GEOPOLITICS, GEOECONOMICS, AND SOVEREIGN RISK: DIFFERENT SHOCKS, DIFFERENT CHANNELS*

VERSION OF MARCH 20TH, 2026

Alvaro Ortiz
BBVA Research & CRIW(NBER)
alvaro.ortiz@bbva.com

Tomasa Rodrigo
BBVA Research
tomasa.rodrigo@bbva.com

Pablo Saborido
BBVA Research
pablo.saborido@bbva.com

ABSTRACT

Geopolitical and geoeconomic shocks reprice sovereign credit risk through different transmission channels. Using a daily panel of 42 advanced and emerging economies over 2018–2025, we show that geopolitical shocks raise sovereign CDS spreads primarily through direct sovereign repricing, while the Global Financial Cycle (GFC) channel moves in the opposite direction and partly offsets that increase—a "scissors pattern." Geoeconomic shocks, by contrast, transmit mainly through financial conditions, policy uncertainty, and domestic amplification, with only a limited direct repricing component. A semistructural framework provides sign benchmarks for four transmission channels, and a Shapley–Taylor decomposition of nonlinear machine-learning predictions partitions each observation's spread into Direct, GFC, Uncertainty, and Local components. Narrative local projections around four dated crisis events recover the scissors pattern for Russia–Ukraine and support the broader channel taxonomy in the remaining episodes. Additional scorecard, placebo, and sign-restricted SVAR evidence corroborates the taxonomy beyond the baseline ML decomposition. Geopolitical direct effects decay with distance from the conflict zone in a gravity-style pattern ($R^2 = 0.35$ for Russia–Ukraine), while policy-uncertainty shocks activate the Uncertainty channel more globally. The taxonomy implies that liquidity provision can mitigate GFC-driven spread widening, but not direct geopolitical sovereign repricing.

Keywords Geopolitics · Geoeconomics · Sovereign Credit Risk · Global Financial Cycle · Transmission Channels · Shapley–Taylor Decomposition · Narrative Identification · Machine Learning · Text-as-Data

1 Introduction

Geopolitical and geoeconomic shocks both widen sovereign credit spreads, but they do not do so through the same transmission channels. Using a daily panel of 42 advanced and emerging economies over 2018–2025, we show that geopolitical shocks raise sovereign CDS spreads primarily through direct sovereign repricing, while the Global Financial Cycle (GFC) channel moves in the opposite direction and partly offsets that increase—a "scissors pattern." Geoeconomic shocks, by contrast, load only weakly on direct repricing and transmit mainly through financial conditions, policy uncertainty, and domestic amplification.

This distinction matters for both theory and policy. When spreads widen because the sovereign itself is repriced, the effect is persistent, concentrated among exposed countries, and decays with distance from the conflict zone in a gravity-style pattern. When widening is mediated by common financial conditions or policy uncertainty, it is more diffuse and potentially more reversible. In policy terms, liquidity tools, swap lines, and backstop facilities can address

*The authors thank, without implicating, Stephen Hansen (UCL), Daniel J. Lewis (UCL), Andreas Joseph (Bank of England), Juri Marcucci (Banca d'Italia), Santiago E. Alvarez-Blaser (Bank of Spain), Marina Diakonova (Bank of Spain), Javier Perez (Bank of Spain), Martin Saldias (Bank of Portugal) and participants in the internal seminar of Bank of Spain and BBVA research as well as participants 13th ECB Conference on Forecasting Techniques and ECB Internal DGE Seminar . We especially thank Buket Begun Boga, Patricia Soroa and Ismael Frutos for their contribution to build the database. The news-based indicators used in this paper are constructed and publicly distributed by BBVA Research and all three authors are employed by BBVA Research. The authors have no sources of external research funding to disclose and declare no conflicts of interest.

GFC-driven widening, but they do not eliminate the persistent component of direct geopolitical sovereign repricing. An originator penalty also emerges: geoeconomic-shock originators face spread widening through the Local channel even as monetary easing compresses spreads abroad.

The paper combines a semistructural benchmark with a three-layer empirical design. The semistructural framework maps four standard mechanisms—sovereign default, global pricing, policy-regime uncertainty, and domestic amplification—into four empirical channels: Direct, GFC, Uncertainty, and Local. It provides sign and dominance benchmarks for short-horizon, panel-mean responses to geopolitical and geoeconomic shocks. The empirical strategy then proceeds in three layers.

First, we estimate a nonlinear forecasting model for sovereign CDS spreads using high-frequency news-based indicators of geopolitical risk, economic-policy uncertainty, trade-policy uncertainty, and domestic political sentiment. A systematic model-comparison exercise shows that these predictors add substantial forecast content only when nonlinear interactions and state dependence are allowed, with gradient-boosted trees outperforming linear alternatives. This first layer establishes that the mapping from news and exposures to sovereign spreads is not well approximated by a globally linear specification.

Second, we use a Shapley–Taylor decomposition of the nonlinear predictions to recover channel-level contributions for each country-date observation. Guided by the semistructural sign benchmarks, we partition predicted spreads into Direct, GFC, Uncertainty, and Local components. The resulting taxonomy is economically interpretable: geopolitical episodes are characterized by strong direct repricing and offsetting GFC compression, whereas geoeconomic episodes transmit primarily through GFC, Uncertainty, and Local channels, with only limited direct loading.

Third, we validate this taxonomy using identification strategies that do not rely on the baseline ML decomposition. Narrative local projections around four dated crisis events—Russia–Ukraine, Hamas–Israel, the U.S. Presidential Election, and Liberation Day tariffs—recover the scissors pattern at the 1% significance level for Russia–Ukraine and confirm 15 of 16 event–channel sign predictions. Evaluated against the semistructural benchmarks, the point-estimate scorecard matches all 16 predictions, with a cluster bootstrap mean of 15.9/16 (sd = 0.3) and 90% confidence interval $[15,\ 16]$. A placebo falsification based on driver-specific decompositions shows that all four episodes exceed at least 83% of random non-event dates, with the channels exiting the placebo envelope differing across shock types in the direction the taxonomy predicts. A complementary narrative sign-restricted SVAR estimated on raw observable composites—entirely bypassing the Shapley–Taylor decomposition—supports all eight sign and dominance predictions. Because this exercise does not rely on the ML decomposition, it provides independent evidence that the channel assignments reflect properties of transmission rather than artifacts of model fitting.

The empirical design is intentionally layered. The semistructural framework generates ex ante channel predictions, the nonlinear ML stage captures threshold-dependent transmission, the Shapley–Taylor decomposition constructs the channel objects, and the local projection and narrative SVAR exercises provide complementary validation. Because the channel predictions are stated before estimation, the scorecard evaluates pre-specified benchmarks rather than ex post pattern matching.

The paper proceeds as follows. Section 2 describes the data. Section 3 develops the semistructural framework and dominance benchmarks. Section 4 presents the ML framework, the Shapley–Taylor channel decomposition, and the local-projection design. Section 5 evaluates predictive performance and Shapley-based results. Section 6 develops the four-channel taxonomy across the four crisis episodes. Section 7 provides econometric validation via local projections. Section 8 evaluates the taxonomy against the semistructural benchmarks. Section 9 concludes.

1.1 Related Literature

This paper connects to four main literatures.

First, we contribute to the literature that uses news-based indicators to measure geopolitical risk, economic-policy uncertainty, and trade-policy uncertainty (Baker et al., 2016; Caldara et al., 2020; Caldara and Iacoviello, 2022; Bondarenko et al., 2024; Alonso-Alvarez et al., 2025; Gentzkow et al., 2019). We show that the predictive content of these measures for sovereign risk is strongly nonlinear and depends on interactions with country exposures and domestic vulnerabilities.

Second, we contribute to the literature on sovereign risk, geopolitics, and geoeconomic statecraft. Existing work shows that geopolitical risk widens spreads (Fernández-Villaverde et al., 2024; Boubaker et al., 2023), that trade tensions amplify uncertainty (Ahn and Ludema, 2020; Aiyar et al., 2023; Fernández-Villaverde et al., 2025), and that sovereign CDS spreads reflect both default-related and global risk-premium components (Longstaff et al., 2011; Augustin et al., 2022). We extend this literature by showing that geopolitical and geoeconomic shocks operate through qualitatively

different transmission channels, providing a sovereign-risk counterpart to recent work on geoeconomic vulnerability and leverage (Clayton et al., 2026).

Third, we connect to the literature on the Global Financial Cycle and international asset pricing (Rey, 2013; Bruno and Shin, 2015; Miranda-Agrippino and Rey, 2020). In sovereign markets, global financial conditions shape the risk-premium component of CDS spreads and interact with dollar funding conditions and safe-haven flows (Longstaff et al., 2011; Du et al., 2018; Bahaj and Reis, 2022; Maggiori et al., 2020). We show that the GFC channel does not respond uniformly across shock families: under geopolitical shocks it generates the scissors pattern, whereas under geoeconomic shocks it becomes a primary transmission margin.

Fourth, we speak to the growing literature combining machine learning with economic identification (Athey and Imbens, 2019; Mullainathan and Spiess, 2017; Gu et al., 2020; Chernozhukov et al., 2024). Our contribution is not to replace causal designs with machine learning, but to assign each tool a clear role: the nonlinear predictor flexibly captures threshold effects and interactions; the Shapley–Taylor decomposition restores interpretability at the observation level; and the narrative local projections and sign-restricted SVAR provide external validation of the resulting channel taxonomy.

2 Data and Measurement

We assemble a daily panel of sovereign credit risk and its potential drivers for 42 advanced and emerging economies over January 2018 to July 2025.[2] The dependent variable is the 5-year sovereign credit default swap (CDS) spread—the annualized premium a protection buyer pays to insure against a credit event on sovereign debt.[3] A large literature documents that sovereign spreads respond to both macroeconomic fundamentals and shifts in global investor sentiment (Longstaff et al., 2011; Pan and Singleton, 2008; Aizenman et al., 2013; Augustin et al., 2022); our channel decomposition formalizes this distinction.

All variables are smoothed using a 28-day geometrically weighted moving average:

$$\tilde{X}_t = \sum_{k=0}^{27} w_k \, X_{t-k} \,, \tag{1}$$

where the weight vector $\mathbf{w} = (w_0, \ldots, w_{27})$ follows a discrete geometric distribution with $\sum_k w_k = 1$, assigning higher importance to more recent observations. The scheme is strongly front-loaded: approximately 69% of the total weight falls within the most recent week and over 90% within the last two weeks, ensuring high responsiveness to recent changes while spanning four complete weeks to equalize day-of-week composition [4]. Once the weighted moving average is applied, the data is standardized to zero mean and unit variance within each country. We do not perform outlier treatment, as extreme events are informative signals of exceptional shifts in risk perception. The panel is unbalanced; country-specific data availability and percentile distributions are detailed in Appendix A.

Our main contribution on the measurement side is the construction of daily, country-level news-based indicators from the *Global Database of Events, Language, and Tone* (GDELT), an open-source platform that monitors broadcast, print, and online media in more than 100 languages (Leetaru and Schrodt, 2013).[5] Following Bondarenko et al. (2024) and Alonso-Alvarez et al. (2025), we construct all indicators from *local-language* media rather than foreign or global outlets, because local newspapers more accurately reflect how geopolitical risks and policy uncertainties are perceived domestically.[6]

We construct two types of news-based indicators. *Coverage-based* indicators measure uncertainty solely through the volume of relevant news. For Economic Policy Uncertainty (EPU) and Trade Policy Uncertainty (TPU), we compute

[2]The sample includes Argentina, Australia, Austria, Belgium, Brazil, Canada, Chile, China, Colombia, Czech Republic, Denmark, Egypt, Finland, France, Germany, Hungary, India, Indonesia, Israel, Italy, Japan, Jordan, Malaysia, Mexico, Morocco, Netherlands, Norway, Peru, Philippines, Poland, Qatar, Russia, Saudi Arabia, Sweden, Spain, Thailand, Turkey, Ukraine, United Kingdom, United States, Uruguay, and Vietnam.

[3]The 5-year tenor is selected because it is the most liquid and widely traded segment of the CDS curve, establishing it as the standard benchmark for pricing default probability. Unlike government bond yield spreads, CDS spreads isolate default expectations from funding costs, interest rate risk, and bond-specific supply dynamics.

[4]The exact weight vector is provided upon request in the replication package.

[5]An alternative source is the Dow Jones Factiva database, a premium archive widely used in academic research. The historical component of the Economic Policy Uncertainty index, for example, is constructed from Factiva news archives.

[6]Bondarenko et al. (2024) show that local-source measures embed heterogeneity in national perspectives and geographic proximity to conflict that global coverage smooths away. Alonso-Alvarez et al. (2025) formalize *bilateral geopolitical risk* and document that local-source shocks have significant effects on domestic financial markets, whereas Anglosphere-media indicators systematically understate local impacts.

the daily share of articles containing predefined uncertainty-related keywords following Baker et al. (2016), divided by total daily news volume. Complete keyword lists are provided in Appendix A. *Tone-weighted* indicators combine volume with sentiment. For Economic Sentiment (ECO), Interest Rate Sentiment (INT), Geopolitical Risk (GPR), and Political Tensions (POL), we score relevant articles using over 40 GDELT sentiment dictionaries, yielding a tone score from -10 to $+10$, and multiply it by normalized coverage. The product is inverted so that higher values consistently correspond to greater risk.[7]

Table 1 summarizes the full set of variables.[8] The explanatory variables fall into three groups: global financial conditions, domestic macroeconomic sentiment, and political–geopolitical risk.

Table 1: Variable, Sources, and Economic Interpretation

Variable	**Abbr.**	**Type / Source**	**Description**
Dependent variable			
Sovereign CDS spread	CDS	Market (5Y)	Annualized premium on sovereign default insurance
Global financial variables			
U.S. 2-year Treasury yield	FED	Market	Proxy for the monetary policy stance (Swanson, 2021)
CBOE Volatility Index	VIX	Market	Implied S&P 500 volatility; global risk aversion
Domestic macroeconomic sentiment (GDELT, local-language media)			
Economic Sentiment	ECO	Tone $\times$ coverage	Narrative framing of the domestic economy
Interest Rate Sentiment	INT	Tone $\times$ coverage	Expectations on monetary policy and borrowing costs
Econ. Policy Uncertainty	EPU	Coverage-based	Ambiguity regarding economic policy decisions
Trade Policy Uncertainty	TPU	Coverage-based	Uncertainty on trade rules, tariffs, and disputes
Political and geopolitical risk (GDELT, local-language media)			
Geopolitical Risk	GPR	Tone $\times$ coverage	International conflict, military disputes, terrorism
Political Tensions	POL	Tone $\times$ coverage	Domestic instability, unrest, political contestation

The use of local-language sources is not merely a data-collection choice; it reflects a substantive measurement decision. English-language indices such as the global GPR of Caldara and Iacoviello (2022)—constructed from ten Anglosphere newspapers—measure when a country attracts international media attention, not how geopolitical risk is perceived domestically. Bondarenko et al. (2024) demonstrate this distinction sharply: geopolitical risk shocks identified from local Russian-language newspapers have significant adverse effects on the Russian economy, whereas shocks from English-language sources do not. Alonso-Alvarez et al. (2025) formalize bilateral geopolitical risk and show that Anglosphere-media indicators systematically understate local financial-market impacts. Because our decomposition attributes sovereign spread movements to country-specific news indicators, the accuracy of these indicators at the country level is a first-order concern: mismeasured local risk would contaminate the Direct and Local channels and bias the transmission taxonomy.

3 Geopolitics & Geoeconomics Transmission Channels and Dominance Benchmarks

This section develops a semistructural framework to distinguish how geopolitical and geoeconomic shocks transmit to sovereign risk. Rather than imposing a fully structural model, the framework isolates four economically interpretable channels (Direct, GFC, Uncertainty, and Local) and derives sign and dominance predictions for their short-run effects on sovereign spreads. The key implication is that geopolitical shocks should load primarily on direct sovereign repricing, with an offsetting role for global financial conditions, whereas geoeconomic shocks should transmit mainly through financial conditions, policy uncertainty, and domestic amplification. The five subsections that follow develop these mechanisms and summarize the corresponding benchmarks.

[7] Detailed construction formulas for all indicators are provided in Appendix A. The complete set of daily and weekly indicators is publicly available at `https://bigdata.bbvaresearch.com/en/`.

[8] Our focus on high-frequency news-based indicators complements existing work emphasizing structural determinants of external debt dynamics.

3.1 Block 1: Sovereign Default and the Direct Channel

Following Arellano (2008), the direct component of the spread rises with the perceived probability of sovereign distress:

$$s_{i,t}^{\text{DIR}} \;\propto\; \Pr\Big(y_{i,t+1} < \underline{y}_i(b_i) \;\Big|\; \mathcal{F}_t\Big), \tag{2}$$

where $y_{i,t}$ denotes fiscal capacity or repayment capacity, b_i is outstanding debt, and $\underline{y}_i(b_i)$ is a debt-dependent distress threshold.

The key distinction is how the two shock families affect repayment capacity on impact. Geopolitical shocks can reduce expected fiscal capacity through destruction, trade disruption, sanctions on state capacity, regional spillovers, or capital flight from exposed sovereigns:

$$\frac{\partial y_{i,t}}{\partial \varepsilon_t^G} < 0 \quad \text{for sufficiently exposed countries.} \tag{3}$$

By contrast, geoeconomic shocks typically do not impair repayment capacity as directly at short horizons. They may change growth expectations, fiscal expectations, or discount rates, but their impact on sovereign default risk is usually more indirect and weaker than under geopolitical shocks:

$$\left|\frac{\partial y_{i,t}}{\partial \varepsilon_t^E}\right| \;\ll\; \left|\frac{\partial y_{i,t}}{\partial \varepsilon_t^G}\right| \quad \text{on impact and at the panel mean.} \tag{4}$$

Election-type or fiscal-regime episodes may still generate modest direct repricing—particularly for the originator sovereign, where policy uncertainty feeds back into fiscal credibility through the Local channel—but not the large and systematic activation expected under geopolitical shocks.

> **Prediction D.** $\partial s^{\text{DIR}}/\partial \varepsilon^G > 0$; $|\partial s^{\text{DIR}}/\partial \varepsilon^E| \ll |\partial s^{\text{DIR}}/\partial \varepsilon^G|$ at short horizons and at the panel mean.

3.2 Block 2: Global Pricing and the GFC Channel

Let the global-financial component of sovereign spreads load on a common discount-rate factor:

$$s_{i,t}^{\text{GFC}} = \beta_i^{\text{GFC}}\, q_t, \tag{5}$$

where q_t summarizes global financing conditions or the common risk-price component embedded in sovereign spreads, and β_i^{GFC} measures country exposure to that common factor. This reduced-form representation is consistent with intermediary-asset-pricing environments such as He and Krishnamurthy (2013) and international-finance frameworks such as Gabaix and Maggiori (2015).

For geopolitical shocks, the benchmark sign comes from safe-haven reallocation and benchmark-rate compression. When conflict risk rises, capital tends to move toward safe assets, lowering the safe-rate component relevant for the pricing of many sovereign claims. For most non-targeted borrowers, this compresses the global discount-rate component of spreads even as the targeted sovereigns experience a rise in the Direct channel. At the panel mean, this generates:

$$\frac{\partial q_t}{\partial \varepsilon_t^G} < 0. \tag{6}$$

The coexistence of $\partial s^{\text{DIR}}/\partial \varepsilon^G > 0$ and $\partial s^{\text{GFC}}/\partial \varepsilon^G < 0$—sovereign-specific repricing moving in the opposite direction from the common financial factor—is what we term the *scissors condition*. It requires that the portfolio rebalancing effect dominates any general risk-aversion increase for most panel members, which is satisfied when the geopolitical shock triggers safe-haven flows without impairing the balance sheet of global intermediaries directly.[9]

For geoeconomic shocks, the sign is similar but the mechanism is different. Tariffs, policy-regime changes, and election outcomes can alter expected growth and the expected path of monetary policy. At short horizons, those shocks often lower expected policy rates or otherwise compress the common discount-rate factor:

$$\frac{\partial q_t}{\partial \varepsilon_t^E} < 0. \tag{7}$$

[9]This mechanism connects to the distinction between vulnerability and leverage in the geoeconomic framework of Clayton et al. (2026): a geopolitical shock creates vulnerability for the targeted sovereign but generates a partial financial offset for bystander sovereigns through the intermediary's portfolio adjustment. Our scissors condition is the sovereign-risk counterpart of their structural insight.

Thus both shock families can load negatively on the GFC channel, but for different reasons: geopolitical shocks through safe-haven and benchmark-rate effects, geoeconomic shocks through expected policy and discount-rate revaluation. The GFC channel therefore discriminates between shock types not through its own sign—which is negative in both cases—but through its *relative role in the channel constellation*: under geopolitical shocks, the GFC is the secondary, offsetting channel (the scissors); under geoeconomic shocks, it is the dominant transmission channel. This benchmark is intentionally local to the event windows studied here; a sufficiently systemic financial crisis need not satisfy it.

Prediction G. $\partial s^{\text{GFC}}/\partial \varepsilon^G < 0$; $\partial s^{\text{GFC}}/\partial \varepsilon^E < 0$. The distinguishing feature is relative dominance: $|\varphi^{\text{GFC}}| < |\varphi^{\text{DIR}}|$ under geopolitical shocks, $|\varphi^{\text{GFC}}| > |\varphi^{\text{DIR}}|$ under geoeconomic shocks.

3.3 Block 3: Policy-Regime Uncertainty and the Uncertainty Channel

Following Pástor and Veronesi (2012, 2013), let the uncertainty component depend on beliefs over domestic economic-policy regimes. Let π_t denote the perceived probability of one policy regime relative to another; then a convenient reduced-form representation is

$$s_{i,t}^{\text{UNC}} = \alpha_i^{\text{UNC}} H(\pi_t), \tag{8}$$

where $H(\pi_t)$ is an uncertainty index, such as Shannon entropy, and α_i^{UNC} measures country sensitivity to policy-regime ambiguity.

The relevant distinction is between economic-policy uncertainty and generic uncertainty. EPU and TPU are designed to capture uncertainty about taxes, trade barriers, regulation, fiscal priorities, and related policy choices. Geoeconomic shocks naturally activate this margin because they directly alter beliefs about policy regimes:

$$\frac{\partial H(\pi_t)}{\partial \varepsilon_t^E} > 0. \tag{9}$$

Geopolitical shocks need not do so. A military conflict, territorial attack, or security shock may sharply raise aggregate uncertainty, but it does not automatically change the distribution over domestic economic-policy regimes. At the short horizons of our event windows (1–3 months), geopolitical shocks have typically not yet cascaded into domestic policy-regime uncertainty—that cascade may materialise at longer horizons as sanctions and fiscal responses unfold. Accordingly, the short-horizon benchmark for the Uncertainty channel under geopolitical shocks is substantially smaller than under geoeconomic shocks:

$$\left|\frac{\partial H(\pi_t)}{\partial \varepsilon_t^G}\right| \ll \left|\frac{\partial H(\pi_t)}{\partial \varepsilon_t^E}\right|. \tag{10}$$

Within geoeconomic shocks, the activation of the Uncertainty channel is expected to be strongest for policy-regime shocks such as elections and tariff-policy episodes. Sanctions or energy-policy episodes may load more weakly on this channel and more strongly on the GFC or Local channels.

Prediction U. $|\partial s^{\text{UNC}}/\partial \varepsilon^G| \ll |\partial s^{\text{UNC}}/\partial \varepsilon^E|$ at short horizons, especially for EPU/TPU-type shocks.

3.4 Block 4: Domestic Amplification and the Local Channel

The Local channel captures the interaction between external shocks and domestic vulnerabilities. Let

$$s_{i,t}^{\text{LOC}} = \psi_i(z_{i,t})\, \varepsilon_t, \quad \psi_i'(z_{i,t}) > 0, \tag{11}$$

where $z_{i,t}$ collects domestic state variables and $\psi_i(\cdot)$ is increasing in vulnerability. Two components of $z_{i,t}$ are particularly relevant for sovereign risk. Local economic-activity sentiment captures shifts in perceived growth and fiscal capacity: a deterioration lowers expected primary surpluses and moves the sovereign closer to its debt-sustainability boundary. Local interest-rate sentiment captures shifts in domestic borrowing costs and rollover conditions: a tightening raises the effective cost of servicing outstanding debt and amplifies the fiscal impact of any external shock. Together, these two margins determine how much a given external impulse translates into sovereign-spread widening through the domestic channel. The remaining elements of $z_{i,t}$—trade exposure, external dependence, domestic political conditions, and regional sensitivity—modulate the strength of this amplification across countries.

