sum_{t\in w}\sigma_t^2$, and model log realized volatility $y_w = \ln\sqrt{RV_w}$.

2.2. Frequency and non-synchronous trading

A daily panel of banks trading in Asia, Europe and the Americas is non-synchronous: each series is measured over a different local window, so a market that closes later can spuriously appear to Granger-cause those that closed earlier. Aggregating to the calendar week makes the windows effectively overlap and removes this one-day lead–lag (Diebold and Yilmaz, 2009). We report both the weekly (synchronization-robust) and the daily network and treat their comparison as a direct test of whether non-synchronicity drives the directional results.

2.3. The connectedness framework

For the N-vector $y_t = (y_{1t}, \ldots, y_{Nt})'$ of log realized volatilities we estimate a covariance-stationary VAR(p), $y_t = \sum_{i=1}^{p}\Phi_i y_{t-i} + \varepsilon_t$, $\varepsilon_t \sim (0,\Sigma)$, with moving-average form $y_t = \sum_{h=0}^{\infty} A_h\varepsilon_{t-h}$. Connectedness is read off the generalized forecast-error variance decomposition (Koop et al., 1996; Pesaran and Shin, 1998), which is invariant to the variable ordering:

$$\theta_{ij}(H) = \frac{\sigma_{jj}^{-1}\sum_{h=0}^{H-1}(e_i' A_h \Sigma e_j)^2}{\sum_{h=0}^{H-1}(e_i' A_h \Sigma A_h' e_i)}, \qquad \tilde{\theta}_{ij}(H) = \frac{\theta_{ij}(H)}{\sum_{k=1}^{N}\theta_{ik}(H)}, \tag{2}$$

so that $\sum_j \tilde{\theta}_{ij}(H) = 1$. The directional measures (in %) are

$$C_{i\leftarrow\bullet} = \sum_{j\neq i}\tilde{\theta}_{ij}, \qquad C_{\bullet\leftarrow j} = \sum_{i\neq j}\tilde{\theta}_{ij}, \qquad C_j^{NET} = C_{\bullet\leftarrow j} - C_{j\leftarrow\bullet}, \tag{3}$$

denoting variance received from others, transmitted to others, and the net of the two; the system-wide Total Connectedness Index is $TCI = \frac{1}{N}\sum_i C_{i\leftarrow\bullet}$. A bank with $C_j^{NET} > 0$ is a net transmitter of risk; with $C_j^{NET} < 0$, a net receiver.

2.4. Time-varying connectedness

To let the network evolve we estimate the TVP-VAR of Antonakakis et al. (2020), in which coefficients and the covariance matrix follow random walks, $y_t = \Phi_t y_{t-1} + \varepsilon_t$, $\varepsilon_t \sim (0,\Sigma_t)$, $\mathrm{vec}(\Phi_t) = \mathrm{vec}(\Phi_{t-1}) + \nu_t$, estimated by a Kalman filter with forgetting factors $\kappa_1 = \kappa_2 = 0.99$. This yields TCI and net measures at every date without an arbitrary rolling window.

2.5. High-dimensional VAR

With N banks the VAR has $N^2 p$ coefficients, so estimates can be noisy. Following Gabauer et al. (2020) and Demirer et al. (2018) we re-estimate with an Elastic-Net penalty, minimising $\sum_t \|y_t - \sum_i \Phi_i y_{t-i}\|_2^2 + \lambda\left[(1-\alpha)\|\Phi\|_2^2 + \alpha\|\Phi\|_1\right]$, with λ chosen by 10-fold cross-validation. The shrinkage yields a sparse, stable network and checks that the structure is not an over-parameterisation artefact.

