VOLATILITY REGIME SHIFTS AND STRESS TRANSMISSION IN EUROPEAN CAPITAL MARKETS: A DYNAMIC CONNECTEDNESS APPROACH

Lucian-Liviu Albu[1], Iulia Lupu[2], Radu Lupu[1,3]*

[1)] Romanian Academy, Institute for Economic Forecasting, Bucharest, Romania.

[2)] Romanian Academy, "Victor Slăvescu" Centre for Financial and Monetary Research, Bucharest, Romania.

[3)] Bucharest University of Economic Studies, Bucharest, Romania.

Abstract

This paper investigates whether volatility regime shifts in European equity markets are associated with changes in cross-market stress transmission. Using daily returns for seven major European equity indices (1998-2026), plus two extension panels adding Sweden, Denmark and Poland, we estimate EGARCH (1,1)-Student-t volatility models, classify markets into low-, medium- and high-volatility regimes via a hidden Markov model, and measure dynamic connectedness with a rolling-window VAR-based variance-decomposition approach. High-volatility regimes cluster strongly across markets (the same-day coincidence rate of regime-entry dates is 24-33% across panels, against a 1% placebo benchmark), and total connectedness rises significantly before, during and for several weeks after system-wide regime shifts, even after controlling for connectedness persistence. Net-transmission concentration increases during high-volatility periods even as gross transmission becomes more diffuse; France, Germany and the Netherlands are persistent net transmitters of stress, while Switzerland, the United Kingdom and Spain are persistent net receivers. A European Market Stress Transmission Index (EMSTI), combining regime probability, connectedness, and network concentration, peaks at well-documented stress episodes: the 2015 Greek/China selloff and the 2022 Russia-Ukraine energy shock. The findings hold up in nine of ten robustness checks. The one substantive qualification: once global risk factors (VIX, US equity returns, EUR/USD and US interest rates) are controlled for, part of the apparent pre- and post-shift rise in connectedness reflects global risk sentiment rather than a purely European phenomenon; only the contemporaneous shift-window effect is unambiguously European. The results support regime-sensitive connectedness indicators as a complement to, not a replacement for, existing financial-stability monitoring tools.

* Corresponding author, **Radu Lupu** – radulupu.ase@gmail.com.

Keywords
volatility regimes; dynamic connectedness; stress transmission; EGARCH; hidden Markov model; European equity markets; financial stability.

JEL Classification
C32, C58, G01, G15

Introduction

European capital markets are highly integrated: national equity indices share listed multinational constituents, overlapping investor bases, common monetary and regulatory frameworks, and a long history of correlated responses to region-wide and global shocks such as the global financial crisis, the euro-area sovereign debt crisis, the 2016 Brexit referendum, the COVID-19 shock, and the 2022 Russia-Ukraine energy crisis, all of which produced simultaneous volatility spikes across the continent's major markets. This integration is precisely what makes cross-market stress transmission a first-order concern for market participants and prudential supervisors alike: a shock that originates in, or is amplified by, one market can propagate quickly to others through common ownership, correlated risk appetite, and direct financial linkages.

During periods of financial stress, two things typically happen at once. First, volatility rises, often sharply and asymmetrically, with negative shocks raising conditional volatility more than positive shocks of the same size (the well-documented leverage effect). Second, and less often studied alongside the first, the structure of cross-market transmission may change: markets that transmit little stress in calm periods can become dominant transmitters during turmoil, and the network of relationships between markets can become more or less concentrated. Traditional conditional-volatility models (ARCH/GARCH-family) capture the first phenomenon well but say nothing about the second, because they are estimated market by market. Dynamic connectedness models, built on forecast-error variance decompositions of vector autoregressions, capture the second phenomenon directly – total, directional and net stress transmission across markets, but are typically estimated without explicit reference to each market's own volatility regime, leaving the link between regime shifts and network reconfiguration implicit at best.

This paper explicitly combines the two literatures. We first estimate each market's own volatility regime using an EGARCH-based hidden Markov model, then ask whether transitions into system-wide high-volatility states are associated with measurable changes in the connectedness network: does total connectedness rise around regime shifts; does the concentration of net stress transmission change; do some markets consistently act as transmitters ahead of others; and does a high-volatility regime in one market raise the probability that other markets enter a high-volatility regime themselves. Building on these tests, we propose a European Market Stress Transmission Index (EMSTI) that combines regime probability, total connectedness and network concentration into a single, interpretable supervisory indicator.

The paper's contribution is therefore a regime-sensitive view of European equity market stress transmission: rather than treating volatility regimes and connectedness as parallel

but separate descriptions of market stress, we test explicitly whether they move together, in which direction, and with what lead-lag structure, and we check the robustness of these findings along ten dimensions recommended in the connectedness literature, including, importantly, whether the results survive controlling for global (largely US-driven) risk factors.

Consistent with established practice in this literature (Forbes and Rigobon, 2002), we are careful throughout to distinguish stress transmission and connectedness (measurable statistical relationships) from the stronger and harder-to-establish claim of contagion in the strict sense of excess co-movement beyond normal interdependence; we use the language of "stress transmission" and "connectedness" throughout, and reserve "contagion-like behaviour" for instances where the evidence specifically supports an abnormal increase in transmission during stress.

The remainder of the paper is organised as follows. Section 1 reviews the relevant literature on volatility modelling, regime switching, connectedness and network measures of systemic risk. Section 2 describes the data. Section 3 sets out the methodology in four layers: volatility modelling, regime detection, dynamic connectedness, and the joint regime-connectedness tests. Section 4 presents the empirical results. Section 5 reports the robustness analysis. Section 6 discusses the financial-stability implications of the findings, and the paper concludes with a summary, limitations and directions for future research.

1. Review of the scientific literature

1.1 Volatility modelling and asymmetry

The modern literature on conditional volatility begins with the autoregressive conditional heteroskedasticity (ARCH) model of Engle (1982) and its generalisation, GARCH, due to Bollerslev (1986), both motivated by the empirical observation that financial-return volatility clusters over time. Because GARCH imposes symmetric responses to positive and negative shocks and requires non-negative parameters, subsequent work developed asymmetric extensions. Nelson (1991) proposed the exponential GARCH (EGARCH) model, which specifies the log of conditional variance directly, avoiding non-negativity constraints, and allows a separate parameter to capture the leverage effect, whereby negative shocks raise future volatility more than positive shocks of equal magnitude. Glosten, Jagannathan and Runkle (1993) proposed an alternative asymmetric specification (GJR-GARCH) based on a threshold indicator for negative shocks. Both specifications are now standard tools in the equity-volatility literature and are used jointly in this paper, with EGARCH as the baseline and GJR-GARCH as a robustness check.