Geoeconomic shocks have a natural affinity with domestic amplification because they operate through heterogeneous bilateral exposures. Tariffs affect countries differently according to trade structure; election shocks affect countries differently according to policy alignment and financial exposure; sanctions and energy measures affect countries differently according to import dependence, sectoral links, and domestic politics. Crucially, these shocks also shift local activity and interest-rate sentiment differentially: a tariff shock depresses growth expectations in export-dependent

economies while tightening borrowing conditions in fiscally exposed ones, activating both margins of $z_{i,t}$ simultaneously. These heterogeneous effects can survive aggregation and therefore appear at the panel mean:

$$\mathbb{E}_i\left[\frac{\partial s_{i,t}^{\mathrm{LOC}}}{\partial \varepsilon_t^E}\right] > 0. \tag{12}$$

Geopolitical shocks can also generate substantial local responses for individually exposed countries (e.g., Poland's $\varphi^{\mathrm{LOC}} = +0.06$ during Russia–Ukraine). However, at the panel mean these effects are secondary to the Direct and GFC channels, because geopolitical shocks simultaneously produce a broad common repricing that dominates the cross-sectional average. By contrast, the bilateral heterogeneity of geoeconomic shocks means that local amplification—operating through activity and interest-rate sentiment—*is* the primary differentiator across countries:

$$\mathbb{E}_i\left[\frac{|\partial s_{i,t}^{\mathrm{LOC}}/\partial \varepsilon_t^G|}{|\partial s_{i,t}^{\mathrm{DIR}}/\partial \varepsilon_t^G|}\right] \ll \mathbb{E}_i\left[\frac{|\partial s_{i,t}^{\mathrm{LOC}}/\partial \varepsilon_t^E|}{|\partial s_{i,t}^{\mathrm{DIR}}/\partial \varepsilon_t^E|}\right]. \tag{13}$$

Prediction L. Under geopolitical shocks, $\mathbb{E}_i[|\varphi^{\mathrm{LOC}}|] \ll \mathbb{E}_i[|\varphi^{\mathrm{DIR}}|]$; under geoeconomic shocks, $\mathbb{E}_i[\partial s^{\mathrm{LOC}}/\partial \varepsilon^E] > 0$ and φ^{LOC} is a primary transmission channel.

3.5 Dominance Benchmarks and Nonlinearities

Table 2 summarizes the resulting benchmarks. The framework implies a *scissors pattern* for geopolitical shocks: the Direct channel rises while the GFC channel falls, with the Uncertainty and Local channels playing a secondary role at the panel mean. Geoeconomic shocks display a different constellation: the Direct channel is small relative to its geopolitical counterpart, and the response operates primarily through the GFC, Uncertainty, and Local channels. The benchmarks are stated as sign predictions for strong channels and *dominance inequalities* for secondary channels—a formulation that is directly testable by comparing magnitudes across episodes.

Three features deserve emphasis. First, the four blocks are modular: each channel is tied to a distinct mechanism, so a miss in one margin does not invalidate the others. Second, the scissors pattern is not imposed mechanically; it follows from the coexistence of sovereign-specific repricing and common-factor compression under geopolitical shocks. Third, the dominance inequalities yield a clear pass/fail criterion: comparing $|\varphi^{\mathrm{DIR}}|$ or $|\varphi^{\mathrm{UNC}}|$ across geopolitical and geoeconomic episodes is sufficient. The scorecard in Section 8 evaluates these benchmarks across all crisis episodes, counting both sign confirmations ($\uparrow$, $\downarrow$) and dominance confirmations ($\ll$).

The framework also motivates the empirical specification. Default risk is threshold-based, policy-regime uncertainty is nonlinear in beliefs, and domestic amplification arises from interactions between external shocks and country-specific vulnerabilities. The reduced-form mapping from observables to sovereign spreads is therefore unlikely to be well captured by a globally linear model, motivating the boosted-tree ensemble, Shapley–Taylor decomposition, and panel local projections that follow.

4 Econometric Framework

Our empirical framework has three steps. First, we compare a broad set of forecasting models and select the specification with the best pseudo-real-time out-of-sample performance (Section 4.1). Second, holding that specification fixed, we use Shapley–Taylor interaction values to decompose fitted forecasts into four economically interpretable transmission channels (Section 4.2). Third, we assess whether this taxonomy receives independent support in panel local projections under two identification strategies (Section 4.3).

4.1 Forecasting Model Comparison

We model sovereign risk as a potentially nonlinear function of market and news-based predictors. For country i on date t, model class m delivers the one-step-ahead forecast

$$y_{i,t+1} = f_m(\boldsymbol{X}_{i,t}, \boldsymbol{\lambda}_i) + \varepsilon_{i,t+1}, \tag{14}$$

where $y_{i,t+1}$ is the standardized sovereign CDS spread, $\boldsymbol{X}_{i,t}$ is the information set available at date t, and $\boldsymbol{\lambda}_i$ denotes country fixed effects. We compare fifteen model classes spanning linear regression, penalized linear models (Hoerl and Kennard, 1970; Tibshirani, 1996; Zou and Hastie, 2005), tree-based ensembles (Breiman, 2001; Friedman, 2001), and neural networks (LeCun et al., 2015).

Table 2: Dominance Benchmarks by Shock Family

	Channel response			
Shock family	φ^{DIR}	φ^{GFC}	φ^{UNC}	φ^{LOC}
Geopolitical (ε^G: GPR)	$\uparrow$	$\downarrow$ [(a)]	$\ll \varphi^{\text{UNC}}_{\text{GEOECO}}$	$\ll \varphi^{\text{DIR}}$
Geoeconomic (ε^E: EPU/TPU)	$\ll \varphi^{\text{DIR}}_{\text{GEO}}$	$\downarrow$ [(b)]	$\uparrow$	$\uparrow$
Mechanism	Default risk	Global pricing (scissors)	Policy-regime uncertainty	Domestic amplification
Canonical reference	Arellano	He–K / G–M	P–V	Threshold / exposure

Notes: Signs refer to short-horizon, panel-mean responses around identified events. $\ll$ denotes a dominance inequality: the channel response is predicted to be substantially smaller in magnitude than the comparator. The framework is a benchmark for signs and relative activation, not a fully estimated structural model. "Geoeconomic" encompasses policy-regime shocks (elections, tariff episodes), which are expected to load most strongly on the Uncertainty channel, and sanctions- or energy-related episodes, which may load relatively more on the GFC and Local channels. [(a)] Under geopolitical shocks, $|\varphi^{\text{GFC}}| < |\varphi^{\text{DIR}}|$ (the GFC is the offsetting, secondary channel). [(b)] Under geoeconomic shocks, $|\varphi^{\text{GFC}}| > |\varphi^{\text{DIR}}|$ (the GFC is the dominant channel). Arellano: Arellano (2008); He–K: He and Krishnamurthy (2013); G–M: Gabaix and Maggiori (2015); P–V: Pástor and Veronesi (2012, 2013). The scissors condition refers to the coexistence of $\partial s^{\text{DIR}}/\partial \varepsilon^G > 0$ and $\partial s^{\text{GFC}}/\partial \varepsilon^G < 0$ under geopolitical shocks (see Section 3).

Forecasts are generated recursively in pseudo-real time using an expanding estimation window and a 28-day exclusion buffer around each train–test split to prevent leakage from overlapping moving-average windows.[10] Model selection is based exclusively on pooled out-of-sample MAE and RMSE. Once the preferred architecture and hyperparameters are selected, they are held fixed. We then re-estimate that specification on the full estimation sample and use the fitted model only for the interpretability exercise in Section 4.2. This separation ensures that model selection is not influenced by the decomposition stage.

To quantify the incremental predictive value of news-based indicators, we compare two predictor sets. The *Markets-Only* benchmark contains the VIX and the U.S. two-year Treasury yield, providing a parsimonious baseline centered on the Global Financial Cycle (Calvo et al., 1996; Rey, 2013; Miranda-Agrippino and Rey, 2020). The augmented specification adds six news-based indicators: geopolitical risk (GPR), economic policy uncertainty (EPU), trade policy uncertainty (TPU), and three domestic macro-political indicators (ECO, INT, POL). For each model class, the change in out-of-sample loss between the two predictor sets measures the marginal predictive contribution of news and reveals which forecasting technologies are most effective in exploiting it. Section 5 reports results; the core finding is that news variables improve all model classes, but gains are substantially larger for nonlinear methods (15–19%) than for linear specifications (5–9%), implying that the predictive content of news operates through interactions and threshold effects (Varian, 2014; Gu et al., 2020).[11]

A methodological caveat applies: the full-sample re-estimation means that the Shapley–Taylor channels have observed the crisis episodes used in the narrative LP. The LP validation therefore tests whether the ML-constructed channels respond to identified shocks in theory-consistent ways that the model was not optimised to produce. The narrative sign-restricted SVAR, which uses raw observables, provides fully independent cross-validation.

4.2 Shapley-Based Channel Decomposition

After selecting the preferred forecasting specification, we freeze its architecture and use it solely for interpretability—the Shapley analysis attributes predictions, not causal effects. Although the target variable is $y_{i,t+1}$, the decomposition is evaluated on the predictor vector observed at date t. Accordingly, the object being decomposed is the fitted forecast $\hat{y}_{i,t+1|t} = f^\star(\boldsymbol{X}_{i,t})$, and all channel contributions are indexed by the information date t.

Economic logic of the four channels. We partition the predictors into four blocks based on their economic role: a *global financial cycle* block $\mathcal{G} = \{\text{VIX}, \text{US2Y}\}$ containing the two variables common to all 42 countries; an

[10] Hyperparameters are selected by cross-validation within the training sample and held fixed during out-of-sample evaluation. The evaluation period runs from February 2021 to July 2025, so all four crisis episodes are assessed out of sample.

[11] The preferred specification is the Multilayer Random Forest (Two Stages), which combines near-frontier predictive performance with a substantive architectural advantage: the first stage separates advanced-economy from emerging-market panels, and the second refines predictions with region-specific dynamics

economic and trade uncertainty block $\mathcal{U} = \{\text{EPU}, \text{TPU}\}$; a *domestic macro-political* block $\mathcal{L} = \{\text{ECO}, \text{INT}, \text{POL}\}$; and a *regional structure* block $\mathcal{R} = \{\text{REG}\}$, where REG denotes regional group indicators that encode geographic, institutional, and income-group heterogeneity.[12] We absorb $\mathcal{R}$ into the Local channel rather than introducing a fifth channel because regional structure operates as a conditioning factor that modulates the domestic amplification of external shocks, not as an independent transmission mechanism.[13]

These blocks motivate four channels, each isolating a distinct transmission mechanism. The *Direct channel* measures a driver's own contribution orthogonal to all other predictors. The *GFC channel* aggregates interactions with VIX and US2Y, capturing transmission through the global risk-taking and monetary environment (Rey, 2013; Miranda-Agrippino and Rey, 2020). The *Uncertainty channel* captures cross-reinforcement among policy-uncertainty measures (Bloom, 2009; Baker et al., 2016; Caldara et al., 2020). The *Local channel* collects interactions with domestic sentiment, political tensions, and regional structure—daily-frequency proxies for the core determinants of sovereign creditworthiness (growth prospects, interest rate–growth differential, institutional quality)—capturing the degree to which local vulnerabilities amplify or absorb external shocks.

Formal decomposition. For each observation we compute pairwise Shapley–Taylor interaction values (Shapley, 1953; Lundberg and Lee, 2017). These attributions satisfy local accuracy, so they decompose the fitted forecast as

$$\hat{y}_{i,t+1|t} \;=\; \varphi_0 \;+\; \sum_{j=1}^{M} \varphi_{jj,i,t} \;+\; \sum_{j<k} \varphi_{jk,i,t}\,, \tag{15}$$

where φ_0 is the baseline prediction, $\varphi_{jj,i,t}$ the main effect of predictor j, and $\varphi_{jk,i,t}$ the pairwise interaction between predictors j and k.[14] The pairwise structure is essential for the channel decomposition that follows. Standard first-order Shapley values assign a single aggregate attribution to each predictor, conflating a variable's own contribution with its interactions with other predictors. The Shapley–Taylor extension separates the main effect φ_{jj} from each pairwise interaction φ_{jk}, which is what allows us to distinguish a driver's *direct* effect on sovereign risk from its *indirect* transmission through the global financial cycle, uncertainty, or domestic conditions. Without this second-order decomposition, the four-channel partition would not be identified. For any driver Z, we partition its total contribution into the four channels:

$$\underbrace{\varphi_{i,t}^{Z,\text{tot}}}_{\text{total}} \;=\; \underbrace{\varphi_{ZZ,i,t}}_{\varphi_{i,t}^{Z,\text{dir}}} \;+\; \underbrace{\sum_{k\,\in\,\mathcal{G}} \varphi_{Zk,i,t}}_{\varphi_{i,t}^{Z,\text{GFC}}} \;+\; \underbrace{\sum_{\substack{k\,\in\,\mathcal{U}\\ k\,\neq\,Z}} \varphi_{Zk,i,t}}_{\varphi_{i,t}^{Z,\text{UNC}}} \;+\; \underbrace{\sum_{\substack{k\,\in\,\mathcal{L}\,\cup\,\mathcal{R}\\ k\,\neq\,Z}} \varphi_{Zk,i,t}}_{\varphi_{i,t}^{Z,\text{LOC}}} \tag{16}$$

By construction, these four terms sum to $\varphi_{i,t}^{Z,\text{tot}}$ for every observation.[15] Section 6 uses these channel series to characterize how the transmission of different shock types varies over time, across countries, and across the cross-sectional distribution of sovereign risk.

4.3 Econometric Validation: Local Projections with Identified Shocks

The Shapley decomposition yields constructed channel series, but it does not by itself establish that these series capture economically meaningful transmission mechanisms. To assess whether the proposed taxonomy receives independent support, we estimate panel local projections under two identification strategies. Section 7 reports the results.

[12] Regional indicators enter the forecasting model as categorical features and generate Shapley interactions with all other predictors.

[13] Figure 4 supports this design choice: under the log-distance specification of equation (20), the Direct channel decays with distance to the conflict epicenter ($R^2 = 0.35$ for Russia–Ukraine, $p < 0.001$; $R^2 = 0.22$ for Hamas–Israel, $p < 0.01$), while the regional component of the Local channel shows no systematic relationship with distance ($R^2 = 0.003$, $p = 0.72$; Appendix Figure 10).

[14] Shapley values are the unique attribution satisfying local accuracy (attributions sum to the prediction) and consistency (if a feature's marginal impact increases, its attribution cannot decrease). In practice we compute pairwise interaction values using TreeSHAP. The algorithm produces a negligible numerical residual $\eta_{i,t}$ with median magnitude below $10^{-4}\sigma$; we treat the decomposition as exact throughout.

[15] When $Z \in \mathcal{U}$, the Uncertainty channel contains only the single cross-interaction between EPU and TPU. When $Z \in \mathcal{L}$, the Local channel excludes Z's own main effect, which enters the Direct channel. Regional indicators in $\mathcal{R}$ contribute to the Local channel for all drivers Z.

4.3.1 Panel Local Projection Specification

For each horizon $h = 0, 1, \ldots, 90$ days, we estimate:

$$\Delta_h Y^{ch}_{i,t+h} \;=\; \alpha^h_i \;+\; \beta^h S_{i,t} \;+\; \sum_{p=1}^{5} \gamma^h_p Y^{ch}_{i,t-p} \;+\; \delta^{h\prime} \mathbf{X}_{i,t} \;+\; u^h_{i,t+h}\,, \tag{17}$$

where $\Delta_h Y^{ch}_{i,t+h} \equiv Y^{ch}_{i,t+h} - Y^{ch}_{i,t-1}$ is the cumulative change in the outcome; Y^{ch} denotes either the raw standardized CDS spread or one of the four channel series $\{\varphi^{\mathrm{dir}}, \varphi^{\mathrm{GFC}}, \varphi^{\mathrm{UNC}}, \varphi^{\mathrm{LOC}}\}$; α^h_i are country fixed effects; $\mathbf{X}_{i,t}$ collects contemporaneous controls; and $S_{i,t}$ is the identified shock. This generic notation accommodates both identification strategies: under Strategy I, $S_{i,t}$ is country-specific; under Strategy II, it is common across countries for a given event date. Standard errors use the Driscoll and Kraay (1998) estimator, robust to heteroskedasticity, serial correlation up to $h+1$ lags, and cross-sectional dependence. We report 68% and 90% confidence bands throughout. The estimation covers 75,650 country–day observations.

4.3.2 Strategy I: Full-Sample Innovation Identification

We identify shocks as innovations from country-specific AR(5) processes:

$$Z_{i,t} \;=\; \mu_i + \sum_{p=1}^{5} \rho_{i,p}\, Z_{i,t-p} + \varepsilon^Z_{i,t}\,. \tag{18}$$

The standardized residual $S^Z_{i,t} = \hat{\varepsilon}^Z_{i,t} \,/\, \hat{\sigma}_{\varepsilon,i}$ serves as the shock in equation (17).[16] This strategy estimates the average marginal response to a one-standard-deviation country-level innovation, pooling across the entire sample. Its strength is broad coverage. Its limitation is that it averages over routine and crisis periods alike, which attenuates episode-specific transmission patterns—an attenuation that the scissors mechanism makes especially severe, since the offsetting Direct and GFC channels tend to cancel in aggregate.

4.3.3 Strategy II: Narrative Identification

We construct event dummies equal to one in a ± 3-day window around each of four dated events:

$$D^e_t \;=\; \mathbb{1}\big(|t - t_e| \le 3\big), \qquad e \in \{E1, E2, E3, E4\}\,, \tag{19}$$

where t_{E1} = 24 February 2022 (Russia–Ukraine invasion), t_{E2} = 7 October 2023 (Hamas–Israel attack), t_{E3} = 5 November 2024 (U.S. presidential election), and t_{E4} = 2 April 2025 ("Liberation Day" tariff announcement). A separate local projection is estimated for each event, with $S_{i,t} = D^e_t$. Each event is paired with its primary news indicator: GPR for $E1$ and $E2$, EPU for $E3$, TPU for $E4$.

The narrative approach offers two advantages. First, event timing is plausibly exogenous to day-to-day sovereign CDS movements. Second, the taxonomy generates event-specific predictions—geopolitical events should primarily activate the Direct channel, whereas geoeconomic events should largely bypass it—that the narrative local projections evaluate directly.

5 The Layered Structure of Sovereign Risk

This section establishes three facts that motivate the transmission analysis. First, news-based indicators add material predictive content beyond global financial variables, but primarily through nonlinear models. Second, in the selected specification, the largest average contributions come from global financial conditions, while domestic sentiment differentiates countries within that common environment. Third, geopolitical and policy-uncertainty indicators matter primarily through nonlinear interactions rather than through large average main effects. Together, these results motivate the four-channel decomposition developed in Section 6.

5.1 Model Selection

All predictors and the dependent variable are constructed as 28-day moving averages, so adjacent observations mechanically share underlying daily information. To avoid leakage, we evaluate all model classes in a recursive

[16] Qualitatively we obtain similar results under panel AR(5) specifications with country fixed effects or from simple first differences $\Delta Z_{i,t}$. See Appendix E.

pseudo-real-time design with a 28-day exclusion buffer around each forecast origin. This buffer ensures that no training observation overlaps with the test target, so any improvement from adding news-based indicators can be interpreted conservatively as genuine incremental predictive content.[17] Table 3 reports out-of-sample forecast accuracy under the

Table 3: **News Indicators Improve Nonlinear Forecasts of Sovereign Risk**

Machine Learning Model	Benchmark Market Only		News Extended Market + News		Difference (News Extended vs Benchmark) RMSE		MAE	
	RMSE	MAE	RMSE	MAE	Diff	% Var	Diff	% Var
Linear Regression	1.09	0.92	1.03	0.86	-0.06	-5.7%	-0.06	-6.4%
Lasso	1.03	0.87	0.95	0.80	-0.08	-8.0%	-0.08	-8.7%
Ridge	1.09	0.92	1.02	0.85	-0.07	-6.4%	-0.07	-7.2%
Elastic Net	1.04	0.88	0.95	0.79	-0.09	-8.8%	-0.08	-9.5%
Quantile Linear Regression	0.93	0.77	0.90	0.74	-0.03	-3.7%	-0.03	-4.5%
Principal Components (PCR)	1.09	0.92	0.99	0.82	-0.09	-8.6%	-0.09	-10.3%
Factor Models (FAR)	0.92	0.75	0.88	0.72	-0.04	-4.4%	-0.03	-3.6%
Gradient Boosting	1.00	0.81	0.85	0.67	-0.14	-14.2%	-0.14	-17.9%
Bagging	1.03	0.83	0.84	0.67	-0.19	-18.3%	-0.17	-20.1%
Random Forest	0.99	0.75	0.82	0.65	-0.17	-17.1%	-0.11	-14.3%
Extremely Randomized Trees	0.98	0.74	0.80	0.60	-0.18	-18.5%	-0.14	-19.0%
Multilayer Random Forest (1S)	0.97	0.75	0.84	0.62	-0.13	-13.6%	-0.13	-17.2%
Multilayer Random Forest (2S)	1.01	0.77	0.85	0.65	-0.16	-15.9%	-0.12	-15.5%
Shallow CNN	1.05	0.84	0.97	0.77	-0.09	-8.3%	-0.07	-8.8%
Deep CNN	1.04	0.77	0.89	0.66	-0.15	-14.1%	-0.11	-14.3%

Notes: Out-of-sample RMSE and MAE for one-day-ahead forecasts of standardized sovereign CDS spreads under two information sets: Markets-Only (VIX and U.S. two-year Treasury yield) and Markets+News (augmented with GPR, EPU, TPU, ECO, INT, POL). All variables are 28-day moving averages. Models are estimated recursively with a 28-day exclusion buffer. "Diff" denotes the change relative to the Markets-Only benchmark; negative values indicate lower forecast loss.

Markets-Only and Markets+News information sets. Three findings stand out. Flexible nonlinear methods outperform linear specifications, with Extremely Randomized Trees achieving the lowest errors (MAE = 0.60). News variables improve every model class, but the gains are substantially larger for nonlinear methods (15–19% RMSE reduction) than for linear specifications (5–9%), indicating that much of the predictive content of news enters through interactions and threshold effects that linear frameworks understate (Varian, 2014; Gu et al., 2020). The magnitude of improvement rises monotonically with model flexibility (Mullainathan and Spiess, 2017; Athey and Imbens, 2019).

Although Extremely Randomized Trees delivers the lowest raw forecast loss, the remainder of the paper uses the Multilayer Random Forest (Two Stages) as the baseline for interpretation. Its performance remains close to the frontier, while its architecture preserves the advanced-economy/emerging-market split and region-specific structure that are central to the subsequent decomposition.[18]

5.2 Sovereign Risk Drivers

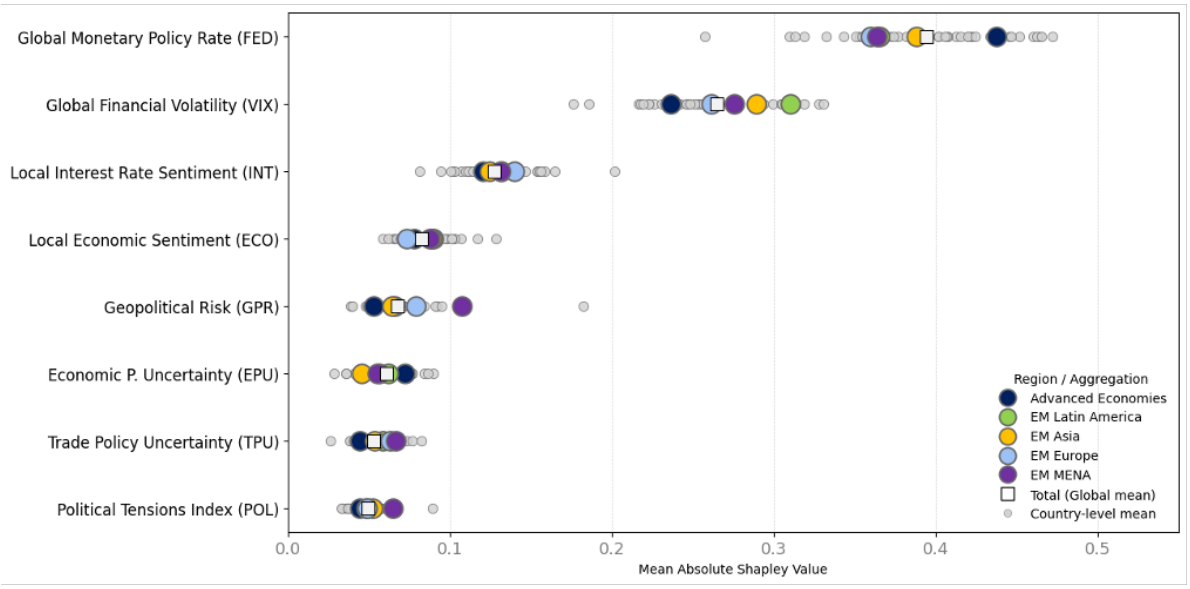
Figure 1 summarizes the fitted model's feature importance using mean absolute Shapley values at the global, regional, and country levels.

The Shapley ranking points to a layered structure. The first layer is global. The U.S. two-year yield and the VIX are the largest contributors in almost every region and for nearly every country, with mean absolute Shapley values exceeding those of any other variable— consistent with the view that sovereign spreads are anchored by common financial conditions (Rey, 2013; Miranda-Agrippino and Rey, 2020), with the VIX corroborating its role as a barometer of global risk aversion whose spikes precede capital-flow reversals and spread widening (Gelos et al., 2011).

[17] A simple AR(1) benchmark performs worse than the Markets-Only specification, confirming that the latter provides a stricter baseline.

[18] The Multilayer Random Forest extends the standard Random Forest by stacking two sequential ensemble layers (Breiman, 2001). The first layer is trained separately on advanced- and emerging-economy panels; the second combines intermediate predictions with the full set of covariates.