3. Data

We take the banks on the Financial Stability Board G-SIB list and use the USD-traded instrument for each – the ordinary share for US, UK, Canadian and Swiss names and an ADR otherwise – with split- and dividend-adjusted daily OHLC from Yahoo! Finance. Because a range estimator requires reliable intraday highs and lows, instrument liquidity matters; the sample keeps the 19 G-SIBs with an exchange-listed instrument and clean OHLC (Table 1). Thinly traded OTC ADRs (BNP Paribas, Crédit Agricole, Société Générale, UniCredit, Standard Chartered and the mainland-China banks) are excluded on this transparent, *ex ante* liquidity criterion, replacing the original informal "currency" exclusion. Credit Suisse was absorbed by UBS in 2023 and delisted; it cannot enter an updated panel and is treated narratively – a G-SIB that exited through the very channel we measure. Groupe BPCE has no USD instrument. The resulting network is best read as that of the USD-liquid G-SIB core; extending it to the mainland-China and thinly traded continental-European names would require local-market data of uneven quality, a trade-off we revisit, for the banks that do have a liquid home listing, in Section 5.2.

The sample runs from 3 January 2007 to 12 June 2026, giving 4,887 daily and roughly 1,010 weekly log-volatility observations.[1] Series are merged on the date index, banks with less than 60% coverage are dropped, and the panel is restricted to common dates; at the weekly frequency the no-trade days that motivate ad-hoc zero filters are absorbed by aggregation.

Table 1. The G-SIB sample (19 USD-listed instruments).

Region	n	Banks (ticker)
United States	8	JPMorgan (JPM), Bank of America (BAC), Citigroup (C), Wells Fargo (WFC), Goldman Sachs (GS), Morgan Stanley (MS), BNY Mellon (BK), State Street (STT)
Europe	6	HSBC, Barclays (BCS), Deutsche Bank (DB), UBS, Santander (SAN), ING
Japan	3	Mitsubishi UFJ (MUFG), Mizuho (MFG), Sumitomo Mitsui (SMFG)
Canada	2	Royal Bank of Canada (RY), Toronto-Dominion (TD)

4. Results

4.1. A highly connected, risk-sharing system

Table 2 reports the Total Connectedness Index across the three estimators and the two frequencies. Connectedness is uniformly high – about 91% at the weekly and 87% at the daily frequency – meaning that the large majority of a bank's volatility forecast-error variance originates in other banks. The level is stable across VAR, TVP-VAR and LASSO, so it is not an artefact of a particular estimator. The 4-point weekly–daily gap is the expected consequence of frequency: higher-frequency data carry more idiosyncratic noise.

Table 2. Total Connectedness Index (%), by model and frequency.

Model	Weekly	Daily
Generalized VAR	90.9	86.6
TVP-VAR	90.0	83.0
LASSO/Elastic-Net VAR	89.6	83.8

4.2. The hub structure: who transmits and who absorbs

The system-wide index conceals a sharp directional structure. As Figure 2 makes clear, reception is homogeneous – every bank imports about 80–92% of its variance, so the points cluster tightly on the horizontal axis – while transmission ranges from about 57% to 110% of own variance. The net position is therefore governed almost entirely by transmission. Table 3 ranks banks by net directional connectedness at both frequencies, and Figure 1 displays the weekly ranking.

[1] The exact weekly count is written by the script to `results/sample_info_weekly.csv`.

Figure 1. Net directional connectedness C^{NET} (%) by bank, weekly generalized VAR, coloured by region. Positive bars are net transmitters of volatility.