1.2 Regime shifts in financial markets

A separate body of literature models financial time series as switching between distinct latent states or "regimes" rather than following a single stationary process throughout. Hamilton (1989) introduced the Markov-switching framework that underlies most subsequent work, including hidden Markov models (HMMs) applied to volatility or returns, in which the latent state evolves according to a first-order Markov chain and each state has its own conditional distribution. Regime-switching tools of this kind extend naturally beyond volatility modelling to bubble detection and crash-risk assessment:

Călin and Lupu (2026a) combine regime-switching with log-periodic power-law signatures to show that precious-metals markets in an identified speculative regime face markedly higher odds of a subsequent crash, illustrating how a market's current regime can itself serve as a risk signal rather than a purely descriptive label. An alternative, non-probabilistic approach detects discrete structural breaks directly in a time series; Bai and Perron (2003) provide the classical framework for multiple structural-change estimation, and change-point algorithms such as PELT (Killick, Fearnhead and Eckley, 2012) offer a computationally efficient, non-parametric alternative to a fitted Markov-switching model. This paper uses a Gaussian hidden Markov model on log conditional volatility as its baseline regime-detection method and PELT change-point detection as a robustness check, following the recommendation that regime probabilities, and not only hard classifications, be used wherever possible.

1.3 Connectedness, spillovers and contagion

The dynamic-connectedness literature originates with Diebold and Yilmaz (2009, 2012), who use forecast-error variance decompositions from a vector autoregression, in the order-invariant, generalised form of Pesaran and Shin (1998), to construct total, directional and net measures of return or volatility spillovers across a system of markets. The original framework uses a rolling estimation window; subsequent extensions, notably the time-varying-parameter VAR (TVP-VAR) connectedness approach of Antonakakis, Chatziantoniou and Gabauer (2020), estimate the underlying VAR with a Kalman filter, producing a daily connectedness series without a window-length choice. An alternative route to dynamic co-movement is Engle's (2002) dynamic conditional correlation (DCC) GARCH model, which estimates time-varying pairwise correlations directly rather than through a variance-decomposition network. This paper uses the rolling-window VAR/generalised-FEVD approach as its baseline and implements both TVP-VAR and DCC-GARCH as robustness checks, reporting whether the main findings are estimator-specific.

A closely related but conceptually distinct strand of the literature studies financial contagion, excess co-movement across markets during a crisis, over and above what normal interdependence would predict. Forbes and Rigobon (2002) show that correlation-based contagion tests can be severely biased by the heteroskedasticity that naturally accompanies market turmoil, and argue for a stricter distinction between contagion and interdependence; Bekaert, Harvey and Ng (2005) develop a related framework for market integration and contagion testing. This paper follows their caution explicitly: it documents changes in measured connectedness around regime shifts, a form of interdependence, without claiming to have identified contagion in the strict, excess-co-movement sense, since doing so would require an explicit adjustment for the normal (non-crisis) level of interdependence that this design does not attempt.

1.4 Network indicators and systemic risk

Further literature translates pairwise or system-wide connectedness measures into indicators of systemic risk intended for financial-stability monitoring. Billio, Getmansky, Lo and Pelizzon (2012) use Granger-causality networks among banks, insurers, broker-dealers and hedge funds to show that measured interconnectedness rose ahead of the

2007-2009 crisis. Adrian and Brunnermeier (2016) propose CoVaR, a conditional value-at-risk measure of one institution's contribution to system-wide tail risk, and Brownlees and Engle (2017) propose SRISK, a capital-shortfall measure combining leverage and expected equity-return shortfall in a systemic crisis. These contributions share a common motivation with the present paper, translating a statistical measure of interconnectedness into a practically interpretable monitoring tool, though they are built at the institution level rather than the national-index level used here.

1.5 Positioning of this paper

This paper differs from standard spillover studies by linking dynamic connectedness explicitly to market-specific volatility regimes, and by testing whether regime transitions are systematically associated with network reconfiguration – rather than simply documenting that both connectedness and volatility are high during known crisis episodes. It further differs from most applications of the Diebold-Yilmaz framework to European markets by (i) combining a regime layer (HMM, cross-checked against PELT change-point detection) with the connectedness layer; (ii) building a single supervisory index (EMSTI) that combines the two; and (iii) subjecting the central findings to an unusually extensive set of ten robustness checks, including an explicit test of whether the regime-connectedness relationship survives controlling for global (US-centred) risk factors – a check that, as reported below, materially qualifies one part of the paper's main result.

2. Data

The empirical analysis uses daily adjusted-close prices for major European equity indices, obtained from Yahoo Finance. The core empirical panel ("core7") comprises seven markets – Germany (DAX), France (CAC 40), Italy (FTSE MIB), Spain (IBEX 35), the Netherlands (AEX), the United Kingdom (FTSE 100) and Switzerland (SMI), chosen because they have the longest continuous history among the sampled markets and jointly span the global financial crisis, the euro-area sovereign debt crisis, the Brexit referendum, the COVID-19 shock and the 2022 Russia-Ukraine energy shock. Two shorter extension panels add Sweden (OMX Stockholm 30) and Denmark (OMX Copenhagen 25) from 2016-12-19, and Poland (WIG20, proxied by a total-return ETF because the native index has no usable pre-2019 history on Yahoo Finance) from 2019-01-07, allowing the main results to be cross-checked against a broader and more recent market set. Consistent with the paper's focus on broader European capital-market dynamics, we exclude the Romanian market from the empirical sample; for direct multi-model evidence on speculative-bubble episodes in the Bucharest stock market specifically, see Călin and Lupu (2026b). Table no. 1 summarises the sample.

Table no. 1. Market sample and data description

Market/index	Country	Currency	Sample start	Sample end	Obs.	Source
DAX	Germany	EUR	1987-12-30	2026-09-28	9,798	Yahoo Finance
CAC 40	France	EUR	1990-03-01	2026-09-28	9,290	Yahoo Finance
FTSE MIB	Italy	EUR	1997-12-31	2026-09-28	7,341	Yahoo Finance
IBEX 35	Spain	EUR	1993-07-12	2026-09-28	8,417	Yahoo Finance
AEX	Netherlands	EUR	1992-10-12	2026-09-28	8,666	Yahoo Finance
FTSE 100	United Kingdom	GBP	1984-01-03	2026-09-28	10,796	Yahoo Finance
SMI	Switzerland	CHF	1990-11-09	2026-09-28	9,010	Yahoo Finance
OMX Stockholm 30	Sweden	SEK	2008-11-20	2026-09-28	4,477	Yahoo Finance
OMX Copenhagen 25	Denmark	DKK	2016-12-19	2026-09-28	2,376	Yahoo Finance
WIG20 (ETF proxy)	Poland	PLN	2019-01-07	2026-09-28	1,932	Yahoo Finance (ETFBW20TR.WA)

Source: Authors' calculations from Yahoo Finance daily adjusted-close data.

