Journal of International Law & Human Rights
ISSN (Print): 3007-0120 ISSN (Online): 3007-0139
Volume 5, No.9, 2026, 29-56
DOI: https://doi.org/ 10.59075/jilhr.v5i9.319
Journal Homepage: http://journals.centeriir.org/index.php/jilhr

Collateral Disruption Reconsidered: A Deep Machine Learning Framework for Modeling Third-Party Sanctions Spillovers, Trade Diversion Dynamics, and De-Risking Behavior in Pakistan's Cross-Border Banking Compliance Regime

Zia Ur Rehman[1], Dr. Muhammad Usman Malik[2], Dr. Zulfqar Khan[3]

[1]National Institute of Intelligence & Security Studies, Quaid-i-Azam University, Islamabad 45320, Pakistan,
Email: ziawanianziawanian@gmail.com
[2]National Institute of Intelligence & Security Studies, Quaid-i-Azam University, Islamabad 45320, Pakistan,
Email: usmanmalik836@gmail.com
[3] Prof. National Institute of Intelligence & Security Studies, Quaid-i-Azam University, Islamabad 45320, Pakistan,
Email: dirniiss@qau.edu.pk

Abstract

Purpose. Pakistan is not the target of any comprehensive sanctions programme, yet over two decades its banks, payment channels and trade routes have absorbed the costs of measures aimed at others, chiefly Iran and Russia, while three spells on the Financial Action Task Force (FATF) grey list and two US penalties on Pakistani banks raised compliance costs at home. This paper asks how large these third-party, or "collateral", disruptions are, through which channels they travel, and whether deep learning adds anything that conventional econometrics misses.

Design/methodology/approach. We build an annual dataset for 2005–2024 that combines World Bank macro-banking indicators, UN Comtrade bilateral trade with six partners, and an author-coded sanctions-compliance pressure index. A two-stage framework first identifies spillovers with bootstrap ARDL bounds tests, nonlinear ARDL, dynamic OLS, Jordà local projections, a Minnesota-prior Bayesian VAR with Diebold–Yilmaz connectedness, Hansen threshold regression, double/debiased machine learning and Toda–Yamamoto causality. It then trains LSTM, GRU, two-layer attention-LSTM and deep multilayer-perceptron networks, a deep ensemble and a denoising autoencoder, evaluates them out of sample with Diebold–Mariano tests and a model confidence set, and explains them with exact group-Shapley values.

Findings. Compliance pressure is associated with a persistent contraction of private credit: the dynamic-OLS long-run semi-elasticity is −25.3 percentage points of GDP per index unit, the bootstrap bounds test finds a levels relationship at the 10% level, and pressure Granger-causes credit but not the reverse. A one-standard-deviation pressure shock cuts reserves by about 0.7 months of imports on impact and lowers credit by up to 0.45 points of GDP after two years. Iran sanctions intensity shifts recorded trade away from Iran and towards UAE intermediation. Remittances follow a threshold regime: above a pressure index of about 0.46 they become sensitive to currency depreciation, consistent with leakage to informal channels. Deep models reduce forecast error by up to 22% relative to a prevailing-mean benchmark, but no gain is statistically significant in eleven annual hold-out years. Their value lies in attribution: sanctions channels account for 32–46% of the deep ensemble's explanatory mass, with Iran sanctions dominating trade diversion and US enforcement and the Russia regime dominating the 2022–2024 remittance shortfall.

Originality/value. The paper is the first to integrate a full small-sample econometric identification strategy with deep sequence models and exact Shapley attribution for a third-country sanctions spillover problem, and it reports, rather than conceals, the limits of deep learning on short macro samples.

Keywords: secondary sanctions; de-risking; correspondent banking; FATF grey list; trade diversion; deep learning; LSTM; Shapley values; Bayesian VAR; nonlinear ARDL; Pakistan

© 2026 The Authors. Published by Center of Innovation in Interdisciplinary Research (CIIR). This is an Open Access Article under the Creative Common Attribution Non-Commercial 4.0

Article History: Received: 05-09-2026 Accepted: 16-10-2026 Published: 03-11- 2026

Corresponding Author Email: ziawanianziawanian@gmail.com

JEL classification: F14; F38; F51; G21; G28; C32; C45; C53

Introduction

Sanctions are designed to hurt their targets. Their reach, however, rarely stops at the target's border. When the United States, the European Union or the United Nations restrict dealings with a country, the cost of enforcing those restrictions is shifted onto every bank, shipper and payment intermediary that might touch a prohibited counterparty. For a neighbouring economy with porous borders, a large diaspora and a banking system that depends on a thin layer of correspondent relationships, the result can be a form of collateral disruption: higher compliance costs, withdrawn correspondent accounts, payments pushed into informal channels and trade re-routed through third hubs. None of these outcomes is the aim of the sanctions, and few are measured.

Pakistan offers an unusually clear setting for studying this problem. It shares a 900-kilometre border with Iran, the most heavily sanctioned large economy of the past two decades. It hosts one of the world's largest remittance corridors, with inflows that rose from US$4.3 billion in 2005 to almost US$35 billion in 2024, much of it originating in the Gulf. It spent close to ten of the twenty years between 2005 and 2024 on the FATF list of jurisdictions under increased monitoring, in three separate stints (February 2008–June 2010, February 2012–February 2015 and June 2018–October 2022). Two of its largest banks were penalized by US regulators: Habib Bank Limited paid US$225 million to the New York State Department of Financial Services in 2017 and closed its New York branch, and National Bank of Pakistan paid US$55.4 million to the Federal Reserve and the same regulator in 2022. Since 2022 Pakistani banks have also had to screen trade and payments against the expanding Russia sanctions regime. Pakistan is the target of none of the sanctions programmes in question, yet each of them reaches its financial system through the correspondent-banking network.

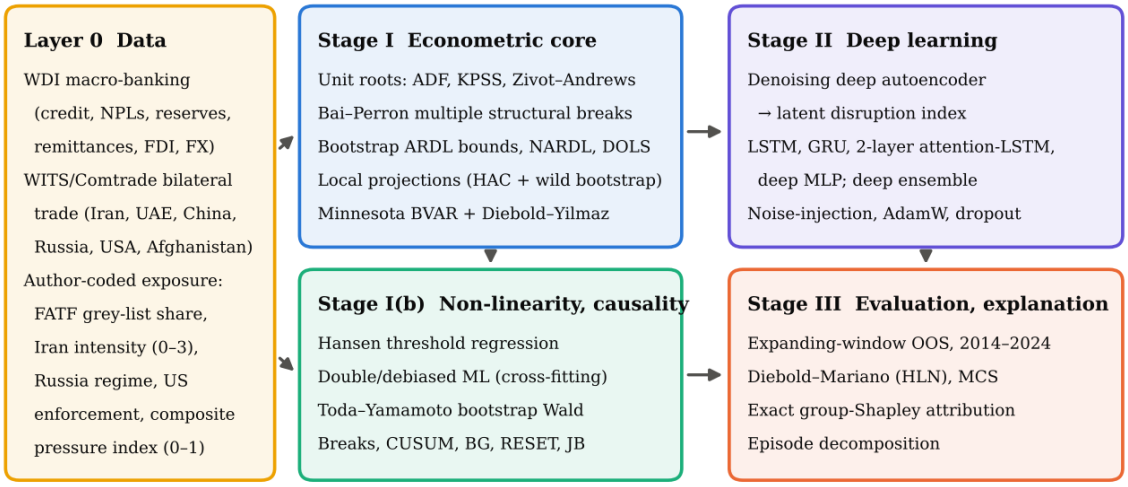
The economics literature has documented two related phenomena in some detail. Work on the direct effects of sanctions shows large trade and output losses for target countries and substantial trade deflection through third countries (Hufbauer et al., 2007; Haidar, 2017; Felbermayr et al., 2020; Morgan et al., 2023). A separate literature on "de-risking" shows that global banks have cut correspondent relationships with smaller and higher-risk jurisdictions since around 2011, with measurable effects on payment flows and trade (World Bank, 2015; Erbenova et al., 2016; Rice et al., 2020). Kida and Paetzold (2021), using machine learning on a cross-country panel, estimate that FATF grey-listing reduces capital inflows by several percentage points of GDP. What is largely missing is an integrated account of how third-party sanctions, AML/CFT listing and enforcement shocks combine to disrupt a single non-target economy, and of which channels (credit, reserves, remittances, foreign investment or trade routing) carry the disruption.

This paper fills that gap with three contributions. The first is a measurement contribution. We assemble an annual dataset for Pakistan for 2005–2024 that links official macro-banking indicators from the World Bank, Pakistan-reported bilateral merchandise trade with Iran, the United Arab Emirates (UAE), China, the United States, Afghanistan and Russia, and a set of exposure variables coded from dated public events: the share of each year spent on the FATF grey list, an ordinal measure of Iran sanctions intensity, a Russia-regime indicator, US enforcement actions against Pakistani banks, and a composite sanctions-compliance pressure index.

The second is a methodological contribution. Short annual samples are where both conventional econometrics and machine learning are most fragile, so we combine them deliberately. The econometric stage uses procedures chosen for their behaviour in small samples and under uncertain integration orders: unit-root tests that allow an endogenous break (Zivot and Andrews, 1992), Bai–Perron multiple-break dating, ARDL bounds testing with bootstrap rather than asymptotic critical values (McNown et al., 2018), nonlinear ARDL to separate tightening from easing (Shin et al., 2014), dynamic OLS, local projections with wild-bootstrap inference, a Bayesian VAR with a Minnesota prior and Diebold–Yilmaz connectedness, Hansen threshold regression, double/debiased machine learning and Toda–Yamamoto causality. The deep-learning stage trains long short-term memory

(LSTM), gated recurrent unit (GRU), attention-augmented LSTM and deep feed-forward networks, combines them in a deep ensemble, builds a denoising autoencoder index of latent disruption, and opens the black box with exact group-Shapley values. Every network was implemented from first principles with verified gradients, and every hyperparameter was fixed before the hold-out evaluation. Fig. 1 summarizes the design.

Fig. 1**. The two-stage econometric and deep-learning framework**.



The third is substantive. We find that compliance pressure is associated with a persistent contraction of private credit; that pressure shocks deplete foreign-exchange reserves on impact; that Iran sanctions intensity shifts recorded trade from Iran to UAE intermediation; that remittances behave differently above a pressure threshold in a way that is consistent with leakage into informal channels when the rupee depreciates; and that tightening appears to transmit faster than easing. The deep models do not deliver statistically significant forecasting gains over simple benchmarks on eleven annual hold-out observations, a result we report in full because it is what the data show and because it is consistent with broader evidence on machine learning in small macroeconomic samples (Makridakis et al., 2018). Their contribution is explanatory: they attribute between a third and nearly half of the variation they capture to the four sanctions channels, and they locate when and through which channel the collateral disruption occurred.

The title echoes Hufbauer et al.'s *Economic Sanctions Reconsidered*, whose case-based assessment of sanctions effectiveness set the agenda for the field. Our reconsideration shifts the focus from the target to the bystander. The rest of the paper is organised as follows. Section 2 reviews the literature and states the hypotheses. Section 3 describes the institutional background. Section 4 presents the data and the construction of the exposure variables. Section 5 sets out the econometric and deep-learning framework. Section 6 reports the results, Section 7 the robustness and causality checks, and Section 8 discusses implications and limitations. Section 9 concludes.

Literature review and hypotheses

Sanctions, Spillovers and Trade Deflection

The modern empirical literature on sanctions grew out of case studies (Hufbauer et al., 2007) and has moved towards large panel datasets such as the Global Sanctions Data Base (Felbermayr et al., 2020). Two findings recur. First, sanctions reduce trade between sender and target substantially, though the size of the effect depends

on the scope of the measure and on the availability of substitutes (Crozet et al., 2021; Morgan et al., 2023). Second, a large share of the lost trade is deflected rather than destroyed. Haidar (2017) shows, using Iranian customs records, that roughly two-thirds of non-oil exports that sanctions would have blocked were redirected to non-sanctioning destinations. Early (2015) documents how third-country "sanctions busters", with the UAE prominent among them, take over trade that sender governments seek to cut off. Studies of Iran's macro economy find sizeable output losses (Dizaji and van Bergeijk, 2013; Gharehgozli, 2017), and cross-country evidence finds that UN and US sanctions lower target-country growth (Neuenkirch and Neumeier, 2015). The Russia sanctions since 2014, and especially since 2022, have prompted modelling of their trade and welfare effects on senders and third countries (Dreger et al., 2016; Mahlstein et al., 2022).

