WORKING PAPERS

N° 1804

November 2026

"Geopolitical Oil Risk and the Maritime Geography of Trade"

Stéphane Auray, Patrick Fève and Zakaria Moussa

Toulouse School of Economics

Geopolitical Oil Risk and the Maritime Geography of Trade

Stéphane Auray* Patrick Fève† Zakaria Moussa‡

November 5, 2026

Abstract

How does oil-related geopolitical risk reshape international trade? Combining a Bayesian structural VAR with a maritime-exposure network covering 37 countries, we document a persistent trade contraction following a transitory Oil-GPR innovation. Trade falls more than industrial production over the medium run, while trade prices respond modestly. Bilateral local projections reveal a larger short-run export contraction along high-Route-Risk corridors among exporters dependent on strategic chokepoints. This differential survives structural-shock uncertainty and country-specific monthly conditions, and conditional estimates indicate stronger transmission when supply chains are strained. Complementary CIF–FOB evidence suggests differential valuation adjustment, but becomes imprecise after shock uncertainty is propagated. Event studies identify episode-specific benefits from geographical alternatives rather than uniform protection. The results connect aggregate trade adjustment to maritime geography without attributing it to a single shipping-cost mechanism.

JEL classification: F14; F40; Q41; Q43; E32.

Keywords: Geopolitical oil shocks; World trade; Route risk; Maritime chokepoints; CIF–FOB wedge.

*CREST-Ensai and Rennes School of Business; stephane.auray@ensai.fr

†Toulouse School of Economics (TSE) and Université Toulouse 1; patrick.feve@tse-fr.eu. Patrick Fève acknowledges funding from the French National Research Agency (ANR) under the Investments for the Future (*Investissements d'Avenir*) program, grant ANR-17-EURE-0010, and under project ANR-23-CE26-0011-01.

‡LEMNA and University of Nantes; zakaria.moussa@univ-nantes.fr

1 Introduction

Oil-specific geopolitical disturbances can disrupt trade not only through energy prices and production, but also through the reliability, availability, and cost of maritime connections. Their effects may differ sharply between corridors with feasible alternatives and those concentrated around strategic passages.[1] This paper examines the aggregate trade response, its geographical incidence, and the valuation frictions and adjustment possibilities underlying that incidence.

We combine monthly world-trade and oil-market indicators, a Bayesian structural VAR, bilateral local projections, and a new Maritime Trade Exposure Network. Our trade-volume index disaggregates annual World Trade Organization benchmarks using monthly indicators while preserving annual totals; our trade-price index aggregates national export and import price changes using trade weights and reporting-coverage criteria.

The paper makes three contributions. First, we document the aggregate transmission of innovations in the Oil Geopolitical Risk index of Iacoviello and Tong (2026), which builds on the geopolitical-risk measurement framework of Caldara and Iacoviello (2022). The benchmark system contains Oil GPR, world oil production, trade volume, trade prices, the real oil price, and inventories. Ordering Oil GPR first permits contemporaneous effects on the other variables while excluding within-month feedback to the index. An augmented system adds industrial production and computes the trade–production response differential at each posterior draw.

The identified innovation is transitory in Oil GPR but produces persistent trade effects: oil production declines, the real oil price rises initially, inventories tend to fall, and world trade contracts more strongly than industrial production over the medium run, while trade prices respond modestly. Matched local projections and an event-based measure following Pinchetti (2026) corroborate the broad aggregate profile. The trade–production differential motivates our maritime analysis, but does not itself identify shipping frictions or exclude demand-side mechanisms.

Second, we examine the maritime geography of transmission. Multiple gateways and plausible routes generate passage-specific exposures and three composite indices: Maritime Vulnerability, Route Risk, and Oil Chokepoint Exposure. Unlike binary shortest-route indicators, these measures accommodate partial chokepoint dependence and feasible alternatives. Their geographical content is cross-validated using independently reconstructed routes from the National Geospatial-Intelligence Agency's World Port Index and an external routing algorithm.