Daily log returns are computed as $r_{i,t} = 100 \times [\log(P_{i,t}) - \log(P_{i,t-1})]$, where $P_{i,t}$ is the adjusted closing price of market i on day t. Because the ten markets have different listing histories and national holiday calendars, three calendar-aligned panels are constructed by trading-day intersection: core7 (1998 onward, seven markets), a nine-market panel adding Sweden and Denmark (2016 onward), and a ten-market panel adding Poland via the ETF proxy (2019 onward). No automatic deletion or interpolation of extreme returns is performed; each series' ten largest absolute daily returns were cross-checked against known European and global stress episodes as a data-quality validation, and every case aligns with a documented event, for example, the region-wide crash of 12 March 2020 (COVID-19 and the oil-price war), the 6 October 2008 global equity rout, the 24 June 2016 Brexit-referendum reaction, and a distinctive Swiss franc shock on 15 January 2015 (the Swiss National Bank's removal of the EUR/CHF floor) that is specific to the SMI series.

Descriptive statistics (Section 4.1) confirm the standard stylised facts of daily equity returns (fat tails, negative skewness and strong volatility clustering) that motivate the conditional-volatility modelling in Section 3.1.

3. Research methodology

The empirical strategy proceeds in four layers: (i) market-specific volatility modelling, (ii) volatility regime detection, (iii) dynamic connectedness estimation, and (iv) joint regime-connectedness tests, culminating in a single supervisory index.

3.1 Return and volatility modelling

For each market i, the mean equation is $r_{i,t} = \mu_i + \varepsilon_{i,t}$, with $\varepsilon_{i,t} = \sigma_{i,t} z_{i,t}$ and $z_{i,t} \sim$ Student-t(0,1). The baseline conditional-variance specification is EGARCH(1,1):

$$\log(\sigma^2_{i,t}) = \omega_i + \beta_i \log(\sigma^2_{i,t-1}) + \alpha_i (|z_{i,t-1}| - E|z_{i,t-1}|) + \gamma_i z_{i,t-1}$$

where β captures volatility persistence, α captures the magnitude effect of shocks, and γ captures asymmetry: a negative γ means negative shocks raise conditional volatility more than positive shocks of the same size (the leverage effect). Each market's EGARCH(1,1)-Student-t model is estimated on its own full native return history (not the calendar-aligned panel, since univariate volatility estimation should use the maximum available sample). As a robustness check, we estimate a GJR-GARCH(1,1)-Student-t model in parallel for each market, selecting the model based on the Akaike and Bayesian information criteria (AIC, BIC).

3.2 Volatility regime detection

Each market's estimated log conditional volatility from the EGARCH model is used to classify days into a latent regime $s_{i,t} \in \{\text{Low, Medium, High}\}$ using a three-state Gaussian hidden Markov model (HMM), with states relabelled by their fitted mean so that state 2 is always "High". Because single-seed HMM estimation converges to a poorly separated local optimum for some markets – confirmed by a 30-seed comparison on one market, where 28 of 30 seeds recovered the same well-separated global optimum, and two converged to a degenerate, near-random-switching solution, every market's HMM is fit with 20 random restarts, keeping the highest-log-likelihood solution. The key output is a continuous regime-stress measure, $\text{HighVolProb}_{i,t} = \Pr(s_{i,t} = \text{High})$, together with a binary indicator $\text{HighVol}_{i,t} = 1\{\text{HighVolProb}_{i,t} \geq 0.50\}$, with thresholds of 0.40 and 0.60 used as a robustness check. As an alternative, model-free regime-detection method, PELT change-point detection (Killick, Fearnhead and Eckley, 2012) is applied directly to each market's log conditional volatility series, with the resulting segments tercile-classified into Low/Medium/High, and used to re-run the paper's central synchronisation test as a robustness check (Section 5, item 2). System-wide regime stress is measured by AvgHighVolShare_t, the cross-sectional average of $\text{HighVol}_{i,t}$ across markets on day t.

3.3 Dynamic connectedness estimation

Dynamic connectedness is estimated using the generalised forecast-error variance decomposition (GFEVD) of Pesaran and Shin (1998), applied to a rolling-window VAR(2) fitted to daily returns, following Diebold and Yilmaz (2012). The baseline specification uses a 250-trading-day rolling window, a one-day step, a two-lag VAR (the standard, parsimonious choice in the connectedness literature, used in preference to the much higher lag order favoured by AIC alone, which would leave few degrees of freedom per rolling window), and a ten-day-ahead forecast horizon. Each day, the GFEVD produces a Total Connectedness Index (TCI, the share of forecast-error variance

explained by cross-market spillovers) and market-level Received, Transmitted, and Net (= Transmitted - Received) connectedness measures. As robustness checks, the same GFEVD machinery is re-estimated using (i) a Kalman-filter time-varying-parameter VAR (TVP-VAR, following Antonakakis, Chatziantoniou and Gabauer, 2020, with forgetting factors $\kappa_1 = 0.99$ for the state equation and $\kappa_2 = 0.96$ for the observation-error covariance), which produces a daily connectedness series without any rolling-window choice, and (ii) a two-step DCC-GARCH(1,1)-Student-t model (Engle, 2002), which measures average pairwise dynamic conditional correlation as an alternative, non-variance-decomposition connectedness concept.

3.4 Regime-connectedness interaction tests

The paper's core empirical layer tests five hypotheses linking the regime and connectedness layers, H1–H5 below. H5, a global-risk-control check on the H2 result, is reported as part of the robustness analysis (Section 5) rather than as a separate empirical subsection:

- H1: volatility regimes are synchronised across markets: high-volatility regime entries cluster across markets more than a within-market-preserving placebo benchmark would predict, tested with coincidence windows of 0, ±3, ±5 and ±10 trading days against 1,000 circular-shift placebo draws per window.
- H2: connectedness rises around regime shifts: an event-study regression of TCI_t on Pre-shift, Shift-window and Post-shift dummies around system-wide regime-shift dates (AvgHighVolShare_t crossing 0.50 from below, after at least 10 quiet days, with at least 20 trading days between events), controlling for TCI persistence (TCI_t−1) and using Newey-West (HAC) standard errors.
- H3: some markets act as regime leaders: a panel logit of next-day high-volatility status on a market's own lagged regime state, the average lagged regime state of all other markets, and its own lagged net transmission (market fixed effects, market-clustered standard errors), together with pairwise Granger-causality tests on regime probabilities.
- H4: the network becomes more centralised during stress: total connectedness and three candidate network-concentration measures – a Herfindahl index of gross transmission shares, a Herfindahl index of |net transmission| shares, and the share of gross transmission from the two largest transmitters – are compared across system-wide Low/Medium/High volatility terciles.
- H5: the regime-connectedness relationship found in H2 is not solely a reflection of global risk sentiment: the pre-shift, shift-window and post-shift connectedness effects are re-tested with the VIX, US equity returns, the EUR/USD rate and US interest-rate changes added as controls to the H2 event-study regression (Section 5).