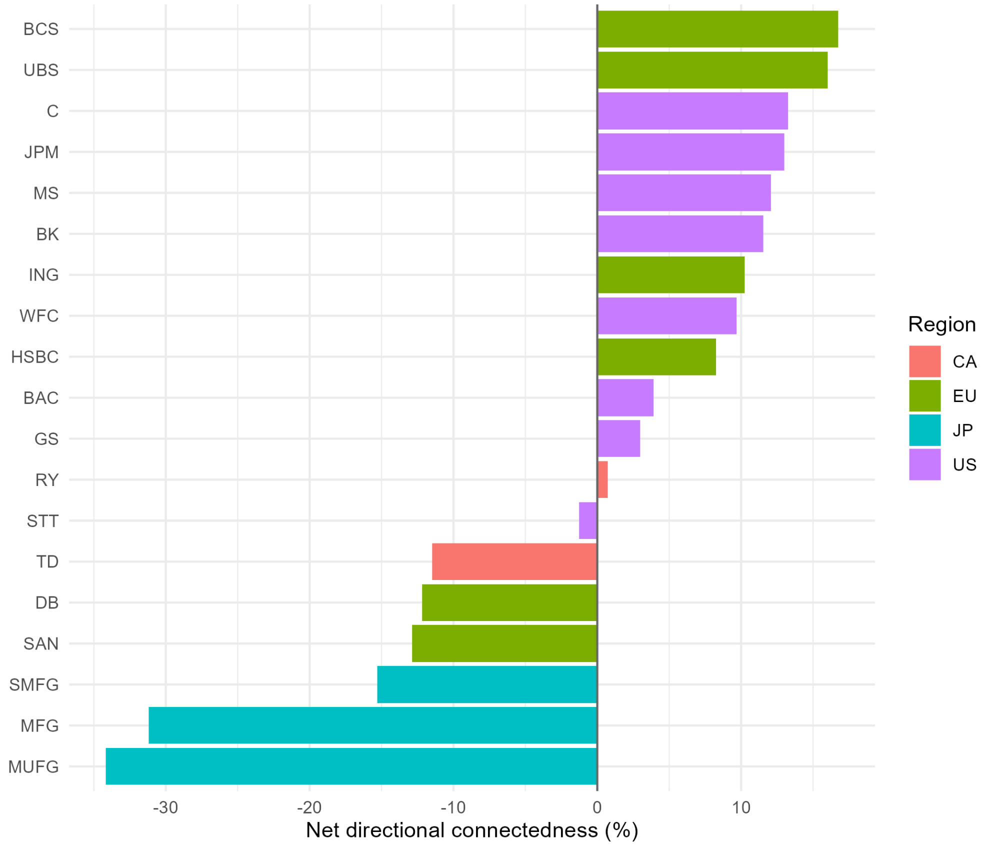

Figure 2. Directional connectedness to versus from others (%), weekly generalized VAR. Banks above the 45-degree line are net transmitters; the tight vertical clustering shows that reception is far more homogeneous than transmission.