Much less is known about the economies that sit between senders and targets without being either. For them, the relevant mechanism is not a legal prohibition but the extraterritorial reach of secondary sanctions, which threaten foreign banks with loss of access to the US financial system if they deal with designated parties. Ahn and Ludema (2020) show that targeted sanctions shift activity away from designated firms; Besedeš et al. (2017) and Efing et al. (2023) show that financial sanctions cause banks to reduce cross-border exposures well beyond the sanctioned entities themselves. These are the micro foundations of collateral disruption.

De-Risking, Correspondent Banking and AML/CFT Listing

De-risking refers to the termination or restriction of business relationships with whole categories of clients or jurisdictions to avoid, rather than manage, compliance risk (World Bank, 2015; Durner and Shetret, 2015). Correspondent relationships, which give domestic banks access to foreign currencies and payment systems, fell by about a fifth worldwide between 2011 and 2018, with the steepest declines in smaller, remittance-dependent and higher-risk jurisdictions (Rice et al., 2020). Survey and supervisory evidence attributes the retreat to rising enforcement penalties, the cost of know-your-customer checks and the low profitability of small correspondent accounts (Artingstall et al., 2016; Erbenova et al., 2016). Collin et al. (2016), using SWIFT data, find that the strengthening of AML regulation is associated with lower payment flows to affected countries.

FATF grey-listing is a particularly visible risk signal. Kida and Paetzold (2021) estimate with machine-learning methods on a large country panel that grey-listing reduces capital inflows by around 7.6% of GDP. For banks, listing triggers enhanced due diligence by foreign correspondents, raising the cost of trade finance and remittance processing and, at the margin, closing accounts. The literature therefore suggests that compliance shocks, whatever their origin, should work through banking intermediation and the external liquidity position.

Machine Learning in Macroeconomic and Sanctions Research

Machine learning has entered empirical macroeconomics mainly as a forecasting tool. Gains are most reliable when many predictors and long samples are available and when the relationship is nonlinear (Medeiros et al., 2021; Goulet Coulombe et al., 2022). In short samples, simple statistical benchmarks are hard to beat, and complex learners frequently overfit (Makridakis et al., 2018). Recurrent networks such as the LSTM (Hochreiter and Schmidhuber, 1997) and GRU (Cho et al., 2014), and attention mechanisms (Bahdanau et al., 2015), were designed for long sequences; applying them to annual data requires strong regularization and, ideally, ensembles (Lakshminarayanan et al., 2017). The second role of machine learning in economics is causal: double/debiased machine learning (Chernozhukov et al., 2018) uses flexible learners to partial out confounders while retaining valid inference on a target parameter (Mullainathan and Spiess, 2017; Athey and Imbens, 2019). The third role, increasingly important for policy work, is explanation: Shapley-value attribution (Shapley, 1953; Lundberg and Lee, 2017) decomposes a model's prediction into additive contributions from its inputs.

We are not aware of a study that brings these three roles together for a sanctions-spillover question, or that pairs them with a full small-sample econometric identification strategy so that the machine-learning results can be checked against well-understood estimators.

Hypotheses

The literature implies six testable propositions for a non-target economy such as Pakistan.

- **H1 (collateral credit contraction).** Higher sanctions-compliance pressure is associated with a lower long-run level of private credit relative to GDP.
- **H2 (external-liquidity channel).** Pressure shocks reduce foreign-exchange reserves in the short run, as correspondent-banking frictions raise the cost of external financing and trade settlement.
- **H3 (trade diversion).** Greater Iran sanctions intensity reduces recorded trade with Iran and raises trade intermediated through the UAE.
- **H4 (asymmetric pass-through).** Tightening of compliance pressure transmits faster or more strongly than easing, because relationships cut during de-risking are slow to be restored.
- **H5 (regime-dependent payment channels).** The behaviour of formal remittance inflows differs above and below a critical level of pressure.
- **H6 (predictive content of deep architectures).** Deep sequence models that allow for nonlinearity and interaction among sanctions channels forecast spillover outcomes better than linear and naive benchmarks.

Institutional Background

Three overlapping regimes shape the compliance environment of Pakistani banks.

Iran Sanctions. The United States has maintained unilateral sanctions on Iran throughout the sample. UN Security Council Resolution 1737 (December 2006) added multilateral measures. The US Comprehensive Iran Sanctions, Accountability, and Divestment Act (CISADA) of July 2010 introduced the threat that foreign banks dealing with designated Iranian entities could lose US correspondent access, the defining feature of secondary sanctions. In March 2012 Iranian banks were disconnected from SWIFT, and in July 2012 the EU oil embargo took effect. The Joint Comprehensive Plan of Action (JCPOA) brought relief from January 2016, but the United States withdrew in May 2018 and re-imposed sanctions in August and November 2018, followed by the "maximum pressure" campaign from 2019. For Pakistan the formal payment channel to Iran effectively closed in 2012, and trade migrated to barter, informal border trade and routing through the UAE.

FATF Listing. Pakistan was placed under increased monitoring in February 2008 and removed in June 2010, re-listed in February 2012 and removed in February 2015, and listed again in June 2018 under a 27-point action plan, later extended to 34 items, before its removal in October 2022. Each listing raised the due-diligence burden imposed by foreign correspondents on Pakistani banks.

Enforcement and the Russia Regime. In 2016 National Bank of Pakistan signed a written agreement with US supervisors over anti-money-laundering deficiencies. In September 2017 the New York regulator fined Habib Bank Limited US$225 million and the bank closed its New York branch, an important dollar-clearing node for the domestic system. In February 2022 the Federal Reserve and the New York regulator penalized National Bank of Pakistan a total of US$55.4 million. In the same month the G7 and EU began imposing sanctions on Russia after its invasion of Ukraine, adding another screening layer for Pakistani trade and payments. Finally, in September 2025, after our sample ends, UN "snapback" sanctions on Iran were re-imposed, which makes the questions studied here newly relevant.

Data

Sources and Coverage

The dataset is annual and covers 2005–2024 for macro-banking variables and 2005–2023 for bilateral trade. Macro, external-sector and banking indicators come from the World Bank World Development Indicators (WDI):

GDP, real growth, CPI inflation, the official PKR/US$ exchange rate, exports and imports of goods and services, personal remittances received, net FDI inflows, total reserves including gold, the current-account balance, domestic credit to the private sector and the bank non-performing loan (NPL) ratio. Derived ratios include reserves in months of imports, remittances and FDI as shares of GDP and the year-on-year depreciation of the rupee. Bilateral merchandise trade with Iran, the UAE, China, the United States, Afghanistan, Russia and the world is from World Bank WITS/UN Comtrade, as reported by Pakistan, in thousands of US dollars. Table 1 lists the variables, units, sources and summary statistics. The figures are calendar-year values and therefore do not match Pakistan's July–June fiscal year used by the State Bank of Pakistan.

Table 1. Variables, Units, Sources and Descriptive Statistics, Pakistan 2005–2024

Variable	Unit	N	Mean	S.D.	Min	Max	Source
Composite sanctions-compliance pressure index	0–1	20	0.38	0.20	0.08	0.96	Author-coded
FATF grey-list exposure	share of year	20	0.48	0.46	0.00	1.00	Author-coded (FATF)
Iran sanctions intensity	0–3	20	2.35	0.75	1.00	3.00	Author-coded
US penalties on Pakistani banks	US$ mn	20	14.02	51.18	0.00	225.00	Fed/NYDFS orders
Private credit	% GDP	20	16.83	3.73	11.46	24.17	WDI
Bank NPLs	% loans	12	8.97	2.30	5.75	12.99	WDI
Personal remittances received	% GDP	20	5.99	1.94	2.95	9.39	WDI
Remittance growth	% (log diff.)	19	11.05	10.03	-12.77	27.36	WDI, derived
FDI net inflows	% GDP	20	1.03	0.81	0.34	3.04	WDI
Total reserves	months of imports	20	3.81	1.24	1.41	5.44	WDI
PKR depreciation vs US$	% y/y	19	8.92	10.63	-0.63	36.85	WDI, derived
Real GDP growth	%	20	3.78	2.25	-1.27	7.28	WDI
CPI inflation	%	20	10.82	6.62	2.53	30.77	WDI
Current account balance	% GDP	20	-2.49	2.03	-7.74	0.12	WDI
Iran share of total trade	%	19	0.92	0.61	0.32	2.46	WITS
UAE share of total trade	%	19	10.48	1.85	7.90	14.74	WITS
UAE/Iran trade ratio	ratio	19	17.04	11.65	4.05	41.35	WITS, derived

Notes: WDI = World Bank World Development Indicators; WITS = World Bank WITS/UN Comtrade (Pakistan-reported). Bilateral trade ends in 2023; NPLs begin in 2013. Coded variables follow the rules in Section 4.2.

Pakistan does not publish bank-level counts of correspondent relationships, and the Bank for International Settlements reports them only in aggregate. We therefore measure de-risking through its macro footprint: private-credit depth, reserves, remittances, FDI and, for the banking system's asset quality, the NPL ratio (available from 2013). Trade with Iran and Afghanistan is substantially informal, so the recorded flows understate true activity; we treat the recorded Iran share and the UAE/Iran ratio as measures of the formal routing of trade, which is precisely the margin on which sanctions-compliance frictions operate.

Coding the Exposure Variables

Five exposure measures are coded from dated public events. FATF exposure is the number of months in each calendar year spent on the grey list divided by twelve. Iran sanctions intensity is an ordinal score: 1 for US unilateral sanctions only; 2 when UN Security Council sanctions are in force (from December 2006) or US measures are partially re-imposed; 3 for SWIFT disconnection and the EU oil embargo (2012–2015) or US "maximum pressure" (from 2019). The JCPOA years 2016–2017 are coded 1 and 2018 is coded 2. The Russia-regime indicator equals 1 from 2022. The enforcement indicator equals 1 in years in which a US penalty was imposed on a Pakistani bank (2017 and 2022). The composite pressure index is the equal-weighted mean of four components scaled to the unit interval:

$$P_t = ¼ (FATF_t + IRAN_t/3 + RUS_t + ENF_t) \quad (1)$$

The index ranges from 0.08 in 2005 and 2016 to 0.96 in 2022, when Pakistan was simultaneously grey-listed, subject to US penalties and screening against the new Russia regime under maximum-pressure Iran sanctions. Because the index is author-constructed, we report every result that uses it alongside results for its components, and we note that equal weights are a transparent default rather than an estimate.

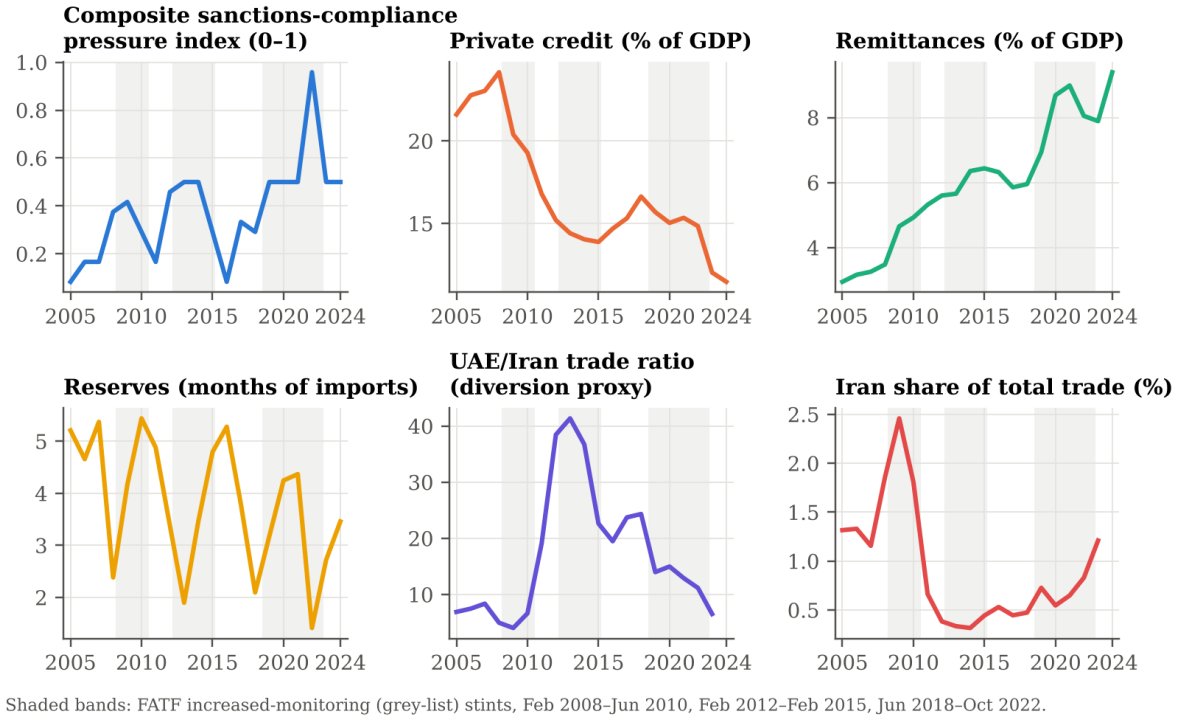


Fig. 2. Exposure and outcome series, 2005–2024. The shaded bands mark FATF grey-list stints. Private credit falls in two steps, after 2008 and after 2021; the UAE/Iran ratio jumps when the formal payment channel to Iran closes in 2012.