3.5 European Market Stress Transmission Index

A single supervisory indicator is constructed as:

$$EMSTI_t = AvgHighVolProb_t \times NormalizedTCI_t \times NetworkConcentration_t$$

where AvgHighVolProb_t is the cross-sectional average high-volatility probability, NormalizedTCI_t is the Total Connectedness Index rescaled to the unit interval, and NetworkConcentration_t is the (min-max rescaled) Herfindahl index of gross

transmission shares. The gross-transmission Herfindahl index, rather than the net-transmission one, is deliberately chosen for this multiplicative construction because – as Section 4.5 shows – it is the concentration measure that does not itself mechanically rise with AvgHighVolProb_t, so EMSTI is not triple-counting the same underlying signal.

4. Results and discussions

4.1 Descriptive evidence

Table no. 2 reports descriptive statistics of daily log returns for the core7 panel. All seven markets display strong excess kurtosis (8.4–11.0), negative skewness (−0.18 to −0.55), and Jarque-Bera normality tests that reject at essentially any conventional significance level – the standard fat-tail, negative-skew signature of daily equity returns. Return autocorrelation at lag 1 is small throughout, but the ARCH-LM test rejects the null of no ARCH effects at $p \approx 0$ in every market, confirming strong volatility clustering and motivating the conditional-volatility modelling in Section 3.1. The two extension panels (not tabulated for space, available on request) display the same qualitative pattern over their shorter windows.

Table no. 2. Descriptive statistics of daily log returns, core7 panel, 1998–2026

Market	Obs.	Mean	Std. dev.	Skewness	Kurtosis	Min	Max	ARCH-LM p
DAX	7,039	0.025	1.453	-0.181	8.377	-13.05	10.80	0.000
CAC 40	7,039	0.014	1.400	-0.195	8.944	-13.10	10.59	0.000
FTSE MIB	7,039	0.010	1.520	-0.554	10.996	-18.55	10.87	0.000
IBEX 35	7,039	0.014	1.446	-0.342	10.345	-15.15	13.48	0.000
AEX	7,039	0.014	1.364	-0.228	9.795	-11.38	10.03	0.000
FTSE 100	7,039	0.010	1.149	-0.324	10.400	-11.51	9.38	0.000
SMI	7,039	0.011	1.144	-0.275	10.058	-10.13	10.79	0.000

Source: Authors' calculations.

4.2 Volatility model estimates

Table no. 3 reports EGARCH(1,1)-Student-t estimates for the core7 markets (each fit on its own full native history). Volatility persistence (β) is high and narrow across markets (0.969–0.982), consistent with the volatility clustering documented above. The asymmetry parameter γ is negative for every market, the classic leverage effect, and the same qualitative conclusion holds under the GJR-GARCH robustness specification, whose asymmetry parameter is positive for all ten sampled markets (its own sign convention). EGARCH is preferred over GJR-GARCH by both AIC and BIC in all ten markets, confirming the choice of EGARCH as the baseline model empirically rather than by convention alone.

Table no. 3. EGARCH(1,1)-Student-t estimation results, core7 markets

Market	ω	α	γ (asymmetry)	β (persist.)	ν (d.o.f.)	Log-lik.	AIC	BIC
DAX	0.0075	0.1417	-0.0928	0.9819	7.594	-15033.9	30079.9	30123.0
CAC 40	0.0060	0.1316	-0.1137	0.9813	8.161	-14157.7	28327.5	28370.3
FTSE MIB	0.0100	0.1569	-0.1055	0.9809	7.764	-11835.1	23682.3	23723.7
IBEX 35	0.0090	0.1589	-0.0877	0.9789	8.576	-13115.5	26242.9	26285.1
AEX	0.0028	0.1623	-0.1056	0.9823	8.956	-12286.7	24585.4	24627.8
FTSE 100	-0.0021	0.1575	-0.0849	0.9797	9.251	-13937.3	27886.5	27930.3
SMI	-0.0037	0.1694	-0.1161	0.9694	8.538	-11792.1	23596.3	23638.9

Source: Authors' calculations. Diagnostics (Ljung-Box, ARCH-LM on standardised residuals) are reported in the online replication material; two markets (IBEX 35, SMI) show borderline mean-equation residual autocorrelation that does not affect the (clean) squared-residual and ARCH-LM diagnostics.

4.3 Volatility regimes across European markets

Table no. 4 summarises the HMM regime characteristics. All seven core markets settle into a well-balanced three-regime structure (roughly 32–39% Low, 34–39% Medium, 25–29% High), with high persistence in the High state (P(High→High) ≈ 0.966–0.974) and average high-volatility spell lengths of 34-45 trading days. The probability of transitioning directly from Low to High is essentially zero in every market (0.000–0.002): regimes always pass through the medium state first. Figure no. 1 plots the resulting regime heatmap for the core7 panel and shows a striking visual pattern that motivates the formal H1 test in Section 4.4: the High-regime (red) bands line up almost vertically across all seven markets at every major stress episode, the dot-com bust, the global financial crisis, the euro-area debt crisis, COVID-19, and the 2022 energy shock.

Table no. 4. Volatility regime characteristics, core7 markets

Market	% Low	% Medium	% High	Avg. High spell (days)	Max. High spell	P(Low→High)	P(High→High)
DAX	38.9	34.3	26.8	42.3	302	0.000	0.973
CAC 40	36.9	38.4	24.8	35.4	249	0.001	0.966
FTSE MIB	32.5	38.7	28.8	39.9	220	0.000	0.971

Market	% Low	% Medium	% High	Avg. High spell (days)	Max. High spell	P(Low→High)	P(High→High)
IBEX 35	34.8	36.4	28.8	40.4	243	0.001	0.972
AEX	36.3	35.8	28.0	43.3	305	0.001	0.974
FTSE 100	33.9	39.3	26.8	45.2	286	0.001	0.974
SMI	36.7	36.9	26.5	33.6	274	0.002	0.966

Source: Authors' calculations. Regime states relabelled by fitted mean (0 = Low, 1 = Medium, 2 = High); best-of-20 random-restart HMM estimation (Section 3.2).

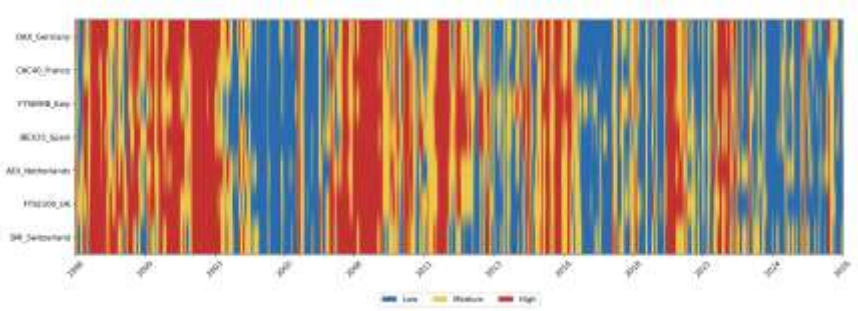


Figure no. 1: Volatility regime heatmap, core7 markets, 1998–2026 (blue = Low, yellow = Medium, red = High)
Source: Authors' calculations.