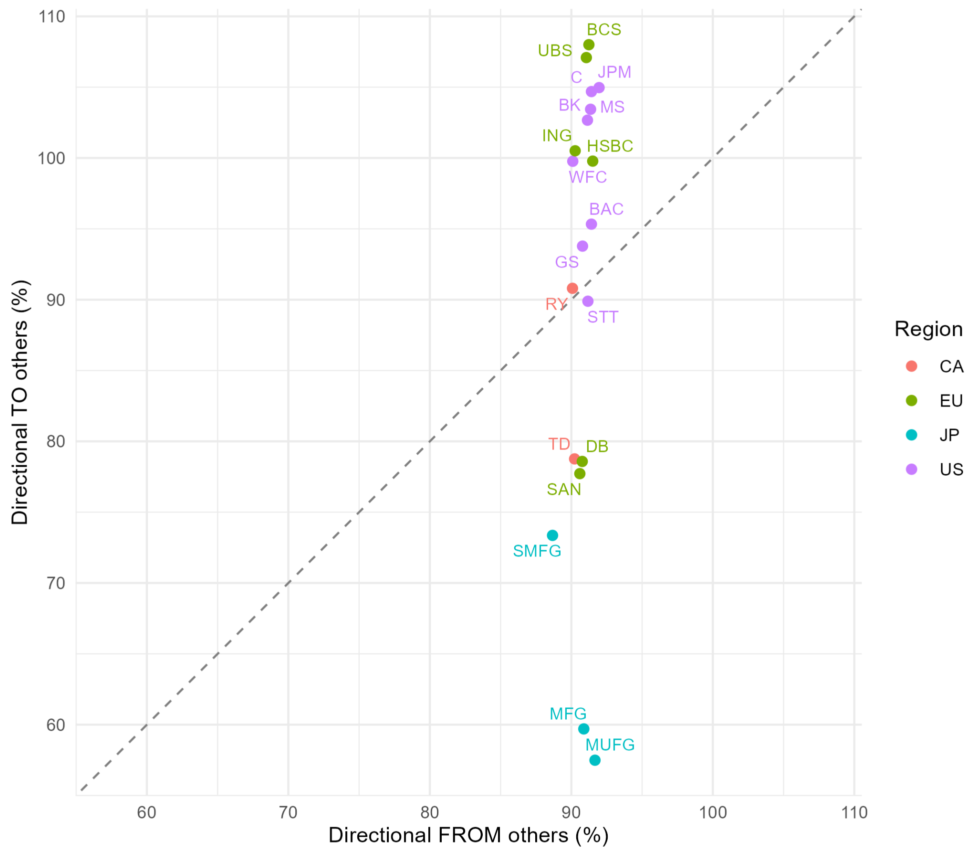

Table 3. Net directional connectedness C^{NET} (%), generalized VAR. Positive: net transmitter; negative: net receiver. Sorted by the weekly value.

Bank	Region	Weekly	Daily
Barclays	EU	16.76	17.30
UBS	EU	16.03	22.30
Citigroup	US	13.26	14.56
JPMorgan	US	13.00	19.66
Morgan Stanley	US	12.08	14.60
BNY Mellon	US	11.53	2.52
ING	EU	10.25	16.27
Wells Fargo	US	9.67	9.29
HSBC	EU	8.27	13.63
Bank of America	US	3.91	8.02
Goldman Sachs	US	2.99	−3.80
Royal Bank Canada	CA	0.73	−6.41
State Street	US	−1.28	−5.49
Toronto-Dominion	CA	−11.47	−16.26
Deutsche Bank	EU	−12.19	−8.11
Santander	EU	−12.88	−16.83
Sumitomo Mitsui	JP	−15.31	−25.19
Mizuho	JP	−31.18	−32.52
Mitsubishi UFJ	JP	−34.17	−23.55

Three findings stand out. First, the net transmitters form an Anglo-American/Swiss/Dutch cluster: UBS, JPMorgan, Barclays, Citigroup, Morgan Stanley, ING and HSBC are positive at both frequencies, joined by Bank of America and Wells Fargo. Second, the net receivers are unambiguously the three Japanese G-SIBs (Mizuho, Mitsubishi UFJ, Sumitomo Mitsui), together with the continental-European universal banks Deutsche Bank and Santander and the Canadian Toronto-Dominion. Third, the ranking generalises rather than overturns the earlier literature: the finding that a Japanese bank is "the most influenced" becomes the sharper statement that all Japanese G-SIBs are the system's risk absorbers, while single-name US dominance gives way to a cluster of Anglo-European hubs.

4.3. Economic interpretation

What accounts for this structure? The net transmitters share a business model built on global wholesale funding and large trading and derivatives books: UBS, JPMorgan, Barclays, Citigroup, Morgan Stanley and HSBC intermediate cross-border capital-market flows, so a shock to their balance-sheet capacity propagates outward through funding and market-making channels – the funding-liquidity mechanism emphasised by Billio et al. (2012). The Japanese megabanks sit at the opposite pole. Mitsubishi UFJ, Mizuho and Sumitomo Mitsui are funded predominantly by a large, stable domestic deposit base and lend heavily into the domestic economy; two decades of near-zero policy rates and yield-curve control further decoupled their marginal funding costs from the global cycle. They import global volatility – they are large, internationally exposed institutions – but transmit little back, because their funding and earnings are anchored at home; hence the large negative net positions that reflect low transmission rather than isolation. Deutsche Bank and Santander are intermediate: continental-European universal banks whose net-receiver position reflects, respectively, a long balance-sheet repair and a retail-funded, geographically concentrated franchise. We caution that the European names' position within this transmitter cluster is sensitive to currency denomination (Section 5.2); the funding-model reading is on firmest ground for the US transmitters and the Japanese absorbers. It is consistent with the exposure-based network evidence in Demirer et al. (2018), and it gives the purely statistical ranking an economic interpretation that the regulatory, indicator-based G-SIB score does not capture directly.