The clearest bilateral finding is a more adverse short-run export response along high-Route-Risk corridors among exporters with high chokepoint dependence. The differential survives structural-shock uncertainty, additional trade dynamics, and exporter- and importer-date fixed effects. Conditional state-dependent estimates indicate stronger adverse transmission when

[1] Oil also moves through pipelines and inland transport, and pipeline connections can bypass particular maritime passages. Their substitution possibilities are geographically specific and capacity constrained. The EIA reports maritime oil flows of 79.7 million barrels per day in 2024, compared with world oil supply of 103.3 million barrels per day, underscoring the scale of maritime transport without implying that all oil trade is seaborne (U.S. Energy Information Administration, 2026). Our network characterizes potential maritime exposure, not the observed modal allocation of each bilateral flow.

pre-shock supply-chain pressures are elevated.[2]

Third, we examine valuation adjustment using mirror CIF and FOB reports. Conditional on the benchmark shock history, high-risk corridors display a more positive valuation-wedge response because FOB values respond more adversely than CIF values. The evidence becomes imprecise after shock uncertainty is propagated and is therefore suggestive rather than decisive mechanism evidence.[3]

Separate disruption event studies suggest that geographical alternatives can facilitate short-run adjustment, but their benefits are episode specific. These exercises use independently dated Hormuz, Suez, Panama, and Red Sea disruptions rather than the Oil-GPR innovation. Panama provides suggestive short-run evidence; Suez and Red Sea estimates are imprecise, while Hormuz faces limited identifying variation and pre-event dynamics. Geographical alternatives need not become usable capacity when congestion, vessel availability, voyage duration, or other network constraints bind.

The contribution is to connect aggregate trade adjustment to its bilateral maritime incidence, rather than to identify a unique shipping mechanism.The analysis complements structural oil-market research (Kilian, 2009; Kilian and Murphy, 2014; Baumeister and Hamilton, 2019a), gravity approaches (Anderson and van Wincoop, 2003; Head and Mayer, 2014), and the vessel-level perspective of Bai et al. (2026). It estimates neither a structural trade-cost elasticity nor an accounting decomposition of the aggregate trade–production differential, and does not observe effective shipping capacity or realized fleet reallocation.

Section 2 reviews the literature; Section 3 presents the aggregate analysis; Section 4 develops the maritime network and bilateral data; Section 5 reports the panel and event-study evidence; and Section 6 concludes. The appendices provide construction details, estimation methods, and robustness exercises.

2 Related literature

Our analysis connects research on oil-market disturbances, geopolitical risk, trade costs, and maritime networks through a common question: how are the trade effects of an oil-specific geopolitical innovation distributed across international corridors?

The structural oil-market literature distinguishes disturbances originating in supply, global activity, and oil-specific demand (Kilian, 2009). Kilian and Murphy (2014) incorporate invento-

[2]At the 75th percentile of exporter dependence, the months-1–6 Route Risk contrast is approximately -1.11 log-percentage points, with a conditional 90% credible interval of $[-1.62, -0.61]$. Integrating posterior shock uncertainty attenuates it to approximately -0.70, with an interval of $[-1.37, -0.03]$. Longer-window estimates are less precise. These are relative corridor responses, not absolute export effects; the other exposure indices yield less uniform patterns.

[3]At high exporter dependence, the conditional Route Risk wedge contrast averages approximately 0.30 log-percentage points over months 1–6 and 0.25 over months 1–12. The short-run CIF and FOB contrasts are approximately -0.84 and -1.13, respectively; the positive wedge contrast does not require a positive CIF response. Insurance, reporting, timing, valuation conventions, and commodity composition also affect mirror values, preventing a direct freight-cost or causal-mediation interpretation.

ries and speculative storage, while Baumeister and Hamilton (2019a) emphasize the identifying restrictions and prior information underlying structural conclusions. Baumeister et al. (2022) develop the Global Economic Conditions index used in our robustness analysis, and Kilian et al. (2009) document different external-adjustment patterns across oil exporters and importers. We complement these studies by examining trade quantities and prices relative to production and tracing bilateral geographical incidence rather than external balances alone.