Fig. 2 shows the main series. Private credit fell from 24.2% of GDP in 2008 to 11.5% in 2024. Remittances rose from 2.9% to 9.4% of GDP. Reserves collapsed in each crisis year (2008, 2013, 2018 and 2022). The ratio of trade with the UAE to trade with Iran rose from about 4 in 2009 to over 41 in 2013, as the formal Iran channel closed and trade moved through Dubai, before declining as direct trade partially recovered.

Methodology

Overall Design

As Fig. 1 shows, Stage I uses econometric estimators with known small-sample properties to establish which relationships exist, their sign and persistence, and whether they are linear or regime-dependent. Stage II uses deep learning to model the same outcomes flexibly. Stage III disciplines the deep models with honest out-of-sample evaluation and decomposes their predictions into channel-level contributions. The two stages are complementary: the econometrics anchors interpretation, and the deep models test whether nonlinear interactions among channels contain information that the linear estimators miss.

Integration Properties and Structural Breaks

We test each series with the augmented Dickey–Fuller (ADF) test (Dickey and Fuller, 1979), using MacKinnon (2010) response-surface critical values adjusted to the effective sample, and the KPSS stationarity test (Kwiatkowski et al., 1992). Because sanctions episodes may induce level shifts that bias conventional tests towards non-rejection, we also apply the Zivot and Andrews (1992) test, which allows one endogenous break in intercept and trend:

$$\Delta y_t = \mu + \beta t + \theta DU_t(\tau) + \gamma DT_t(\tau) + \alpha y_{t-1} + \sum_{j=1}^{k} c_j \Delta y_{t-j} + \varepsilon_t \qquad (2)$$

where DU and DT are level- and trend-shift dummies at break date τ, which is chosen to minimize the t-statistic on α. We date multiple breaks in the mean of each outcome with the Bai and Perron (1998, 2003) procedure, which minimizes the global sum of squared residuals by dynamic programming:

$$(T_1, \ldots, T_m) = arg\ min_{T_1,\ldots,T_m} \sum_{i=1}^{m+1} \sum_{t=T_{i-1}+1}^{T_i} (y_t - \mu_i)^2 \qquad (3)$$

with the number of breaks chosen by the Bayesian information criterion.

Bootstrap ARDL Bounds Testing and Nonlinear ARDL

Our series mix I(0) and I(1) processes, which rules out Johansen-type cointegration. The autoregressive distributed-lag (ARDL) bounds approach of Pesaran et al. (2001) is valid for any such mixture. We estimate the unrestricted error-correction form

$$\Delta y_t = c + \rho y_{t-1} + \theta' x_{t-1} + \varphi' \Delta x_t + \varepsilon_t \qquad (4)$$

and test H_0: $\rho = \theta = 0$ with an F-statistic. With fewer than twenty observations the asymptotic bounds, and even Narayan's (2005) small-sample tables, may be unreliable, so we follow McNown et al. (2018) and generate critical values by a residual bootstrap under the null of no levels relationship (2,000 replications). Long-run coefficients are $-\theta/\rho$, with delta-method standard errors.

To test H4 we decompose the pressure index into partial sums of positive and negative changes (Shin et al., 2014):

$$P_t^+ = \sum_{s=1}^{t} max(\Delta P_s, 0), \quad P_t^- = \sum_{s=1}^{t} min(\Delta P_s, 0) \qquad (5)$$

$$\Delta y_t = c + \rho y_{t-1} + \theta^+ P_{t-1}^+ + \theta^- P_{t-1}^- + \varphi^+ \Delta P_t^+ + \varphi^- \Delta P_t^- + \delta' z_{t-1} + \varepsilon_t \qquad (6)$$

Long-run symmetry ($\theta^+ = \theta^-$) and short-run symmetry ($\varphi^+ = \varphi^-$) are tested with Wald statistics, and we trace cumulative dynamic multipliers for a permanent unit tightening and a unit easing.

Dynamic OLS

Because the ARDL estimates of long-run coefficients can be imprecise when adjustment is slow, we also estimate long-run relationships by dynamic OLS (Stock and Watson, 1993), which adds leads and lags of the differenced regressors to remove endogeneity and serial-correlation bias:

$$y_t = \alpha + \beta' x_t + \sum_{j=-q}^{q} \gamma_j' \Delta x_{t+j} + u_t \qquad (7)$$

with q = 1 and Newey and West (1987) standard errors. DOLS is our main estimator of the trade-diversion elasticity, with the log UAE/Iran ratio and the Iran share of trade regressed on Iran sanctions intensity and FATF exposure.

Local Projections

Impulse responses are estimated directly with local projections (Jordà, 2005), which are robust to misspecification of the data-generating process:

$$y_{t+h} - y_{t-1} = \alpha_h + \beta_h \Delta P_t + \gamma_h \Delta y_{t-1} + \delta_h g_{t-1} + u_{t+h}, \; h = 0, \ldots, 4 \qquad (8)$$

where g is real GDP growth. Standard errors are Newey–West with h + 1 lags. Because HAC inference is unreliable in very small samples, significance is judged from wild (Rademacher) bootstrap-t p-values with 999 replications.

Bayesian VAR and Connectedness

To capture joint dynamics without over-parameterizing, we estimate a five-variable VAR(1) in the pressure index, reserves, private credit, remittances and FDI with a Minnesota prior (Litterman, 1986), implemented through dummy observations (Bańbura et al., 2010) with overall tightness λ = 0.2 and a random-walk prior mean for persistent variables. We draw 4,000 times from the Normal–inverse-Wishart posterior and retain stable draws. Structural pressure shocks are identified recursively with the pressure index ordered first, on the grounds that FATF decisions, foreign sanctions and US enforcement are set abroad and do not respond to Pakistan's macro variables within the year. We then compute the Diebold and Yilmaz (2012, 2014) connectedness measures from the generalized forecast-error variance decomposition (Pesaran and Shin, 1998), which is invariant to ordering:

$$C(H) = (1/N) \sum_{i \neq j} \theta_{ij}(H) \times 100, \; NET_i = \sum_{j \neq i} \theta_{ji}(H) - \sum_{j \neq i} \theta_{ij}(H) \qquad (9)$$

where $\theta_{ij}(H)$ is the row-normalised share of variable i's H-step forecast-error variance attributable to shocks in j. Ten-year rolling windows trace how connectedness evolved.

Threshold Regression

H5 is tested with Hansen's (2000) threshold model, in which the pressure index q_t determines the regime:

$$y_t = \beta_1' x_t \cdot 1(q_t \leq \gamma) + \beta_2' x_t \cdot 1(q_t > \gamma) + e_t \qquad (10)$$

The threshold γ is estimated by least squares over the central 60% of the q distribution. Because γ is not identified under the null of linearity, the p-value of the Sup-F statistic is obtained by a heteroskedasticity-robust bootstrap (Hansen, 1996), and a confidence set for γ is formed by inverting the likelihood-ratio statistic.

Double/Debiased Machine Learning

Linear controls may fail to capture confounding by the macro cycle. We therefore estimate the partially linear model Y = θD + g(W) + ε, D = m(W) + v, where D is the pressure index or FATF exposure and W contains GDP growth, inflation, depreciation, the current account and the lagged outcome. Following Chernozhukov et al. (2018), g and m are learned by random forests or cross-validated ridge regression, and θ is estimated from the orthogonalized residuals with cross-fitting:

$$\theta = \sum_i \tilde{v}_i \tilde{u}_i / \sum_i \tilde{v}_i^2, \ \tilde{u}_i = Y_i - \hat{g}_{(-k)}(W_i), \ \tilde{v}_i = D_i - \hat{m}_{(-k)}(W_i) \qquad (11)$$

We use four folds, repeat the sample split 25 times and aggregate by the median, as recommended for small samples.

Deep-Learning Architectures

Inputs and Targets. Three targets represent the spillover channels: the annual change in private credit (percentage points of GDP), remittance growth (log difference, %) and the change in the log UAE/Iran trade ratio. For each year t the input is a three-year sequence of nine features: the four sanctions-channel exposures at t, four macro conditions (GDP growth, inflation, depreciation and reserves) at t − 1, and the lagged target. Features are standardized with training-window moments.

LSTM. At each step the LSTM (Hochreiter and Schmidhuber, 1997) updates a memory cell through input, forget and output gates:

$$i_t = \sigma(W_i x_t + U_i h_{t-1} + b_i), \ f_t = \sigma(W_f x_t + U_f h_{t-1} + b_f), \ o_t = \sigma(W_o x_t + U_o h_{t-1} + b_o) \quad (12)$$

$$c_t = f_t \odot c_{t-1} + i_t \odot \tanh(W_c x_t + U_c h_{t-1} + b_c), \ h_t = o_t \odot \tanh(c_t) \quad (13)$$

GRU. The GRU (Cho et al., 2014) uses update and reset gates and no separate cell: $h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \tilde{h}_t$.

Attention-LSTM. A two-layer stacked LSTM is followed by additive temporal attention (Bahdanau et al., 2015), which lets the network weight the three years of the window:

$$e_s = w_a' \tanh(V_a h_s), \ \alpha_s = \exp(e_s) / \sum_r \exp(e_r), \ \hat{h} = \sum_s \alpha_s h_s \qquad (14)$$

Deep MLP and Deep Ensemble. A feed-forward network maps the flattened window through three ReLU layers (32-16-8). The deep ensemble averages the MLP, LSTM, GRU and attention-LSTM, following evidence that ensembles of independently trained networks are better calibrated and more robust than single networks (Lakshminarayanan et al., 2017). A hybrid model adds an LSTM correction to the residuals of a linear ridge model, in the spirit of combining econometric and neural components.

Regularization for a Short Sample. Annual data impose severe constraints. We use small networks (six recurrent units), dropout (Srivastava et al., 2014), AdamW with decoupled weight decay (Kingma and Ba, 2015; Loshchilov and Hutter, 2019), and Gaussian noise injection, which is equivalent to Tikhonov regularization (Bishop, 1995). Table 2 gives the full protocol.

Table 2. Deep-Learning Architectures and Training Protocol (Fixed Ex Ante; No Tuning on the Evaluation Window)

Component	Specification
Inputs (per year, 9)	FATF exposure, Iran intensity, Russia dummy, US-enforcement dummy (time t); GDP growth, inflation, depreciation, reserves (t−1); lagged target
Sequence window	L = 3 years (t−2 … t); unavailable early steps zero-padded after standardization
LSTM / GRU	One recurrent layer, 6 hidden units, linear read-out; forget-gate bias initialized at 1
Attention-LSTM	Two stacked LSTM layers (6 units) + additive (Bahdanau-type) temporal attention over the window
Deep MLP	Flattened window → 32-16-8 ReLU → linear; L2 penalty 0.01
Deep ensemble	Equal-weight average of deep MLP, LSTM, GRU and attention-LSTM
Regularization	AdamW (lr 0.01, decoupled decay 0.01), dropout 0.2 on the final state, Gaussian noise injection (σ = 0.08, 40 replicas)
Training	80 epochs, mini-batch 64; three seeds per recurrent model and five per MLP in the OOS exercise; six seeds in the explanation stage
Implementation	Reverse-mode automatic differentiation written in NumPy (gradients verified by finite differences); scikit-learn for MLP and tree benchmarks

Notes: Hyperparameters were fixed before the evaluation exercise and kept identical across targets, which guards against selection on the hold-out years.

All recurrent networks were coded on a reverse-mode automatic-differentiation engine written in NumPy, whose gradients were verified against finite differences.