4.4. Model comparison

The full generalized-VAR connectedness matrices for the weekly and daily samples, together with the TVP-VAR and LASSO/Elastic-Net matrices, deliver the same ordering. The set of net transmitters and receivers, and the magnitude of the TCI, are stable across the three estimators; only Goldman Sachs changes sign between frequencies, and it is close to zero in both. The agreement across models is evidence that the network reflects the data, not the specification.

4.5. Frequency robustness: not a time-zone artefact

Table 4 aggregates net connectedness by region and Figure 3 plots each bank's weekly net position against its daily one. The directional structure is essentially identical across frequencies: the points line up along the 45-degree line, Japan is a strong absorber ($C^{NET} \approx -27\%$) at both, and the transmitter/receiver membership is unchanged. A non-synchronicity artefact would inflate the net transmission of the latest-closing market (the US) and depress that of the earliest-closing market (Japan) at the daily frequency; we observe neither. If anything, moving to the synchronization-robust weekly frequency raises US net transmission relative to Europe (from parity to a 3.7-point gap), the opposite of the artefact prediction. We conclude that volatility connectedness among G-SIBs is not materially contaminated by non-synchronous trading – plausibly because the daily range spans the whole local session and volatility shocks cluster over several days, muting the one-day lead–lag that contaminates returns.

Table 4. Mean net directional connectedness by region (%), generalized VAR.

Region	Weekly	Daily
United States	+8.1	+7.4
Europe	+4.4	+7.4
Canada	−5.4	−11.3
Japan	−26.9	−27.1

Figure 3. Net directional connectedness at the weekly versus the daily frequency (%), generalized VAR. Proximity to the 45-degree line indicates that the transmitter/receiver structure does not depend on the sampling frequency.