Figure 1: **Global Financial Conditions Dominate Sovereign Risk: Shapley Feature Importance**



Notes: Mean absolute Shapley values for all predictors over 2018–2025, shown at the global, regional, and country levels. Country-level values are computed by averaging over all observations for each country; the global value is the unweighted mean across country means.

The second layer is domestic. Local interest-rate sentiment (INT) and economic sentiment (ECO) rank just below the global variables but display much greater cross-country dispersion, indicating that country-specific narratives are critical for differentiating sovereign spreads conditional on the global environment (Eichengreen et al., 2021).

The third layer is conditional. Geopolitical and policy-uncertainty indicators—GPR, EPU, and TPU—are smaller on average but become much more prominent in EM Europe and MENA, where exposure to geopolitical and geoeconomic risk is greater (Baker et al., 2016; Caldara et al., 2020), foreshadowing the state-dependent nonlinearities documented below.

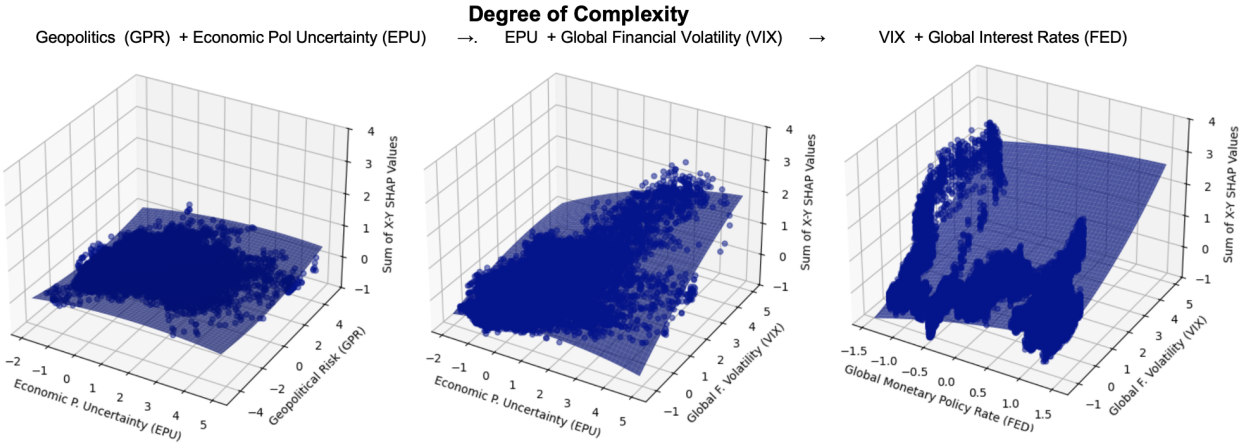
5.3 Nonlinear Interactions and State Dependence

Pairwise Shapley–Taylor interaction values show that global financial conditions are not only the largest main effects but also the dominant sources of nonlinear interaction. The largest and most pervasive interaction terms involve US2Y and the VIX, as well as interactions between US2Y and domestic interest-rate sentiment. By contrast, GPR and EPU rarely appear as large stand-alone drivers across the full sample, but they interact strongly with both global and domestic variables in EM Europe and MENA while remaining much weaker elsewhere. In the fitted model, geopolitical and policy-uncertainty variables matter primarily by changing how sovereign risk responds to the broader financial environment. This pattern provides a direct motivation for the channel grouping introduced in Section 4.2, where the global financial cycle block ($\mathcal{G}$) is defined as the set of variables that generate the strongest and most pervasive interactions.

The interaction surfaces reinforce this state dependence (Figure 2). When global financial conditions are benign, the joint contribution of geopolitical risk and policy uncertainty is relatively muted in most regions. By contrast, uncertainty becomes materially more important in high-volatility states (Bloom, 2009), and the strongest nonlinearities arise when elevated volatility coincides with tight U.S. rates: volatility shocks are powerfully amplified under tight monetary policy (Rey, 2013; Bruno and Shin, 2015) but partially contained under accommodative conditions (Bekaert et al., 2013). Geoeconomic or policy uncertainty, while often manageable in isolation, thus becomes a potent threat when it coincides with financial stress—a finding that motivates the four-channel decomposition developed next.

Cross-country connectedness measures applied to the panel of daily Shapley attributions confirm that global financial variables are the dominant transmitters of cross-country co-movement, while geopolitical shocks generate episodic synchronization that dissipates unless reinforced by tighter financial conditions.

Figure 2: **State Dependence in Sovereign Risk: Two-Factor Interaction Surfaces**



Notes: The figure illustrates the combined impact of key predictors on sovereign CDS spreads for Advanced Economies using two-factor SHAP dependence surfaces. Left panel: joint influence of Geopolitical Risk (GPR) and Economic Policy Uncertainty (EPU)—the surface is nearly flat, indicating low sensitivity when financial conditions are uncontrolled. Centre panel: joint contribution of EPU and Global Financial Volatility (VIX)—high uncertainty in a high-VIX environment is substantially more detrimental than either shock alone (Bloom, 2009). Right panel: joint effect of VIX and global monetary policy (US2Y)—the largest and most nonlinear increases in sovereign risk arise when volatility spikes coincide with tight monetary conditions (Rey, 2013; Bruno and Shin, 2015), but are partially contained under accommodative policy (Bekaert et al., 2013). See Section 5.3 for discussion.

6 Transmission Anatomy of Geopolitical and Geoeconomic Shocks

Applying the four-channel decomposition, we begin by inspecting the raw attribution dynamics around each episode. Country-level Shapley paths around the Russia–Ukraine invasion (Appendix Figure 9) are consistent with staged propagation from geopolitical news to energy-related macro pressures and then to interactions with global financial conditions; the Hamas–Israel war appears more regionally contained; and the U.S. election and tariff episodes transmit less through the Direct channel and more through policy uncertainty, trade exposure, and global financial conditions.

These raw attributions are informative but combine several mechanisms at once. To separate them, we apply the four-channel Shapley–Taylor decomposition to realized episode data, feeding observations through the estimated model—held fixed at its pre-episode parameters—and tracking daily channel contributions for the 42-country panel.[19] Three features organize the results. First, the configuration of channels differs systematically across geopolitical and geoeconomic episodes. Second, the cross-sectional distribution of each channel follows a distinct geographic or economic pattern. Third, the channels decay at different speeds. The remainder of this section develops these results in turn.

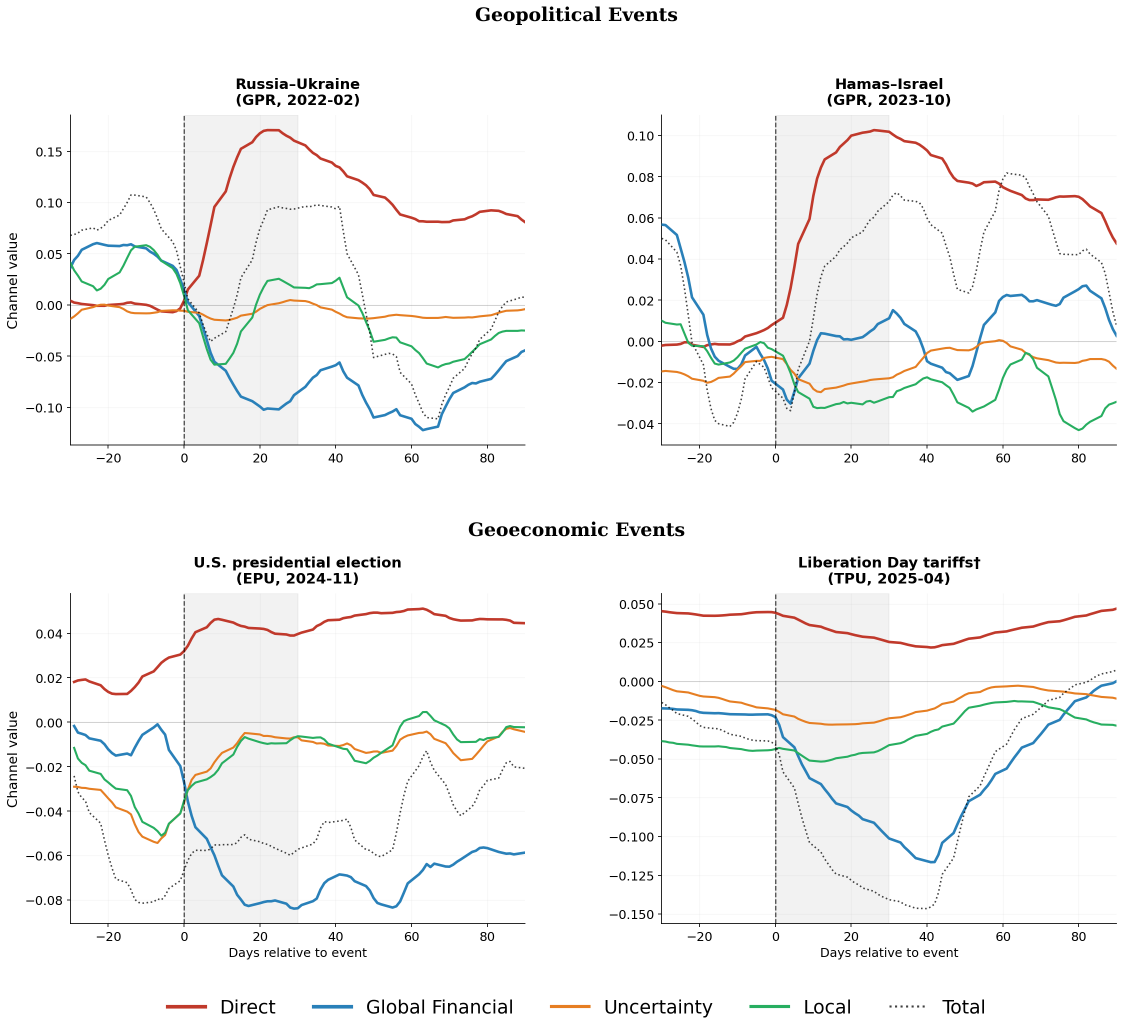
6.1 The Scissors Pattern

Figure 3 plots the cross-sectional average of the four Shapley–Taylor channels across all 42 countries for each episode. Geopolitical episodes generate a scissors pattern between the Direct and GFC channels, whereas geoeconomic episodes display limited direct repricing and transmit mainly through financial and uncertainty channels.

Geopolitical episodes. In panel (a), the Russia–Ukraine invasion opens a clear "scissors": the Direct channel (red) jumps to roughly $+0.10$ within days as conflict exposure reprices default probability (Caldara and Iacoviello, 2022), while the GFC channel (blue) drops to approximately -0.05 as the global financial system partially absorbs the shock through the risk-taking channel (Rey, 2013; Bruno and Shin, 2015). The two move in opposite directions, so the total Shapley contribution is slightly negative at the panel mean ($\varphi_{1m}^{\text{tot}} = -0.07$)—the GFC and Local channels more than offset the Direct component, implying that bystander sovereigns experience net spread compression even as conflict-exposed countries are sharply repriced. Panel (b) shows a similar but more muted pattern for Hamas–Israel: the Direct channel spikes, but the GFC response is much smaller in magnitude, consistent with a conflict whose spillovers remain more regionally contained. In both geopolitical episodes, the Uncertainty and Local channels are comparatively muted in cross-country mean terms at impact, even though country-level responses can be substantial for exposed sovereigns.

[19] Country-level Shapley attribution dynamics for each episode are reported in Appendix D.

Figure 3: **Geopolitical vs. Geoeconomic Shocks: The Scissors Pattern**



Notes: Each panel plots the cross-sectional average (42 countries) of the four Shapley–Taylor transmission channels for the indicator most directly associated with each episode: GPR for geopolitical shocks (top row); EPU/TPU for geoeconomic shocks (bottom row). Direct (red): own Shapley–Taylor term ϕ^{dir}. Global Financial (blue): interactions with VIX and US2Y. Uncertainty (orange): interactions with EPU and TPU. Local (green): interactions with ECO, INT, POL, and REG. Dotted black line: total Shapley contribution. Vertical dashed lines mark the event date. All series are 7-day moving averages with a trailing smoother.

Geoeconomic episodes: transmission through financial and uncertainty channels. The geoeconomic episodes invert the picture. In panel (c), the Direct channel remains economically negligible throughout—the U.S. election does not reprice default probability. Instead, the GFC channel begins declining three months *before* the event as markets price in looser monetary policy (Miranda-Agrippino and Rey, 2020), and continues falling to roughly -0.08 by December, while the Uncertainty channel also contributes positively. Panel (d) is the sharpest illustration: the GFC plunges from near zero to approximately -0.12 within two weeks of the April 2025 tariff announcement. The Direct channel remains near zero—the tariff shock shifts expected monetary policy and risk appetite, not default probability.

Quantitative summary. Table 4 summarizes these patterns at one- and three-month horizons. The Direct channel exceeds the GFC at the panel level in both geopolitical episodes ($\varphi^{\text{dir}}_{1\text{m}} = +0.11$ vs. $\varphi^{\text{GFC}}_{1\text{m}} = -0.11$ for Russia–Ukraine; $+0.08$ vs. -0.01 for Hamas–Israel); the GFC dominates the Direct channel in both geoeconomic episodes. Taken together, geopolitical shocks load primarily on the Direct channel while provoking a partial GFC offset; geoeconomic shocks largely bypass the Direct channel and transmit through financial and uncertainty channels. Appendix Tables 8 and 9 report the complete country-level decomposition for all 42 economies, confirming that these patterns hold broadly across the panel.

6.2 Country Heterogeneity

The scissors pattern describes the average response across countries. At the country level, transmission depends on the relative importance of two additional forces: exposure to the global financial cycle (φ^{GFC}) and the strength of domestic amplification (φ^{LOC}). By construction, these channels depend on different sets of predictors and need not move together.

Table 4: Four-Channel Decomposition across Crisis Episodes: Impact and Persistence

		Short-Term Impact ($\Delta_{1\mathrm{m}}$)					Persistence ($\Delta_{3\mathrm{m}}$)				
	Nature of Shock	φ^{dir}	φ^{GFC}	φ^{UNC}	φ^{LOC}	φ^{tot}	φ^{dir}	φ^{GFC}	φ^{UNC}	φ^{LOC}	φ^{tot}
Panel A : Russia–Ukraine (GPR, Feb. 2022)											
Panel avg.	**Geopolitics** (Panel A)	+.10	-.12	.00	-.03	-.04	+.07	-.09	.00	-.03	-.05
Ukraine		+.13	+.66	-.02	+.48	+1.25	+.13	+.66	-.02	+.48	+1.25
Poland		+.15	-.10	+.04	+.06	+.14	+.11	-.02	+.04	+.06	+.20
Germany		+.25	-.24	+.02	+.06	+.09	+.15	-.11	+.01	+.05	+.10
Panel B : Hamas–Israel (GPR, Oct. 2023)											
Panel avg.	**Geopolitics** (Panel B)	+.08	-.01	+.01	-.02	+.06	+.06	-.02	.00	-.01	+.02
Israel		+.21	+.32	+.19	+.30	+1.02	+.21	+.34	+.16	+.33	+1.04
Egypt		+.27	+.08	+.08	+.07	+.51	+.21	-.07	+.01	-.06	+.09
Germany		+.09	-.15	-.02	+.02	-.06	+.09	-.18	-.04	.00	-.13
Panel C : U.S. Presidential Election (EPU, Nov. 2024)											
Panel avg.	**Economic Uncertainty** (Panel C)	+.03	-.06	+.03	+.01	+.01	+.04	-.06	+.03	+.01	+.01
USA		+.04	-.08	+.01	+.02	-.01	+.06	-.10	-.02	+.02	-.04
Mexico		+.03	-.03	+.01	.00	+.01	+.05	-.06	-.03	-.04	-.08
Germany		+.04	.00	+.06	+.04	+.15	+.05	-.04	+.05	+.04	+.10
Panel D : U.S. Tariffs / Liberation Day (TPU, Apr. 2025)											
Panel avg.	**Trade Uncertainty** (Panel D)	-.01	-.01	+.01	+.02	+.01	-.01	-.01	+.01	+.02	+.01
USA		-.01	-.02	-.03	+.10	+.03	-.01	-.02	-.03	+.10	+.03
India		-.02	-.08	-.03	+.02	-.10	-.02	-.08	-.03	+.02	-.10
Germany		-.01	-.05	.00	-.01	-.07	-.01	-.05	.00	-.01	-.07

Notes: Each cell reports $\Delta \equiv \bar{\varphi}^{\mathrm{post}} - \bar{\varphi}^{\mathrm{pre}}$ in standardized CDS units. Short-Term Impact: post-window = 1 month after event onset. Persistence: post-window = 3 months after event onset. Pre-event window = 3 months in both cases. Cell shading: $\Delta > +0.05$, $0.005 < \Delta \leq +0.05$, $|\Delta| \leq 0.005$, $-0.05 \leq \Delta < -0.005$, $\Delta < -0.05$. Values in σ-CDS units. *Panel avg.* = unweighted cross-country mean (42 countries). Germany appears in all panels as a common benchmark. 28-day moving averages applied to all series. φ^{LOC} includes interactions with local macro-political variables and regional indicators (ECO, INT, POL, REG; see equation 16).

Table 4 illustrates three distinct transmission modes within the geopolitical episodes. *Ukraine's* response is dominated by the GFC channel: $\varphi^{\mathrm{GFC}}_{1\mathrm{m}} = +0.67$, the largest single-channel value in the sample. Sanctions and market exclusion converted a geopolitical shock into a severe financial-intermediation event (Fernández-Villaverde et al., 2024), compounded by collapsing domestic sentiment ($\varphi^{\mathrm{LOC}} = +0.62$). *Israel* presents a different balance: $\varphi^{\mathrm{GFC}} = +0.35$ and $\varphi^{\mathrm{LOC}} = +0.46$ contribute with broadly similar magnitudes, consistent with the joint activation of global risk-appetite sensitivity and domestic amplifiers (Gorodnichenko et al., 2025). The contrast is economically meaningful: Ukraine's crisis is primarily one of financial exclusion; Israel's reflects domestic strain in equal measure. *Poland* illustrates partial GFC compensation: the Direct-channel rise of $+0.18$ is materially offset by $\varphi^{\mathrm{GFC}} = -0.15$, consistent with the idea that deep integration into European capital markets can cushion part of the initial shock.

The geoeconomic episodes reveal a different asymmetry. At the panel level, the tariff shock compresses spreads through the GFC as easier expected global financial conditions benefit most sovereigns. But the United States—the originator—is a partial exception. Its GFC contribution is a muted -0.02, while its Local channel registers $+0.10$: the policy uncertainty the originator created raises its own risk premia through the interaction of TPU with domestic sentiment and political tensions. The originator does not receive the same financial offset as the rest of the panel—consistent with evidence that geoeconomic instruments generate domestic risk premia (Clayton et al., 2025) and with the theoretical prediction of Pástor and Veronesi (2013) that uncertainty-generated risk premia are largest for the source country.

6.3 Differential Persistence

If the Direct channel is associated with more persistent sovereign-specific repricing, while the GFC channel reflects more transitory financial adjustment, the two should fade at different speeds. Table 4 is consistent with that prediction. The Direct channel retains roughly two-thirds of its one-month value at three months in both geopolitical episodes, while the GFC decays faster and can reverse sign (Bekaert et al., 2013; Bloom, 2009).

Egypt illustrates this temporal separation. At impact, all four channels contribute positively ($\varphi^{\text{tot}}_{1\text{m}} = +0.56$). By three months, the GFC has reversed ($+0.07 \to -0.08$) and the Local channel has followed, pulling the total to $+0.10$. But the Direct channel barely moved ($+0.27 \to +0.21$): Egypt's fundamental exposure persisted even after the financial amplifiers faded.

6.4 Cross-Sectional Fingerprints

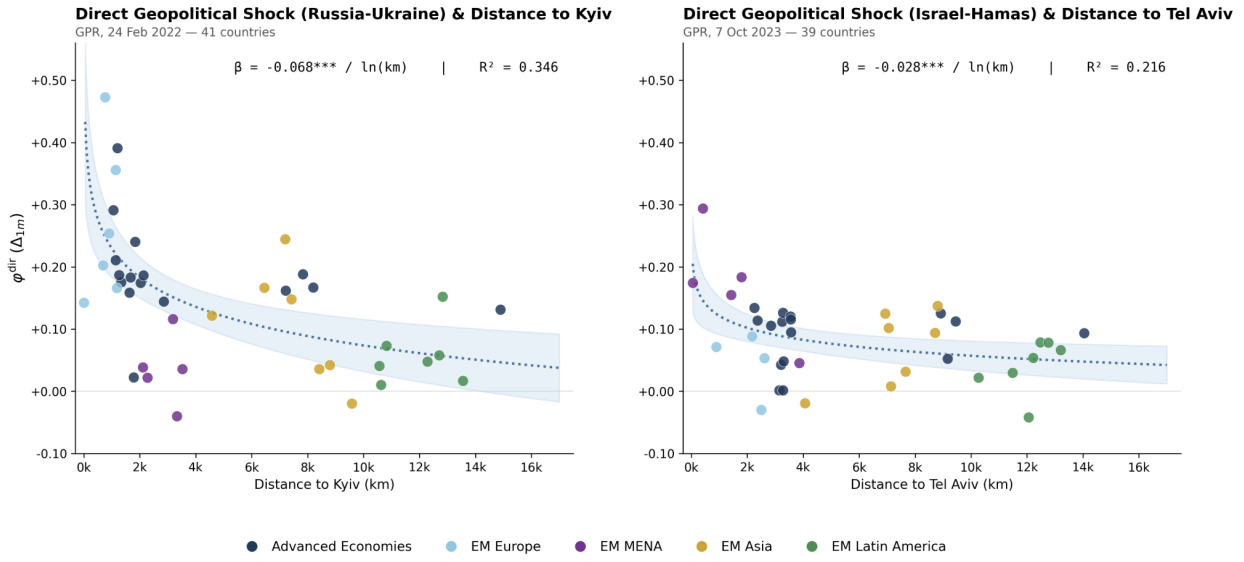
The preceding subsections established that different shock types activate different channels over time. We now show that each activated channel also exhibits a distinct *cross-sectional* distribution, completing the transmission taxonomy.

Geopolitical shocks: gravity-like distance decay. If geopolitical shocks reprice default probability in proportion to conflict exposure, countries closer to the conflict should register larger Direct-channel effects—a gravity-style prediction familiar from the trade literature (Head and Mayer, 2014) and from work linking geographic proximity to financial contagion (Glick and Rose, 1999). Figure 4 evaluates this prediction by regressing the short-term direct effect on log great-circle distance:

$$\varphi_i^{\text{dir}} \;=\; \alpha \;+\; \beta \,\ln d_i \;+\; u_i \,, \tag{20}$$

where d_i is distance to Kyiv (left panel) or Tel Aviv (right panel).

Figure 4: **Direct Geopolitical Effects Decline with Distance from Conflict**



Notes: Each panel plots the one-month Direct-channel response φ^{dir} ($\Delta_{1\text{m}}$) against great-circle distance (km) from the national capital to the conflict epicenter (Kyiv, left panel; Tel Aviv, right panel). The decomposition follows equation (16). Left panel: 41 sovereigns (Russia–Ukraine, 24 February 2022; Ukraine at $d = 0$ plotted but excluded from log-distance OLS). Right panel: 39 sovereigns (Hamas–Israel, 7 October 2023). Dotted lines: OLS fitted values from equation (20); shaded bands: 95% confidence intervals. All values in within-country standard deviations of CDS spreads (σ-CDS). Jordan excluded due to insufficient pre-2023 CDS data. $^{***}\,p < 0.01$, $^{**}\,p < 0.05$, $^{*}\,p < 0.10$.

The Russia–Ukraine shock produces a steep, concave distance gradient ($\hat{\beta} = -0.068^{***}$ per log-km; $R^2 = 0.35$, $n = 40$). Countries within 1,200 km of Kyiv—Czech Republic ($+0.32$), Germany ($+0.28$), Austria ($+0.23$)—register Direct-channel responses two to three times the cross-sectional mean ($+0.11$); most Latin American sovereigns beyond 10,000 km fall below $+0.08$, though notable exceptions such as Argentina ($+0.15$) and Australia ($+0.13$) suggest that non-geographic linkages also play a role. The Hamas–Israel shock reveals a different geography: the slope is roughly half as steep ($\hat{\beta} = -0.028^{***}$; $R^2 = 0.22$, $n = 39$), with the four largest responses all belonging to MENA

sovereigns within 2,000 km—Egypt ($+0.27$), Qatar ($+0.20$), Israel ($+0.19$), Saudi Arabia ($+0.14$). Both coefficients are significant at 1% under the log specification, which captures the concave decay pattern—steep attenuation near the epicenter, flattening at greater distances—better than the linear alternative. The cross-sectional footprint of the Direct channel is gravity-like in both cases, but much more pronounced for the more systemic of the two geopolitical shocks.

Appendix Figure 10 confirms that the regional component of the Local channel shows no distance relationship ($R^2 = 0.003$, $p = 0.72$), validating its absorption into LOC rather than creating a fifth channel.