Denoising Autoencoder. To summarize the outcome side of disruption without imposing weights, we train a denoising autoencoder (Hinton and Salakhutdinov, 2006; Vincent et al., 2008) with an 8-6-3-1-3-6-8 architecture on eight standardized indicators of banking depth, external liquidity and trade routing. The one-dimensional code is a latent collateral-disruption index (LCDI):

$$z_t = f_\theta(x_t + \eta_t),\ \hat{x}_t = g_\phi(z_t),\ min_{\theta,\phi} \sum_t \|x_t - \hat{x}_t\|^2 \qquad (15)$$

5.10 Evaluation and Explanation

Out-of-Sample Design. Each model is re-estimated every year on all earlier data and forecasts the following year, giving hold-out years 2014–2024 (2014–2023 for trade). Benchmarks are the prevailing mean, an AR(1), ridge regression, a random forest and gradient boosting. Accuracy is compared with the Diebold and Mariano (1995) test with the Harvey et al. (1997) small-sample correction, on the loss differential $d_t = e_{1t}^2 - e_{2t}^2$, and jointly with the model confidence set (MCS) of Hansen et al. (2011), which identifies the set of models that cannot be distinguished from the best at the 10% level.

Exact Group-Shapley Attribution. The prediction f(x) of the deep ensemble is decomposed among six feature groups G (FATF, Iran sanctions, Russia regime, US enforcement, macro-financial conditions and own dynamics). With only six groups, Shapley values can be computed exactly over all 64 coalitions rather than approximated:

$$\phi_j = \sum_{S \subseteq G \setminus j} [|S|!\,(K - |S| - 1)!/K!]\,[v(S \cup j) - v(S)],\ K = 6 \qquad (16)$$

where v(S) is the ensemble prediction when groups outside S are set to their sample means across all three steps of the input window. The four sanctions-channel values summed give the sanctions-attributable component of each year's outcome relative to a counterfactual of average exposure.

Results

Integration Properties and Structural Breaks

Table 3 reports the unit-root evidence. Reserves, GDP growth, FATF exposure and the composite pressure index reject a unit root in levels at the 10% level or better; remittances (% of GDP) reject only around a deterministic trend. Private credit, FDI, inflation, depreciation, Iran sanctions intensity and the Iran trade share are best treated as I(1), since their first differences reject comfortably. The log UAE/Iran ratio is borderline: neither its level nor its first difference rejects at 10% with the ADF test, while KPSS does not reject stationarity. No series is I(2). This mixture is exactly the case for which the bounds approach was designed. The Zivot–Andrews test places the most important breaks in 2009 (private credit, FDI and reserves, coinciding with the first grey-listing and the global financial crisis), 2015–2016 (FATF exposure, the pressure index and Iran intensity, coinciding with the second exit and the JCPOA) and 2017–2018 (remittances, inflation and depreciation).

Table 3. Unit-Root and Stationarity Tests, with an Endogenous Structural Break

Series	ADF (c)	ADF (c,t)	KPSS (c)	ADF (c)	Δ	Z–A (C)	Z–A break	Decision
Composite pressure index	-2.74	-3.38	0.462	-5.61*		-4.48	2016	I(0)
FATF grey-list exposure	-4.11*	-3.86*	0.094	-4.27*		-4.22	2015	I(0)
Iran sanctions intensity	-2.52	-3.09	0.316	-3.78*		-8.08*	2016	I(1)
Private credit (% GDP)	-0.66	-2.81	0.583*	-3.51*		-3.20	2009	I(1)
Remittances (% GDP)	-0.34	-4.51*	0.733*	-4.19*		-7.06*	2017	I(0)
FDI (% GDP)	-1.32	-1.70	0.439	-3.83*		-7.65*	2009	I(1)
ln(UAE/Iran trade ratio)	-2.04	-1.63	0.225	-2.64		-4.13	2015	I(1)?
Iran share of trade (%)	-1.26	-2.34	0.321	-3.58*		-4.03	2011	I(1)
Reserves (months of imports)	-3.94*	-6.78*	0.380	-7.50*		-5.89*	2009	I(0)
Real GDP growth	-3.77*	-3.65	0.175	-5.95*		-4.16	2013	I(0)
CPI inflation	-2.38	-2.39	0.169	-4.03*		-3.49	2018	I(1)
PKR depreciation	-0.87	-1.34	0.292	-9.80*		-12.24*	2018	I(1)

Notes: ADF lag order chosen by AIC (max 2); critical values from MacKinnon (2010) response surfaces adjusted to the effective sample. KPSS null is stationarity (5% c.v. 0.463). Zivot–Andrews model C (break in intercept and trend), 5% c.v. −5.08. * rejects the respective null at 5%. Decision uses the 10% level; "I(1)?" marks a series where neither level nor first-difference ADF rejects at 10%; the bounds framework is used because it is valid for any I(0)/I(1) mixture.

The Bai–Perron dating in Table 4 gives a consistent picture. Private credit shifts down twice, from a mean of 21.9% of GDP before 2011 to 15.2% in 2011–2021 and 12.8% from 2022. Remittances shift up in 2009 and 2020, from 3.2% to 5.8% and then 8.6% of GDP. The log UAE/Iran ratio rises sharply in 2011 and falls back in 2019, which in levels means that the ratio of UAE trade to Iran trade averaged about 6 before 2011, about 27 in 2011–2018 and about 11.5 thereafter. The Iran share of Pakistan's trade peaks in 2008–2010 and collapses from 2011.

The pressure index shifts up in 2019 from an average of 0.29 to 0.58. The break dates line up with the sanctions and FATF chronology in Section 3: the formal Iran channel closed between 2010 (CISADA) and 2012 (SWIFT disconnection), and the third grey-list stint and the maximum-pressure campaign coincide in 2018–2019.

Table 4. Bai–Perron Multiple Structural-Break Tests (Mean-Shift Model)

Series	sup-F(1|0)	sup-F(2|0)	Breaks (BIC)	Break years	Regime means
Private credit (% GDP)	86.42**	62.00**	2	2011, 2022	21.88 → 15.18 → 12.77
Remittances (% GDP)	33.33**	86.64**	2	2009, 2020	3.21 → 5.82 → 8.60
ln(UAE/Iran ratio)	24.06**	40.47**	2	2011, 2019	1.83 → 3.30 → 2.44
Iran trade share (%)	42.50**	45.35**	2	2008, 2011	1.27 → 2.04 → 0.58
FDI (% GDP)	90.12**	50.87**	1	2009	2.47 → 0.67
Pressure index	13.41**	10.20**	1	2019	0.29 → 0.58

Notes: Global minimization of the sum of squared residuals by dynamic programming; trimming h = 3 observations; up to two breaks. ** exceeds the Bai and Perron (2003) 5% critical values (8.58 for one break; 7.22 for two breaks, ε = 0.15). Break year = first year of the new regime. Because the series are persistent, the sup-F statistics are best read as dating devices rather than formal tests.

Long-Run Relationships: Credit Contraction and Trade Diversion

Table 5 reports the ARDL bounds tests. For private credit, with the pressure index and inflation as forcing variables, the F-statistic of 3.77 exceeds the bootstrap 10% critical value (3.59) but not the 5% value (4.66), giving a bootstrap p-value of 0.092. We therefore find evidence of a levels relationship at the 10% level. The implied long-run semi-elasticity of credit to pressure is −23.9 percentage points of GDP per index unit (standard error 12.5, p = 0.078), so a 0.1-point rise in the index, roughly the difference between a year off the grey list and a year spent half on it, is associated with a long-run fall in credit depth of about 2.4 points of GDP. The error-correction coefficient (−0.17) implies slow adjustment, with a half-life of around four years. For the trade-diversion ratio, remittances and FDI the bounds tests do not reject the null of no levels relationship (bootstrap p = 0.80, 0.11 and 0.28). The diagnostics in Table 6 are satisfactory for the credit and remittance models; the UAE/Iran and FDI models show first-order serial correlation (Breusch–Godfrey p = 0.005 and 0.010), which is one reason we turn to DOLS with HAC inference for long-run estimation. CUSUM statistics stay within their 5% bounds in all four models.

Table 5. ARDL Bounds Tests with Bootstrap Critical Values and Long-Run Coefficients

Model (y ← x)	n	F_PSS	Boot c.v. 10% / 5%	Boot p	ECT ρ (t)	Long-run coefficients (s.e.)
Private credit ← Pressure, Inflation	19	3.77*	3.59 / 4.66	0.091	-0.168 (-1.48)	Pressure -23.89* (12.48); Inflation -0.35 (0.61)
ln(UAE/Iran) ← Iran int., FATF	18	0.61	3.75 / 4.85	0.797	-0.118 (-0.63)	Iran int. -1.19 (3.96); FATF -0.44 (4.00)
Remittances ← Pressure,	18	2.58	2.72 / 3.63	0.113	-0.133 (-1.42)	Pressure -0.08 (9.60); Depreciation 0.27 (0.27)

Model (y ← x)	n	F_PSS	Boot c.v. 10% / 5%	Boot p	ECT ρ (t)	Long-run coefficients (s.e.)
Depreciation						
FDI ← Pressure, GDP growth	19	2.34	3.86 / 4.89	0.280	-0.257 (-1.72)	Pressure -2.61 (2.65); GDP growth 0.43 (0.42)

Notes: Unrestricted error-correction model, ARDL(1,1,1), case III (unrestricted intercept, no trend). Critical values and p-values for F come from 2,000 residual-bootstrap replications under the null of no levels relationship (McNown et al., 2018), which is preferable to asymptotic Pesaran et al. (2001) bounds at n < 20. Long-run coefficients are −θ/ρ with delta-method standard errors. ***, **, * = 1%, 5%, 10%.

Table 6. ARDL Residual Diagnostics

Dependent variable	R^2	Adj. R^2	BG-LM(1) [p]	Jarque–Bera [p]	RESET [p]	CUSUM (5%)
Private credit	0.498	0.304	0.51 [0.474]	0.85 [0.655]	1.37 [0.264]	Stable
ln(UAE/Iran)	0.146	-0.209	7.72 [0.005]	1.19 [0.551]	0.71 [0.416]	Stable
Remittances	0.670	0.533	0.88 [0.348]	0.38 [0.828]	0.26 [0.622]	Stable
FDI	0.379	0.140	6.61 [0.010]	0.06 [0.969]	0.61 [0.450]	Stable

Notes: Breusch–Godfrey LM test for first-order serial correlation; Jarque–Bera normality; Ramsey RESET with squared fitted values; CUSUM of recursive residuals with 5% significance bounds (Brown et al., 1975).

The DOLS estimates in Table 7 sharpen these findings. For private credit, the long-run coefficient on the pressure index is −25.3 (standard error 6.0, p = 0.003), very close to the ARDL estimate but much more precisely estimated, and inflation enters positively, consistent with nominal-valuation effects. This supports H1. For trade routing, a one-notch increase in Iran sanctions intensity lowers the Iran share of Pakistan's total recorded trade by 2.0 percentage points (p = 0.017) and raises the log UAE/Iran ratio by 2.1 (p = 0.080), while FATF exposure has no significant effect on either. The magnitudes are large relative to the sample mean Iran share of 0.92% because they are identified mainly by the collapse of the formal Iran channel between 2009 and 2012, when the share fell from 2.5% to 0.4% and the UAE/Iran ratio rose roughly tenfold. The evidence therefore supports H3 in the sense that sanctions intensity, not FATF listing, drives the re-routing of recorded trade towards UAE intermediation. The absence of a bounds-test levels relationship for the ratio cautions that this is a regime shift rather than a stable equilibrium elasticity.

Table 7. Dynamic OLS Long-Run Elasticities: Trade Diversion and Credit Contraction

Dependent variable	n	Coefficient 1 (s.e.)	Coefficient 2 (s.e.)	R^2
ln(UAE/Iran ratio)	16	2.121* (1.036)	-2.945 (2.009)	0.514
Iran share of trade (%)	16	-2.011** (0.642)	2.514 (1.452)	0.650
Private credit (% GDP)	17	-25.334*** (5.969)	0.450** (0.135)	0.817

Notes: Stock–Watson (1993) DOLS with one lead and one lag of the differenced regressors; Newey–West HAC standard errors. Rows 1–2: coefficient 1 = Iran sanctions intensity, coefficient 2 = FATF exposure. Row 3: coefficient 1 = composite pressure index, coefficient 2 = CPI inflation. ***, **, * = 1%, 5%, 10%.