4.4 Dynamic connectedness and stress transmission

Table no. 5 reports the rolling-window connectedness summary for the core7 panel. The full-sample Total Connectedness Index averages 79.8% (range 71.4–83.4%); these seven European equity markets share, on average, close to four-fifths of their ten-day-ahead forecast-error variance with one another at any point in time, a level consistent with the high integration reported in the connectedness literature for developed European markets. Figure no. 2 plots the daily TCI series together with the system-wide high-volatility share; the correlation between the two series is 0.233 over the full sample, an informal first look at hypothesis H2 that is tested formally below.

Table no. 5. Connectedness summary, core7 panel (full sample averages)

Market	Avg. received	Avg. transmitted	Avg. net	% days net transmitter
CAC 40	81.65	90.21	+8.56	100%
AEX	80.66	84.09	+3.43	81%
DAX	80.73	84.13	+3.40	87%
FTSE MIB	79.83	79.57	-0.27	51%
FTSE 100	79.11	75.68	-3.43	16%

Market	Avg. received	Avg. transmitted	Avg. net	% days net transmitter
IBEX 35	79.14	75.45	-3.69	14%
SMI	77.38	69.37	-8.01	3%

Source: Authors' calculations. Net = Transmitted - Received; positive net = net transmitter.

France (CAC 40), Germany (DAX) and the Netherlands (AEX) are persistent net transmitters of stress throughout the sample; Switzerland (SMI) is a persistent net receiver, consistent with its relatively defensive, safe-haven-adjacent market role, and the United Kingdom and Spain are milder, less consistent net receivers. Italy (FTSE MIB) is close to net-neutral on average but swings between being a strong transmitter and a strong receiver over time (Figure no. 3), a pattern consistent with its role flipping specifically during sovereign-stress episodes, exactly the kind of regime-dependence tested formally in the next subsection.

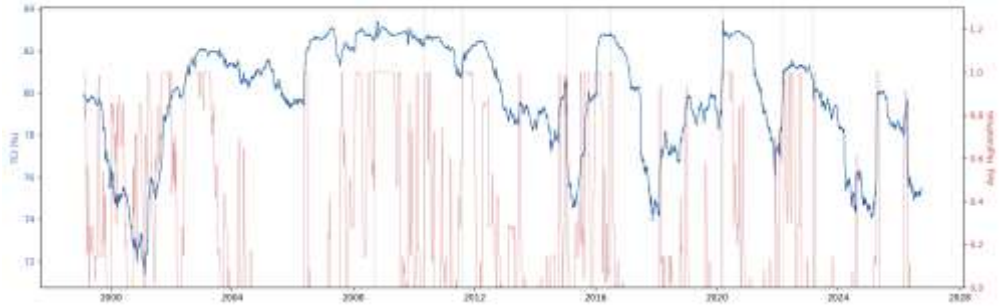


Figure no. 2: Total connectedness index (TCI) and system-wide high-volatility share, core7 panel

Source: Authors' calculations.



Figure no. 3: Net transmission by market over time, core7 panel

Source: Authors' calculations.

4.5 Regime shifts and network reconfiguration

Table no. 6 reports the formal H1 synchronisation test. Regime-entry dates (a market transitioning into the HMM High state) coincide across the core7 markets far more often than a within-market-preserving placebo benchmark, each market's own entry series circularly shifted by a random offset, 1,000 draws per window, would predict, at every coincidence window and with $p < 0.001$ throughout. Same-day coincidence alone is 24.0%, against a placebo mean of 0.7%; within a ±10-trading-day window, 59.8% of one market's regime entries have another market entering High within the same window, against a 15.0% placebo benchmark. H1 is strongly supported: European volatility regimes behave as a genuinely integrated stress system rather than as independent national events, and this conclusion is unchanged when regime entries are instead dated with PELT change-point detection rather than the HMM (Section 5, item 2).

Table no. 6. Cross-market synchronisation of volatility-regime entries (H1), core7 panel

Coincidence window	Observed rate	Placebo mean	Placebo 95th pct.	p-value
Same day (±0d)	0.240	0.007	0.012	< 0.001
±3 trading days	0.430	0.050	0.062	< 0.001
±5 trading days	0.505	0.079	0.094	< 0.001
±10 trading days	0.598	0.150	0.170	< 0.001

Source: Authors' calculations. Placebo distribution from 1,000 circular-shift simulations per window.

Table no. 7 compares connectedness and network-concentration measures across system-wide Low/Medium/High volatility terciles (H4). Total connectedness is significantly higher in High-regime days than in Low-regime days (+0.67 points, $p < 0.001$), as H2 implies. The picture on network concentration is genuinely mixed depending on which of the three candidate measures is used, and all three are reported rather than selecting the one that best fits the hypothesis: the Herfindahl index of |net transmission| shares (HHI_net) rises significantly in High-volatility periods (+0.012, $p < 0.001$), supporting H4, a smaller subset of markets accounts for a larger share of the net stress-transmission imbalance during stress. But the Herfindahl index of gross transmission (HHI_transmitted) and the two-largest-transmitters' share (Top2Share) both move slightly in the opposite direction: gross transmission becomes marginally more evenly spread across markets, not less, during high-volatility periods. Together, these results suggest that during stress every market transmits more broadly and diffusely to every other market (consistent with the integrated-system picture from H1/H2), while the net balance of who ends up a transmitter versus a receiver becomes more polarised. H4 is therefore supported for net-transmission concentration and not supported (mildly reversed) for gross-transmission concentration; this nuance is carried explicitly into the EMSTI construction (Section 3.5), which uses HHI_transmitted precisely because it is the measure that does not mechanically co-move with AvgHighVolProb.

Table no. 7. Connectedness and network concentration by system-wide volatility regime (H4), core7 panel

Metric	Low	Medium	High	High − Low	p-value
TCI	79.84	79.00	80.51	+0.67	< 0.001
HHI (gross transmission)	0.1447	0.1447	0.1443	-0.0005	< 0.001
HHI (|net transmission|)	0.1974	0.1991	0.2095	+0.0122	< 0.001
Top-2 share (gross)	0.3207	0.3199	0.3128	-0.0079	< 0.001
Avg. |net transmission|	5.69	5.61	4.56	-1.13	< 0.001

Source: Authors' calculations. Terciles of AvgHighVolShare_t; p-values from two-sided t-tests, High vs. Low.

Table no. 8 reports the formal H2 event study: total connectedness is regressed on Pre-shift, Shift-window and Post-shift dummies around 46 system-wide regime-shift events (core7, 1999-2026), controlling for TCI persistence and using HAC standard errors. TCI is significantly elevated in the pre-shift window ($p < 0.001$), the shift week itself ($p < 0.01$) and for roughly a month afterwards ($p < 0.05$), even after controlling for connectedness's own persistence. Connectedness is already rising before the regime shift is formally dated, consistent with an early-warning interpretation, and remains elevated for weeks afterwards. H2 is strongly supported in the uncontrolled specification; Section 5 reports an important qualification once global risk factors are added as controls.