Policy-uncertainty shocks: global and simultaneous. The election episode leaves a different signature. Its most distinctive feature is the broad-based activation of the Uncertainty channel φ^{UNC}, which turns positive for a large share of countries at roughly the same time. At the panel level it registers $+0.03$ at both horizons, with a cross-country interquartile range of $[+0.01, +0.06]$; the appendix tables confirm positive values for France, Brazil, China, Germany ($+0.06$), and most other sovereigns.[20] This broad-based activation is much less visible in the other three episodes. In the tariff shock, φ^{UNC} is comparatively small; in the geopolitical episodes, it is limited mainly to the most exposed countries. The election episode therefore stands out not only by which channel activates, but also by how broadly and synchronously it does so—consistent with models in which political uncertainty generates correlated risk premia across all assets simultaneously (Pástor and Veronesi, 2012; Baker et al., 2016).

A three-way fingerprint. Together, these results complete the cross-sectional dimension of the taxonomy. Geopolitical shocks activate φ^{dir} with gravity-like structure. Trade shocks operate through φ^{GFC} in proportion to GFC exposure. Policy-uncertainty shocks activate φ^{UNC} broadly across countries. Each shock type has a distinct first-responder channel *and* a distinct cross-sectional distribution. This combination provides a useful empirical diagnostic for distinguishing the nature of a shock in real time.

In sum, the transmission anatomy rests on four findings: (i) geopolitical shocks display a scissors pattern while geoeconomic shocks load on financial and uncertainty channels; (ii) country-level heterogeneity depends on GFC exposure and domestic amplification; (iii) channels decay at different speeds; and (iv) each shock type leaves a distinct cross-sectional fingerprint.

7 Econometric Validation via Local Projections

This section asks whether the four-channel taxonomy developed in Section 6 receives separate econometric support in panel local projections under the two identification strategies introduced in Section 4.3. The logic is simple. If the proposed channels capture economically meaningful differences in transmission, then they should respond differently to identified shocks even outside the machine-learning decomposition stage. Table 5 summarizes the main evidence at the 30-day horizon. The left panel reports the full-sample innovation LP; the right panel reports the narrative LP.

7.1 Full-Sample Innovation LP

The left panel of Table 5 reports the full-sample innovation LP. For GPR, EPU, and TPU innovations, the 30-day responses are small in magnitude and statistically indistinguishable from zero across all four channels and all three control sets. No coefficient exceeds $\pm 0.005\sigma$ in absolute value, and none is significant at the 10% level.[21]

This null is consistent with the full-sample design averaging over many routine news days on which channel responses are weak or offsetting. Three results favour attenuation over absence of transmission: the narrative LP recovers large and significant channel responses for precisely the episodes where the taxonomy predicts they should be strongest; the placebo falsification confirms that these episodes are exceptional relative to random dates; and a narrative sign-restricted SVAR estimated on raw observables recovers the scissors pattern with theory-consistent signs (section 8.1 and further details in Appendix F).

7.2 Narrative LP: Four Crisis Episodes

In sharp contrast to the full-sample null, the narrative LP detects large and statistically significant aggregate CDS responses: Russia–Ukraine generates $+0.54\sigma$ at $h = 30$ ($p < 0.01$), rising to $\sim 1.5\sigma$ by $h = 90$; Hamas–Israel produces

[20] The cross-country standard deviation of $\varphi^{\mathrm{UNC}}_{1\mathrm{m}}$ is 0.02 for the election episode—roughly one-third of the corresponding dispersion for φ^{dir} in the Russia–Ukraine episode (0.07). The narrow spread confirms that the Uncertainty channel activates broadly rather than selectively.

[21] Despite the aggregate null, channel-level IRFs produce qualitative sign patterns consistent with four of the taxonomy's sign conditions at the 68% level, though none reaches conventional significance. See Appendix E for channel-level detail at $h = 5, 30, 60$.

Table 5: **Local-Projection Validation of the Channel Taxonomy**

		Full-Sample Innovation LP			Narrative LP		
		(A) Baseline	(B) Global	(C) Extended	(A) Baseline	(B) Global	(C) Extended
Episode	**Channel**	$\hat{\beta}$ Sig.	$\hat{\beta}$ Sig.	$\hat{\beta}$ Sig.	$\hat{\beta}$ Sig.	$\hat{\beta}$ Sig.	$\hat{\beta}$ Sig.
Russia–Ukraine GPR, Feb. 2022 *Geopolitical*	φ^{dir}	$+.005$ –	$+.005$ –	$+.005$ –	$+.122$ ***	$+.121$ ***	$+.125$ ***
	φ^{GFC}	$-.001$ –	$-.001$ –	$-.001$ –	$-.102$ ***	$-.104$ ***	$-.102$ ***
	φ^{UNC}	$-.000$ –	$-.000$ –	$-.000$ –	$-.009$ *	$-.010$ *	$-.009$ *
	φ^{LOC}	$-.004$ –	$-.004$ –	$-.004$ –	$-.018$ –	$-.023$ –	$-.023$ –
Hamas–Israel GPR, Oct. 2023 *Geopolitical*	φ^{dir}	$+.005$ –	$+.005$ –	$+.005$ –	$+.071$ ***	$+.072$ ***	$+.068$ ***
	φ^{GFC}	$-.001$ –	$-.001$ –	$-.001$ –	$+.003$ –	$+.000$ –	$-.000$ –
	φ^{UNC}	$-.000$ –	$-.000$ –	$-.000$ –	$+.006$ –	$+.006$ –	$+.006$ –
	φ^{LOC}	$-.004$ –	$-.004$ –	$-.004$ –	$-.017$ –	$-.018$ –	$-.019$ –
U.S. Election EPU, Nov. 2024 *Geoeconomic*	φ^{dir}	$+.002$ –	$+.002$ –	$+.003$ –	$+.013$ ***	$+.013$ ***	$+.015$ ***
	φ^{GFC}	$+.002$ –	$+.002$ –	$+.004$ –	$-.044$ ***	$-.046$ ***	$-.046$ ***
	φ^{UNC}	$+.001$ –	$+.001$ –	$+.001$ –	$+.021$ ***	$+.021$ ***	$+.024$ ***
	φ^{LOC}	$+.000$ –	$+.000$ –	$+.000$ –	$+.016$ ***	$+.015$ ***	$+.017$ ***
"Liberation Day" TPU, Apr. 2025 *Geoeconomic*	φ^{dir}	$-.001$ –	$-.001$ –	$-.002$ –	$-.008$ ***	$-.009$ ***	$-.016$ ***
	φ^{GFC}	$-.001$ –	$-.001$ –	$-.001$ –	$-.015$ –	$-.019$ *	$-.028$ ***
	φ^{UNC}	$+.001$ –	$+.001$ –	$+.001$ –	$+.026$ ***	$+.026$ ***	$+.034$ ***
	φ^{LOC}	$+.000$ –	$+.000$ –	$-.001$ –	$+.034$ **	$+.032$ **	$+.023$ *

Notes: Each cell reports the cumulative impulse response at horizon $h = 30$ days. Coefficients rounded to three decimal places. *Full-sample innovation LP*: panel local projection with country-specific AR(5) innovations; Driscoll–Kraay SE with bandwidth $\max(20, h)$, country and time FE, 5 lags. *Narrative LP*: panel local projection with $\pm$3-day event dummies; Driscoll–Kraay SE with country FE. Controls: (A) none; (B) VIX, US2Y; (C) VIX, US2Y, ECO, INT, POL, EPU, TPU. $^{***}\,p < 0.01$, $^{**}\,p < 0.05$, $^{*}\,p < 0.10$; – = not significant at 10%. Cell shading indicates significance level of the Narrative LP estimate: $p < 0.01$, $p < 0.05$, $p < 0.10$. Because the full-sample innovation LP is estimated by shock type rather than by narrative episode, the GPR-based coefficients are repeated in the two geopolitical rows to facilitate comparison with the narrative panel. GPR = Geopolitical Risk, EPU = Economic Policy Uncertainty, TPU = Trade Policy Uncertainty, US2Y = U.S. two-year Treasury yield.

-0.18σ ($p < 0.05$), consistent with flight-to-quality compression for non-MENA sovereigns; the U.S. election raises CDS by $+0.11\sigma$ ($p < 0.10$) with a gradual 60-day build-up; and "Liberation Day" tariffs produce a hump-shaped response peaking at $+0.15\sigma$ ($p < 0.01$) before partially reversing.[22] The contrast with the innovation null resolves an apparent puzzle: geopolitical shocks have large effects on sovereign spreads, but these effects are concentrated in discrete crisis episodes that are washed out when averaged over routine news-innovation days.

Channel decomposition. Figure 5 plots the full horizon profiles for the four channels under the narrative design. Table 5 reports the corresponding cumulative responses at $h = 30$ under the three control sets. For interpretation, specification (B) is our preferred baseline because it conditions on global financial variables while remaining parsimonious; specifications (A) and (C) show that the results are stable to omitting or extending the control set. Four results stand out.

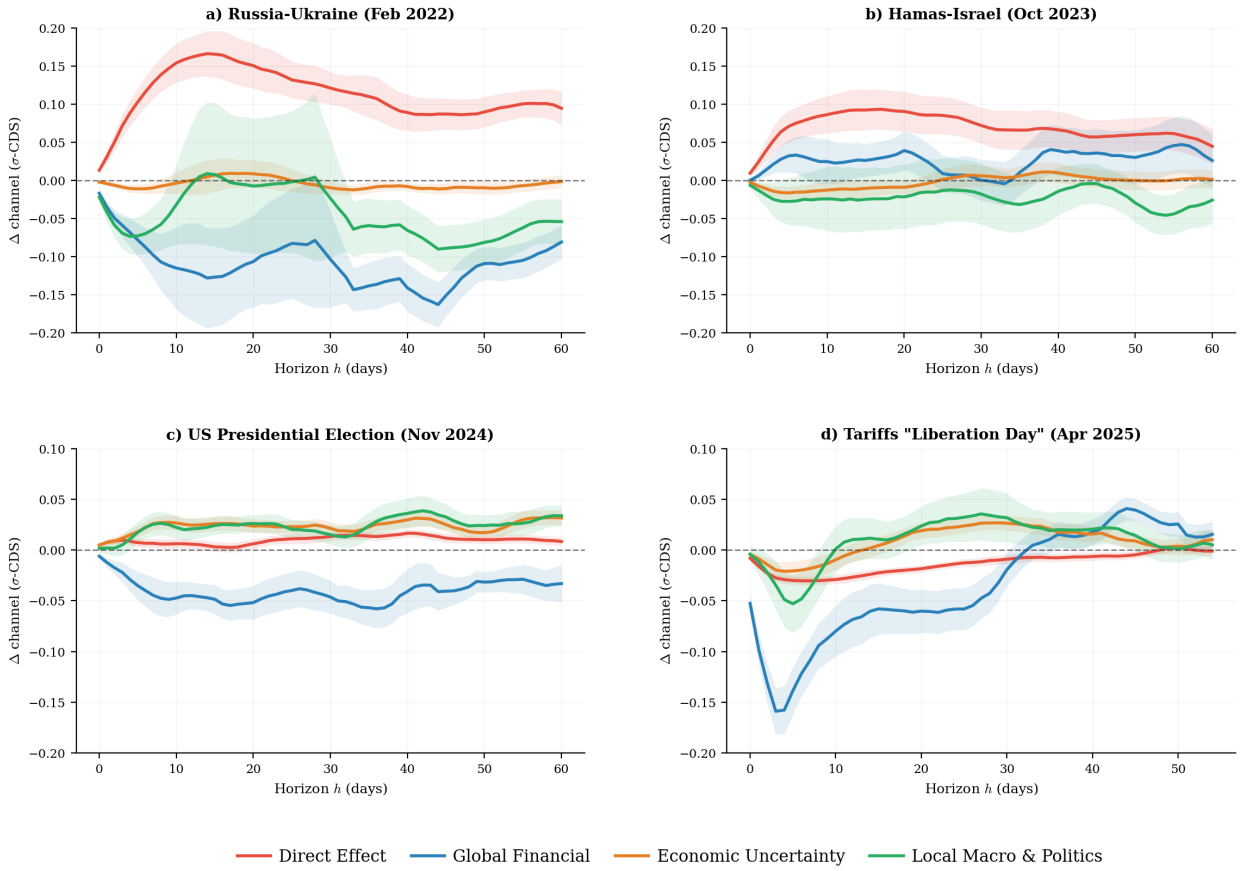
E1—Russia–Ukraine (GPR shock). The Direct channel rises sharply ($\hat{\beta}^{h=30} = +0.121$, $p < 0.01$) while the GFC channel moves in the opposite direction (-0.104, $p < 0.01$), recovering the scissors pattern at the 1% significance level. The horizon profiles in Figure 5 suggest that the Direct and GFC channels decay at different speeds, with the Direct response more persistent.[23]

E2—Hamas–Israel (GPR shock). The Direct channel activates ($+0.072$, $p < 0.01$) at roughly half the E1 magnitude, while the GFC channel is indistinguishable from zero ($+0.000$). The E1–E2 contrast supports a distinction between

[22] Full aggregate impulse-response paths are reported in Appendix E.

[23] We also estimate a threshold specification and a smooth-transition variant following Auerbach and Gorodnichenko (2012) and Ramey and Zubairy (2018). Neither recovers the scissors using full-sample innovation shocks.

Figure 5: **Channel Impulse Responses from Narrative Local Projections**



Notes: Each panel plots cumulative impulse responses of the four Shapley–Taylor channels—Direct (φ^{dir}, red), GFC (φ^{GFC}, blue), Uncertainty (φ^{UNC}, orange), and Local (φ^{LOC}, green)—from narrative local projections with $\pm$3-day event dummies at horizons $h = 0, \ldots, 30$ days. Top row: geopolitical episodes; bottom row: geoeconomic episodes. Specification (B): VIX and US2Y controls; country FE; Driscoll–Kraay SE. Shaded bands: 90% confidence intervals..

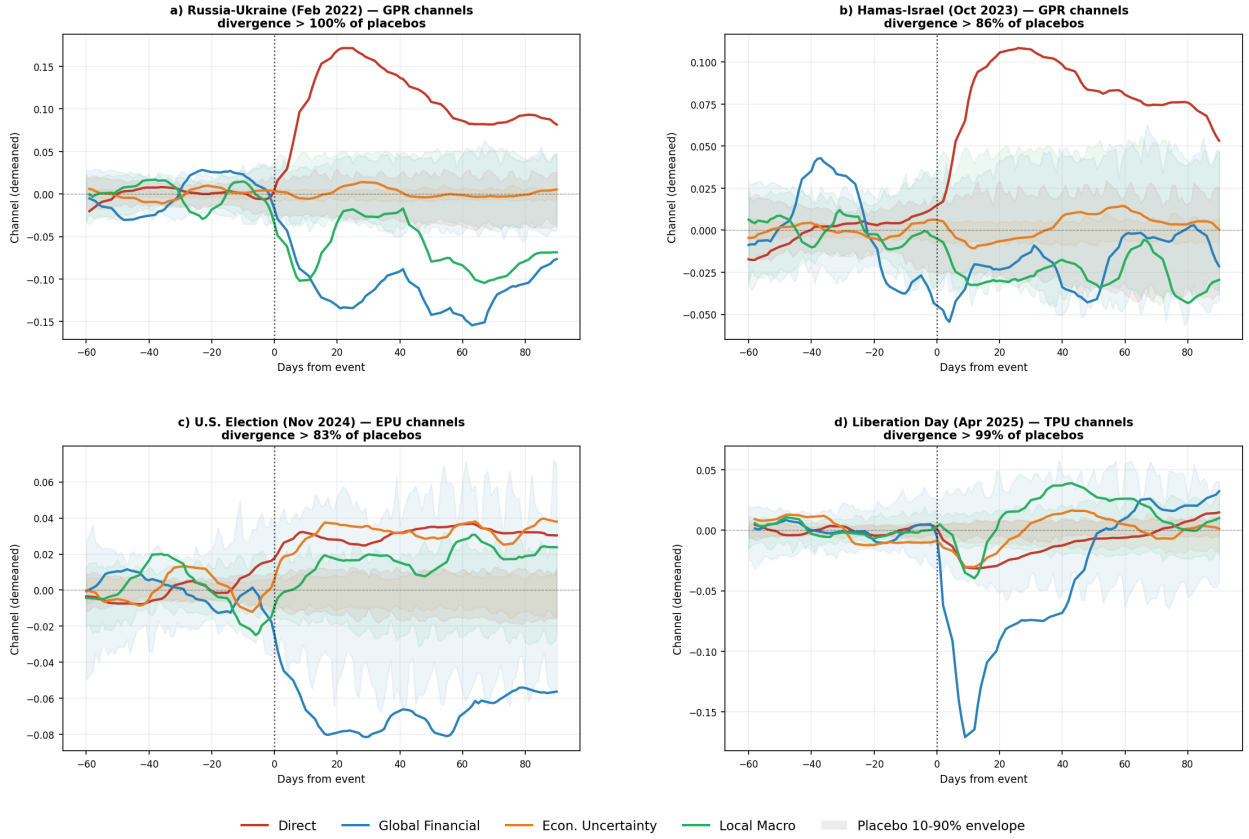
globally systemic geopolitical shocks (textbook scissors) and regionally contained ones (no GFC offset)—the single unconfirmed prediction among the sixteen event–channel tests.

E3—U.S. Presidential Election (EPU shock). The Uncertainty channel is the largest positive response ($+0.021$, $p < 0.01$), the GFC is significantly negative (-0.046, $p < 0.01$), and both Direct ($+0.013$) and Local ($+0.015$) channels are positive and significant. The key point is not that the Direct channel is literally absent, but that it is secondary to the Uncertainty and GFC responses. This multi-channel activation—with Uncertainty as first responder—is the broad-based fingerprint that distinguishes policy-uncertainty shocks from both geopolitical and trade-policy episodes.

E4—"Liberation Day" Tariffs (TPU shock). The Uncertainty ($+0.026$, $p < 0.01$) and Local ($+0.032$, $p < 0.05$) channels are the dominant positive contributors, while the GFC drops on impact (-0.019, $p < 0.10$). The Direct channel response (-0.009, $p < 0.01$) is economically small relative to the dominant channels—an order-of-magnitude difference from Russia–Ukraine ($\varphi^{\mathrm{dir}} = +0.121$), establishing that geopolitical and geoeconomic shocks generate qualitatively different Direct-channel responses.

Stability and scorecard. The qualitative pattern is stable across the three control sets. Moving from specification (A) to specification (C) changes point estimates only modestly (e.g., the E1 Direct channel moves from $+0.122$ to $+0.125$; the E1 GFC from -0.102 to -0.102) and leaves the main sign pattern intact. Overall, the narrative LP supports 15 of 16 event–channel predictions: φ^{dir} activates for both geopolitical events (E1, E2) but is economically negligible for the geoeconomic event (E4); φ^{UNC} dominates for the political-uncertainty event (E3); and the scissors is recovered at the 1% significance level for E1. Because this validation exercise is separate from model selection and from the Shapley decomposition itself, it provides complementary econometric support for the proposed interpretation of the channels.

Figure 6: **Placebo Falsification: Actual Channel Responses vs. Random-Date Envelope.**



Notes: Each panel plots the cross-country mean of the four demeaned channel series (φ^{dir}, φ^{gfc}, φ^{unc}, φ^{loc}) in a $[-60, +90]$-day window around the event date (solid lines) against the 10–90% envelope from 100 randomly drawn non-event dates (shaded bands). Each episode uses its driver-specific Shapley–Taylor decomposition: GPR for Russia–Ukraine and Hamas–Israel; EPU for the U.S. Presidential Election; TPU for the Liberation Day tariff shock. Panel titles report the percentile of the placebo distribution exceeded by the actual event's maximum post-event channel divergence. Channel series smoothed with a 7-day trailing moving average and demeaned over the pre-event window.

Placebo falsification. To verify that the scissors pattern is event-specific, we draw 100 random non-event dates and construct the 10–90% envelope of channel responses in a $[-60, +90]$-day window around each placebo date, using the driver-specific Shapley–Taylor decomposition for each episode (GPR for geopolitical events; EPU and TPU for geoeconomic events). All four episodes produce channel divergence well outside this envelope: Russia–Ukraine exceeds 100% of placebos, Liberation Day 99%, Hamas–Israel 86%, and the U.S. Presidential Election 83%. Crucially, the channels that exit the placebo bands differ across shock types in the direction the taxonomy predicts. Figures 11–14 in Appendix E.1 decompose each episode channel by channel. For geopolitical episodes, the Direct channel exits the 90% band upward while the GFC channel exits downward, and the Uncertainty channel remains inside the placebo range. For geoeconomic episodes, the pattern reverses: the GFC channel dominates the departure, the Uncertainty channel exits the envelope (upward for the Election, consistent with EPU-driven repricing), while the Direct channel stays within or moves opposite to the geopolitical case.

Block bootstrap validation. Because the channel decomposition treats Shapley–Taylor values as observed regressors, the Driscoll–Kraay standard errors in Table 5 do not account for estimation uncertainty in the first-stage ML model. We address this with a generated-regressor block bootstrap ($B = 500$) that, in each replication, jointly re-estimates the Multilayer Random Forest tree, recomputes the Shapley–Taylor decomposition, and re-runs the panel local projection. Appendix Figure 15 overlays the resulting 90% confidence intervals on the baseline Driscoll–Kraay bands.

The graph yields three findings. First, for the raw CDS spread—a directly observed variable unaffected by the first-stage model—bootstrap and Driscoll–Kraay confidence intervals are comparable in width, confirming that the two methods agree when no generated regressor is involved. Second, for the four channel series, bootstrap bands are substantially tighter than the Driscoll–Kraay bands. This reflects two distinct uncertainty sources: the channel decomposition is highly stable across re-estimations of the random forest trained on approximately 75,000 observations, so the bootstrap

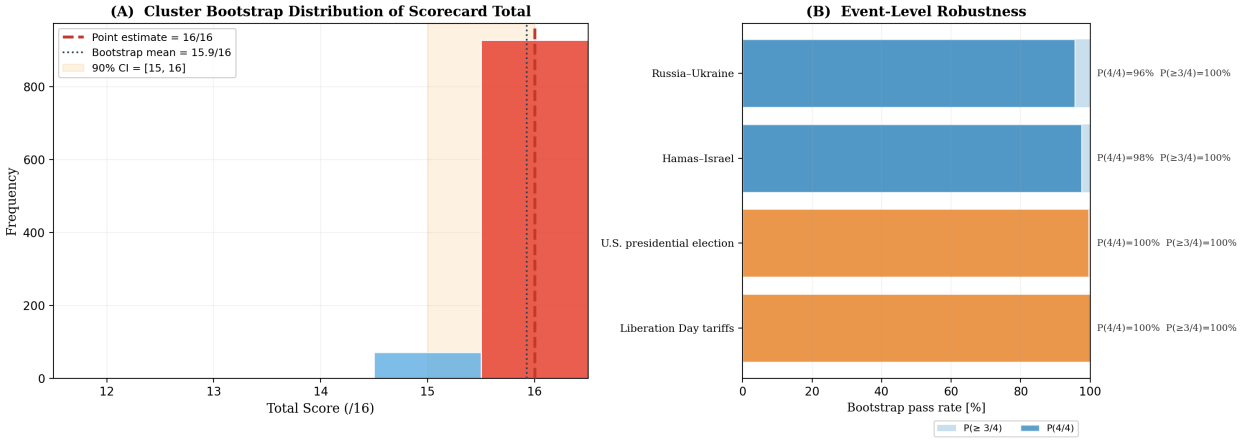
correctly reports that first-stage estimation noise is small; the Driscoll–Kraay estimator additionally corrects for cross-sectional dependence across 42 sovereigns exposed to common shocks—a source of uncertainty that the block bootstrap, which resamples temporal blocks within a fixed cross-section, does not fully replicate. The two methods therefore bound different margins: the bootstrap captures first-stage ML instability, while Driscoll–Kraay captures second-stage panel dependence. We use the latter as the conservative baseline throughout. Third, no channel that is significant under Driscoll–Kraay loses significance under the bootstrap; if anything, several impact responses ($h = 0$) that fall short of conventional thresholds under DK become significant once generated-regressor uncertainty (Pagan, 1984) is properly internalized. Together, these results confirm that the narrative LP evidence supporting the transmission taxonomy—the scissors pattern, the event–channel sign predictions, and the placebo rankings documented in Sections 7–8—is not an artifact of a single ML partition.

8 Evaluating the Channel Taxonomy

The semistructural framework in Section 3 generates four dominance benchmarks—Predictions D, G, U, and L—each mapping a shock family to a specific channel configuration (Table 2). Evaluated jointly across the four crisis episodes and four channels, the taxonomy produces $4 \times 4 = 16$ event–channel predictions. We assess each cell as a pass if the sign and relative magnitude of the estimated $\Delta^c_{1m} \equiv \bar{\varphi}^c_{\text{post},1m} - \bar{\varphi}^c_{\text{pre}}$ conform to the predicted direction, using the calibrated thresholds from the Shapley–Taylor decomposition (Table 2).

Figure 7: **Cluster Bootstrap Validation of Taxonomy of Events**

(a) 1000 replications, resampling 42 countries with replacement



Notes: Each of B = 1,000 replications resamples 42 countries with replacement from the original panel, recomputes the cross-country daily panel mean of each Shapley-Taylor channel, and re-scores the 1-event taxonomy. Point estimate: 16/16. Bootstrap mean: 15.9/16 (sd = 0.3). 90% CI: [15, 16]. Panel (B) dark bars show P(1/4); light bars show P($\geq$ 3/4). Values report P(4/4) | P($\geq$ 3/4).