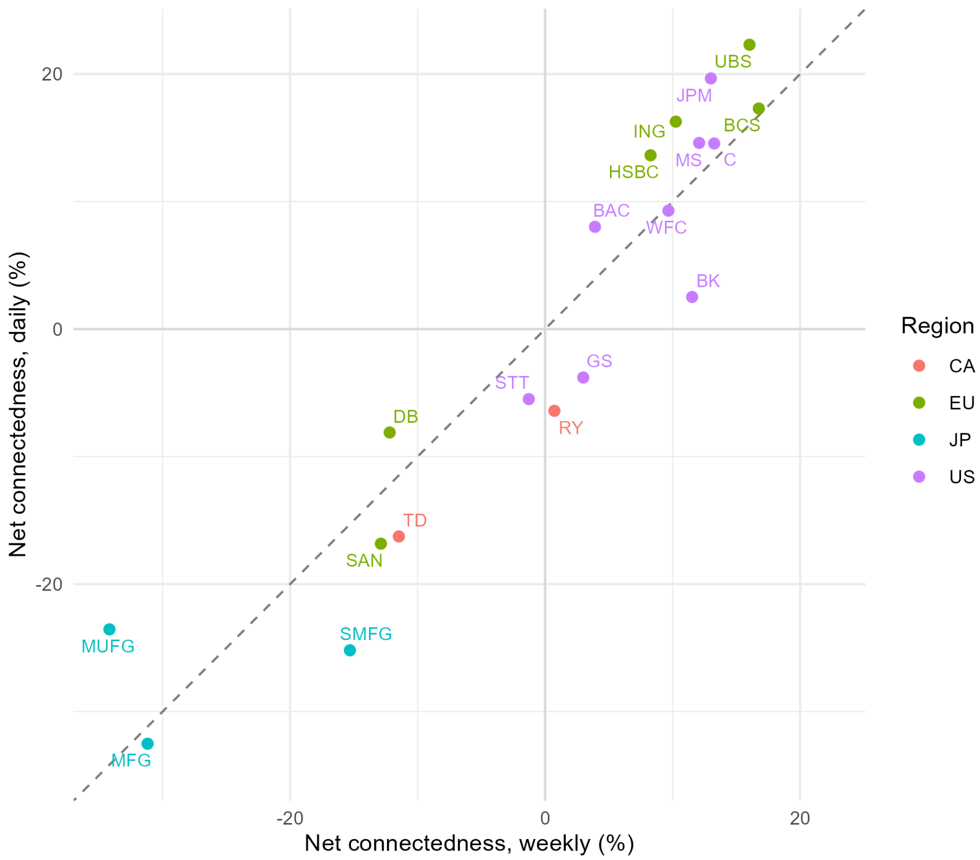

4.6. Dynamic connectedness

Figure 4 plots the TVP-VAR total connectedness index with markers at Lehman (September 2008), the COVID-19 shock (March 2020) and the SVB/Credit Suisse episode (March 2023). Connectedness is high throughout and tightens in each systemic episode; the rise around March 2023 – an event the model never used for estimation – is an out-of-sample validation: the network closes ranks precisely when contagion is observed, supporting its use as a monitoring tool.

Figure 4. Dynamic (TVP-VAR) total volatility connectedness among G-SIBs, weekly. Dashed lines mark Lehman, COVID-19 and the March-2023 SVB/Credit Suisse stress.

5. Robustness

5.1. Lag and horizon

Table 5 reports the average TCI over VAR lags p and FEVD horizons H. The index is stable within each frequency – 89–91% (weekly) and 79–87% (daily) – and the transmitter/receiver ranking is unchanged across the grid. Together with the agreement across estimators (Section 4) and across frequencies (Section 4.5), this indicates that the network's central features are not sensitive to the modelling choices.

Table 5. Average TCI (%) over VAR lag p and horizon H.

		$H = 5$	$H = 10$	$H = 20$
Weekly	$p = 1$	90.2	90.9	91.0
	$p = 2$	89.5	90.3	90.7
	$p = 3$	89.3	90.1	90.6
	$p = 4$	89.1	89.9	90.4
Daily	$p = 1$	85.2	86.6	86.9
	$p = 2$	82.3	85.0	86.3
	$p = 3$	80.5	83.6	85.6
	$p = 4$	79.2	82.4	84.8

5.2. Currency denomination

For the non-US banks the USD-denominated ADR embeds a currency leg, raising the concern that common dollar/FX volatility – rather than bank-specific risk – contributes to the measured spillovers, especially within the euro-area cluster. We address this with a same-sample orthogonalisation, the cleanest control. Keeping the USD ADR panel – and hence the sample, calendar and microstructure – fixed, we regress each non-US bank's weekly log realized volatility on the realized volatility of its home-currency/USD exchange rate (GBP/USD for HSBC and Barclays, EUR/USD for Deutsche Bank, Santander and ING, CHF/USD for UBS, JPY/USD for the Japanese G-SIBs, CAD/USD for the Canadian banks) and re-estimate the network on the residuals, which are purged of the currency component. Figure 5 plots the orthogonalised net positions against the raw ADR positions. On the full weekly sample the two rankings correlate at Spearman $\rho = 0.83$, and the reshuffling is systematic: the Japanese G-SIBs remain net absorbers and the US banks remain – indeed become more pronounced – net transmitters, but the cross-listed European banks move from net transmitters to net receivers (HSBC, Barclays, UBS, ING). Two caveats temper the magnitude. First, only the non-US banks are orthogonalised – the US ADRs carry no currency leg – which mechanically raises the US banks' relative share. Second, for the euro bloc the exchange-rate and bank-risk volatilities are partly co-determined (both spiked in the 2010–12 sovereign crisis), so the orthogonalisation strips some genuine common bank risk and is best read as an upper bound on the currency contribution. Even so, the direction is unambiguous and agrees with the home-listing check below: the European net-transmitter result is substantially a dollar-overlay effect.

A complementary check re-estimates the network on the banks' home-market listings (HSBA.L, BARC.L, DBK.DE, UBSG.SW, SAN.MC, INGA.AS, 8306.T, 8411.T, 8316.T, RY.TO, TD.TO), which also strips the FX leg but at the cost of changing the common sample. It points the same way (Figure 6): the US and Japanese poles are unchanged, while the cross-listed European banks (Barclays, UBS, ING and especially HSBC) move toward net-receiver positions, with an overall rank correlation of only $\rho = 0.55$. Because that re-estimation also alters the sample – the local exchanges have different calendars and histories, which is why even the unchanged US instruments shift – it cannot isolate FX by itself; but with the same-sample orthogonalisation pointing the same way, the two controls together bound the European banks' true standing between their raw-ADR (net-transmitter) and FX-stripped (net-receiver) positions. We therefore treat the US-transmitter and Japanese-absorber backbone as the denomination-invariant finding and flag the European names as denomination-sensitive.