Asymmetric Pass-Through

Table 8 and panel (a) of Fig. 3 report the nonlinear ARDL estimates. For private credit the long-run effects of tightening and easing are almost identical ($L^+ = -22.6$, $L^- = -24.3$; Wald p = 0.80), so in the long run credit depth responds symmetrically. The short run is different. A unit tightening lowers credit by 6.0 points on impact, whereas a unit easing does not raise it on impact (the contemporaneous coefficient on easing has the opposite sign), and the cumulative response to easing turns positive only after the first year. The short-run symmetry test has p = 0.11, so the evidence is suggestive rather than conclusive, but the pattern is the one predicted by H4: relationships cut during de-risking are restored more slowly than they were withdrawn. The asymmetry line in Fig. 3 narrows from about −10 points on impact to near zero after six years. For FDI the asymmetry tests are not significant, and for remittances the error-correction coefficient is indistinguishable from zero, so long-run effects are not identified.

Table 8. Nonlinear ARDL: Asymmetric Pass-through of Sanctions-Compliance Pressure

Dependent variable	F_PSS	ECT ρ (t)	L^+ (s.e.)	L^- (s.e.)	W_LR [p]	W_SR [p]
Private credit	3.13	-0.245 (-1.16)	-22.65 (14.78)	-24.32 (21.66)	0.07 [0.796]	2.56 [0.110]
Remittances	2.62	0.007 (0.02)	124.74 (4963.98)	92.58 (3771.03)	0.04 [0.844]	3.20 [0.073]
FDI	2.05	-0.238 (-1.31)	-2.01 (3.73)	-2.95 (5.26)	0.26 [0.612]	1.34 [0.246]

Notes: Pressure is decomposed into partial sums of tightening (P^+) and easing (P^-) changes (Shin et al., 2014). L^+ and L^- are long-run coefficients $-\theta\pm/\rho$. W_LR tests $\theta^+ = \theta^-$; W_SR tests equality of the contemporaneous coefficients; both are $\chi^2(1)$. Controls: inflation (credit), depreciation (remittances), GDP growth (FDI). The remittance model has $\rho \approx 0$, so its long-run coefficients are not identified and are omitted from interpretation.

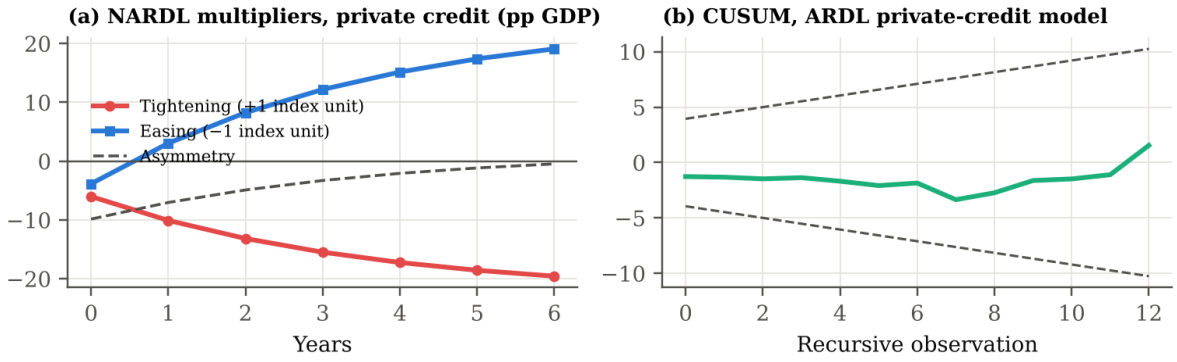


Fig. 3. (a) Nonlinear ARDL cumulative dynamic multipliers of private credit to a permanent unit tightening and a unit easing of the pressure index, and their sum (asymmetry). (b) CUSUM of recursive residuals for the ARDL private-credit model with 5% bounds.

Short-Run Dynamics: Local Projections, Bayesian VAR and Connectedness

The local projections in Table 9 and Fig. 4 isolate the external-liquidity channel. A unit change in the pressure index lowers reserves by 4.3 months of imports on impact (wild-bootstrap p = 0.013) and by 3.9 months a year later (p = 0.081), with the effect fading by the third year. Scaled to the sample standard deviation of the change in pressure (0.21), this is a loss of about 0.9 months of import cover on impact. Private credit falls by about 6.2–6.7 points of GDP per unit at horizons two to four, in line with the long-run estimates, but the local-projection

bands are wide and the bootstrap p-values (0.23–0.38) do not reach conventional levels. FDI falls on impact (−0.54 points of GDP, p = 0.11). Responses of remittances and of the UAE/Iran ratio are statistically indistinguishable from zero at every horizon.

Table 9. Local-Projection Responses to a Unit Change in the Pressure Index

Outcome	h = 0	h = 1	h = 2	h = 3	h = 4
Private credit (pp GDP)	1.06 (1.60)	-1.39 (2.20)	-6.22 (3.64)	-6.68 (4.66)	-6.64 (5.19)
Remittances (pp GDP)	0.10 (0.53)	-0.19 (1.15)	1.08 (1.62)	2.30 (1.35)	0.47 (0.65)
FDI (pp GDP)	-0.54 (0.26)	-0.83 (0.65)	0.09 (0.96)	-0.81 (1.05)	-0.79 (1.09)
ln(UAE/Iran)	0.10 (0.48)	-0.85 (1.24)	-1.06 (1.74)	0.73 (1.67)	2.44 (1.66)
Reserves (months)	-4.35** (1.07)	-3.92* (1.75)	-2.12 (1.99)	-0.33 (0.89)	0.33 (1.88)

Notes: y(t+h) − y(t−1) regressed on ΔP(t), Δy(t−1) and lagged GDP growth (Jordà, 2005). Newey–West standard errors with h + 1 lags in parentheses; stars from wild (Rademacher) bootstrap-t p-values, 999 replications. The sample standard deviation of ΔP is 0.21 index points.

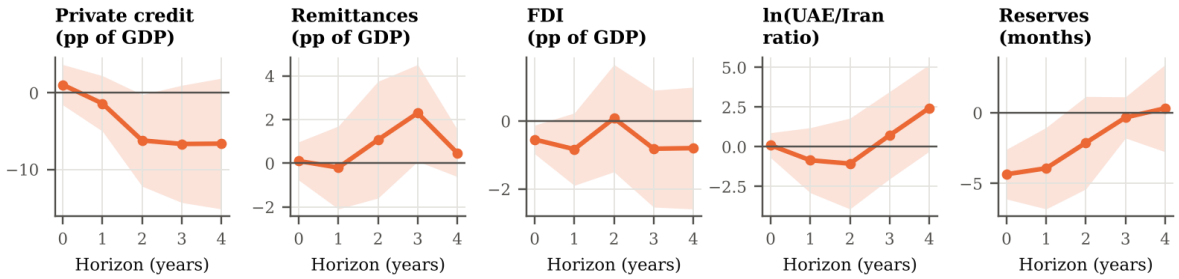


Fig. 4. Local-projection responses to a unit change in the pressure index, with 90% HAC confidence bands.

The Minnesota-prior BVAR imposes more structure and so delivers sharper responses (Fig. 5). A one-standard-deviation pressure shock, about 0.16 index points on impact, lowers reserves by 0.69 months of imports on impact (68% credible band −0.96 to −0.42) and private credit by up to 0.45 points of GDP after two years (68% band −0.75 to −0.16). The credit response is persistent, still around −0.3 points after six years. The response of FDI is negative but its 68% band includes zero, and remittances respond slightly positively but insignificantly. Together, the local projections and the BVAR support H2: compliance shocks drain external liquidity immediately, and the credit contraction follows with a lag of one to two years, consistent with banks rationing credit after correspondent and trade-finance lines tighten. Only 1,523 of 4,000 posterior draws were stable, reflecting the near-unit-root behaviour of private credit; results are reported for the stable draws.

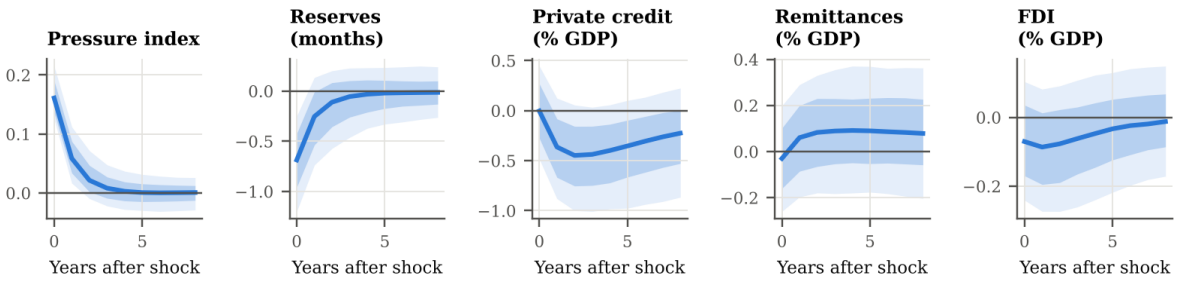


Fig. 5. Bayesian VAR impulse responses to a one-standard-deviation pressure-index shock (recursive identification, pressure ordered first). Lines: posterior medians; dark and light bands: 68% and 90% credible intervals.

Table 10 reports the connectedness decomposition. Evaluated at the posterior-mean parameters, the total connectedness index is 32.5%: about a third of the ten-year forecast-error variance of the five variables is due to shocks to other variables in the system. Integrating over parameter uncertainty raises the figure, with a posterior median of 51.3% across draws (68% band 43.2–60.0). The most robust feature of the decomposition is the position of private credit, which takes 55% of its forecast-error variance from other variables, 13.4% of it from the pressure index and 27.2% from FDI, and is the largest net receiver both at the posterior mean (−34.1 points) and at the posterior median (−20.3). The pressure index accounts for 13.3% of the forecast-error variance of reserves. The transmitter roles of the other variables are not sharply identified: their 68% bands for net spillovers all include zero, which is unsurprising with eighteen annual observations. The rolling estimates in Fig. 6, computed at posterior-mean parameters, show how the system evolved. Total connectedness fell from about 44% in the window ending 2016 to 20% in the window ending 2018 and then rose steadily to 39% in the window ending 2024, and the pressure index, a net receiver in the windows ending 2019 and 2020, became a net transmitter from the window ending 2021, transmitting 10.4 points net in the window ending 2024. Systemic interdependence between compliance pressure and the macro-banking system has thus grown since the third grey-list stint.

Table 10. Diebold–Yilmaz Connectedness from the Minnesota-Prior BVAR (Generalized FEVD, H = 10)

From \ To	Pressure	Reserves	Priv. credit	Remittances	FDI	From others
Pressure	69.1	17.2	1.0	10.3	2.4	30.9
Reserves	13.3	67.0	6.5	13.0	0.2	33.0
Priv. credit	13.4	8.7	45.0	5.7	27.2	55.0
Remittances	4.4	9.5	1.5	83.5	1.1	16.5
FDI	3.7	1.3	11.9	10.2	72.8	27.2
To others	34.9	36.7	20.9	39.1	30.9	32.5
Net (posterior median)	-4.6	7.5	-20.3	1.8	1.4	
Net 68% band	-27.9, 22.1	-16.6, 42.7	-45.9, 8.8	-33.0, 40.9	-26.6, 34.8	

Notes: Rows 1–6: row i, column j = share (%) of the 10-step forecast-error variance of i due to shocks in j, evaluated at the posterior-mean parameters of the BVAR(1), λ = 0.2; the bottom-right cell is the total connectedness index at those parameters. Rows 7–8: net spillovers (to − from) computed draw by draw, posterior median and 68% band; positive = net transmitter. Posterior median of the total connectedness index across draws = 51.3% (68% band 43.2–60.0); 1,523 stable draws out of 4,000.

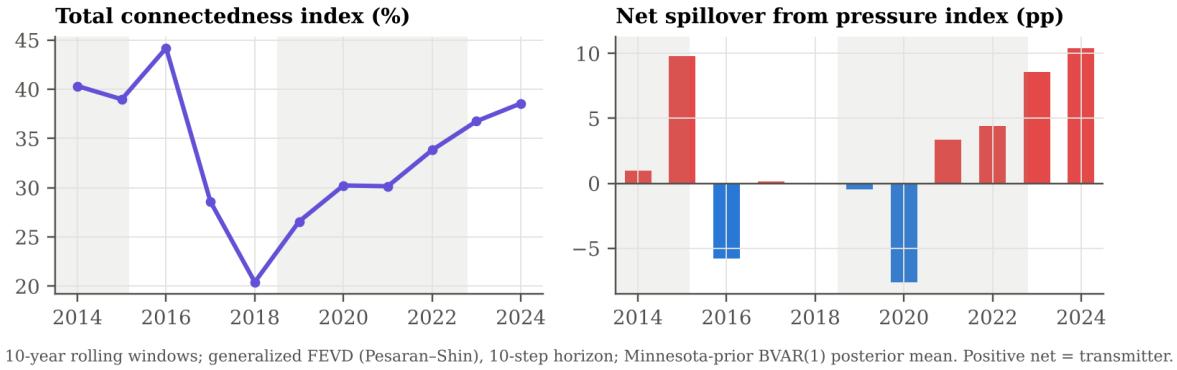


Fig. 6. Rolling Diebold–Yilmaz connectedness (ten-year windows): (a) total connectedness index; (b) net directional spillover from the pressure index to the rest of the system.