The point estimate is **16 out of 16**: every event–channel cell matches its predicted sign or dominance inequality. The two geopolitical episodes activate the Direct channel ($\varphi^{\text{dir}} > 0$, significant) while the GFC channel moves in the opposite direction, recovering the scissors condition from Prediction G. The Uncertainty and Local channels remain secondary, consistent with Predictions U and L. The two geoeconomic episodes display the inverted configuration: the Direct channel is small relative to its geopolitical counterpart ($|\varphi^{\text{dir}}_{\text{geoeco}}| \ll |\varphi^{\text{dir}}_{\text{geo}}|$), while the GFC, Uncertainty, and Local channels carry the transmission, with $|\varphi^{\text{GFC}}| > |\varphi^{\text{dir}}|$ confirming GFC dominance under geoeconomic shocks.

To assess whether this result depends on the particular composition of the cross-sectional panel, we perform a cluster bootstrap with $B = 1{,}000$ replications, resampling 42 countries with replacement. Figure 7 reports the results. The bootstrap mean is $15.9/16$ (sd $= 0.3$), with a 90% confidence interval of $[15, 16]$; over 95% of replications score 16/16. Panel (B) decomposes robustness by episode: the two geoeconomic episodes pass all four predictions in 100% of replications; Hamas–Israel at $P(4/4) = 98\%$; Russia–Ukraine at $P(4/4) = 95\%$, with the only residual fragility in the marginally significant Uncertainty channel—consistent with geopolitical shocks not activating policy-regime uncertainty at short horizons (Prediction U).

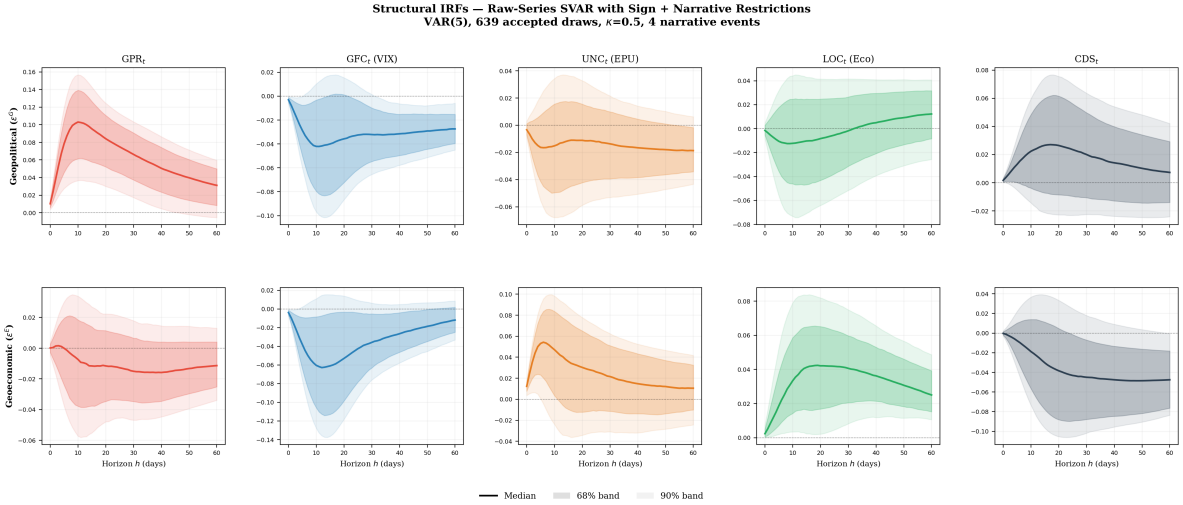
8.1 Independent Validation via Sign-Restricted SVARs

The scorecard and bootstrap validation above rely on the Shapley–Taylor channel decomposition. A natural concern is that the channel assignments may reflect properties of the machine-learning model rather than properties of transmission. To address this, we estimate narrative sign-restricted SVARs on *raw observable composites*—entirely bypassing the Shapley–Taylor decomposition—following the identification logic of Antolín-Díaz and Rubio-Ramírez (2018).

For each country i, we estimate a five-variable VAR(5) on observables that map directly into the four transmission channels: the sovereign CDS spread, the country-level geopolitical risk index (GPR, proxying the Direct channel), a global financial-conditions composite (VIX and U.S. two-year Treasury yield, proxying the GFC channel), an uncertainty composite (EPU and TPU, proxying the Uncertainty channel), and a domestic macro-political composite (ECO, INT, POL, proxying the Local channel). Identification combines sign restrictions derived from the semistructural benchmarks in Table 2, relative-magnitude restrictions encoding the scissors logic ($|\Theta^{\text{GFC}}| < |\Theta^{\text{GPR}}|$ under geopolitical shocks; $|\Theta^{\text{GFC}}| > |\Theta^{\text{GPR}}|$ under geoeconomic shocks), and narrative restrictions on the four crisis dates.[24]

Figure 8 reports the mean-group structural impulse responses from 639 accepted draws. Under the geopolitical shock, GPR rises sharply (approximately $+0.10$ around $h = 10$, with the 68% credible band above zero) while the financial-conditions composite falls (approximately -0.06, with the 68% band below zero), recovering the scissors pattern with $|\Theta^{\text{GFC}}| < |\Theta^{\text{GPR}}|$. The Uncertainty and Local composites remain secondary, consistent with Predictions U and L. Under the geoeconomic shock, the GPR response is muted while the financial-conditions channel is the dominant response (approximately -0.12), with the Uncertainty ($+0.06$) and Local ($+0.04$) composites carrying the remaining transmission and $|\Theta^{\text{GFC}}| > |\Theta^{\text{GPR}}|$. All eight sign and dominance predictions from Table 2 are supported at the posterior median.

Figure 8: **Narrative Sign-Restricted SVAR: Structural Impulse Responses on Raw Observables**



Notes: Mean-group structural impulse responses from country-level VAR(5) models estimated on raw observable composites (equation 29), identified with sign restrictions (Table 12), relative-magnitude restrictions, and narrative restrictions on four crisis dates following Antolín-Díaz and Rubio-Ramírez (2018). Top row: geopolitical shock (ε^G); bottom row: geoeconomic shock (ε^E). Columns: GPR (Direct proxy), GFC composite (VIX-based), Uncertainty composite (EPU/TPU), Local composite (ECO/INT/POL), and sovereign CDS spread. Solid lines: posterior median across 639 accepted draws ($\kappa = 0.5$). Dark shaded bands: 68% pointwise credible intervals; light bands: 90%. Under the geopolitical shock, GPR rises sharply while GFC falls, recovering the scissors pattern ($|\Theta^{\text{GFC}}| < |\Theta^{\text{GPR}}|$). Under the geoeconomic shock, the GFC channel dominates ($|\Theta^{\text{GFC}}| > |\Theta^{\text{GPR}}|$), with Uncertainty and Local channels carrying the remaining transmission. All eight sign and dominance predictions from Table 2 are supported at the posterior median. Because neither the estimation nor the identification uses Shapley–Taylor values, this exercise provides independent validation of the channel taxonomy.

When the same SVAR-identified shocks are used as regressors in full-sample panel local projections, the responses attenuate sharply toward zero—economically small and statistically indistinguishable from zero across all horizons and control sets (Appendix Figures 17 and 18). This attenuation mirrors the innovation-based null in Section 7.1 and is consistent with state-dependent transmission: the time-invariant SVAR weights average over states in which transmission is active and states in which it is dormant (Plagborg-Møller and Wolf, 2021).

[24] Full specification details—including the restriction design, estimation algorithm, and companion SVAR on Shapley–Taylor channel series—are provided in Appendix F.

Table 6 summarises how the four empirical frameworks relate to one another. The Shapley–Taylor scorecard, the narrative LP, the sign-restricted SVAR on raw observables, and the full-sample LP with SVAR shocks form a coherent body of evidence: the taxonomy is recoverable from raw observables without the ML decomposition, and transmission is state-dependent rather than pervasive and linear.

Table 6: Cross-Validation Summary Across Methods

Figure / section	Method	Variables	Result	Relation to baseline
Main text	Shapley Δ_{1m} + narrative LP	ML channels	16/16 event–channel predictions	Baseline
Appendix Figure 16	Sign-restricted SVAR	ML channels	8/8 sign and magnitude relations	Internal coherence
Figure 8	Sign-restricted mean-group SVAR	Raw observables	8/8 sign and magnitude relations	Independent of ML decomposition
Appendix Figure 17	Full-sample LP with identified SVAR shocks	Raw observables	Null average responses	Consistent with state dependence

Notes: "Result" reports the number of predicted sign and relative-magnitude relations supported by the baseline posterior median. "Independent of ML decomposition" means that the raw-observable SVAR neither estimates nor identifies shocks using Shapley–Taylor values. "Null average responses" means that the full-sample LP responses are economically small and statistically indistinguishable from zero at conventional levels.

9 Conclusion

This paper shows that geopolitical and geoeconomic shocks transmit to sovereign credit risk through systematically different channels. The central result is a scissors pattern: geopolitical shocks raise spreads primarily through direct sovereign repricing while the Global Financial Cycle channel moves in the opposite direction, whereas geoeconomic shocks bypass the Direct channel and transmit mainly through financial conditions, policy uncertainty, and domestic amplification. This taxonomy is consistently recovered across three independent empirical frameworks—Shapley–Taylor decomposition, narrative local projections, and sign-restricted SVARs on raw observables—providing a unified characterization of how different shock families map into sovereign risk.

This distinction has direct implications for policy design. When spread widening is driven by global financial conditions, central-bank liquidity provision and swap lines can help stabilize funding markets (Bahaj and Reis, 2022). By contrast, when spreads reflect sovereign-specific geopolitical repricing, liquidity tools are unlikely to address the underlying risk premium. In such cases, the relevant margins lie in diplomatic de-escalation, fiscal adjustment, or institutional credibility. The originator penalty documented for geoeconomic shocks adds a further dimension: the country that initiates a tariff or sanctions episode may face spread widening through the Local channel even as monetary-easing expectations compress spreads abroad, implying that geoeconomic instruments carry a sovereign-risk cost for the originator that standard trade-policy analysis does not capture. More broadly, the results highlight a risk of policy misdiagnosis: similar movements in spreads may require fundamentally different responses depending on the underlying transmission channel. High-frequency news-based indicators are therefore valuable not only as measures of risk, but as tools for identifying the dominant channel in real time.

Several limitations deserve acknowledgment. The channel measures remain model-based decompositions, and the semistructural framework serves as a benchmark rather than a fully estimated structural model. The Shapley–Taylor attribution is an accounting exercise; the narrative validation confirms theory-consistent signs across crisis episodes but does not close the gap between accounting and structural causation over the full sample. The analysis is also silent on welfare and optimal policy: we characterize *how* shocks transmit, but not *what* the optimal central-bank or fiscal response should be conditional on the active channel.

These limitations point to natural extensions. First, embedding the four-channel decomposition within a structural geoeconomic framework such as Clayton et al. (2026) would link empirical channel assignments to underlying policy trade-offs and move from sign predictions to quantitative welfare analysis. Second, augmenting the episode set with pure global financial shocks—such as the SVB/Credit Suisse stress of March 2023 or the yen carry-trade unwind of August 2024—would complete a three-way taxonomy (geopolitical, geoeconomic, financial) and test whether financial shocks generate a third distinct channel configuration in which the GFC channel dominates *positively*, rather than through the offsetting or transmitting roles documented here. Third, modeling sovereign spreads as state-dependent functions of the four channels would formalize the threshold and interaction effects that the gradient-boosted trees capture nonparametrically, providing a bridge between the machine-learning and structural approaches.

References

Ahn, D. P. and Ludema, R. D. (2020). The sword and the shield: The economics of targeted sanctions. *European Economic Review*, 130:103578.

Aiyar, S., Presbitero, A. F., and Ruta, M. (2023). Geoeconomic fragmentation and the future of multilateralism. Staff Discussion Note SDN/2023/001, International Monetary Fund.

Aizenman, J., Hutchison, M., and Jinjarak, Y. (2013). What is the risk of European sovereign debt defaults? Fiscal space, CDS spreads and market pricing of risk. *Journal of International Money and Finance*, 34:37–59.

Alonso-Alvarez, I., Diakonova, M., and Pérez, J. J. (2025). Rethinking GPR: The sources of geopolitical risk. Documentos de Trabajo 2522, Banco de España.

Antolín-Díaz, J. and Rubio-Ramírez, J. F. (2018). Narrative sign restrictions for SVARs. *American Economic Review*, 108(10):2802–2829.

Arellano, C. (2008). Default risk and income fluctuations in emerging economies. *American Economic Review*, 98(3):690–712.

Athey, S. and Imbens, G. W. (2019). Machine learning methods that economists should know about. *Annual Review of Economics*, 11:685–725.

Auerbach, A. J. and Gorodnichenko, Y. (2012). Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy*, 4(2):1–27.

Augustin, P., Subrahmanyam, M. G., Tang, D. Y., and Wang, S. Q. (2022). Sovereign credit risk. *Journal of Financial Economics*, 146(2):431–453.

Bahaj, S. and Reis, R. (2022). Central banking in challenging times: Central bank liquidity lines. *American Economic Journal: Macroeconomics*, 14(2):110–150.

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4):1593–1636.

Bekaert, G., Hoerova, M., and Lo Duca, M. (2013). Risk, uncertainty and monetary policy. *Journal of Monetary Economics*, 60(7):771–788.

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3):623–685.

Bondarenko, Y., Lewis, V., Rottner, M., and Schüler, Y. (2024). Geopolitical risk perceptions. *Journal of International Economics*, 152:104005.

Boubaker, S., Goodell, J. W., Kumar, S., and Sureka, R. (2023). Geopolitical risk and energy stock performance: Evidence from the Russia–Ukraine conflict. *Journal of Commodity Markets*, 30:100325.

Breiman, L. (2001). Random forests. *Machine Learning*, 45(1):5–32.

Bruno, V. and Shin, H. S. (2015). Capital flows and the risk-taking channel of monetary policy. *Journal of Monetary Economics*, 71:119–132.

Caldara, D. and Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, 112(4):1194–1225.

Caldara, D., Iacoviello, M., Molligo, P., Prestipino, A., and Raffo, A. (2020). The economic effects of trade policy uncertainty. *Journal of Monetary Economics*, 109:38–59.

Calvo, G. A., Leiderman, L., and Reinhart, C. M. (1996). Inflows of capital to developing countries in the 1990s. *Journal of Economic Perspectives*, 10(2):123–139.

Chernozhukov, V., Hansen, C., Kallus, N., Spindler, M., and Syrgkanis, V. (2024). Applied causal inference powered by ml and ai.

Clayton, C., Coppola, A., Maggiori, M., and Schreger, J. (2025). Geoeconomic pressure. Working Paper 34020, National Bureau of Economic Research.

Clayton, C., Maggiori, M., and Schreger, J. (2026). A framework for geoeconomics. *Econometrica*, 94(1):105–136.

Driscoll, J. C. and Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics*, 80(4):549–560.

Du, W., Im, J., and Schreger, J. (2018). The dollar, bank funding and global financial conditions. *Journal of Finance*, 73(2):537–586.

Eichengreen, B., Hausmann, R., and Panizza, U. (2021). Original sin redux: A modern assessment of emerging-market vulnerabilities. *Journal of International Money and Finance*, 110:102280.

Farhi, E. and Tirole, J. (2018). Deadly embrace: Sovereign and financial balance sheets doom loops. *Review of Economic Studies*, 85(3):1781–1823.

Fernández-Villaverde, J., Li, Y., Xu, L., and Zanetti, F. (2025). Charting the uncharted: The (un)intended consequences of oil sanctions and dark shipping. Working Paper 33486, National Bureau of Economic Research.

Fernández-Villaverde, J., Mineyama, T., and Song, D. (2024). Are we fragmented yet? measuring geopolitical fragmentation and its causal effects. Working Paper 32638, National Bureau of Economic Research.

Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5):1189–1232.

Gabaix, X. and Maggiori, M. (2015). International liquidity and exchange rate dynamics. *Quarterly Journal of Economics*, 130(3):1369–1420.

Gelos, R. G., Sahay, R., and Sandleris, G. (2011). Sovereign borrowing by developing countries: What determines the bond spread? *Journal of Development Economics*, 96(2):273–282.

Gentzkow, M., Kelly, B., and Taddy, M. (2019). Text as data. *Journal of Economic Literature*, 57(3):535–574.

Glick, R. and Rose, A. K. (1999). Contagion and trade: Why are currency crises regional? *Journal of International Money and Finance*, 18(4):603–617.

Gorodnichenko, Y., Georgarakos, D., Kenny, G., and Coibion, O. (2025). The impact of geopolitical risk on consumer expectations and spending. Working Paper 34195, National Bureau of Economic Research.

Gu, S., Kelly, B. T., and Xiu, D. (2020). Empirical asset pricing via machine learning. *Review of Financial Studies*, 33(5):2223–2273.

He, Z. and Krishnamurthy, A. (2013). Intermediary asset pricing. *American Economic Review*, 103(2):732–770.

Head, K. and Mayer, T. (2014). Gravity equations: Workhorse, toolkit, and cookbook. In Gopinath, G., Helpman, E., and Rogoff, K., editors, *Handbook of International Economics*, volume 4, pages 131–195. Elsevier.

Hoerl, A. E. and Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67.

Jordà, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.

LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553):436–444.

Leetaru, K. and Schrodt, P. A. (2013). GDELT: Global data on events, location, and tone, 1979–2012. Technical Report Version 1.0, GDELT Project. Presented at the International Studies Association Annual Convention, San Francisco, CA.

Longstaff, F. A., Pan, J., Pedersen, L. H., and Singleton, K. J. (2011). How sovereign is sovereign credit risk? *American Economic Journal: Macroeconomics*, 3(2):75–103.

Lundberg, S. M. and Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, volume 30, pages 4765–4774. Curran Associates, Inc.

Maggiori, M., Neiman, B., and Schreger, J. (2020). International currencies and capital allocation. *Journal of Political Economy*, 128(6):2019–2066.

Miranda-Agrippino, S. T. and Rey, H. (2020). U.S. monetary policy and the global financial cycle. *Review of Economic Studies*, 87(6):2754–2776.

Mullainathan, S. and Spiess, J. (2017). Machine learning: An applied econometric approach. *Journal of Economic Perspectives*, 31(2):87–106.

Pagan, A. (1984). Econometric issues in the analysis of regressions with generated regressors. *International Economic Review*, 25(1):221–247.

Pan, J. and Singleton, K. J. (2008). Default and recovery implicit in the term structure of sovereign CDS spreads. *Journal of Finance*, 63(5):2345–2384.

Pástor, V. and Veronesi, P. (2012). Uncertainty about government policy and stock prices. *Journal of Finance*, 67(4):1219–1264.

Pástor, V. and Veronesi, P. (2013). Political uncertainty and risk premia. *Journal of Political Economy*, 121(1):36–72.

Plagborg-Møller, M. and Wolf, C. K. (2021). Local projections and VARs estimate the same impulse responses. *Econometrica*, 89(2):955–980.

Ramey, V. A. and Zubairy, S. (2018). Government spending multipliers in good times and in bad: Evidence from us historical data. *Journal of Political Economy*, 126(2):850–901.

Rey, H. (2013). Dilemma not trilemma: The global financial cycle and monetary policy independence. In *Proceedings of the Jackson Hole Economic Policy Symposium*, pages 285–333.

Shapley, L. S. (1953). A value for n-person games. In *Contributions to the Theory of Games*, volume 2, pages 307–317. Princeton University Press.

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

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288.

Varian, H. R. (2014). Big data: New tricks for econometrics. *Journal of Economic Perspectives*, 28(2):3–28.

Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320.

A APPENDIX:The data

This Appendix details data availability and percentile distributions by country (Appendix A.1), the construction of news indicators (Appendix A.2) and LOESS Unconditional Scatters.

A.1 Data availability by country

The following table presents country-specific data availability together with descriptive statistics summarizing the distribution of the main variables. For each country, the start and end dates of the available sample are reported. To characterize the distribution, we provide the 25th and 75th percentiles (p25 and p75), which capture the inter-quartile range where the central 50% of observations lie, for each of the described variables in the data section.

Table 7: **Country-specific sample periods and interquartile distributions of variables**