Identification of oil-relevant disturbances builds on historical and narrative approaches (Hamilton, 1983). Pinchetti (2026) use geopolitical-risk and oil-price comovement around events to distinguish contractionary, inflationary geopolitical energy shocks from contractionary, deflationary geopolitical macro shocks. Zanetti and Verduzco-Bustos (2026) combine oil futures and geopolitical threats in an external instrument, documenting heterogeneous output and inflation responses associated with oil-import exposure and energy inventories. Brignone et al. (2026) further show that geopolitical transmission depends on shock size and type: larger positive risk shocks have disproportionately adverse U.S. effects, while threats and acts operate through different uncertainty and demand channels.

Our benchmark instead recursively identifies innovations in the Oil GPR index of Iacoviello and Tong (2026), building on Caldara and Iacoviello (2022). Matched local projections assess dynamic propagation, while the Pinchetti (2026) exercise uses alternative identifying information. These approaches are complementary, not interchangeable. Relative to Zanetti and Verduzco-Bustos (2026), our emphasis is world trade and its maritime incidence rather than cross-country output and inflation. Likewise, Känzig and Raghavan (2026) identify shipping-related supply-chain shocks from Suez and Panama disruptions and high-frequency shipping-rate changes, finding higher costs and prices, lower activity, and little geopolitical-risk response. We study the converse transmission question: how trade responds across maritime corridors following an oil-specific geopolitical innovation.

The trade literature establishes the importance of economic size, distance, and barriers (Anderson and van Wincoop, 2003; Head and Mayer, 2014). Chen and Hsu (2013) examine oil-price volatility and trade, while Hummels and Schaur (2013) show that time in transit is a substantial barrier, especially for parts and components. Our exposure interactions address route-dependent adjustment without estimating structural gravity or a trade-cost elasticity. Pair fixed effects absorb persistent bilateral characteristics; date effects absorb common disturbances, leaving relative corridor responses rather than absolute route effects or a decomposition of the aggregate trade–production differential.

Research on maritime access provides the geographical foundation. Feyrer (2011) use changes in Suez route availability to study trade, while U.S. Energy Information Administration (2017) and UNCTAD (2020) document concentrated flows and disruption risks at strategic passages. Network studies examine closure-induced changes in distances, time, costs, and congestion (Pratson, 2023), ports' systemic criticality (Verschuur et al., 2022), and losses from disruption hazards, delays, rerouting, insurance, freight increases, and interruption (Verschuur et al., 2025). We complement this work by linking a probabilistic bilateral network to dynamic responses to identified Oil-GPR innovations. Multiple gateways and routes capture partial passage dependence and feasible alternatives, with geographical cross-validation that does not equate modelled exposure with observed vessel movements.

A closely related operational perspective is provided by Bai et al. (2026). Using AIS vessel records for U.S.-bound container routes, they distinguish potential from effective capacity and model endogenous fleet allocation, congestion, markups, and sourcing. Their framework shows how adjustment can transmit disruption beyond the initially affected route. Our international corridor measures instead describe ex ante geography: we neither observe fleet reallocation nor estimate congestion externalities or welfare. The event studies accordingly distinguish feasible alternatives from effective capacity constrained by vessel availability, congestion, and sailing time.[4]

3 A structural VAR for geopolitical risk and global trade

This section documents the aggregate effects of oil-specific geopolitical risk that motivate the bilateral analysis. We estimate a monthly Bayesian structural vector autoregression (SVAR) containing the Oil Geopolitical Risk (Oil GPR) index of Iacoviello and Tong (2026), global oil-market variables, and measures of world trade quantities and prices. The benchmark exercise asks whether an unexpected increase in oil-specific geopolitical risk contracts international trade and traces the associated adjustment in oil production, traded-goods prices, the real price of oil, and oil inventories. A separate augmented system introduces world industrial production and permits a comparison between the responses of trade and production at each posterior draw.