Table no. 8. Event-study results around system-wide regime shifts (H2), core7 panel

Window	Coefficient	Std. error	p-value	Significance
Pre-shift	0.0244	0.0073	0.0008	***
Shift window	0.0447	0.0161	0.0053	***
Post-shift	0.0143	0.0069	0.0373	**

*Source: Authors' calculations. Dependent variable: TCI_t; regression includes TCI_t−1 and a constant (not shown); HAC (Newey-West) standard errors, 10 lags. *** p<0.01, ** p<0.05, * p<0.10.*

4.6 Market transmitters, receivers, and regime leadership

Table no. 9 reports the H3 panel-logit results: a market's own lagged high-volatility status and the lagged average high-volatility status of every other market both strongly predict its own next-day high-volatility status (own-lag odds ratio ≈ 891, foreign-lag odds ratio ≈ 16.4, both $p < 0.001$), clear evidence of cross-market spillover in regime transitions. A market's own lagged net transmission adds no further predictive power once the regime-state variables are included ($p = 0.58$), suggesting that a market's directional connectedness role does not carry information about its own next-day regime beyond what the regime states themselves already capture. Pairwise Granger-causality tests on regime probabilities (lags 1–5, 5% threshold) find that, within the tightly integrated core7 panel, essentially every market Granger-causes every other market; leadership scores are uniformly zero, i.e. the core markets are too synchronised to distinguish individual

leaders, itself consistent with the H1 clustering result. Once the shorter, more peripheral nine- and ten-market panels are used, genuine differentiation emerges: the United Kingdom and Spain are consistent net followers, while the Netherlands, Germany, Denmark and Poland are consistent net leaders in their respective panels. H3 is therefore supported for cross-market spillover in general, and mixed, informative only outside the tightly integrated core, for individual market leadership specifically.

Table no. 9. Lead-lag and spillover tests (H3), core7 panel

Variable	Coefficient	Odds ratio	p-value	Significance
Own lagged high-vol. status	6.792	890.7	< 0.001	***
Foreign lagged high-vol. status (avg.)	2.797	16.4	< 0.001	***
Own lagged net transmission	-0.011	0.99	0.576	

Source: Authors' calculations. Panel logit, dependent variable HighVol_i,t; market fixed effects; market-clustered standard errors.

4.7 European Market Stress Transmission Index

Figure no. 4 plots the resulting EMSTI series for the core7 panel. Without being told where the known stress episodes are, EMSTI's historical peak falls on 24 August 2015, in the midst of the Greek bailout referendum aftermath and the August 2015 China-led global equity selloff; the correlation between EMSTI and AvgHighVolProb is 0.627, confirming that the index tracks system-wide volatility stress while also incorporating the connectedness and concentration information from the other two layers. In the shorter, more recent nine- and ten-market panels, EMSTI's peak instead falls in March-May 2022, coinciding with Russia's invasion of Ukraine and the ensuing European energy shock, and the correlation with AvgHighVolProb rises to 0.90–0.96, reflecting the more concentrated crisis history of the shorter samples. EMSTI is not a crisis-prediction model; it is a market-monitoring framework that identifies periods when volatility stress and cross-market transmission reinforce each other, intended to complement – not replace – existing risk indicators such as implied-volatility indices.

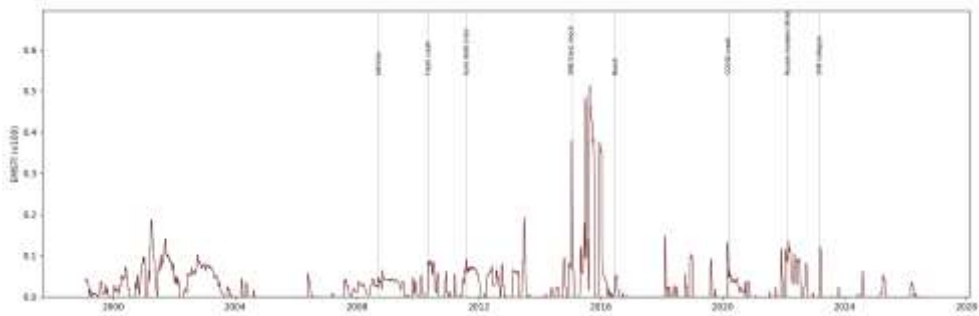


Figure no. 4: European Market Stress Transmission Index (EMSTI), core7 panel
Source: Authors' calculations.

Figure no. 5 (panels a-b) illustrates the underlying network directly for two of these dates, a calm pre-crisis snapshot (250-day window ending 15 June 2006) and the COVID-19 crisis snapshot (window ending 20 March 2020); the same network measures are also computed, but not plotted here, for two further dates, the euro-area debt-crisis (window ending 25 November 2011) and the 2022 energy-shock (window ending 15 September 2022) snapshots. Switzerland and the United Kingdom are net receivers in each of these four snapshots, a persistent safe-haven/periphery role that holds across four very different market regimes, while France, Spain and Italy, specifically during the crisis snapshots, are the persistent net transmitters.

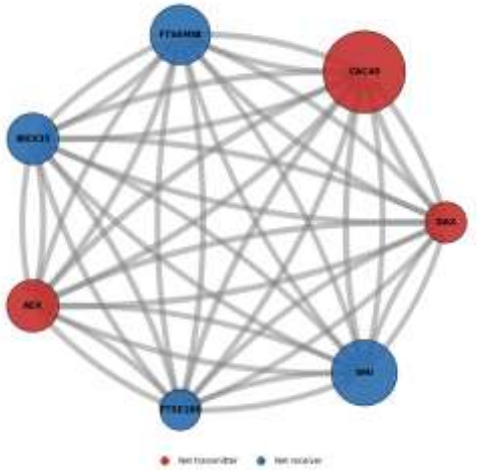


Figure no. 5a: Network snapshot, calm period (250-day window ending 2006-06-15)

Source: Authors' calculations.

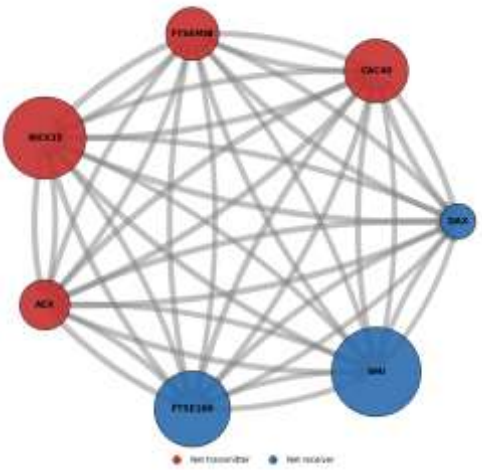


Figure no. 5b: Network snapshot, COVID-19 crisis (250-day window ending 2020-03-20)

Source: Authors' calculations. Edge width proportional to directional connectedness; node size proportional to net transmission.