Country	Sample		CDS		Global financial var.				Macroeconomic sentiment var.								Political and Geop. var.			
					FED		VIX		ECO		INT		EPU		TPU		GPR		POL	
	Start	End	p25	p75	p25	p75	p25	p75	p25	p75	p25	p75	p25	p75	p25	p75	p25	p75	p25	p75
Argentina	11/08/2020	13/06/2025	-0.78	0.77	-1.30	0.79	-0.72	0.60	-0.72	0.77	-0.39	0.68	-0.68	0.29	-0.77	0.66	-0.82	0.70	-0.76	0.52
Australia	01/01/2018	13/06/2025	-0.64	0.66	-1.03	0.99	-0.71	0.41	-0.34	0.61	-0.49	0.70	-0.65	0.30	-0.63	0.41	-0.69	0.58	-0.72	0.58
Austria	01/01/2018	13/06/2025	-0.95	0.64	-1.03	0.99	-0.71	0.41	-0.39	0.61	-0.39	0.66	-0.64	0.27	-0.56	0.21	-0.57	0.33	-0.69	0.61
Belgium	01/01/2018	13/06/2025	-0.82	0.63	-1.03	0.99	-0.71	0.41	-0.34	0.64	-0.36	0.67	-0.69	0.48	-0.71	0.53	-0.58	0.44	-0.67	0.45
Brazil	01/01/2018	13/06/2025	-0.73	0.66	-1.03	0.99	-0.71	0.41	-0.37	0.67	-0.38	0.67	-0.76	0.57	-0.76	0.52	-0.70	0.57	-0.72	0.47
Canada	21/02/2019	14/06/2025	-0.82	0.78	-1.19	0.96	-0.68	0.44	-0.41	0.68	-0.36	0.68	-0.54	0.16	-0.44	-0.05	-0.72	0.61	-0.75	0.57
Chile	01/01/2018	13/06/2025	-0.65	0.43	-1.03	0.99	-0.71	0.41	-0.62	0.72	-0.52	0.55	-0.73	0.54	-0.69	0.60	-0.67	0.54	-0.85	0.61
China	01/01/2018	13/06/2025	-0.79	0.57	-1.03	0.99	-0.71	0.41	-0.73	0.79	-0.50	0.76	-0.64	0.37	-0.58	0.45	-0.75	0.58	-0.55	0.52
Colombia	01/01/2018	13/06/2025	-0.91	0.67	-1.03	0.99	-0.71	0.41	-0.29	0.58	-0.21	0.50	-0.75	0.54	-0.67	0.52	-0.69	0.49	-0.74	0.66
CzechRep	01/01/2018	13/06/2025	-0.78	0.80	-1.03	0.99	-0.71	0.41	-0.45	0.54	-0.04	0.34	-0.68	0.45	-0.69	0.69	-0.70	0.44	-0.68	0.57
Denmark	01/01/2018	13/06/2025	-0.74	0.49	-1.03	0.99	-0.71	0.41	-0.50	0.71	-0.29	0.60	-0.72	0.46	-0.70	0.51	-0.64	0.45	-0.62	0.41
Egypt	01/01/2018	13/06/2025	-0.74	0.30	-1.03	0.99	-0.71	0.41	-0.44	0.56	-0.33	0.51	-0.68	0.54	-0.67	0.58	-0.75	0.49	-0.66	0.47
Finland	01/01/2018	13/06/2025	-0.79	0.80	-1.03	0.99	-0.71	0.41	-0.46	0.69	-0.57	0.73	-0.68	0.47	-0.70	0.49	-0.70	0.48	-0.58	0.51
France	01/01/2018	13/06/2025	-0.78	0.55	-1.03	0.99	-0.71	0.41	-0.29	0.60	-0.36	0.61	-0.49	0.10	-0.67	0.47	-0.60	0.44	-0.79	0.59
Germany	01/01/2018	13/06/2025	-0.69	0.34	-1.03	0.99	-0.71	0.41	-0.41	0.66	-0.38	0.66	-0.67	0.36	-0.72	0.58	-0.61	0.40	-0.69	0.59
Hungary	01/01/2018	13/06/2025	-0.75	0.52	-1.03	0.99	-0.71	0.41	-0.50	0.64	-0.15	0.56	-0.69	0.55	-0.71	0.48	-0.57	0.31	-0.60	0.49
India	01/01/2018	13/06/2025	-0.59	0.47	-1.03	0.99	-0.71	0.41	-0.46	0.63	-0.42	0.71	-0.67	0.38	-0.69	0.47	-0.66	0.46	-0.75	0.59
Indonesia	01/01/2018	13/06/2025	-0.72	0.42	-1.03	0.99	-0.71	0.41	-0.64	0.63	-0.28	0.62	-0.62	0.25	-0.61	0.38	-0.66	0.58	-0.67	0.60
Israel	01/01/2018	13/06/2025	-0.75	0.17	-1.03	0.99	-0.71	0.41	-0.49	0.74	-0.38	0.70	-0.69	0.51	-0.52	-0.06	-0.74	0.57	-0.86	0.59
Italy	01/01/2018	13/06/2025	-0.83	0.58	-1.03	0.99	-0.71	0.41	-0.43	0.74	-0.17	0.61	-0.79	0.58	-0.54	0.09	-0.62	0.44	-0.63	0.37
Japan	01/01/2018	13/06/2025	-0.71	0.47	-1.03	0.99	-0.71	0.41	-0.42	0.68	-0.40	0.68	-0.75	0.58	-0.75	0.45	-0.73	0.52	-0.70	0.38
Jordan	27/10/2023	13/06/2025	-0.63	0.55	-0.83	0.87	-0.76	0.41	-0.64	0.49	-0.36	0.64	-0.65	0.26	-0.73	0.65	-0.78	0.59	-0.63	0.63
Malaysia	01/01/2018	13/06/2025	-0.78	0.57	-1.03	0.99	-0.71	0.41	-0.66	0.63	-0.43	0.66	-0.66	0.30	-0.72	0.55	-0.84	0.69	-0.68	0.68
Mexico	01/01/2018	13/06/2025	-0.66	0.34	-1.03	0.99	-0.71	0.41	-0.35	0.64	-0.27	0.60	-0.72	0.66	-0.59	0.19	-0.69	0.58	-0.64	0.53
Morocco	01/01/2018	13/06/2025	-0.56	0.17	-1.03	0.99	-0.71	0.41	-0.56	0.52	-0.16	0.42	-0.74	0.61	-0.69	0.62	-0.68	0.52	-0.66	0.58
Netherlands	01/01/2018	13/06/2025	-0.79	0.35	-1.03	0.99	-0.71	0.41	-0.70	0.72	-0.32	0.64	-0.74	0.51	-0.74	0.52	-0.64	0.49	-0.67	0.43
Norway	01/01/2018	13/06/2025	-0.71	0.44	-1.03	0.99	-0.71	0.41	-0.45	0.73	-0.27	0.62	-0.70	0.38	-0.63	0.14	-0.65	0.57	-0.80	0.66
Peru	01/01/2018	13/06/2025	-0.53	0.37	-1.03	0.99	-0.71	0.41	-0.39	0.60	-0.31	0.62	-0.59	0.22	-0.64	0.45	-0.60	0.43	-0.64	0.36
Philippines	01/01/2018	13/06/2025	-0.78	0.69	-1.03	0.99	-0.71	0.41	-0.52	0.75	-0.20	0.54	-0.61	0.35	-0.71	0.51	-0.76	0.68	-0.69	0.45
Poland	01/01/2018	13/06/2025	-0.52	0.00	-1.03	0.99	-0.71	0.41	-0.45	0.69	-0.37	0.65	-0.74	0.51	-0.64	0.40	-0.71	0.50	-0.73	0.67
Qatar	01/01/2018	13/06/2025	-0.65	0.49	-1.03	0.99	-0.71	0.41	-0.65	0.49	-0.12	0.47	-0.64	0.25	-0.68	0.46	-0.75	0.36	-0.54	0.40
Russia	01/01/2018	06/04/2022	-0.21	-0.11	-1.08	1.02	-0.67	0.36	-0.15	0.61	0.01	0.48	-0.76	0.51	-0.67	0.45	-0.50	0.36	-0.62	0.30
SaudiA.	01/01/2018	13/06/2025	-0.65	0.46	-1.03	0.99	-0.71	0.41	-0.42	0.56	-0.14	0.55	-0.76	0.53	-0.73	0.55	-0.73	0.57	-0.55	0.44
Spain	01/01/2018	13/06/2025	-0.76	0.38	-1.03	0.99	-0.71	0.41	-0.35	0.65	-0.27	0.58	-0.65	0.39	-0.52	0.27	-0.63	0.41	-0.67	0.51
Sweden	01/01/2018	13/06/2025	-0.71	0.83	-1.03	0.99	-0.71	0.41	-0.53	0.69	-0.41	0.68	-0.71	0.45	-0.45	0.11	-0.64	0.52	-0.76	0.60
Thailand	01/01/2018	13/06/2025	-0.55	0.31	-1.03	0.99	-0.71	0.41	-0.44	0.72	-0.25	0.50	-0.60	0.36	-0.42	0.19	-0.68	0.43	-0.28	0.59
Turkey	01/01/2018	13/06/2025	-0.77	0.73	-1.03	0.99	-0.71	0.41	-0.46	0.64	-0.20	0.53	-0.66	0.47	-0.68	0.50	-0.74	0.59	-0.47	0.56
UK	01/01/2018	13/06/2025	-0.71	0.72	-1.03	0.99	-0.71	0.41	-0.39	0.69	-0.43	0.69	-0.65	0.44	-0.55	0.34	-0.62	0.39	-0.66	0.41
USA	01/01/2018	13/06/2025	-0.79	0.76	-1.03	0.99	-0.71	0.41	-0.17	0.58	-0.45	0.70	-0.62	0.28	-0.69	0.41	-0.63	0.50	-0.72	0.51
Ukraine	17/01/2020	06/04/2022	-0.33	-0.14	-0.54	0.09	-0.73	0.28	-0.41	0.70	0.15	0.47	-0.81	0.59	-0.72	0.47	-0.69	0.29	-0.67	0.67
Uruguay	01/01/2018	13/06/2025	-0.75	0.45	-1.03	0.99	-0.71	0.41	-0.49	0.65	-0.48	0.69	-0.71	0.51	-0.62	0.59	-0.72	0.54	-0.79	0.65
Vietnam	01/01/2018	13/06/2025	-0.69	0.40	-1.03	0.99	-0.71	0.41	-0.70	0.64	-0.60	0.66	-0.68	0.35	-0.65	0.26	-0.78	0.75	-0.23	0.56

Notes: The table reports country-specific sample periods and interquartile ranges (25th and 75th percentiles) for all variables used in the empirical analysis. The sample covers 42 sovereigns over January 2018–July 2025 (unbalanced panel); start dates vary by CDS data availability. All variables are expressed as 28-day moving averages, demeaned and standardized to zero mean and unit variance within each country. **CDS**: 5-year sovereign credit default swap spread (basis points, standardized). **Global financial variables**: FED = 2-year U.S. Treasury yield (proxy for global monetary policy); VIX = CBOE Volatility Index (proxy for global financial volatility). **Macroeconomic sentiment variables**: ECO = local economic sentiment index; INT = local interest-rate sentiment index; EPU = Economic Policy Uncertainty index; TPU = Trade Policy Uncertainty index. **Political and geopolitical variables**: GPR = Geopolitical Risk index; POL = Political Tensions index. All news-based indicators (ECO, INT, EPU, TPU, GPR, POL) are constructed by BBVA Research from GDELT data as described in Appendix A.2. FED and VIX are global (non-country-specific) variables; their interquartile ranges differ across countries only because of variation in sample periods. Russia and Ukraine have truncated samples ending in June 2022 and January 2020, respectively, reflecting CDS market disruptions following the onset of the conflict. Jordan enters the sample in October 2023.

A.2 Media sentiment indicators development process

All news-based indicators are constructed by BBVA Research using the Global Database of Events, Language, and Tone (Leetaru and Schrodt, 2013). Each index is computed at the country–day level, smoothed with a 28-day moving average, and standardized to zero mean and unit variance over the country-specific sample period. No outlier treatment is applied, as extreme events are informative signals of shifts in risk perception. Four indices are used.[25]

Economic Policy Uncertainty (EPU). Following Baker et al. (2016), the EPU index measures the relative coverage of news articles that jointly mention a country name, terms related to economic uncertainty, and policy-relevant institutions or instruments. The generic query template is:

```
{COUNTRY_NAME} (uncertain OR uncertainty) (economic OR economy) (policy OR tax OR
spending OR regulation OR central bank OR budget OR deficit)
```

Country-specific queries adapt this template to local institutional vocabulary. For instance, the United States query adds references to Congress, the White House, and the Federal Reserve; Mexico includes Banxico, BdeM, and the Chamber of Deputies; Spain references Hacienda, parliament, and specific fiscal terms.[26] The index equals the ratio of EPU-related articles to total published articles for each country–day.

Trade Policy Uncertainty (TPU). Following Caldara et al. (2020), the TPU index captures news coverage related to trade-policy instruments. The query requires the co-occurrence of a trade term and a policy-instrument term:

```
(tariff OR tariffs OR import OR imports OR export OR exports OR trade OR dumping OR
antidumping OR GATT OR WTO) (duty OR duties OR barrier OR barriers OR ban OR bans OR
tax OR taxes OR subsidy OR subsidies)
```

The index equals the ratio of TPU-related articles to total published articles. Unlike the EPU, the TPU query is uniform across countries.

Geopolitical Risk (GPR). Following Caldara and Iacoviello (2022), the GPR index combines tone (sentiment) and relative coverage of articles classified under two groups of GDELT conflict-related taxonomy themes. Group 1 covers military and political-violence categories (e.g., `WAR`, `CONFLICT`, `TERROR`, `MILITARY`, `REBELLION`, `PEACEKEEPING`). Group 2 covers threat and disruption categories (e.g., `SANCTIONS`, `BLOCKADE`, `ARMEDCONFLICT`, `CRISIS`, `RAID`, `KILL`). An article must match at least one theme from each group. The index is the product of tone and relative coverage, multiplied by -1 for interpretability (higher values = greater geopolitical risk).

Political Tensions (POL). The Political Tensions Index captures tone and coverage of news classified under the GDELT taxonomy theme `USPEC_POLITICS_GENERAL1` (elections, government institutions, legislative activity, political scandals, civil rights, national security, among others) that additionally contain at least one of the following keywords: *political instability, political uncertainty, political crisis, political polarization, political extremism, political turmoil, political conflict*. As with GPR, the index is the product of tone and relative coverage, multiplied by -1.

[25] Complete country-specific keyword lists and GDELT taxonomy codes for all 42 countries are available in the replication package at `https://bigdata.bbvaresearch.com/en/`.

[26] Representative examples: *United States*: "United States (uncertainty OR uncertain) (economic OR economy) (congress OR legislation OR white house OR regulation OR federal reserve OR deficit)"; *Mexico*: "Mexico (economic OR economy) (uncertain OR uncertainty) (regulation OR deficit OR budget OR Bank OR BdeM OR Banxico OR congress OR senate OR deputies OR legislation OR taxes OR Federal Reserve)"; *Spain*: "Spain (uncertainty OR uncertain OR instability OR risk) (economic OR economy) (parliament OR government OR Hacienda OR deficit OR budget OR expenditure OR debt OR taxes OR law OR reform OR regulation OR Bank)".

B APPENDIX: Semistructural Microfoundations

This appendix provides micro-founded mechanisms for the dominance benchmarks summarized in Table 2. Each channel is given minimal internal structure—enough to show that the sign predictions, the dominance inequalities, and the key empirical patterns (gravity decay, the scissors, broad uncertainty activation, domestic amplification) arise from primitive economic forces. Readers who accept the benchmarks of Section 3 on economic grounds may skip this appendix without loss of continuity.

Setup. Let $s_{i,t}$ denote the sovereign CDS spread of country i at time t. We write

$$s_{i,t} \;=\; \mu_i + \tau_t + D_{i,t} + G_{i,t} + U_{i,t} + L_{i,t} + \varepsilon_{i,t}\,, \tag{21}$$

where μ_i and τ_t are country and time effects, $\varepsilon_{i,t}$ is an idiosyncratic disturbance, and the four channel terms capture distinct transmission mechanisms.

Channel I: Direct sovereign repricing. A geopolitical shock of intensity $g_t \geq 0$ reduces expected output in proportion to conflict proximity. Building on the endogenous-default tradition (Arellano, 2008), we model the Direct channel as

$$D_{i,t} = h\big(\alpha(d_i)\, g_t,\; F_{i,t}\big)\,, \tag{22}$$

where d_i is geodesic distance from country i to the conflict epicenter, $\alpha(\cdot)$ is a decreasing and concave exposure function ($\alpha' < 0$, $\alpha'' > 0$), and $F_{i,t}$ is a fragility index summarising the sovereign's proximity to its default boundary. The function $h(\cdot)$ is increasing in both arguments and exhibits a *threshold*: when $F_{i,t}$ is low, moderate exposure has a small effect on spreads; when $F_{i,t}$ exceeds a critical level $\bar{F}$, the same exposure triggers a nonlinear jump as the sovereign approaches its default boundary. The concavity of $\alpha(\cdot)$ in distance implies steep attenuation near the epicenter and flattening at greater distances—the functional form captured by the log-distance specification estimated in Section 6.4.

This channel provides the microfoundation for Prediction D of Section 3: geopolitical shocks generate large $D_{i,t}$ for exposed sovereigns, while geoeconomic shocks—which do not directly impair repayment capacity—produce $|D_{i,t}^{\text{geoeco}}| \ll |D_{i,t}^{\text{geo}}|$.

Channel II: Global Financial Cycle. Sovereign CDS are priced by global financial intermediaries whose risk-bearing capacity fluctuates with balance-sheet constraints (He and Krishnamurthy, 2013; Gabaix and Maggiori, 2015). Let W_t denote intermediary-sector wealth and $\bar{W}$ its leverage constraint. The GFC channel is

$$G_{i,t} = \lambda_i \cdot \xi(W_t)\,, \qquad \xi' < 0\,,\;\; \xi'' > 0\,, \tag{23}$$

where $\xi(W_t)$ is a convex amplification function that rises steeply as the intermediary approaches its constraint, and λ_i is country i's loading on the global financial factor, determined by its integration into dollar funding markets, reliance on portfolio flows, and exchange-rate regime (Rey, 2013; Miranda-Agrippino and Rey, 2020). When W_t is high (accommodative conditions), $\xi(\cdot)$ is flat and a volatility spike has moderate effects; when W_t is low (tight conditions), $\xi(\cdot)$ is steep and the same spike is powerfully amplified (Bruno and Shin, 2015).

The scissors condition. The intermediary framework delivers the scissors pattern. A large geopolitical shock ($g_t > 0$) raises $D_{i,t}$ for exposed sovereigns. Simultaneously, two offsetting forces act on W_t: mark-to-market losses on exposed debt reduce it, while flight-to-quality flows and central-bank liquidity provision (Bahaj and Reis, 2022) replenish it. When intermediary capacity is sufficient ($W_t \gg \bar{W}$) and the central-bank response is aggressive, the net effect is $\Delta W_t > 0$, so $\xi(W_t)$ *falls*:

$$\underbrace{\Delta D_{i,t} > 0}_{\text{Direct rises}} \qquad \text{while} \qquad \underbrace{\Delta G_{i,t} < 0}_{\text{GFC falls}} \tag{24}$$

This is the scissors. Russia–Ukraine (large shock, robust intermediary, aggressive ECB/Fed action) produces a textbook scissors; Hamas–Israel (smaller shock, limited global spillover) produces a muted GFC offset. When the shock overwhelms intermediary capacity ($\Delta W_t < 0$), the GFC *reinforces* rather than offsets the Direct channel—the financial-exclusion regime exemplified by Ukraine ($\varphi_{1m}^{\text{GFC}} = +0.66$).

This mechanism delivers the dominance component of Prediction G: under geopolitical shocks the GFC is the secondary, offsetting channel ($|\varphi^{\text{GFC}}| < |\varphi^{\text{DIR}}|$), whereas under geoeconomic shocks—which compress the discount-rate factor without activating the Direct channel—the GFC becomes the dominant transmission channel ($|\varphi^{\text{GFC}}| > |\varphi^{\text{DIR}}|$).

Channel III: Geoeconomic Uncertainty. We depart from Pástor and Veronesi (2012, 2013) in two respects. First, in PV, political uncertainty resolves when the government announces policy. In a geoeconomic setting, uncertainty is *self-reinforcing*: a tariff announcement generates uncertainty about retaliation, escalation, and regime change. We model the underlying state as a hidden Markov geoeconomic regime $\theta_t \in \{\theta_L, \theta_H\}$, with investors updating beliefs $\pi_{i,t} = \Pr(\theta_t = \theta_H \mid \mathcal{I}_{i,t})$ from noisy signals. Second, PV operates in a closed-economy equity setting with a representative agent; our setting requires a multi-country sovereign credit market with a global intermediary.

The Uncertainty channel is:

$$U_{i,t} = \kappa_i \cdot \mathcal{H}(\pi_{i,t}) + \delta_i \cdot \pi_{i,t}\,, \tag{25}$$

where $\mathcal{H}(\pi) = -\pi \ln \pi - (1-\pi)\ln(1-\pi)$ is the Shannon entropy of the posterior and κ_i is country i's sensitivity to geoeconomic ambiguity. The first term captures the pure *uncertainty premium*, maximised at $\pi = \frac{1}{2}$ and vanishing at the extremes. The second captures the *level effect*: higher $\pi_{i,t}$ raises spreads because the high-tension regime implies worse expected fundamentals.

The entropy function is concave, generating *convexity in the uncertainty premium*: marginal uncertainty increases matter more at moderate levels than at extremes—the nonlinearity documented in the EPU × VIX interaction surface of Figure 2. Crucially, when a common geoeconomic signal arrives (an election outcome, a tariff announcement), $\pi_{i,t}$ updates simultaneously for all countries, producing correlated uncertainty premia with narrow cross-country dispersion—the pattern documented for the election episode (IQR $[+0.01, +0.06]$; $\sigma = 0.02$).

This structure delivers both Prediction U and the temporal distinction underlying its dominance formulation: geopolitical shocks do not directly shift the distribution over $\{\theta_L, \theta_H\}$ at short horizons, so $|\varphi^{\text{UNC}}_{\text{geo}}| \ll |\varphi^{\text{UNC}}_{\text{geoeco}}|$. The cascade from geopolitical tension into policy-regime uncertainty (sanctions → fiscal response → regime ambiguity) operates at longer horizons, outside our event windows.

The channel also generates the *originator penalty* of Pástor and Veronesi (2013), but through a different mechanism: the originator faces a dual burden as its posterior $\pi_{US,t}$ rises from domestic policy uncertainty while the interaction with domestic sentiment amplifies spread widening through the Local channel—even as the GFC compresses spreads abroad.

Channel IV: Local amplification. Domestic fundamentals amplify or absorb external shocks through a threshold process:

$$L_{i,t} = \beta_L\, X^{\text{macro}}_{i,t} + \gamma_L\, X^{\text{macro}}_{i,t} \cdot \mathbf{1}\{F_{i,t} > \bar{F}\}\,, \tag{26}$$

where $X^{\text{macro}}_{i,t}$ stacks domestic sentiment, interest-rate expectations, political tensions, and regional structure—daily-frequency proxies for the core determinants of debt sustainability. Below the fragility threshold, shocks pass through with modest amplification (β_L). Above it, amplification jumps to $\beta_L + \gamma_L$, capturing sovereign–bank doom loops, capital flight, or loss of market access (Farhi and Tirole, 2018).

The threshold structure delivers Prediction L: under geopolitical shocks, the Direct and GFC channels dominate the panel mean, so $\mathbb{E}_i[|\varphi^{\text{LOC}}|] \ll \mathbb{E}_i[|\varphi^{\text{DIR}}|]$. Under geoeconomic shocks, bilateral heterogeneity in $X^{\text{macro}}_{i,t}$ and the threshold indicator make the Local channel a primary differentiator across countries.

The nonlinear pricing function. Combining the four channels, the spread is

$$s_{i,t} = m\Big(\underbrace{\alpha(d_i)\, g_t,\, F_{i,t}}_{\text{Direct}},\, \underbrace{\xi(W_t),\, \lambda_i}_{\text{GFC}},\, \underbrace{\mathcal{H}(\pi_{i,t}),\, \kappa_i}_{\text{Uncertainty}},\, \underbrace{X^{\text{macro}}_{i,t},\, \mathbf{1}\{F_{i,t} > \bar{F}\}}_{\text{Local}}\Big) + \mu_i + \tau_t + \varepsilon_{i,t}\,, \tag{27}$$

where $m(\cdot)$ is an unknown function encoding three sources of nonlinearity: (i) threshold effects through $\mathbf{1}\{F_{i,t} > \bar{F}\}$; (ii) interaction effects through the dependence of $G_{i,t}$ on $\xi(W_t)$, which interacts with g_t via intermediary losses; (iii) convexity through $\mathcal{H}(\pi_{i,t})$. The illustrative expansion

$$s_{i,t} = \mu_i + \tau_t + \beta_D\, \alpha(d_i)\, g_t + \beta_G\, \xi(W_t) + \beta_U\, \mathcal{H}(\pi_{i,t}) + \beta_L\, X^{\text{macro}}_{i,t} + \gamma_1\, \alpha(d_i)\, g_t\, \mathbf{1}\{F_{i,t} > \bar{F}\} + \gamma_2\, g_t\, \xi(W_t) + \gamma_3 \left[\mathcal{H}(\pi_{i,t})\right]^2 + u_{i,t} \tag{28}$$

shows that even a parsimonious approximation generates interaction and higher-order terms that a linear panel model cannot capture, motivating the use of flexible ML estimators to approximate $m(\cdot)$.

Summary of predictions. The framework generates four predictions that correspond to the dominance benchmarks of Table 2:

P1. Gravity decay (Prediction D). The concavity of $\alpha(d_i)$ implies that $\varphi^{\text{DIR}}_{i,t}$ decays with log-distance to the conflict epicenter. *Test*: Figure 4.

P2. Scissors and GFC dominance (Predictions D and G). Equation (24) delivers $\varphi^{\text{DIR}} > 0$ and $\varphi^{\text{GFC}} < 0$ under geopolitical shocks, with $|\varphi^{\text{GFC}}| < |\varphi^{\text{DIR}}|$. Geoeconomic shocks bypass the Direct channel and transmit primarily through the GFC, so $|\varphi^{\text{GFC}}| > |\varphi^{\text{DIR}}|$. *Test*: Table 5, Figure 5.

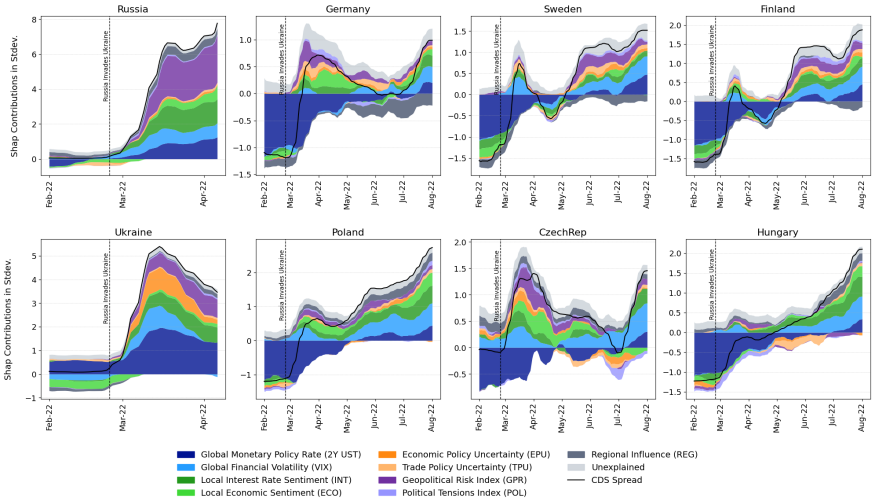
P3. Broad-based uncertainty activation (Prediction U). Because $\pi_{i,t}$ updates on a common geoeconomic signal, φ^{UNC} activates broadly under geoeconomic shocks. Under geopolitical shocks, $|\varphi^{\text{UNC}}_{\text{geo}}| \ll |\varphi^{\text{UNC}}_{\text{geoeco}}|$ at short horizons. *Test*: cross-country dispersion across episodes.

P4. Nonlinear amplification and originator penalty (Prediction L). The threshold $\bar{F}$ produces qualitatively different transmission for high- vs. low-fragility sovereigns. Under geopolitical shocks, $\mathbb{E}_i[|\varphi^{\text{LOC}}|] \ll \mathbb{E}_i[|\varphi^{\text{DIR}}|]$; under geoeconomic shocks, φ^{LOC} is a primary channel, and the originator faces widening spreads through it even as the GFC compresses spreads abroad. *Test*: Table 4.

The framework nests Pástor and Veronesi (2012, 2013) as special cases—our Direct, Uncertainty, and GFC channels correspond to PV's impact component, political uncertainty premium, and the intermediary response absent in their representative-agent setting—and connects to Clayton et al. (2026), whose vulnerability–leverage model provides the structural foundation for the strategic interaction sustaining θ_H in Channel III.