3.1 Data and transformations

The benchmark SVAR is estimated on monthly data from January 1984 to December 2024. Oil-specific geopolitical risk is measured by the Oil GPR index developed by Iacoviello and Tong (2026) and obtained from the companion dataset provided by the authors. Building on the text-based geopolitical-risk measurement framework of Caldara and Iacoviello (2022), the index uses a two-layer large-language-model classification of international newspaper articles. The first layer evaluates the geopolitical-risk intensity of each article, while the second identifies references to actual or potential disruptions to oil or energy supply, including conflict, sanctions, threats to energy infrastructure, and risks to critical energy transportation routes. We standardize the monthly index over the benchmark estimation sample.

Oil GPR should be interpreted as a text-based measure of the intensity of oil-related geopolitical risk rather than as a measure of realized physical disruptions. The recursively identified shock is an unexpected innovation in this index conditional on the lagged joint history of the oil market and world trade. It may therefore capture changes in expected scarcity and oil-supply risk even when no contemporaneous physical disruption is directly observed.

World oil production (WOP) is obtained from the U.S. Energy Information Administration's *Monthly Energy Review*. The real price of oil is the refiners' acquisition cost of imported

[4]Following Hummels and Lugovskyy (2006), we interpret CIF–FOB ratios as broad valuation discrepancies affected by insurance, timing, reporting, and composition, not direct freight-cost measures. Joint valuation decompositions and independently dated events provide complementary evidence without constructing exposure from realized post-event rerouting.

crude oil deflated by the U.S. consumer price index for all urban consumers. World crude-oil inventories are constructed following Kilian and Murphy (2014): total OECD petroleum inventories are multiplied by the ratio of U.S. crude-oil inventories to total U.S. petroleum inventories. The resulting series approximates the global stock of crude oil available to absorb oil-market disturbances.

World trade volume (WTRADE) is constructed by temporally disaggregating the annual World Trade Organization (WTO) merchandise-trade volume index using the Chow–Lin procedure (Chow and Lin, 1971). The candidate monthly indicators include the measure of world industrial production used by Baumeister and Hamilton (2019b), which aggregates industrial production across OECD economies and six major emerging economies, together with the real price of oil, U.S. real imports, world trade shares, and an aggregate export indicator constructed from IMF Direction of Trade Statistics. The preferred specification is selected according to its ability to reproduce the dynamics of the CPB World Trade Monitor over their common sample, while preserving the annual WTO benchmark exactly. Appendix A reports the disaggregation procedure and validation diagnostics. Because global industrial production is among the candidate indicators, Section A.2 verifies the trade–production comparison using the CPB World Trade Monitor directly.

The World Trade Price Index (WTPI) is constructed from country-level export and import price indices obtained from the IMF *International Financial Statistics*. Monthly national price changes are aggregated using contemporaneous trade weights from the IMF Direction of Trade Statistics. The resulting export and import price indices are chain-linked and combined using their respective shares in world trade. Reporting-coverage and concentration restrictions limit distortions from missing observations, and the final index is normalized to an average of 100 in 2015. Appendix A provides the complete construction method.

World trade volume and world oil inventories (WSTOCKS) are seasonally adjusted using X-13. World oil production, seasonally adjusted world trade volume, the World Trade Price Index, and seasonally adjusted inventories enter in logarithms after removal of a quadratic deterministic trend. A tilde denotes the resulting cyclical component. The real oil price enters in logarithmic levels. Responses of variables expressed in logs or log cycles[5] are multiplied by 100 and can be interpreted approximately as percentage deviations from their pre-shock paths.