5. Robustness Analysis

The two central findings - H1 (regime synchronisation) and H2 (connectedness rises around regime shifts) – are checked against ten robustness variations, summarised in Table no. 10 and reported in full in the accompanying replication material. Nine of the ten checks leave the main conclusions essentially unchanged.

The alternative volatility model check (EGARCH versus GJR-GARCH) leaves the results unchanged: the two models agree completely on the model-fit ranking and on the leverage-effect finding in all ten sampled markets. Substituting PELT change-point detection for the hidden Markov model as the regime-detection method also leaves H1 intact: full-label agreement between the two methods is moderate (46-57%, as expected for two conceptually different detection methods, one probabilistic and one based on discrete breaks), but H1 re-run on PELT-dated regime entries still holds strongly, with a same-day coincidence rate of 56.5% against an approximately zero placebo benchmark ($p < 0.001$). Re-estimating connectedness with two alternative estimators – the Kalman-filter TVP-VAR and the two-step DCC-GARCH model – produces series that correlate strongly with the baseline rolling-VAR TCI (0.69 and 0.71, respectively); H2 re-estimated on the TVP-VAR series confirms the pre-shift and shift-week rise in connectedness, though the post-shift effect is weaker under this smoother, window-free estimator.

Restricting the market set (excluding the United Kingdom; using a eurozone-only subset) leaves the core findings unchanged: TCI correlations with the core7 baseline are 0.99 and 0.91, respectively, and 10 of the 11 market roles (net-transmitter or net-receiver sign) examined across the two variants are preserved. Truncating the sample to start in 2006 or

2010 instead of the full history leaves mean TCI and the H2 pre-shift/shift-window effects essentially unchanged, ruling out that the results are driven by the pre-2006 portion of the sample. Varying the per-market high-volatility threshold to 0.40 or 0.60 instead of the baseline 0.50 produces H1 and H2 results that are nearly identical to the baseline at both thresholds, and dropping the COVID-19 window (2020-02-01 to 2020-08-31) entirely leaves both H1 and H2 nearly unchanged, ruling out that either finding is a COVID-19-specific artefact.

The eighth check, whether H2 survives controlling for global risk factors (H5), is the one genuine qualification, and it is built from real market data rather than assumed: daily series for the VIX (CBOE volatility index), the S&P 500, the EUR/USD exchange rate and the US 10-year Treasury yield were downloaded from Yahoo Finance (tickers ^VIX, ^GSPC, EURUSD=X and ^TNX respectively) over the full available history, converted to log(VIX), the US equity daily log return, the EUR/USD daily log return and the daily change in the US 10-year yield, and added as controls to the event-study regression (sample necessarily restricted to 2003-12-02 onward by EUR/USD's Yahoo Finance history, 37 regime-shift events; a VSTOXX control was sought under three candidate tickers but could not be sourced from Yahoo Finance in this environment, and is disclosed here as a data limitation rather than silently dropped). Adding these controls changes the picture substantially. Three of the four – log(VIX), the US equity returns, and the US 10-year yield change are themselves highly significant ($p < 0.01$); the EUR/USD return control is not statistically significant over this sample ($p = 0.16$). That the significant controls are specifically the US/global-risk ones, and not the exchange-rate control, is itself consistent with European connectedness co-moving with global, largely US-driven risk sentiment rather than with currency movements per se. Once the four controls are included, the pre-shift effect flips sign and becomes significantly negative (-0.0071, $p = 0.026$), and the post-shift effect becomes insignificant ($p = 0.929$); only the contemporaneous shift-window effect survives, smaller in magnitude but still significant ($p = 0.021$). A meaningful part of what appears, in the uncontrolled specification, to be "connectedness rises for weeks before and after a European regime shift" is therefore at least partly a reflection of concurrent global risk-factor movements rather than a purely European phenomenon; only the contemporaneous jump in connectedness during the shift window itself is unambiguously attributable to the European regime shift once global conditions are held constant. This qualification is discussed further in Section 6.

Converting the two non-euro markets (FTSE 100 and SMI) into EUR terms using GBP/EUR and CHF/EUR exchange rates produces a TCI series with a 0.97 correlation to the local-currency baseline over the same restricted sample, and 6 of the 7 markets, including both converted markets, keep their net-transmitter/receiver sign. Finally, re-estimating connectedness with 125-day and 500-day rolling windows instead of the baseline 250 days leaves the mean level of TCI essentially unchanged; the H2 pre/shift/post pattern is present at the shorter window and progressively smooths out at the longer window, exactly as expected of a slower-moving rolling estimator.

Table no. 10. Robustness checks: summary

Check	Main conclusion supported?	Effect direction
1. Alternative volatility model	Yes	Unchanged
2. Alternative regime detection	Yes	Unchanged
3. Alternative connectedness estimator	Yes	Unchanged
4. Alternative market set	Yes	Unchanged
5. Alternative sample start	Yes	Unchanged
6. Alternative high-vol. threshold	Yes	Unchanged
7. Crisis (COVID-19) exclusion	Yes	Unchanged
8. Global risk controls (H5)	Partially / No	Weakened; pre-shift sign-flips
9. Currency robustness	Yes	Unchanged
10. Window-length robustness	Yes	Unchanged (level); pattern smooths

Source: Authors' calculations; see the accompanying replication code and table9_robustness_summary.csv for full coefficients and commentary on every check.

6. Discussion and Financial Stability Implications

The results support four practically oriented uses of a regime-sensitive connectedness framework for financial-stability monitoring. First, the strong regime-synchronisation result (H1) means that a supervisor observing one European market enter a high-volatility state has genuine, statistically grounded reason to watch the others: the same-day coincidence rate of 24% against a 1% placebo baseline indicates that European volatility stress is a system property, not a sequence of independent national events. Second, the event-study evidence (H2) that connectedness rises before regime shifts are formally dated, even after the H5 qualification is taken into account for the pre- and post-shift windows, suggests that a rising TCI can act as one useful early signal that system-wide stress is building, complementing rather than substituting for market-specific volatility indicators. Third, the persistent transmitter/receiver roles documented in Section 4.4 (France, Germany and the Netherlands as transmitters; Switzerland, and more mildly the United Kingdom and Spain, as receivers) and the network-concentration results (Section 4.5) together suggest that certain markets deserve closer monitoring as potential amplification points, while the finding that gross transmission becomes more diffuse during stress, even as net-transmission concentration rises, is itself informative: during a systemic episode, most markets are transmitting to most other markets simultaneously, and it is the net imbalance, not the raw transmission pattern, that concentrates around a smaller set of markets. Fourth, EMSTI offers a single, interpretable summary of these three layers that could complement existing risk indicators such as implied-volatility indices; EMSTI is presented here explicitly as a market-monitoring framework rather than a crisis-prediction model.