C Appendix. Country-Level Transmission Dynamics

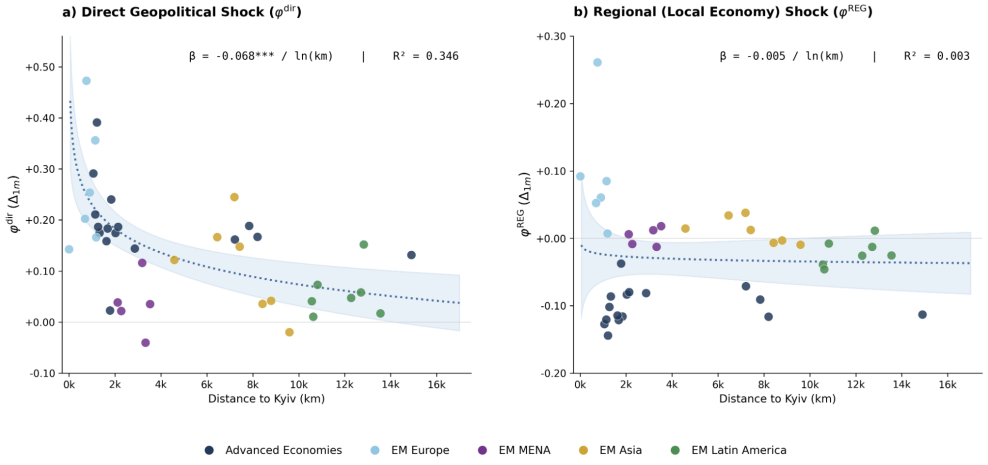
Figure 9: **Russia–Ukraine Invasion: Country-Level Shapley Dynamics**



Notes: Each panel plots the evolution of Shapley contributions to sovereign CDS spreads for selected countries before and after February 2022. Stacked areas represent individual predictor contributions; black line is the observed CDS spread. For Russia and Ukraine, contributions from GPR, global financial volatility, and domestic conditions surge simultaneously. For European economies, the initial geopolitical impulse is followed by a cascade through domestic inflation, local rate tightening, and global financial tightening. All series are 28-day moving averages.

Country-level Shapley attribution dynamics for the remaining three episodes (Hamas–Israel, U.S. Election, Liberation Day) are available in the replication package.

Figure 10: **Gravity Validation: Direct Channel vs. Regional Component — Russia–Ukraine (Feb. 2022)**



Notes: Panel (a) plots the one-month Direct-channel response $\varphi_i^{\text{dir}}(\Delta_{1m})$ and panel (b) the REG-within-LOC Shapley value $\varphi_i^{\text{REG}}(\Delta_{1m})$ against geodesic distance from each country's capital to Kyiv. The Direct channel exhibits significant gravity-like decay ($\hat{\beta} = -0.068^{***}$, $R^2 = 0.346$), consistent with geopolitical shocks repricing default risk in proportion to conflict proximity. The regional component—which captures the Shapley–Taylor attribution of the regional group indicators (Advanced, EM Europe, EM MENA, EM Asia, EM Latin America) absorbed into the Local channel—shows no distance relationship ($\hat{\beta} = -0.005$, $R^2 = 0.003$, $p = 0.72$), confirming that it does not proxy for geographic proximity. This contrast validates the absorption of REG into the Local channel rather than the creation of a fifth transmission channel. Dotted lines: OLS log-distance fit from eq. (20); shaded bands: 95% confidence intervals. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. OLS includes all countries with available CDS data (excluding Ukraine at $d = 0$).

D Appendix. Country-Level Channel Decomposition Tables

Table 8: **Geopolitical Episodes** — Country-level channel decomposition (Δ_{1m})

	Russia–Ukraine (Feb. 2022)					Hamas–Israel (Oct. 2023)				
Country	φ^{dir}	φ^{GFC}	φ^{UNC}	φ^{LOC}	φ^{tot}	φ^{dir}	φ^{GFC}	φ^{UNC}	φ^{LOC}	φ^{tot}
Advanced Economies										
Group avg.	+.15	−.15	+.01	−.17	−.16	+.09	−.02	.00	−.06	+.01
USA	+.14	−.12	.00	−.17	−.14	+.08	−.12	−.03	−.01	−.08
Canada	+.09	−.04	.00	−.20	−.15	+.10	−.09	−.02	−.15	−.15
Australia	+.10	−.18	+.02	−.21	−.27	+.08	−.20	−.04	−.09	−.26
UK	+.15	−.14	.00	−.19	−.17	+.11	+.08	+.02	−.04	+.18
Germany	+.25	−.24	+.02	−.03	.00	+.09	−.15	−.02	.00	−.08
France	+.14	−.08	+.02	−.16	−.07	+.16	−.04	−.02	−.11	.00
Italy	+.12	−.28	+.04	−.39	−.50	+.12	.00	−.02	−.12	−.01
Spain	+.11	−.11	+.01	−.08	−.07	+.15	−.04	.00	−.17	−.07
Netherlands	−.03	−.03	−.01	−.14	−.22	+.06	+.10	+.03	−.04	+.15
Belgium	+.17	−.17	+.01	−.15	−.14	+.14	−.01	+.02	−.12	+.04
Austria	+.23	−.26	+.01	−.12	−.14	+.07	−.09	+.01	−.03	−.04
Denmark	+.17	−.12	+.01	−.08	−.02	−.02	+.03	+.02	+.01	+.04
Finland	+.24	−.12	−.01	−.16	−.06	+.08	+.05	+.02	−.01	+.14
Norway	+.16	−.18	−.02	−.17	−.20	+.07	+.06	+.02	−.09	+.06
Sweden	+.20	−.12	+.01	−.19	−.10	+.03	+.08	+.04	−.01	+.14
Japan	+.14	−.14	−.01	−.23	−.24	+.08	+.04	−.01	−.01	+.10
Emerging Europe										
Group avg.	+.24	+.22	+.05	+.52	+1.03	+.05	−.02	.00	−.03	.00
Russia	+.52	+1.55	+.31	+2.23	+4.61	—	—	—	—	—
Ukraine	+.13	+.66	−.02	+.64	+1.41	—	—	—	—	—
Poland	+.15	−.10	+.04	+.11	+.19	−.03	+.06	+.01	.00	+.05
Czech Rep.	+.32	−.35	−.01	+.12	+.08	+.05	+.04	.00	+.02	+.11
Hungary	+.21	−.27	−.02	+.07	−.01	+.11	−.03	−.01	−.09	−.02
Turkey	+.13	−.20	−.01	−.03	−.12	+.05	−.14	.00	−.06	−.15
Middle East & North Africa										
Group avg.	+.03	−.08	−.02	−.14	−.22	+.17	+.12	+.07	+.13	+.50
Egypt	+.01	−.09	−.02	−.17	−.26	+.27	+.08	+.08	+.16	+.59
Israel	+.08	−.06	−.02	−.10	−.10	+.21	+.32	+.19	+.45	+1.17
Morocco	−.01	−.09	−.03	−.10	−.24	+.06	+.17	+.04	−.06	+.20
Qatar	−.04	−.06	.00	−.19	−.28	+.20	+.16	+.04	.00	+.39
Saudi Arabia	+.08	−.12	−.03	−.17	−.23	+.14	−.11	.00	+.12	+.14
Asia-Pacific										
Group avg.	+.08	−.14	−.01	−.01	−.07	+.06	−.09	.00	−.06	−.10
China	+.12	−.17	.00	+.06	+.01	−.02	−.46	−.05	−.13	−.66
India	+.14	−.16	−.01	−.04	−.07	−.03	−.08	.00	+.02	−.10
Indonesia	−.04	−.02	+.01	−.07	−.13	+.09	−.02	.00	−.09	−.02
Malaysia	.00	−.12	.00	−.09	−.20	.00	−.07	+.01	.00	−.07
Philippines	.00	−.02	+.03	−.03	−.03	+.20	+.02	+.02	−.19	+.04
Thailand	+.14	−.14	−.01	+.07	+.06	+.10	+.06	+.01	−.05	+.11
Vietnam	+.19	−.33	−.05	+.04	−.15	+.05	−.10	+.02	−.01	−.04
Latin America										
Group avg.	+.03	−.11	−.01	−.04	−.13	+.02	−.01	−.01	−.01	−.01
Argentina	+.10	−.14	−.03	−.03	−.10	+.04	−.02	−.01	−.05	−.04
Brazil	−.03	−.14	−.03	−.11	−.31	−.04	−.04	−.03	+.03	−.08
Chile	.00	−.04	.00	−.06	−.11	+.08	+.05	.00	−.03	+.09
Colombia	.00	−.12	−.04	−.05	−.21	.00	.00	−.02	−.03	−.05
Mexico	+.05	−.09	+.01	−.11	−.14	+.06	−.10	−.03	−.01	−.09
Peru	+.06	−.09	.00	+.05	+.03	+.05	−.02	−.02	.00	+.01
Uruguay	+.03	−.11	−.02	+.03	−.08	−.01	+.07	+.01	−.01	+.06

$\Delta > +0.05$ $+0.005 < \Delta \leq +0.05$ $|\Delta| \leq 0.005$ $-0.05 \leq \Delta < -0.005$ $\Delta < -0.05$. Russia/Ukraine excl. from Hamas–Israel (no pre-event data). φ^{LOC} includes Region cross-SHAP.

Table 9: **Geoeconomic Episodes** — Country-level channel decomposition (Δ_{1m})

	Trump Election (Nov. 2024)					Liberation Day (Apr. 2025)				
Country	φ^{dir}	φ^{GFC}	φ^{UNC}	φ^{LOC}	φ^{tot}	φ^{dir}	φ^{GFC}	φ^{UNC}	φ^{LOC}	φ^{tot}
USMCA										
Group avg.	+.04	−.07	.00	+.02	−.01	—	—	—	—	—
USA	+.04	−.08	+.01	+.06	+.03	.00	−.02	.00	+.07	+.05
Canada	+.05	−.10	−.01	.00	−.07	+.01	+.11	+.07	+.22	+.42
Mexico	+.03	−.03	+.01	+.01	+.01	+.02	+.06	+.10	+.14	+.32
Advanced Economies										
Group avg.	+.04	−.05	+.04	+.05	+.08	.00	−.05	.00	.00	−.05
Australia	+.04	−.12	+.05	+.03	−.01	−.02	−.07	−.03	+.03	−.10
UK	+.04	−.05	+.03	+.01	+.02	.00	−.07	+.01	+.03	−.03
Germany	+.04	.00	+.06	+.08	+.19	.00	−.05	+.01	−.01	−.05
France	+.03	+.14	+.05	+.16	+.38	+.01	.00	+.01	.00	+.02
Italy	+.04	−.15	+.05	−.01	−.07	+.01	−.05	+.01	−.03	−.07
Spain	+.02	−.11	+.01	+.02	−.06	+.01	−.02	−.01	+.02	.00
Netherlands	+.05	−.10	+.07	+.06	+.08	.00	−.01	−.02	−.02	−.05
Japan	+.02	−.01	+.03	+.03	+.08	−.01	−.09	−.01	+.02	−.09
Emerging Europe										
Group avg.	+.04	−.05	+.02	−.01	.00	.00	.00	−.01	.00	.00
Poland	+.02	+.02	+.02	−.01	+.05	.00	−.01	−.01	−.01	−.03
Czech Rep.	+.03	−.06	+.05	+.03	+.05	+.01	−.01	+.01	−.03	−.02
Hungary	+.04	−.09	+.01	−.02	−.06	−.02	+.02	−.01	+.07	+.05
Turkey	+.05	−.05	.00	−.04	−.04	.00	+.02	.00	−.02	−.01
Asia-Pacific										
Group avg.	+.03	−.05	+.04	+.03	+.04	−.01	−.03	−.01	+.04	−.01
China	+.03	−.03	+.05	+.07	+.12	.00	−.05	−.01	−.01	−.06
India	+.03	−.09	.00	+.01	−.04	.00	−.09	−.02	+.01	−.10
Indonesia	+.03	−.06	+.03	+.03	+.03	+.01	−.02	.00	+.02	+.01
Malaysia	.00	−.01	+.04	+.01	+.04	−.03	−.03	−.02	+.07	.00
Philippines	+.03	−.08	+.04	+.01	+.01	−.02	+.02	.00	.00	+.01
Thailand	+.02	−.03	+.05	+.04	+.09	−.03	+.05	+.01	+.14	+.18
Vietnam	+.04	−.07	+.03	+.03	+.02	−.01	−.11	−.02	+.03	−.12
Latin America										
Group avg.	+.03	−.07	+.04	+.02	+.01	.00	.00	+.01	−.02	.00
Argentina	+.03	−.05	+.05	−.04	−.02	.00	+.02	−.01	.00	+.02
Brazil	+.05	−.10	+.06	+.07	+.09	+.02	−.01	.00	−.08	−.06
Chile	+.03	−.07	+.04	+.03	+.03	−.03	.00	+.01	−.01	−.02
Colombia	+.02	−.07	+.01	+.03	−.02	+.01	−.01	+.02	−.03	.00
Peru	+.04	−.08	+.04	−.01	.00	−.01	.00	+.02	.00	+.01
Uruguay	+.01	−.06	+.03	+.01	−.01	−.01	+.02	+.02	+.01	+.04
Middle East & North Africa										
Group avg.	+.03	−.08	+.03	+.01	−.02	−.01	.00	+.01	+.01	+.01
Egypt	+.02	−.08	.00	+.01	−.05	.00	−.05	+.01	−.05	−.08
Israel	+.04	−.08	+.02	−.06	−.08	−.01	+.06	.00	+.04	+.10
Qatar	+.03	−.12	+.05	+.04	.00	−.02	−.02	−.01	.00	−.04
Saudi Arabia	+.03	−.05	+.05	+.04	+.07	−.02	+.01	+.04	+.03	+.06

$\Delta > +0.05$ $+0.005 < \Delta \leq +0.05$ $|\Delta| \leq 0.005$ $-0.05 \leq \Delta < -0.005$ $\Delta < -0.05$. Liberation Day USMCA group avg. not reported separately; see individual countries. φ^{LOC} includes Region cross-SHAP.

E Appendix. Local Projection Robustness

Table 10: Average Innovation LP : Estimated Channel Impulse Responses at $h = 5$, 30, and 60 Days

Event	Channel	$h = 5$		$h = 30$		$h = 60$	
		$\hat{\beta}$	Sig.	$\hat{\beta}$	Sig.	$\hat{\beta}$	Sig.
Russia–Ukraine GPR, Feb. 2022 *Geopolitical*	φ^{dir}	0.0130	–	0.0050	–	0.0034	–
	φ^{GFC}	-0.0039	–	-0.0007	–	-0.0002	–
	φ^{UNC}	-0.0012	–	-0.0003	–	-0.0002	–
	φ^{LOC}	-0.0077	–	-0.0037	–	-0.0023	–
Hamas–Israel GPR, Oct. 2023 *Geopolitical*	φ^{dir}	0.0130	–	0.0050	–	0.0034	–
	φ^{GFC}	-0.0039	–	-0.0007	–	-0.0002	–
	φ^{UNC}	-0.0012	–	-0.0003	–	-0.0002	–
	φ^{LOC}	-0.0077	–	-0.0037	–	-0.0023	–
U.S. Election EPU, Nov. 2024 *Geoeconomic*	φ^{dir}	0.0034	–	0.0017	–	0.0014	–
	φ^{GFC}	0.0032	–	0.0017	–	0.0010	–
	φ^{UNC}	0.0008	–	0.0006	–	0.0004	–
	φ^{LOC}	-0.0006	–	0.0003	–	0.0003	–
"Liberation Day" TPU, Apr. 2025 *Geoeconomic*	φ^{dir}	-0.0020	–	-0.0008	–	-0.0005	–
	φ^{GFC}	-0.0016	–	-0.0010	–	-0.0000	–
	φ^{UNC}	0.0015	–	0.0009	–	0.0005	–
	φ^{LOC}	-0.0009	–	0.0000	–	0.0002	–

Notes: Estimated impulse responses from panel local projections (Jordà, 2005) using AR(5) residual innovations as structural shocks. Standard errors follow Driscoll and Kraay (1998) with bandwidth $= \max(20, h)$, with country and time fixed effects and 5 lags of Δy. Contemporaneous controls: global VIX (moving average) and sovereign yield spread (moving average).

Table 11: Narrative LP: Estimated Channel Impulse Responses at $h = 5$, 30, and 60 Days

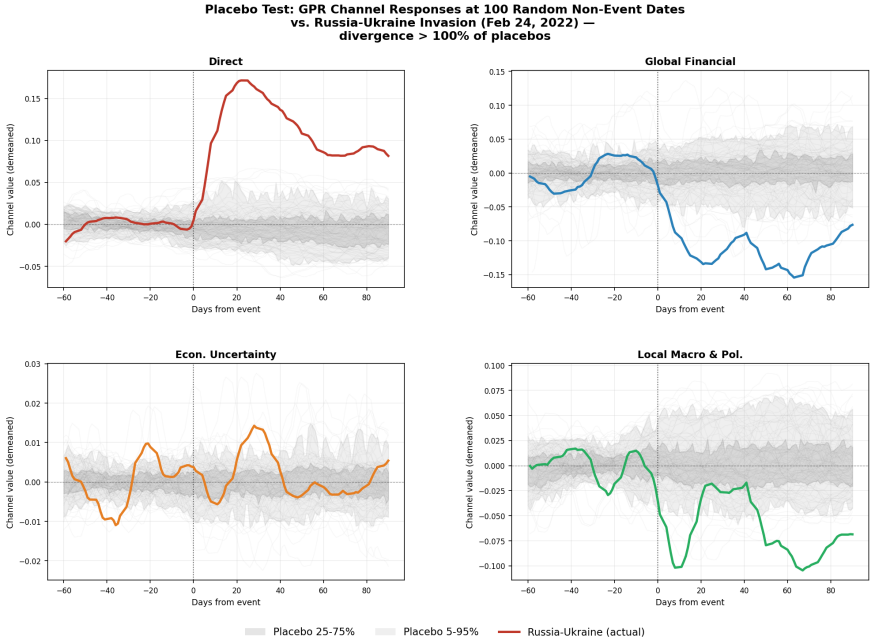
Event	Channel	$h = 5$		$h = 30$		$h = 60$	
		$\hat{\beta}$	Sig.	$\hat{\beta}$	Sig.	$\hat{\beta}$	Sig.
Russia–Ukraine GPR, Feb. 2022 *Geopolitical*	φ^{dir}	0.1031	***	0.1213	***	0.0948	***
	φ^{GFC}	-0.0782	***	-0.1035	***	-0.0808	***
	φ^{UNC}	-0.0110	***	-0.0096	*	-0.0013	–
	φ^{LOC}	-0.0728	***	-0.0234	–	-0.0540	***
Hamas–Israel GPR, Oct. 2023 *Geopolitical*	φ^{dir}	0.0710	***	0.0722	***	0.0448	***
	φ^{GFC}	0.0323	**	0.0003	–	0.0263	+
	φ^{UNC}	-0.0164	***	0.0060	–	0.0014	–
	φ^{LOC}	-0.0275	+	-0.0179	–	-0.0258	+
U.S. Election EPU, Nov. 2024 *Geoeconomic*	φ^{dir}	0.0073	***	0.0132	***	0.0084	***
	φ^{GFC}	-0.0351	***	-0.0463	***	-0.0331	***
	φ^{UNC}	0.0180	***	0.0214	***	0.0316	***
	φ^{LOC}	0.0174	***	0.0151	***	0.0337	***
"Liberation Day" TPU, Apr. 2025 *Geoeconomic*	φ^{dir}	-0.0300	***	-0.0090	***	–	–
	φ^{GFC}	-0.1386	***	-0.0194	*	–	–
	φ^{UNC}	-0.0206	***	0.0264	***	–	–
	φ^{LOC}	-0.0529	***	0.0319	**	–	–

Notes: Estimated impulse responses from narrative local projections using dated event dummies within a ± 3-day window around each crisis episode. Standard errors are clustered by country; country fixed effects are included. Contemporaneous controls: global VIX (moving average) and sovereign yield spread (moving average).
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$, $^{+}p < 0.32$ (68% CI); $--$ = not significant at the 10% level.

E.1 Placebo Falsification: Channel-Level Evidence

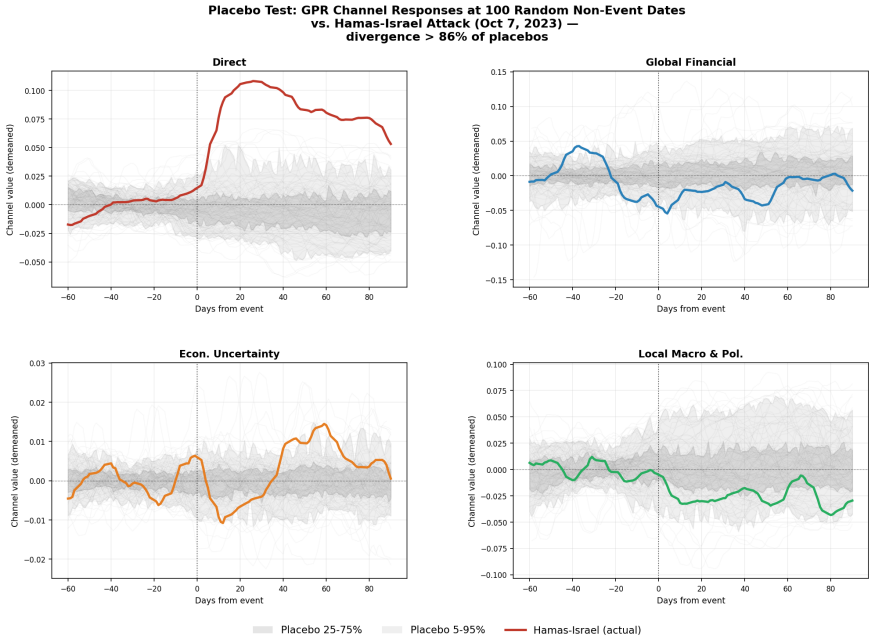
Figures 11–14 present the channel-level placebo tests for all four episodes, using the driver-specific Shapley–Taylor decomposition (GPR for geopolitical events; EPU and TPU for geoeconomic events). Each panel plots a single channel's cross-country mean response against the placebo envelope constructed from 100 randomly drawn non-event dates.

Figure 11: **Placebo Test (GPR): Russia–Ukraine, Feb. 24, 2022**



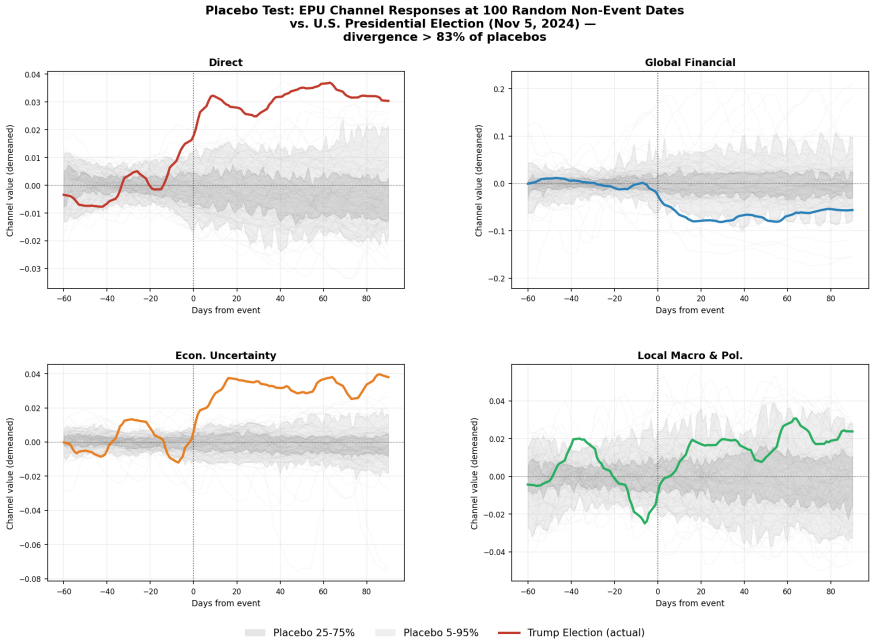
Notes: Channel-level cross-country mean responses vs. 25–75% (dark) and 5–95% (light) placebo bands from 100 random non-event dates. The Direct channel exits the 95th percentile band upward; the GFC channel exits downward. The Uncertainty channel remains within the placebo range, consistent with geopolitical—rather than geoeconomic—transmission. 7-day trailing MA; values demeaned over the pre-event window.

Figure 12: **Placebo Test (GPR): Hamas–Israel, Oct. 7, 2023 — Divergence > 86% of Placebos**



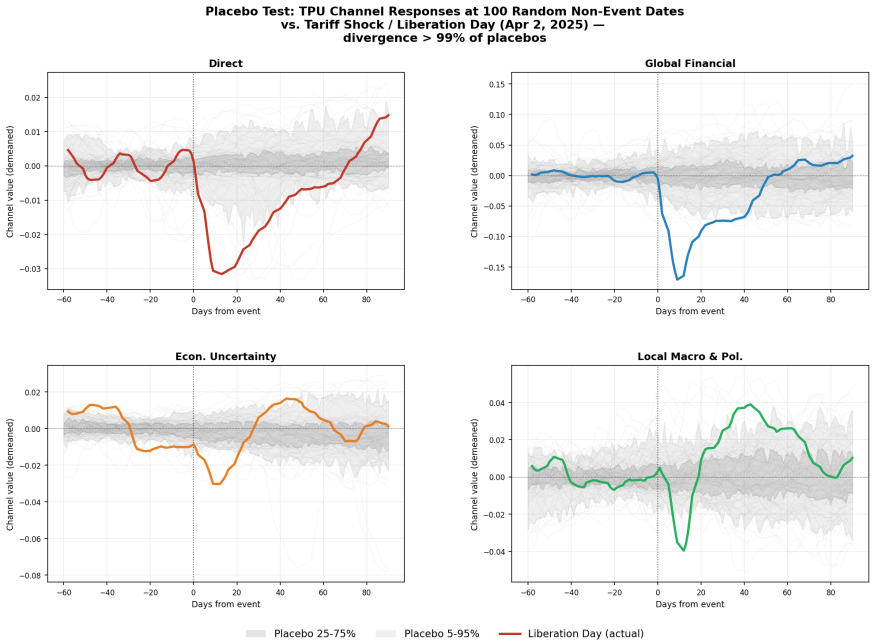
Notes: The Direct channel exits the 95th percentile band upward immediately after day 0, replicating the geopolitical signature observed for Russia–Ukraine. The GFC channel remains largely within the placebo range, consistent with the more regionally contained nature of this conflict. The Uncertainty and Local channels stay inside the bands. 25–75% (dark) and 5–95% (light) placebo bands from 100 random non-event dates. 7-day trailing MA; values demeaned over the pre-event window.