3.2 Specification, identification, and Bayesian estimation

The benchmark six-variable system is

$$y_t = A_0 + \sum_{\ell=1}^{p} A_\ell y_{t-\ell} + u_t, \qquad u_t \sim \mathcal{N}(0, \Sigma_u), \tag{1}$$

[5]We assess the sensitivity of the aggregate responses to alternative model specifications. In addition to the cyclical benchmark, we consider specifications in logarithmic levels and monthly log differences. The latter produces cumulative responses expressed in levels. Across these specifications, the qualitative decline in world trade remains. The models also differ in their treatment of world oil inventories, so the comparison reflects changes in both variable transformations and system composition. The results are reported in Appendix C.1.

where A_0 is a vector of constant terms, p is the number of lags, and the vector of endogenous variables is

$$y_t = \left[GPR_t^{oil} \quad \log \widetilde{WOP}_t \quad \log \widetilde{WTRADE}_t \quad \log \widetilde{WTPI}_t \quad \log RRAC_t \quad \log \widetilde{WSTOCKS}_t \right]'.$$

In practice, we set $p = 12$, allowing the monthly system to capture adjustment over a full annual cycle.[6] World industrial production (WIP) is introduced only in the augmented system considered below, where it is placed after world oil production and before world trade.

Structural innovations ε_t are identified recursively. Let

$$u_t = B\varepsilon_t, \qquad \mathbb{E}(\varepsilon_t \varepsilon_t') = I, \qquad \Sigma_u = BB',$$

where B is lower triangular. Our analysis focuses on the first structural innovation, corresponding to the first column of B, and examines its effects on the global oil market and world trade.

Because Oil GPR is ordered first, its structural innovation is proportional to the reduced-form innovation in the Oil GPR equation and does not combine contemporaneous innovations from the remaining equations. Recovering this first shock therefore does not require imposing a contemporaneous ordering among oil production, trade, prices, and inventories. The identifying restriction instead rules out within-month feedback from those variables to the Oil GPR innovation, while allowing the innovation to affect all remaining variables contemporaneously.

This assumption is motivated by the political and military origins of the underlying geopolitical developments. It nevertheless remains an identifying restriction: the recursive scheme does not rule out all forms of anticipation or omitted common shocks. We consequently complement it below with local projections and an alternative event-based measure of oil-related geopolitical risk.

At each posterior draw, the structural innovation is rescaled so that the impact response of the standardized Oil GPR index equals one. Because Oil GPR has unit variance over the benchmark sample, the reported responses correspond to an impact increase of one sample standard deviation in the index. Responses are traced from horizon zero to horizon 36.

The reduced-form coefficients and covariance matrix are estimated using a block-Minnesota prior. The prior shrinks the system toward parsimonious univariate dynamics while allowing stronger interactions within economic blocks than across blocks. Overall tightness is set to 0.10, with relative tightness parameters of 0.50 within blocks and 0.25 across blocks. Lag decay equals 1.50, and the prior standard deviation of the intercept is 10.

Posterior inference uses 45,000 Gibbs iterations. The first 5,000 are discarded and every twentieth subsequent draw is retained, yielding 2,000 posterior draws. Identification and normalization are applied separately at every retained draw. We report posterior medians and central 68% and 90% credible intervals. Appendix B.5 provides further estimation details.

[6] Specifications with 3, 6, and 9 lags preserve the qualitative response profile; see Appendix C.2, Figure 11.

3.3 Benchmark oil-market and trade responses

Figure 1 reports the benchmark responses. Oil GPR rises sharply on impact and returns close to baseline within a few months. The shock is therefore transitory in the risk indicator itself but has more persistent economic consequences. World oil production declines, while the real price of oil initially rises before returning toward baseline. This combination is consistent with an adverse supply-type disturbance that reduces effective supply, raises expected scarcity, or increases the risk premium attached to future oil availability.[7]

The median response of oil inventories is negative, although it is imprecisely estimated. This pattern is consistent with an initial drawdown of inventories following an unanticipated disruption to current oil supply, as documented in studies of oil supply shocks (Kilian and Murphy, 2014; Baumeister and Hamilton, 2019a). By contrast, Känzig (2021) finds that inventories increase following adverse news about future oil supply, consistent with precautionary stock-building.