The H5 finding is, in our view, the paper's most important qualification and deserves emphasis rather than being smoothed over. Once the analysis controls for global risk

sentiment, the VIX, US equity returns, the EUR/USD rate and US interest-rate changes, the pre-shift rise in European connectedness that appeared so clearly in the uncontrolled H2 test is substantially reduced and, for the pre-shift window specifically, reverses sign. This does not overturn the paper's central claim that connectedness and volatility regimes move together; the contemporaneous shift-window effect survives the controls, but it does mean that a supervisor reading a rising TCI as a purely European early-warning signal should first ask how much of that rise reflects global, rather than region-specific, risk conditions. We regard this as exactly the kind of result that the interpretive caution urged by Forbes and Rigobon (2002) is meant to surface: throughout this paper we use the language of "stress transmission" and "connectedness", and reserve "contagion-like behaviour" for instances where the evidence specifically supports an abnormal, non-fundamental increase in cross-market transmission during stress, which the H5 result shows is, at most, only partially the case for the pre- and post-shift windows of the paper's central event study.

Conclusions

This paper tests whether volatility regime shifts in European equity markets are associated with changes in the structure of cross-market stress transmission. Using daily data for seven, and, in extension panels, up to ten major European equity indices over as much as 1998-2026, combining EGARCH-based hidden Markov regime detection with rolling-window generalised-variance-decomposition connectedness, we find that they are: volatility regimes are strongly synchronized across markets (H1), total connectedness rises significantly around system-wide regime shifts (H2), cross-market regime spillover is pervasive, with individual market leadership identifiable only outside the most tightly integrated core markets (H3), and the net-transmission network becomes more concentrated during stress even as gross transmission becomes more diffuse (H4). France, Germany and the Netherlands are the sample's persistent net transmitters of stress; Switzerland is a persistent net receiver. A composite European Market Stress Transmission Index built from these three layers identifies well-known stress episodes, the 2015 Greek/China-selloff period and the 2022 Russia-Ukraine energy shock, without being told where to look.

These findings are robust to nine of ten recommended robustness variations, alternative volatility and connectedness models, alternative regime-detection methods and market sets, sample-period and threshold sensitivity, COVID-19 exclusion, currency denomination, and rolling-window length. The one substantive qualification (H5) is that part of the connectedness rise observed around European regime shifts, particularly in the weeks before and after the shift itself, reflects concurrent global (largely US-driven) risk sentiment rather than a purely European phenomenon; only the contemporaneous shift-week effect is unambiguously attributable to the European regime shift once global risk factors are held constant.

Several limitations should be kept in mind when interpreting these results. First, and most importantly, connectedness models based on forecast-error variance decompositions capture predictive and variance-attribution relationships, not structural causality; the careful language and lead-lag evidence used throughout this paper are intended to support the more defensible claim that the evidence is "consistent with" stress transmission, not

that it demonstrates contagion in the strict sense of excess co-movement over and above normal interdependence. Second, some sampled indices contain overlapping multinational constituents, which national-index-level analysis such as this cannot fully disentangle; a firm- or sector-level connectedness analysis is a natural extension. Third, the analysis uses daily data throughout and may therefore miss intraday stress-transmission dynamics; a high-frequency extension is left for future work. Fourth, regime detection depends on the chosen method, which is why the paper explicitly cross-checks its central synchronisation finding against an independent, model-free change-point method (Section 5). Fifth, and as already discussed, a downside- or quantile-specific connectedness measure – needed to formally test whether negative shocks propagate more strongly than positive ones – was not estimated for this paper and remains a natural next step, alongside closing the data gap that currently prevents a European-specific (VSTOXX-based) volatility control from being included in the H5 global-risk-control check. Future work along these lines, together with an extension to a longer post-2026 sample as it becomes available, would further strengthen the case for regime-sensitive connectedness measures as a standard component of European financial-stability monitoring.

References

[1] Adrian, T. and Brunnermeier, M.K. (2016) 'CoVaR', American Economic Review, 106(7), pp. 1705–1741.

[2] 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), p. 84.

[3] Bai, J. and Perron, P. (2003) 'Computation and analysis of multiple structural change models', Journal of Applied Econometrics, 18(1), pp. 1–22.

[4] Bekaert, G., Harvey, C.R. and Ng, A. (2005) 'Market integration and contagion', Journal of Business, 78(1), pp. 39–69.

[5] 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), pp. 535–559.

[6] Bollerslev, T. (1986) 'Generalized autoregressive conditional heteroskedasticity', Journal of Econometrics, 31(3), pp. 307–327.

[7] Brownlees, C. and Engle, R.F. (2017) 'SRISK: A conditional capital shortfall measure of systemic risk', Review of Financial Studies, 30(1), pp. 48–79.

[8] Călin, A.C. and Lupu, R. (2026a) 'Not All Bubbles Are Alike: Crash Risk and Speculative Regimes in Precious Metals', International Journal of Finance & Economics, (ahead of print), doi: 10.1002/ijfe.70267.

[9] Călin, A.C. and Lupu, R. (2026b) 'Tracing exuberance: multi-model evidence of bubble episodes in the Bucharest stock market', International Journal of Emerging Markets, (ahead of print), doi: 10.1108/IJOEM-12-2025-2934.

[10] Diebold, F.X. and Yilmaz, K. (2009) 'Measuring financial asset return and volatility spillovers, with application to global equity markets', Economic Journal, 119(534), pp. 158–171.
[11] 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), pp. 57–66.
[12] Engle, R.F. (1982) 'Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation', Econometrica, 50(4), pp. 987–1007.
[13] Engle, R.F. (2002) 'Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models', Journal of Business & Economic Statistics, 20(3), pp. 339–350.
[14] Forbes, K.J. and Rigobon, R. (2002) 'No contagion, only interdependence: Measuring stock market comovements', Journal of Finance, 57(5), pp. 2223–2261.
[15] Glosten, L.R., Jagannathan, R. and Runkle, D.E. (1993) 'On the relation between the expected value and the volatility of the nominal excess return on stocks', Journal of Finance, 48(5), pp. 1779–1801.
[16] Hamilton, J.D. (1989) 'A new approach to the economic analysis of nonstationary time series and the business cycle', Econometrica, 57(2), pp. 357–384.
[17] Killick, R., Fearnhead, P. and Eckley, I.A. (2012) 'Optimal detection of changepoints with a linear computational cost', Journal of the American Statistical Association, 107(500), pp. 1590–1602.
[18] Nelson, D.B. (1991) 'Conditional heteroskedasticity in asset returns: A new approach', Econometrica, 59(2), pp. 347–370.
[19] Pesaran, H.H. and Shin, Y. (1998) 'Generalized impulse response analysis in linear multivariate models', Economics Letters, 58(1), pp. 17–29.