Figure 5. Net directional connectedness before and after FX orthogonalisation (%).

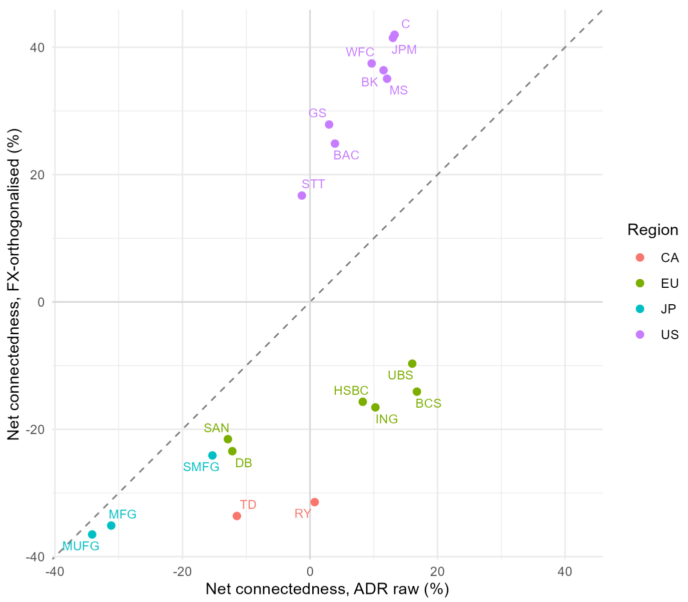

Notes: Same weekly sample, generalized VAR. Each non-US bank's volatility is purged of its home-currency/USD volatility. The European (green) and Canadian (red) banks fall below the 45-degree line – their net transmission shrinks or reverses once the currency component is removed – while the US banks (purple) rise above it and the Japanese absorbers (cyan) are little changed. Distance from the line measures sensitivity to the currency overlay; the asymmetric residualisation (US ADRs carry no FX leg) inflates the US shift, so the figure bounds rather than point-identifies the effect.

Figure 6. Net directional connectedness: ADR prices versus home-market listings (%).

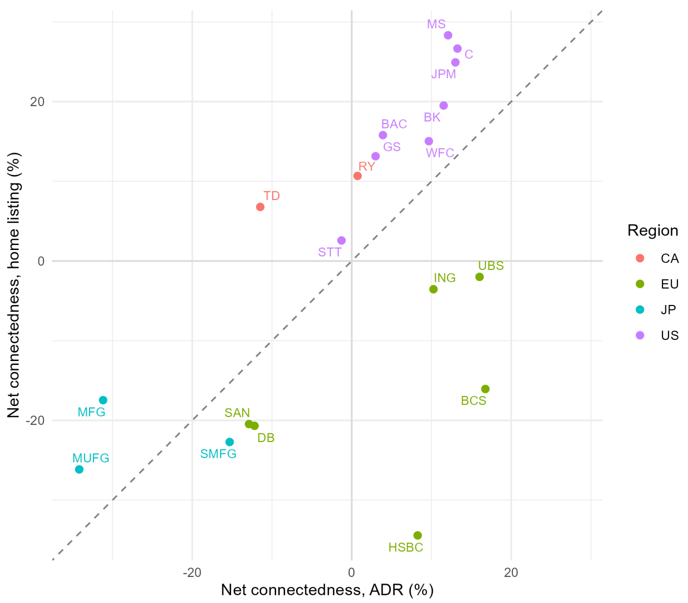

Notes: Complementary check, weekly generalized VAR. The US (purple) and Japanese (cyan) banks stay on the same side of the line; the European banks (green) move from transmitters to receivers. This re-estimation also changes the common sample, which is why the same-sample orthogonalisation of Figure 5 is the cleaner test.