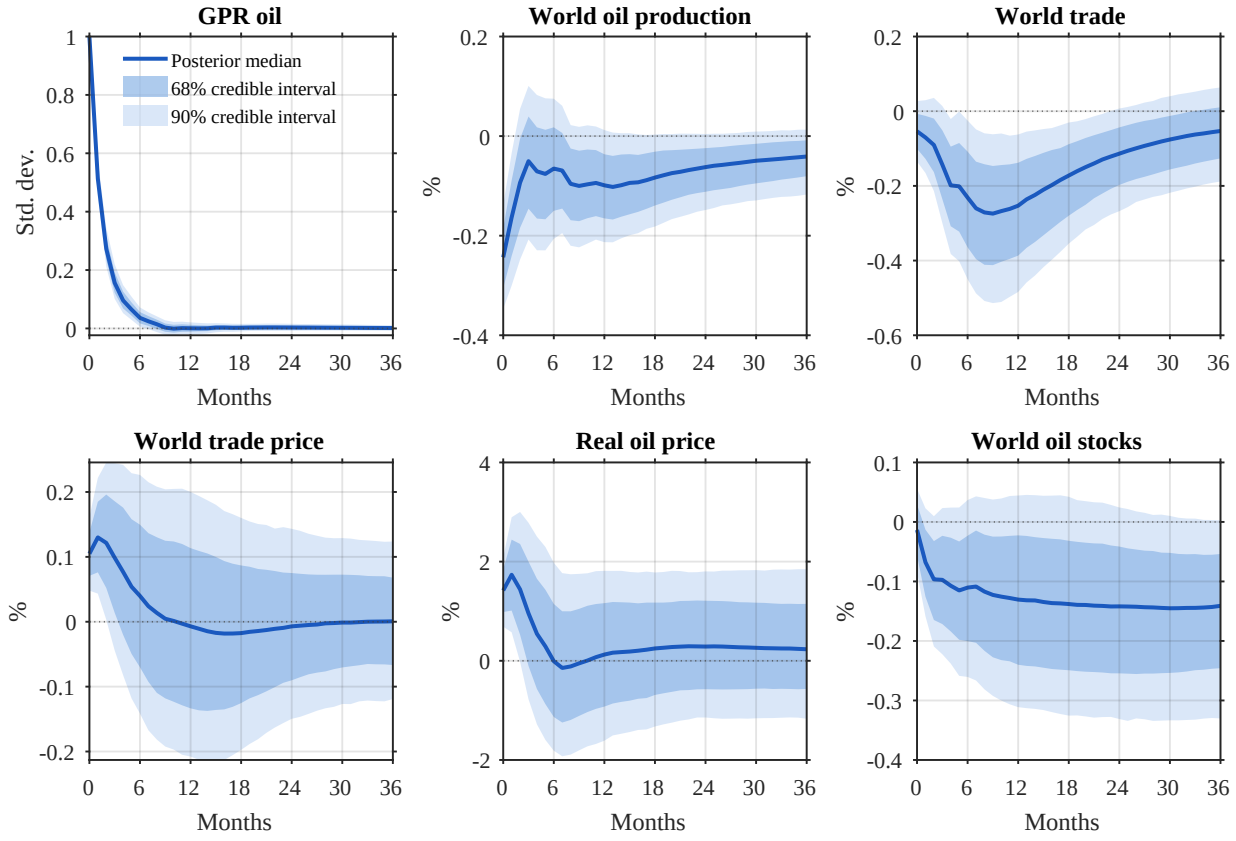
The estimated inventory response therefore suggests that our Oil GPR innovation may partly capture an unanticipated disturbance to current oil supply, rather than only news about future scarcity. However, the wide credible intervals prevent us from drawing a firm conclusion: the shock may combine realized disruptions and changes in expectations, and the inventory response alone cannot disentangle these components.

World trade contracts persistently after the disturbance, whereas the World Trade Price Index increases only modestly before returning toward baseline. This combination is consistent with two offsetting forces. Higher energy costs, insurance premia, route uncertainty, and other cross-border frictions can raise delivery costs, while weaker global demand and trade volumes limit aggregate price pass-through and induce adjustment primarily through quantities. These aggregate responses do not identify the relative contribution of these channels; the bilateral analysis below examines their geographical incidence.