Regime Dependence in the Payment Channel

Table 11 reports the threshold regressions. Linearity is rejected only for the remittance equation (sup-F = 8.49, bootstrap p = 0.042), with an estimated threshold of 0.458 on the pressure index. Below the threshold (eleven years), the intercept is close to zero and the change in remittances (% of GDP) rises weakly with depreciation (slope 0.032). Above it (eight years: 2013–2014 and 2019–2024), remittances rise faster on average (intercept 0.96 points of GDP per year, standard error 0.40), but their response to depreciation turns negative (slope −0.034). This pattern is consistent with H5 and with two features of the high-pressure regime. On one hand, the FATF action plans of 2012–2015 and 2018–2022 pushed transactions into documented channels and coincided with official incentive schemes for formal remittances. On the other hand, when the rupee depreciates sharply under high pressure, as in 2022–2023 when a large gap opened between interbank and open-market exchange rates, formal inflows lose ground to informal hundi/hawala channels. The confidence set for the threshold is wide (0.17–0.46), and with eight observations in the upper regime the estimates are fragile, but they are the only evidence of nonlinearity that survives a bootstrap test. For private credit and reserves the linear model is not rejected (p = 0.47 and 0.94).

Table 11. Hansen Threshold Regressions with the Pressure Index as the Regime Variable

Outcome	Regressor	$\hat{\gamma}$ [95% LR set]	F_1 [boot p]	Low regime: const.; slope	High regime: const.; slope	n low/high
Δ Private credit	GDP growth	0.375 [0.17, 0.46]	4.95 [0.465]	-1.60 (0.95); 0.418 (0.206)	-1.78 (0.65); 0.202 (0.173)	9/10
Δ Remittances (% GDP)	Depreciation	0.458 [0.17, 0.46]	8.49 [0.042]	0.01 (0.15); 0.032 (0.017)	0.96 (0.40); -0.034 (0.022)	11/8
Δ Reserves (months)	GDP growth	0.292 [0.17, 0.46]	1.15 [0.944]	1.36 (1.13); -0.270 (0.233)	0.55 (0.90); -0.254 (0.243)	7/12

Notes: Hansen (2000) least-squares threshold estimation, 20% trimming; F_1 tests linearity against one threshold, p-value from 1,000 heteroskedasticity-robust wild-bootstrap replications (Hansen, 1996). Confidence set for γ inverts the likelihood-ratio statistic at 7.35.

Deep-Learning Results

A latent Collateral-Disruption Index

The denoising autoencoder compresses eight outcome indicators into one latent dimension that reconstructs 93.3% of their standardized variance, compared with 50.4% for the first principal component (Table 12). The index loads negatively on private credit (−0.86), FDI (−0.81), reserves (−0.56) and the Iran trade share (−0.51) and positively on remittances (0.74), depreciation (0.57) and the UAE/Iran ratio (0.42). It therefore captures a syndrome of shrinking banking depth, weak capital inflows, external stress and greater dependence on remittances and intermediated trade. Fig. 7 shows that the LCDI rises from between −1.2 and −1.4 standard deviations in 2006–2008 to its sample maximum of 1.52 in 2023, with local peaks in 2013, 2019 and 2023. Its correlation with the author-coded pressure index is 0.57, higher than the correlation between the first principal component and the pressure index (0.41). The outcome-side index, learned without any knowledge of sanctions, thus tracks the coded exposure, which is independent evidence that the coded index measures something economically meaningful. Reconstruction errors are largest in 2017 and 2020, the years of the HBL penalty and New York branch closure and of the pandemic, when the usual co-movement among indicators broke down.

Table 12. Deep Autoencoder: Structure of the Latent Collateral-Disruption Index (LCDI)

Indicator / statistic	Value
Private credit (% GDP)	-0.86
Remittances (% GDP)	0.74
FDI (% GDP)	-0.81
Reserves (months)	-0.56
PKR depreciation	0.57
Current account (% GDP)	0.39
ln(UAE/Iran ratio)	0.42
Iran trade share	-0.51
Corr(LCDI, coded pressure index)	0.574
Corr(LCDI, first principal component)	0.825
Corr(PC1, coded pressure index)	0.414
Autoencoder reconstruction R^2	0.933
PC1 share of variance	0.504

Notes: Denoising autoencoder 8-6-3-1-3-6-8 (tanh hidden layers, linear code and output), Gaussian input noise σ = 0.05, AdamW, 3,000 epochs, ten seeds averaged after sign alignment (higher = more disruption, i.e. negatively related to reserves). Loadings are correlations between the LCDI and each standardized input, 2006–2023.

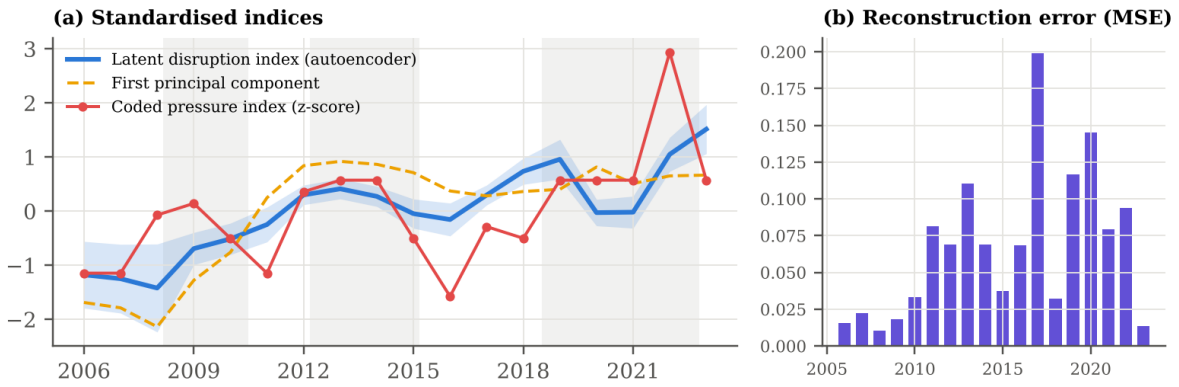


Fig. 7. (a) Latent collateral-disruption index from the denoising autoencoder (mean of ten seeds with 95% band), the first principal component and the coded pressure index, all standardized. (b) Mean squared reconstruction error by year.

Out-of-Sample Performance

Table 13 and Fig. 8 report the expanding-window forecasts. For the change in private credit, the deep MLP has the lowest RMSE, 22% below the prevailing mean, followed by the deep ensemble (15% lower), the random forest (14%) and the LSTM (7%). For the trade-diversion ratio the attention-LSTM is best (17% below the mean), followed by the deep ensemble (12%), the LSTM (10%) and the AR(1) (10%). For remittance growth no model beats the prevailing mean; the deep recurrent models are 21–28% worse. Under the Harvey–Leybourne–Newbold-corrected Diebold–Mariano test, none of the improvements is significant (the smallest p-value among models that beat the benchmark is 0.30), and the 90% model confidence set retains all eleven models for all three targets. The hybrid linear–LSTM model is the worst performer for credit and trade, which illustrates how adding a nonlinear correction to an already noisy linear fit compounds estimation error in short samples.

Table 13. Expanding-Window out-of-Sample Accuracy, 2014–2024

Model	Credit RMSE	rRMSE [DM p]	Remit. RMSE	rRMSE [DM p]	Trade RMSE	rRMSE [DM p]
Prevailing mean	1.249	1.000	12.134	1.000	0.388	1.000
AR(1)	1.227	0.982 [0.67]	14.104	1.162 [0.39]	0.350	0.903 [0.36]
Ridge	1.616	1.293 [0.53]	13.110	1.080 [0.55]	0.498	1.285 [0.35]
Random forest	1.075	0.861 [0.57]	12.239	1.009 [0.84]	0.367	0.946 [0.73]
Gradient boosting	1.255	1.005 [0.99]	14.055	1.158 [0.27]	0.434	1.119 [0.62]
Deep MLP (32-16-8)	0.972	0.778 [0.30]	12.875	1.061 [0.58]	0.408	1.052 [0.76]
LSTM	1.165	0.933 [0.76]	14.704	1.212 [0.49]	0.348	0.897 [0.62]
GRU	1.230	0.985 [0.96]	15.490	1.277 [0.37]	0.384	0.990 [0.96]
Attention-LSTM (2-layer)	1.369	1.096 [0.59]	15.025	1.238 [0.51]	0.322	0.832 [0.47]
Hybrid linear–LSTM	1.670	1.337 [0.50]	13.575	1.119 [0.55]	0.528	1.362 [0.25]
Deep ensemble	1.060	0.849 [0.52]	14.225	1.172 [0.52]	0.342	0.881 [0.51]

Notes: One-step-ahead forecasts re-estimated each year on all prior data (first origin trains on 2007–2013). Credit = Δ private credit (pp GDP), n = 11; Remit. = remittance growth (%), n = 11; Trade = Δ ln(UAE/Iran ratio), n = 10. rRMSE = RMSE relative to the prevailing mean; DM p = two-sided Diebold–Mariano test with the Harvey et al. (1997) small-sample correction against the prevailing mean. The 90% model confidence set (T_max, block bootstrap) retains all eleven models for every target (MCS p = 0.51, 0.73, 0.39).

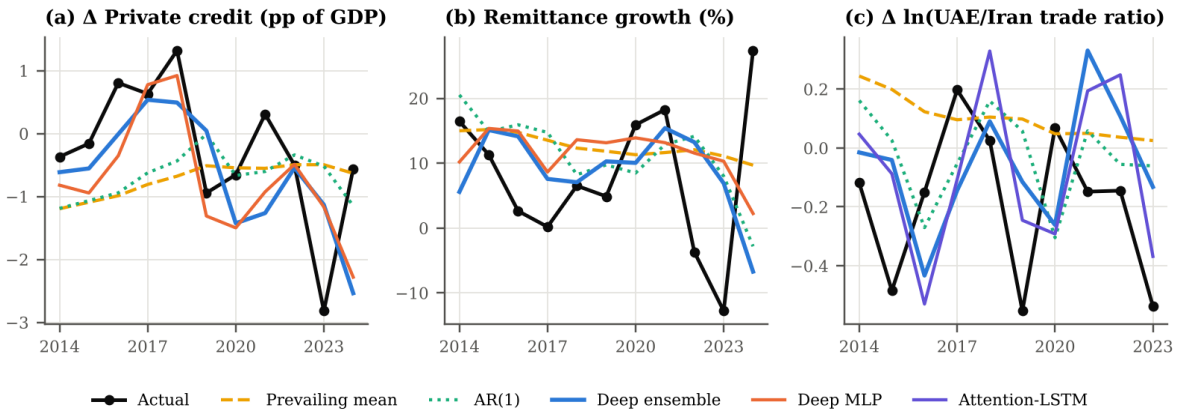


Fig. 8. Expanding-window one-step-ahead forecasts, 2014–2024, against actual outcomes: prevailing mean, AR(1), deep ensemble and the best single deep learner for each target.

H6 is therefore not supported. With eleven hold-out years, a 20% RMSE reduction cannot be distinguished from luck, and the remittance results show that deep networks can be decisively worse than a constant when the target is dominated by shocks outside the information set, such as the 2023 exchange-rate wedge and the 2024 rebound. The gains for credit and trade routing are nevertheless consistent in direction with the econometric evidence of persistent, nonlinear and regime-dependent dynamics in those two channels. The attention weights are informative too. For credit the network places more weight on the oldest and the current year of the window (0.37 and 0.36) than on the middle year (0.27), and for trade the oldest year receives the most weight (0.37). This is consistent with the delayed responses found in the BVAR and NARDL estimates.

Which Channels Drive the Deep Ensemble?