5.3. Limitations

The network is built from market prices, not balance-sheet exposures, so it captures expected transmission as priced by investors rather than realised interbank claims; weekly aggregation trades temporal resolution for synchronisation; and the USD-liquidity sampling criterion leaves out the mainland-China G-SIBs and the thinly traded continental-European names, so the results describe the liquid core of the global bank network rather than its entirety.

6. Conclusion

The world's systemically important banks form a tightly connected volatility network in which roughly 87–91% of a typical bank's risk is imported from its peers. Behind this high average lies a stable hub structure: a wholesale-funded Anglo-American, Swiss and Dutch core (UBS, JPMorgan, Barclays, ING, Citigroup, Morgan Stanley, HSBC) transmits risk, while the deposit-funded Japanese G-SIBs, together with Deutsche Bank, Santander and Toronto-Dominion, absorb it. The structure holds across three estimators, across VAR lags and horizons, and across the weekly and daily frequencies, ruling out a non-synchronous-trading artefact. A currency check shows the US-transmitter and Japanese-absorber poles survive removal of the dollar overlay, while the cross-listed European banks' net positions are denomination-sensitive – a caveat we flag for the euro-area cluster. The time-varying network tightens in every major stress episode, including March 2023, when Credit Suisse left the sample through the very channel measured here.

For regulators, the net-transmitter ranking is a market-based, contagion-oriented complement to the FSB's exposure-based bucketing, and a candidate input for surcharge calibration and real-time monitoring. For investors and risk managers, the hub banks are the natural focus for hedging and stress design, since shocks to them diffuse furthest. Methodologically, the paper shows that measuring bank contagion on volatility – and probing robustness to sampling frequency and to currency denomination – yields a network whose US and Japanese poles are economically interpretable and free of the non-synchronicity artefact that distorts daily, return-based cross-border panels, while flagging where the currency overlay still matters.

References

Adams, Z., Füss, R., and Gropp, R. (2014). Spillover effects among financial institutions: A state-dependent sensitivity value-at-risk approach. *Journal of Financial and Quantitative Analysis*, 49(3):575–598.

Antonakakis, N., Chatziantoniou, I., and Gabauer, D. (2020). Refined measures of dynamic connectedness based on time-varying parameter vector autoregressions. *Journal of Risk and Financial Management*, 13(4):84.

Bank for International Settlements (2018). Global systemically important banks: Revised assessment methodology and the higher loss absorbency requirement. Bank for International Settlements.

Billio, M., Getmansky, M., Lo, A. W., and Pelizzon, L. (2012). Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, 104(3):535–559.

Demirer, M., Diebold, F. X., Liu, L., and Yilmaz, K. (2018). Estimating global bank network connectedness. *Journal of Applied Econometrics*, 33(1):1–15.

Diebold, F. X. and Yilmaz, K. (2009). Measuring financial asset return and volatility spillovers, with application to global equity markets. *Economic Journal*, 119(534):158–171.

Diebold, F. X. and Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1):57–66.

Diebold, F. X. and Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1):119–134.

Forbes, K. J. and Rigobon, R. (2002). No contagion, only interdependence: Measuring stock market comovements. *Journal of Finance*, 57(5):2223–2261.

Fratzscher, M., Lo Duca, M., and Straub, R. (2013). On the international spillovers of us quantitative easing. ECB Working Paper 1557, European Central Bank.

Gabauer, D., Gupta, R., Marfatia, H., and Miller, S. (2020). Estimating us housing price network connectedness: Evidence from dynamic elastic net, lasso, and ridge vector autoregressive models. Working Paper 2020-65, University of Pretoria.

Garman, M. B. and Klass, M. J. (1980). On the estimation of security price volatilities from historical data. *Journal of Business*, 53(1):67–78.

Koop, G., Pesaran, M. H., and Potter, S. M. (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics*, 74(1):119–147.

Pesaran, H. H. and Shin, Y. (1998). Generalized impulse response analysis in linear multivariate models. *Economics Letters*, 58(1):17–29.