Because the quantity variables enter as log cycles, persistence should not be interpreted as a permanent level effect. It indicates instead that adjustment unfolds gradually and remains economically relevant over the medium run. The qualitative response profile is preserved when the system is re-estimated through December 2019, although the shorter pre-pandemic sample produces wider credible intervals; Appendix C.3 reports this exercise.

[7]The SVAR Oil-GPR innovation is positively but imperfectly associated with the adverse oil-supply innovation of Baumeister and Hamilton (2019a). Over January 1985–December 2024, correlations computed separately across posterior SVAR histories have median Pearson correlation 0.228, with a central 90% interval of $[0.205, 0.249]$, but a much smaller median Spearman correlation of 0.076, with an interval of $[0.033, 0.117]$. Median correlations of first differences are 0.263 and 0.117, respectively. These results indicate limited positive comovement, not equivalence of the shocks. The Pearson–rank difference is consistent with an important role for large episodes, although an influence analysis would be needed to establish this. Figure 17 and Table 7 report the comparison. The uncertainty summaries vary the SVAR history conditional on the supplied BH series.

Figure 1: Responses to an oil-specific geopolitical disturbance



Notes: The figure reports posterior median responses to a one-standard-deviation SVAR Oil-GPR innovation. The monthly Bayesian SVAR is estimated from January 1984 to December 2024 and contains Oil GPR, world oil production, world trade volume, the World Trade Price Index, the real price of oil, and world oil inventories, in that recursive order. The model contains 12 monthly lags and a constant. Quantity variables and trade prices use the quadratic-detrended log transformations described in the text. Dark and light shaded regions denote 68% and 90% posterior credible intervals. Responses are shown over 36 months and are normalized so that the impact response of the standardized Oil GPR index equals one.

3.4 The trade–production response differential

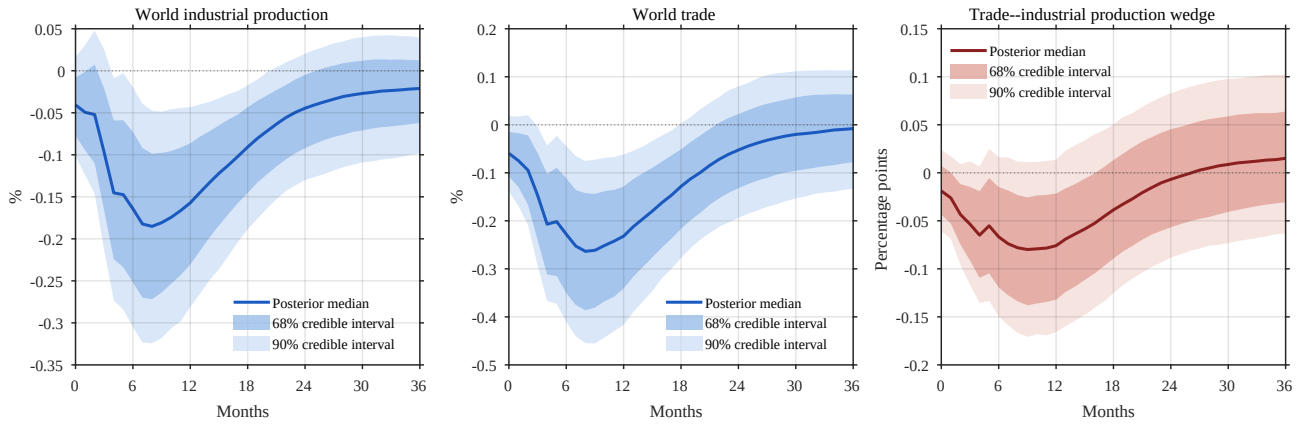
To assess whether the trade contraction exceeds the decline in global goods production, we estimate a seven-variable augmented system ordered as Oil GPR, world oil production, world industrial production, world trade, the World Trade Price Index, the real price of oil, and oil inventories. The lag length, prior, recursive identification, and draw-specific normalization are otherwise identical to the benchmark.