Table 14 decomposes the deep ensemble's attributions into the six feature groups, and Fig. 9 shows them year by year. Macro-financial conditions carry the largest share for every target (45–53%), as expected. The four sanctions channels together account for 31.9% of the attribution mass for credit, 40.2% for remittances and 46.2% for trade diversion. The composition differs by target in a way that matches the econometrics. Iran sanctions are the dominant sanctions channel for trade diversion, with a quarter of the total (25.0%; seed range 21.4–28.6%), consistent with the DOLS evidence. US enforcement penalties are the largest sanctions channel for remittances (13.2%), followed by Iran sanctions (9.8%), the Russia regime (8.8%) and FATF listing (8.3%). For credit, Iran sanctions (12.6%) and FATF listing (9.1%) dominate.

Table 14. Exact Group-Shapley Decomposition of the Deep Ensemble: Share of Mean |Attribution| (%)

Feature group	Δ Private credit	Remittance growth	Δ ln(UAE/Iran)
FATF grey-listing	9.1 [7.8, 10.6]	8.3 [7.1, 11.2]	10.5 [7.6, 12.8]
Iran secondary sanctions	12.6 [10.7, 13.7]	9.8 [7.6, 12.5]	25.0 [21.4, 28.6]
Russia sanctions regime	5.7 [3.2, 9.5]	8.8 [7.8, 10.3]	5.1 [3.5, 6.1]
US enforcement penalties	4.5 [3.6, 5.7]	13.2 [11.3, 15.6]	5.6 [4.1, 7.3]
Macro-financial conditions	53.1 [48.0, 59.0]	47.6 [41.7, 50.9]	45.1 [40.6, 47.9]
Own dynamics	14.9 [9.6, 18.8]	12.2 [9.8, 14.1]	8.6 [6.1, 12.0]
All four sanctions channels	31.9	40.2	46.2
In-sample R^2 of ensemble	0.998	0.998	0.999
Attention weights (t−2, t−1, t)	0.37, 0.27, 0.36	0.34, 0.31, 0.34	0.37, 0.31, 0.32

Notes: Interventional Shapley values computed exactly over all $2^6 = 64$ coalitions of six feature groups, with absent groups set to their sample mean across all three time steps of the input window. Brackets: 5th–95th percentiles across six independently trained seed-ensembles (each averaging LSTM, GRU, attention-LSTM and deep MLP). Attention weights are averaged over years and seeds.

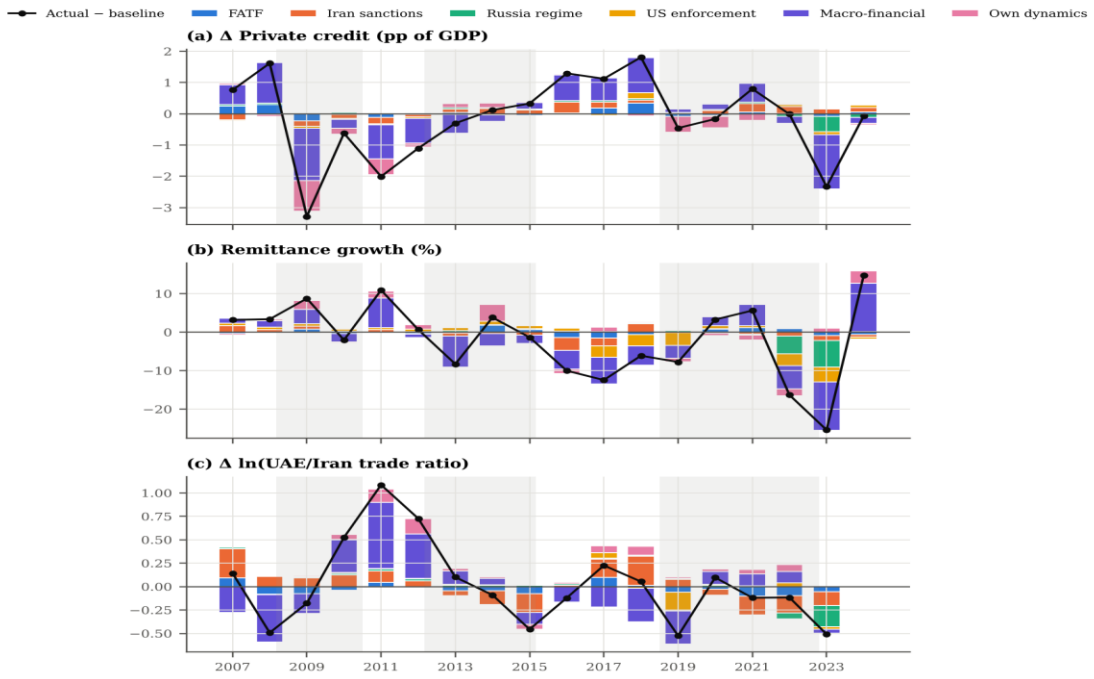


Fig. 9. Exact group-Shapley attributions of the deep ensemble by year, relative to an average-exposure baseline. Stacked bars sum to the prediction minus the baseline; the line shows the actual outcome minus the baseline.

Table 15 aggregates the four sanctions channels by episode. Three episodes stand out. In 2016–2017 the sanctions channels added about 0.4 points of GDP a year to credit growth relative to average exposure, reflecting JCPOA relief and the absence of grey-listing, but subtracted about 5 percentage points a year from remittance growth, with the Iran and FATF groups and, in 2017, the enforcement group (the year of the HBL penalty and the closure of its New York branch) all contributing negatively. In 2022–2024 the sanctions channels subtracted about 7.5 points a year from remittance growth, driven by the Russia-regime and enforcement groups, and, through the Iran and Russia groups, lowered the log UAE/Iran ratio by 0.38 a year. The latter reflects a partial return of recorded trade to direct Iran channels, which coincides with steps taken in 2023 to bring border and barter trade with Iran into documented channels. Because the Russia-regime indicator is collinear with the 2022–2024 macro crisis, its attribution should be read as an upper bound on the role of the Russia regime itself, a limitation we return to in Section 8.

Table 15. Sanctions-Attributable Component by Episode (Sum of the Four Sanctions-Channel Shapley Values, Annual Average)

Episode	Δ Private credit (pp)	Remittance growth (pp)	Δ ln(UAE/Iran)
2008–2010: first grey-list stint	-0.09 [-0.33, 0.17]	1.13 [0.05, 1.87]	0.06 [0.02, 0.10]
2012–2015: peak Iran sanctions; second stint	0.08 [-0.00, 0.16]	0.94 [-0.37, 2.28]	-0.11 [-0.15, -0.06]
2016–2017: JCPOA window; HBL penalty	0.39 [0.22, 0.50]	-5.01 [-6.38, -3.86]	0.19 [0.13, 0.27]
2018–2021: third grey-list stint; maximum pressure	0.27 [0.01, 0.46]	-0.46 [-1.30, 0.29]	-0.05 [-0.11, -0.01]
2022–2024: Russia regime; NBP penalty; exit	-0.05 [-0.21, 0.12]	-7.46 [-9.54, -5.44]	-0.38 [-0.44, -0.32]

Notes: Contribution relative to a counterfactual with every sanctions channel at its 2007–2024 average exposure, in the units of the target. Brackets average the yearly 5th–95th percentiles across seed-ensembles. Trade data end in 2023.

Robustness and Causality

Double/Debiased Machine Learning

Table 16 reports the cross-fitted DML estimates, which control flexibly for the macro cycle. The clearest result concerns reserve: a unit increase in the pressure index lowers the annual change in reserves by 4.3 months of imports with random-forest nuisance functions and by 3.9 months with ridge (both $p < 0.001$), essentially the same as the local-projection impact estimate. FATF exposure lowers the change in FDI by about 0.38 points of GDP ($p = 0.08$ with random forests and 0.06 with ridge) and is associated with a faster rise in remittances relative to GDP (0.60, $p = 0.07$ with random forests), consistent with the formalization effect found in the threshold model. For the change in credit and in the trade-diversion ratio the DML estimates are small and insignificant. This is not in conflict with the long-run results: credit responds slowly and in levels, whereas the DML specification identifies contemporaneous effects on annual changes after partialling out the macro cycle, which itself transmits part of the pressure.

Table 16. Double/Debiased Machine Learning Estimates of Sanctions-Pressure Effects

Outcome	Pressure: RF	Pressure: Ridge	FATF: RF	FATF: Ridge
Δ Private credit	-0.27 (1.16)	-0.52 (0.82)	0.06 (0.56)	-0.17 (0.46)
Δ Reserves	-4.27*** (1.09)	-3.91*** (0.84)	-0.63 (0.63)	-0.78 (0.53)
Δ Remittances	-0.39 (1.15)	0.34 (0.97)	0.60* (0.33)	0.46 (0.31)
Δ FDI	-0.52 (0.45)	-0.68* (0.41)	-0.38* (0.22)	-0.38* (0.20)
Δ ln(UAE/Iran)	-0.13 (0.45)	0.10 (0.45)	-0.10 (0.27)	-0.01 (0.19)

Notes: Partially linear model $Y = \theta D + g(W) + \varepsilon$ with Neyman-orthogonal partialling-out (Chernozhukov et al., 2018); 4-fold cross-fitting repeated 25 times, median aggregation of point estimates and variance. W = GDP growth, inflation, depreciation, current account and the lagged outcome level. Nuisance learners: random forest (100 trees) or cross-validated ridge. ***, **, * = 1%, 5%, 10%.

Toda–Yamamoto Causality

Table 17 reports lag-augmented Granger non-causality tests, which are valid whatever the integration and cointegration properties. The pressure index Granger-causes private credit (Wald = 4.72, asymptotic p = 0.030, bootstrap p = 0.049), and there is no evidence of reverse causality (p = 0.82). No other pair shows significant predictive causality at one lag. The credit result is the strongest single piece of evidence that the association documented in Sections 6.2–6.4 runs from compliance pressure to banking depth rather than the reverse.

Table 17. Toda–Yamamoto Granger Non-Causality Tests (Lag-Augmented VAR, p = 1, d_max = 1)

Variable y	n	W (P → y)	Asy. p	Boot p	W (y → P)	Asy. p
Private credit	18	4.72	0.030	0.049	0.05	0.824
Reserves	18	0.48	0.487	0.491	0.18	0.668
Remittances	18	0.66	0.418	0.433	1.07	0.301
FDI	18	0.00	0.985	0.987	0.02	0.881
ln(UAE/Iran)	17	1.85	0.174	0.198	0.32	0.571
Iran trade share	17	0.21	0.650	0.645	0.32	0.571

Notes: Modified Wald test on the first p lags in a VAR(p + d_max), which is valid irrespective of integration and cointegration (Toda and Yamamoto, 1995). Bootstrap p from 2,000 residual resamples under the null.

Further Checks

Three further checks support the main conclusions. First, the ARDL diagnostics and the CUSUM test (Table 6; Fig. 3, panel b) show no evidence of parameter instability in the credit model. Second, re-estimating the BVAR with a tighter ($\lambda = 0.1$) or looser ($\lambda = 0.5$) prior leaves the impact response of reserves essentially unchanged (−0.69 and −0.67 months) and keeps the credit response negative, with a trough between −0.28 and −0.54 points of GDP. The posterior-median total connectedness ranges from 44.7% to 55.1%, and private credit remains a net receiver; because the generalized decomposition does not depend on variable ordering, these rankings are not an artefact of the recursive identification. Third, the deep-learning attributions are stable across six independently trained seed-ensembles: the seed ranges in Table 14 never reverse the ranking of Iran sanctions as the leading sanctions channel for trade diversion or of macro conditions as the leading group overall.

Discussion

Interpretation

The evidence supports a coherent account of collateral disruption. Third-party sanctions, AML/CFT listing and enforcement shocks reach Pakistan through two fast channels and one slow channel. The fast channels are external liquidity, as reserves fall on impact, and trade routing, as formal trade with Iran gives way to intermediation through the UAE when Iran sanctions tighten. The slow channel is banking depth. Private credit contracts with a lag of one to two years, persistently, and apparently more readily than it recovers. Remittances are the channel where formal and informal behaviour diverge most: compliance drives push flows into documented channels, but under high pressure a depreciating rupee drives them out again.

These results connect the sanctions-deflection literature to the de-risking literature. Studies of target countries emphasize that trade is deflected rather than destroyed (Haidar, 2017; Early, 2015). Our results show the mirror image from the bystander's side: the deflection creates a hub-and-spoke structure in which the UAE captures intermediation rents, and the domestic banking system absorbs the compliance costs. The credit contraction is consistent with the cross-country evidence that grey-listing reduces capital inflows (Kida and Paetzold, 2021) and with bank-level evidence that sanctions cause banks to retreat well beyond sanctioned counterparties (Efing et al., 2023).