Figure 13: **Placebo Test (EPU): U.S. Presidential Election, Nov. 5, 2024 — Divergence > 83% of Placebos**



Notes: The Uncertainty channel exits the envelope upward after day 0, consistent with EPU-driven repricing. The GFC channel drops below the band. The Direct channel rises modestly above the envelope, reflecting tariff-related repricing accompanying the election. 25–75% (dark) and 5–95% (light) placebo bands from 100 random non-event dates. 7-day trailing MA; values demeaned over the pre-event window.

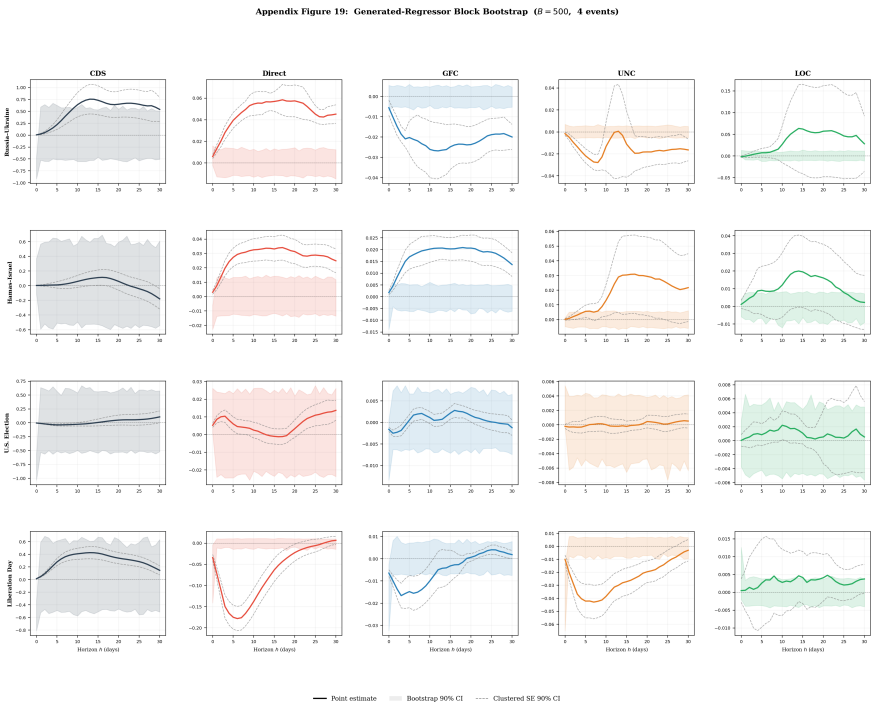
Figure 14: **Placebo Test (TPU): Liberation Day Tariff Shock, Apr. 2, 2025 — Divergence > 99% of Placebos**



Notes: The GFC channel collapses below the envelope immediately after day 0—the largest single-channel departure across all four episodes. The Direct channel moves *negative*, consistent with the taxonomy's prediction that tariff shocks bypass conflict-proximity repricing. The Local channel exits upward, reflecting heterogeneous tariff exposure. 25–75% (dark) and 5–95% (light) placebo bands from 100 random non-event dates. 7-day trailing MA; values demeaned over the pre-event window.

E.2 Block Bootstrap Validation of Channel Decomposition

Figure 15: **Block Bootstrap Validation of Narrative LP Impulse Responses ($B = 500$, 4 narrative events)**



Notes: Each panel plots the narrative LP impulse-response function for the episode indicated in the row label on the variable indicated in the column header: the raw CDS spread and the four Shapley–Taylor transmission channels ($\varphi^{\mathrm{DIR}}, \varphi^{\mathrm{GFC}}, \varphi^{\mathrm{UNC}}, \varphi^{\mathrm{LOC}}$) over horizons $h = 0, \ldots, 30$ days. Solid coloured lines: point estimates from the baseline narrative local projection (Table 5). Dashed grey lines: clustered 90% confidence intervals, robust to heteroskedasticity and cross-sectional dependence. Coloured shaded bands: 90% confidence intervals from $B = 500$ block-bootstrap replications that jointly re-estimate the Multilayer Random Forest, recompute the Shapley–Taylor decomposition via TreeSHAP, and re-run the panel local projection. For the raw CDS spread (first column), bootstrap and clustered intervals are comparable in width, confirming that the two methods agree when the dependent variable is directly observed. For the four channel series, bootstrap bands are substantially tighter because the channel decomposition is highly stable across re-estimations of the random forest trained on approximately 75,000 observations; the clustered estimator is wider because it additionally corrects for cross-sectional dependence across 42 sovereigns—a source of uncertainty the block bootstrap does not fully replicate. No channel that is significant under clustered inference loses significance under the bootstrap. Clustered standard errors are used as the conservative baseline throughout.

F Appendix F. SVAR Specification and Additional Results

This appendix provides the technical specification underlying the sign-restricted SVAR results summarised in Section 8.1, together with additional robustness exercises. The key results—the 8/8 scorecard on raw observables (Figure 8) and the cross-validation summary (Table 6)—are reported in the main text.

F.1 Specification

Observable mapping. For each country i, we estimate a five-variable SVAR on observables that map directly into the four transmission channels without passing through the machine-learning model:

$$\mathbf{Y}_{i,t} = \left(\mathrm{CDS}_{i,t},\ \mathrm{GPR}_{i,t},\ \mathrm{GFC}_{t}^{\mathrm{obs}},\ \mathrm{UNC}_{i,t}^{\mathrm{obs}},\ \mathrm{LOC}_{i,t}^{\mathrm{obs}} \right)'. \tag{29}$$

Here, $\mathrm{GPR}_{i,t}$ is the country-level geopolitical risk index that proxies the Direct channel. $\mathrm{GFC}_{t}^{\mathrm{obs}}$ is a two-variable global financial-conditions composite based on the standardised VIX and U.S. Treasury yield, normalised so that lower values correspond to the flight-to-safety configuration that offsets the Direct channel in the main text. $\mathrm{UNC}_{i,t}^{\mathrm{obs}}$ combines economic and trade policy uncertainty measures and proxies the Uncertainty channel. $\mathrm{LOC}_{i,t}^{\mathrm{obs}}$ combines domestic economic, institutional, and political indicators and proxies the Local channel.[27]

[27] Results are robust to alternative normalisations of the global financial composite, to using the underlying components individually, and to principal-component extraction within each composite.

All series are transformed into 28-day moving averages for comparability with the main text. Country-specific variables are standardised within country; global variables are standardised once over the full sample and then stacked across countries in the panel.

Reduced form. For each country i, we estimate a reduced-form VAR(p) with $p = 5$, matching the lag structure used in the main text:

$$\mathbf{Y}_{i,t} = \boldsymbol{\mu}_i + \sum_{\ell=1}^{5} \mathbf{A}_{i,\ell}\mathbf{Y}_{i,t-\ell} + \mathbf{u}_{i,t}, \qquad \mathbf{u}_{i,t} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_i). \tag{30}$$

Structural representation. Let $\mathbf{u}_{i,t} = \mathbf{B}_i\boldsymbol{\varepsilon}_{i,t}$, where $\boldsymbol{\varepsilon}_{i,t}$ contains five orthogonal structural shocks and $\mathbf{B}_i\mathbf{B}_i' = \boldsymbol{\Sigma}_i$. We label two shocks of interest: a geopolitical shock $\varepsilon^G_{i,t}$ and a geoeconomic shock $\varepsilon^E_{i,t}$. The remaining three orthogonal shocks are left unlabeled. Identification combines sign restrictions, relative-magnitude restrictions, and narrative restrictions derived from the semistructural benchmarks in Section 3.

F.2 Restriction Design

Layer 1: Sign restrictions. The semistructural framework implies sign predictions for the initial response of each observable system. We impose sign restrictions at $h = 0$ and $h = 5$ (approximately one trading week in the daily system):

Table 12: SVAR Sign Restrictions Implied by the Semistructural Framework

	Geopolitical shock (ε^G)		Geoeconomic shock (ε^E)	
Variable	h=0	h=5	h=0	h=5
$\text{CDS}_{i,t}$	+	+	–	–
$\text{GPR}_{i,t}$ (Direct)	+	+	≈ 0	≈ 0
$\text{GFC}^{\text{obs}}_t$	–	–	–	–
$\text{UNC}^{\text{obs}}_{i,t}$	–	–	+	+
$\text{LOC}^{\text{obs}}_{i,t}$	–	–	+	+

Notes: "+" and "−" denote sign restrictions; "–" denotes unrestricted. "≈ 0" for the geoeconomic response of GPR is implemented as a relative-magnitude restriction: $|\Theta^h_{\text{GPR},\varepsilon^E}| 0.5 \cdot |\Theta^h_{\text{GPR},\varepsilon^G}|$ for $h \in \{0, 5\}$. CDS is left unrestricted under ε^E because the identifying content of the taxonomy lies in the relative channel configuration rather than in a universally positive immediate spread response at all horizons.

Layer 2: Relative-magnitude restrictions. The key empirical content of the taxonomy is not only the sign of each response but also the ranking of the dominant channel. We therefore impose two relative-magnitude inequalities over the initial week:

$$\text{Geopolitical shock:} \quad \max_{h\in\{0,\ldots,5\}} |\Theta^h_{\text{GFC},\varepsilon^G}| < \max_{h\in\{0,\ldots,5\}} |\Theta^h_{\text{GPR},\varepsilon^G}| \tag{31}$$

$$\text{Geoeconomic shock:} \quad \max_{h\in\{0,\ldots,5\}} |\Theta^h_{\text{GFC},\varepsilon^E}| > \max_{h\in\{0,\ldots,5\}} |\Theta^h_{\text{GPR},\varepsilon^E}|. \tag{32}$$

The first inequality encodes the scissors logic: under a geopolitical shock, the GFC response moves in the opposite direction but remains secondary to the direct repricing channel. The second inequality captures the geoeconomic benchmark: the financial-conditions channel dominates the muted direct channel.

Layer 3: Narrative restrictions. We impose narrative restrictions on the four crisis dates used in the main text:

$t_{E1} =$ 24 Feb 2022 (Russia–Ukraine)
$t_{E2} =$ 7 Oct 2023 (Hamas–Israel)
$t_{E3} =$ 5 Nov 2024 (U.S. election)
$t_{E4} =$ 2 Apr 2025 (Liberation Day / tariff shock).

Following Antolín-Díaz and Rubio-Ramírez (2018), the restrictions are imposed country by country at these common calendar dates. We use two types of narrative conditions.

Type 1 (shock sign). The target shock must be positive on the corresponding narrative date:

$$\varepsilon^{G}_{i,t_{E1}} > 0, \qquad \varepsilon^{G}_{i,t_{E2}} > 0, \qquad \varepsilon^{E}_{i,t_{E3}} > 0, \qquad \varepsilon^{E}_{i,t_{E4}} > 0. \tag{33}$$

Type 2 (shock ranking among labeled shocks). On each narrative date, the target shock must be larger in absolute value than the other labeled shock:

$$|\varepsilon^{G}_{i,t_{E1}}| > |\varepsilon^{E}_{i,t_{E1}}|, \qquad |\varepsilon^{G}_{i,t_{E2}}| > |\varepsilon^{E}_{i,t_{E2}}|, \tag{34}$$

$$|\varepsilon^{E}_{i,t_{E3}}| > |\varepsilon^{G}_{i,t_{E3}}|, \qquad |\varepsilon^{E}_{i,t_{E4}}| > |\varepsilon^{G}_{i,t_{E4}}|. \tag{35}$$

These narrative conditions anchor the sign-restricted system on the same dated episodes used in the local-projection validation, while leaving the remaining three structural shocks unlabeled.

F.3 Estimation Algorithm and Aggregation

We follow the Bayesian accept–reject procedure in Antolín-Díaz and Rubio-Ramírez (2018). For each country, we: (i) draw reduced-form VAR parameters from the Normal–inverse-Wishart posterior; (ii) draw an orthogonal rotation matrix $\mathbf{Q}$ uniformly from the Haar measure on $O(5)$; (iii) compute structural impulse responses at horizons $h = 0, \ldots, 60$; and (iv) retain the draw if all sign, relative-magnitude, and narrative restrictions are satisfied.

The baseline raw-observable system yields 639 accepted draws. For each accepted draw and each horizon h, we compute a mean-group response as the unweighted cross-country average of the corresponding country-level IRF. Pointwise posterior medians and 68%/90% credible bands are then constructed from the distribution of these mean-group IRFs across accepted draws. This aggregation preserves country-level heterogeneity in the reduced form while reporting a single cross-country response profile.

F.4 Internal Coherence Check: SVAR on Shapley–Taylor Channel Series

As an internal coherence check, we estimate the same sign-restricted system on the Shapley–Taylor channel series $(\varphi^{\text{DIR}}, \varphi^{\text{GFC}}, \varphi^{\text{UNC}}, \varphi^{\text{LOC}})$ together with CDS. This exercise is not independent of the machine-learning decomposition, but it tests whether the channel series admit a structural representation aligned with the semistructural benchmarks. The channel SVAR yields the same qualitative 8/8 scorecard, with somewhat tighter credible bands and 617 accepted draws (Figure 16).

F.5 Full-Sample LP with SVAR-Identified Shocks

We use the identified structural shocks $\hat{\varepsilon}^{G}_{i,t}$ and $\hat{\varepsilon}^{E}_{i,t}$ as regressors in full-sample panel local projections with country fixed effects and Driscoll–Kraay standard errors. In these regressions, the responses of both the raw observables and CDS are economically small and statistically indistinguishable from zero across horizons. This finding is stable across three control sets: (A) lags only, (B) lags plus global financial controls, and (C) the full contemporaneous control set shown in Figure 18.

Taken together, Figures 8 and 17 mirror the main-text contrast between sharp responses in narrative episode designs and weak average responses in full-sample innovation designs. The SVAR evidence therefore corroborates, rather than mechanically restates, the state-dependent interpretation of sovereign-risk transmission.

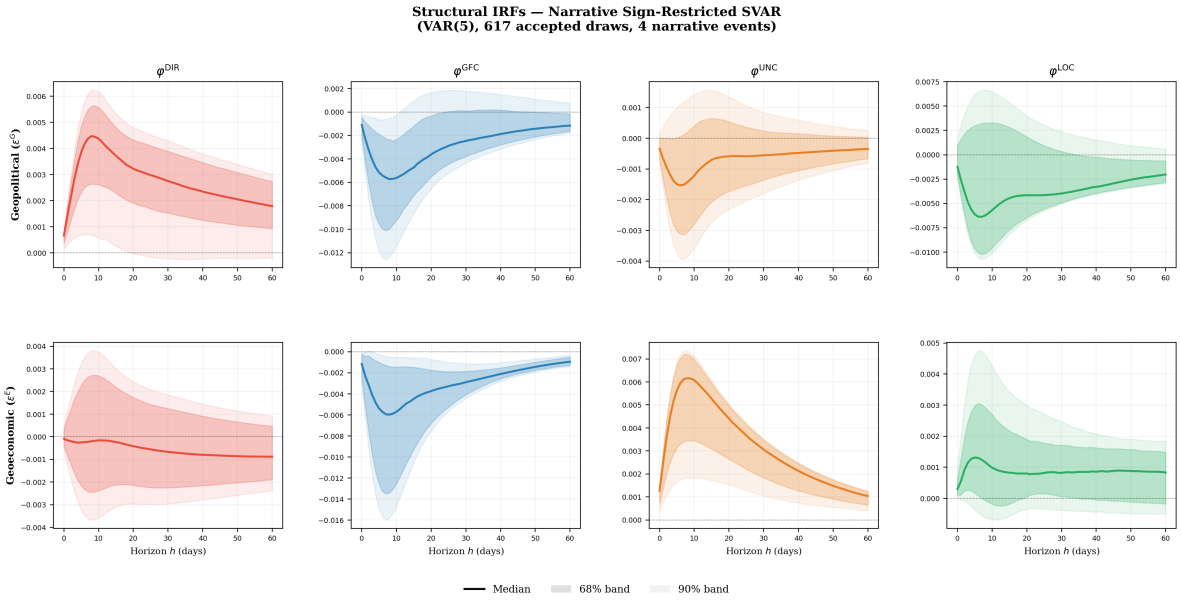
F.6 Interpretation: Why the Full-Sample LP Attenuates

The identified SVAR shocks are linear combinations of reduced-form innovations with time-invariant weights:

$$\hat{\boldsymbol{\varepsilon}}_{i,t} = \mathbf{B}_i^{-1}\mathbf{u}_{i,t}. \tag{36}$$

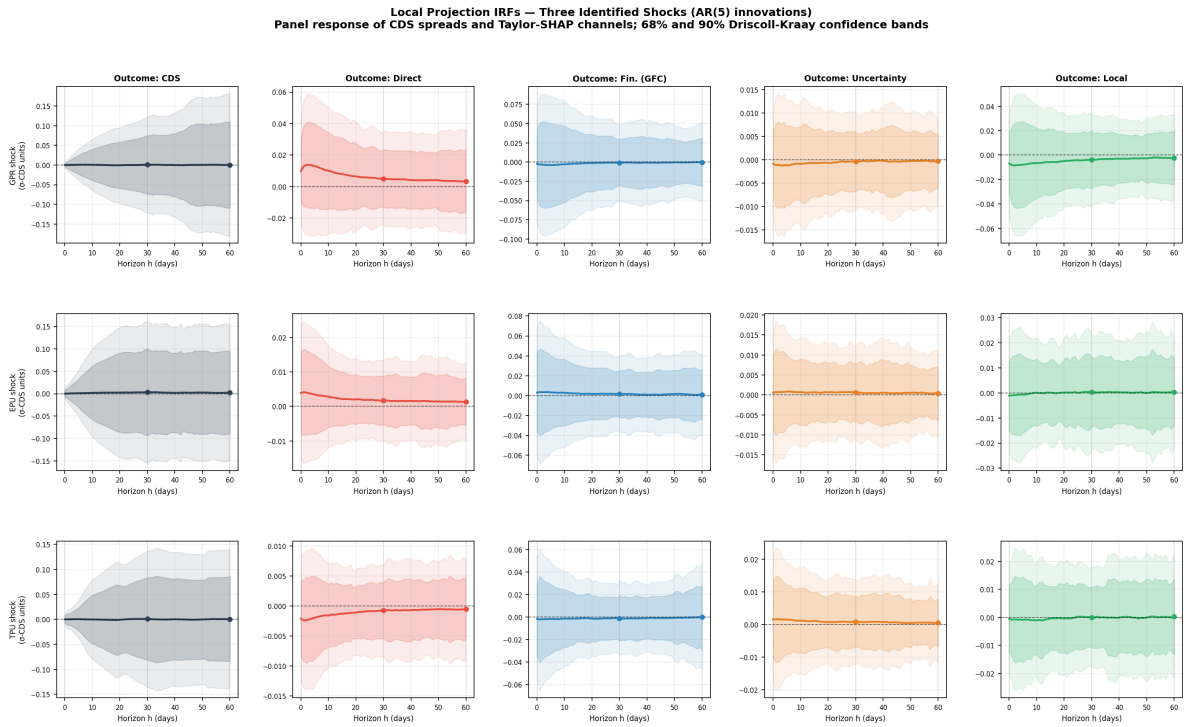
Because these weights do not vary with the state of the financial system, a full-sample LP averages over states in which transmission is active and states in which it is dormant, attenuating the average response toward zero. The SVAR structural IRFs are less affected because the narrative restrictions concentrate identification power on tail events, and the VAR propagates shocks through a cross-equation dynamic system (Plagborg-Møller and Wolf, 2021). The combined pattern—sharp structural IRFs in the narrative SVAR but weak average responses in full-sample LPs—is consistent with state-dependent transmission.

Figure 16: **Narrative Sign-Restricted SVAR: Structural Impulse Responses on Shapley–Taylor Channel Series**



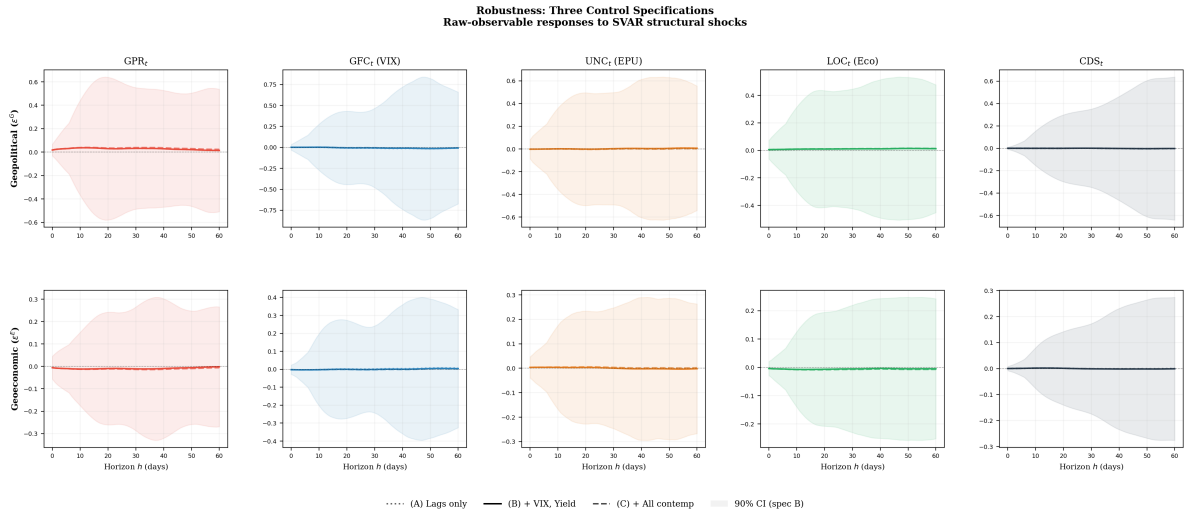
Notes: Mean-group structural impulse responses from country-level VAR(5) models estimated on the four Shapley–Taylor channel series (φ^{DIR}, φ^{GFC}, φ^{UNC}, φ^{LOC}), identified with sign restrictions (Table 12), relative-magnitude restrictions, and narrative restrictions on four crisis dates following Antolín-Díaz and Rubio-Ramírez (2018). Top row: geopolitical shock (ε^{G}); bottom row: geoeconomic shock (ε^{E}). Solid lines: posterior median across 617 accepted draws. Dark shaded bands: 68% pointwise credible intervals; light bands: 90%. The geopolitical shock raises φ^{DIR} while φ^{GFC} falls, reproducing the scissors pattern. The geoeconomic shock produces a muted and eventually negative φ^{DIR}, with the GFC, Uncertainty, and Local channels carrying transmission. All eight sign and dominance predictions are confirmed at the posterior median. Because this system uses Shapley–Taylor values as inputs, it serves as an internal coherence check rather than an independent validation; Figure 8 provides the independent test on raw observables.

Figure 17: **Full-Sample Local Projections with SVAR-Identified Shocks: Attenuation Toward Zero**



Notes: Panel local-projection impulse responses of sovereign CDS spreads and the four Shapley–Taylor channels (φ^{DIR}, φ^{GFC}, φ^{UNC}, φ^{LOC}) to SVAR-identified structural shocks used as regressors over the full sample. Rows: GPR shock (top), EPU shock (middle), TPU shock (bottom). Columns: raw CDS spread and the four channel series. Structural shocks are extracted from country-level sign-restricted SVARs ($\hat{\varepsilon}_{i,t} = B_i^{-1} u_{i,t}$; see Appendix F for identification details); standard errors follow Driscoll and Kraay (1998) with bandwidth $\max(20, h)$, country and time fixed effects, and 5 lags. Dark shaded bands: 68% confidence intervals; light bands: 90%. Dots at $h = 30$ mark the horizon reported in Table 5. All responses are economically small and statistically indistinguishable from zero at conventional levels, consistent with state-dependent transmission concentrated in discrete crisis episodes rather than a pervasive linear relation in the full sample. This attenuation mirrors the innovation-based null documented in Section 7.1 and contrasts with the large and significant responses recovered by the narrative design (Figure 5).

Figure 18: **Robustness of Full-Sample LP Attenuation Across Three Control Sets**



Notes: Raw-observable impulse responses to SVAR-identified structural shocks under three control specifications: (A) lags only, (B) lags plus VIX and U.S. Treasury yield, and (C) the full contemporaneous control set. Top row: geopolitical shock (ε^G); bottom row: geoeconomic shock (ε^E). Columns: GPR (Direct proxy), GFC composite, Uncertainty composite, Local composite, and sovereign CDS spread. Shaded bands: 90% Driscoll–Kraay confidence intervals (specification B). All responses remain economically small and statistically indistinguishable from zero across control sets, confirming that the full-sample attenuation documented in Figure 17 is not driven by the choice of controls.