The deep-learning results deserve a candid reading. They do not show that deep networks out-forecast simple models on twenty annual observations; that claim would not survive the Diebold–Mariano and MCS tests. What they do show is that a flexible model that is free to find interactions assigns a large and channel-specific role to sanctions exposure, that the autoencoder recovers a disruption factor from outcomes alone that tracks the coded exposure, and that the channel ranking agrees with the econometric estimates. Used this way, deep learning serves as a cross-check on structure and a tool for attribution rather than as an oracle.

Policy Implications

Four implications follow. First, because the credit effect is slow and asymmetric, exiting the grey list or obtaining sanctions relief does not by itself restore banking depth; supervisors should pair exits with active efforts to rebuild correspondent relationships, for example through standardized due-diligence information shared with foreign correspondents and regional clearing arrangements. Second, the reserve channel implies that compliance shocks are also balance-of-payments shocks, and reserve-adequacy planning, including the timing of IMF programmes, should treat FATF reviews and enforcement risk as contingent liabilities. Third, the threshold result for remittances suggests that keeping the formal and informal exchange rates close matters most precisely when compliance pressure is high, since that is when a widening premium diverts inflows into informal channels. Fourth, since trade with Iran is re-routed rather than eliminated, policies that bring border and barter trade into documented, screened channels can lower intermediation costs without breaching sanctions, which matters again with the re-imposition of UN sanctions on Iran in September 2025.

Limitations

The study has five limitations. First, twenty annual observations are too few for deep networks to learn reliably from data alone; our regularization, ensembles and noise injection mitigate but cannot remove this, and the in-sample R^2 of the explanatory ensembles (above 0.99) shows that the networks interpolate the training years. The Shapley attributions describe the fitted function and should be read together with the econometric estimates rather than on their own. Second, the exposure variables are author-coded and the composite index uses equal weights; alternative weights could change magnitudes. Third, trade data are Pakistan-reported and miss the informal trade that dominates the Iran and Afghanistan corridors; partner-reported mirror data would be a natural robustness check. Fourth, the Russia-regime and enforcement indicators coincide with the 2022–2023 macroeconomic crisis,

so their separate effects are only weakly identified. Fifth, without bank-level correspondent-relationship data, de-risking is measured through its aggregate footprint. Monthly State Bank of Pakistan series on remittances by corridor and trade by partner, and a cross-country panel of grey-listed economies, would allow the deep models to be trained on samples large enough to test H6 properly.

Conclusion

This paper has reconsidered the costs of sanctions from the point of view of an economy that is not their target. Combining a small-sample econometric identification strategy with deep learning and exact Shapley attribution, we find that third-party sanctions, FATF listing and US enforcement shocks have imposed real costs on Pakistan through external liquidity, trade routing and, most persistently, banking depth. A one-standard-deviation rise in compliance pressure costs about 0.7 months of import cover on impact and lowers private credit by up to 0.45 points of GDP after two years, and the long-run association between pressure and credit depth is large, at about 25 points of GDP per index unit. Iran sanctions, not FATF listing, re-route recorded trade towards the UAE. Remittances switch regime above a pressure threshold. Deep networks reproduce and refine this structure and attribute a third to almost a half of their explanatory power to sanctions channels, but they do not deliver statistically significant forecasting gains on annual data. Collateral disruption is real, measurable and channel-specific, and studying it requires methods that are flexible enough to capture nonlinearity and honest enough to report what the data cannot yet show.

References

1. Ahn, D. P., & Ludema, R. D. (2020). The sword and the shield: The economics of targeted sanctions. *European Economic Review*, 130, 103587.
2. Artingstall, D., Dove, N., Howell, J., & Levi, M. (2016). *Drivers & impacts of derisking: A study of representative views and data in the UK*. John Howell & Co. for the Financial Conduct Authority.
3. Athey, S., & Imbens, G. W. (2019). Machine learning methods that economists should know about. *Annual Review of Economics*, 11, 685–725.
4. Bahdanau, D., Cho, K., & Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. *Proceedings of the International Conference on Learning Representations (ICLR)*.
5. Bai, J., & Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, 66(1), 47–78.
6. Bai, J., & Perron, P. (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, 18(1), 1–22.
7. Bańbura, M., Giannone, D., & Reichlin, L. (2010). Large Bayesian vector auto regressions. *Journal of Applied Econometrics*, 25(1), 71–92.
8. Besedeš, T., Goldbach, S., & Nitsch, V. (2017). You're banned! The effect of sanctions on German cross-border financial flows. *Economic Policy*, 32(90), 263–318.
9. Bishop, C. M. (1995). Training with noise is equivalent to Tikhonov regularization. *Neural Computation*, 7(1), 108–116.
10. Brown, R. L., Durbin, J., & Evans, J. M. (1975). Techniques for testing the constancy of regression relationships over time. *Journal of the Royal Statistical Society: Series B*, 37(2), 149–192.
11. Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., & Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1), C1–C68.
12. Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder–decoder for statistical machine translation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1724–1734.
13. Collin, M., Cook, S., & Soramäki, K. (2016). *The impact of anti-money laundering regulation on payment flows: Evidence from SWIFT data* (Working Paper No. 445). Center for Global Development.

14. Crozet, M., Hinz, J., Stammann, A., & Wanner, J. (2021). Worth the pain? Firms' exporting behaviour to countries under sanctions. *Economic Policy*, 36(105), 97–146.
15. Dickey, D. A., & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74(366), 427–431.
16. Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263.
17. Diebold, F. X., & Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1), 57–66.
18. Diebold, F. X., & Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1), 119–134.
19. Dizaji, S. F., & van Bergeijk, P. A. G. (2013). Potential early phase success and ultimate failure of economic sanctions: A VAR approach with an application to Iran. *Journal of Peace Research*, 50(6), 721–736.
20. Dreger, C., Kholodilin, K. A., Ulbricht, D., & Fidrmuc, J. (2016). Between the hammer and the anvil: The impact of economic sanctions and oil prices on Russia's ruble. *Journal of Comparative Economics*, 44(2), 295–308.
21. Durner, T., & Shetret, L. (2015). *Understanding bank de-risking and its effects on financial inclusion*. Global Center on Cooperative Security and Oxfam.
22. Early, B. R. (2015). *Busted sanctions: Explaining why economic sanctions fail*. Stanford University Press.
23. Efing, M., Goldbach, S., & Nitsch, V. (2023). Freeze! Financial sanctions and bank responses. *The Review of Financial Studies*, 36(11), 4417–4459.
24. Erbenova, M., Liu, Y., Kyriakos-Saad, N., López-Mejía, A., Gasha, G., Mathias, E., Norat, M., Fernando, F., & Almeida, Y. (2016). *The withdrawal of correspondent banking relationships: A case for policy action* (IMF Staff Discussion Note SDN/16/06). International Monetary Fund.
25. Felbermayr, G., Kirilakha, A., Syropoulos, C., Yalcin, E., & Yotov, Y. V. (2020). The global sanctions data base. *European Economic Review*, 129, 103561.
26. Gharehgozli, O. (2017). An estimation of the economic cost of recent sanctions on Iran using the synthetic control method. *Economics Letters*, 157, 141–144.
27. Goulet Coulombe, P., Leroux, M., Stevanovic, D., & Surprenant, S. (2022). How is machine learning useful for macroeconomic forecasting? *Journal of Applied Econometrics*, 37(5), 920–964.
28. Haidar, J. I. (2017). Sanctions and export deflection: Evidence from Iran. *Economic Policy*, 32(90), 319–355.
29. Hansen, B. E. (1996). Inference when a nuisance parameter is not identified under the null hypothesis. *Econometrica*, 64(2), 413–430.
30. Hansen, B. E. (2000). Sample splitting and threshold estimation. *Econometrica*, 68(3), 575–603.
31. Hansen, P. R., Lunde, A., & Nason, J. M. (2011). The model confidence set. *Econometrica*, 79(2), 453–497.
32. Harvey, D., Leybourne, S., & Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, 13(2), 281–291.
33. Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, 313(5786), 504–507.
34. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
35. Hufbauer, G. C., Schott, J. J., Elliott, K. A., & Oegg, B. (2007). *Economic sanctions reconsidered* (3rd ed.). Peterson Institute for International Economics.
36. Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1), 161–182.
37. Kida, M., & Paetzold, S. (2021). *The impact of gray-listing on capital flows: An analysis using machine learning* (IMF Working Paper WP/21/153). International Monetary Fund.
38. Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. *Proceedings of the*

International Conference on Learning Representations (ICLR).
39. Kwiatkowski, D., Phillips, P. C. B., Schmidt, P., & Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, 54(1–3), 159–178.
40. Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in Neural Information Processing Systems*, 30.
41. Litterman, R. B. (1986). Forecasting with Bayesian vector autoregressions: Five years of experience. *Journal of Business & Economic Statistics*, 4(1), 25–38.
42. Loshchilov, I., & Hutter, F. (2019). Decoupled weight decay regularization. *Proceedings of the International Conference on Learning Representations (ICLR)*.
43. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30.
44. MacKinnon, J. G. (2010). *Critical values for cointegration tests* (Queen's Economics Department Working Paper No. 1227). Queen's University.
45. Mahlstein, K., McDaniel, C., Schropp, S., & Tsigas, M. (2022). Estimating the economic effects of sanctions on Russia: An allied trade embargo. *The World Economy*, 45(11), 3344–3383.
46. Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: Concerns and ways forward. *PLoS ONE*, 13(3), e0194889.
47. McNown, R., Sam, C. Y., & Goh, S. K. (2018). Bootstrapping the autoregressive distributed lag test for cointegration. *Applied Economics*, 50(13), 1509–1521.
48. Medeiros, M. C., Vasconcelos, G. F. R., Veiga, Á., & Zilberman, E. (2021). Forecasting inflation in a data-rich environment: The benefits of machine learning methods. *Journal of Business & Economic Statistics*, 39(1), 98–119.
49. Morgan, T. C., Syropoulos, C., & Yotov, Y. V. (2023). Economic sanctions: Evolution, consequences, and challenges. *Journal of Economic Perspectives*, 37(1), 3–30.
50. Mullainathan, S., & Spiess, J. (2017). Machine learning: An applied econometric approach. *Journal of Economic Perspectives*, 31(2), 87–106.
51. Narayan, P. K. (2005). The saving and investment nexus for China: Evidence from cointegration tests. *Applied Economics*, 37(17), 1979–1990.
52. Neuenkirch, M., & Neumeier, F. (2015). The impact of UN and US economic sanctions on GDP growth. *European Journal of Political Economy*, 40, 110–125.
53. Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708.
54. Pesaran, H. H., & Shin, Y. (1998). Generalized impulse response analysis in linear multivariate models. *Economics Letters*, 58(1), 17–29.
55. Pesaran, M. H., Shin, Y., & Smith, R. J. (2001). Bounds testing approaches to the analysis of level relationships. *Journal of Applied Econometrics*, 16(3), 289–326.
56. Rice, T., von Peter, G., & Boar, C. (2020). On the global retreat of correspondent banks. *BIS Quarterly Review*, March, 37–52.
57. Shapley, L. S. (1953). A value for n-person games. In H. W. Kuhn & A. W. Tucker (Eds.), *Contributions to the theory of games II* (pp. 307–317). Princeton University Press.
58. Shin, Y., Yu, B., & Greenwood-Nimmo, M. (2014). Modelling asymmetric cointegration and dynamic multipliers in a nonlinear ARDL framework. In R. C. Sickles & W. C. Horrace (Eds.), *Festschrift in honor of Peter Schmidt* (pp. 281–314). Springer.
59. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15, 1929–1958.
60. Stock, J. H., & Watson, M. W. (1993). A simple estimator of cointegrating vectors in higher order integrated systems. *Econometrica*, 61(4), 783–820.
61. Toda, H. Y., & Yamamoto, T. (1995). Statistical inference in vector autoregressions with possibly integrated processes. *Journal of Econometrics*, 66(1–2), 225–250.
62. Vincent, P., Larochelle, H., Bengio, Y., & Manzagol, P.-A. (2008). Extracting and composing robust

features with denoising autoencoders. *Proceedings of the 25th International Conference on Machine Learning*, 1096–1103.

63. World Bank. (2015). *Withdrawal from correspondent banking: Where, why, and what to do about it*. World Bank Group.
64. Zivot, E., & Andrews, D. W. K. (1992). Further evidence on the great crash, the oil-price shock, and the unit-root hypothesis. *Journal of Business & Economic Statistics*, 10(3), 251–270.