For posterior draw b and horizon h, define

$$\mathcal{D}_h^{(b)} = IRF_{WTRADE,h}^{(b)} - IRF_{WIP,h}^{(b)}. \tag{2}$$

A negative value indicates that world trade contracts more strongly than world industrial production. Computing $\mathcal{D}_h^{(b)}$ at each posterior draw preserves the joint posterior dependence between the two responses and yields credible intervals for their difference.

Figure 2: The trade–production response differential



Notes: The left and center panels report posterior responses of world industrial production and world trade to a one-standard-deviation SVAR Oil-GPR innovation. The right panel reports the trade–production response differential, $\mathcal{D}_h^{(b)} = IRF_{WTRADE,h}^{(b)} - IRF_{WIP,h}^{(b)}$, computed at each posterior draw b. A negative value indicates that trade contracts more strongly than production. The augmented monthly SVAR is estimated from January 1984 to December 2024, contains 12 lags, and orders Oil GPR, world oil production, world industrial production, world trade, the World Trade Price Index, the real price of oil, and oil inventories. Solid lines denote posterior medians; dark and light shaded regions denote 68% and 90% credible intervals.

Figure 2 shows that trade declines more strongly and remains weaker for longer than world industrial production. The differential becomes negative shortly after the shock and remains negative over the medium run.[8] This finding indicates that the international-trade response is not merely a mechanical reflection of the contraction in global production. It does not, by itself, establish that maritime frictions cause the difference or exclude contributions from demand, inventories, composition, or other cross-border adjustment margins.

3.5 Robustness of the aggregate trade response

We perform three additional exercises that assess the contraction in world trade rather than independently re-identifying the trade–production differential. First, we estimate horizon-specific local projections using the same historical SVAR Oil-GPR innovation and transformed variables as the benchmark system. For response variable j and horizon h, the specification is

$$y_{j,t+h} = a_{j,h} + \theta_{j,h} Shock_t^{SVAR} + \sum_{\ell=1}^{p} \Gamma_{j,h,\ell} y_{t-\ell} + e_{j,t+h}. \tag{3}$$

In practice, we set $p = 12$, as in the benchmark VAR, holding the lagged information set fixed when comparing propagation methods.

[8] As a measurement robustness check, we replace our WTO-based trade index with the CPB World Trade Monitor index and re-estimate both specifications over the common January 2000–December 2024 sample. The medium-run trade–production differential remains negative, but is smaller and less precisely estimated, with the 90% credible interval generally including zero. This exercise supports the qualitative ranking of the responses, while highlighting sensitivity to trade measurement; see Appendix C.4, Figure 13.

The LP estimates reproduce the initial decline in oil production, the persistent trade contraction, the short-run increase in the real price of oil, and inventory drawdowns (see Appendix C.5). They are less smooth and less precise at longer horizons, consistent with estimating each horizon separately. Because these LPs use the shock recovered from the benchmark SVAR, they assess sensitivity to the dynamic propagation method rather than to the recursive identification scheme. Their uncertainty intervals condition on the estimated benchmark SVAR shock history and do not propagate first-stage shock uncertainty.

Second, we use the standardized observed Oil GPR index directly in otherwise matched local projections. This measure combines predictable and unexpected variation and therefore has a descriptive rather than structural interpretation. The comparison nevertheless shows that the aggregate trade contraction is not confined to the recursively identified innovation. Third, we construct an alternative event-based measure following the logic of Pinchetti (2026), classifying sharp geopolitical-risk events using their short-window oil-price comovement and introducing the resulting monthly series in a Bayesian VARX. The response profile remains close to the benchmark: oil production and inventories decline, the real price of oil initially rises, and world trade contracts persistently. Appendices C.6 and C.7 report the observed-index LP and event-based VARX comparisons, respectively.

These measures answer related but distinct questions. Observed Oil GPR captures the prevailing intensity of oil-related geopolitical risk; the benchmark SVAR innovation isolates its unexpected component conditional on lagged system dynamics; and the Pinchetti-style indicator selects discrete events using high-frequency oil-price information. Their agreement supports the robustness of the aggregate trade contraction without treating the three objects as interchangeable structural shocks.

4 Bilateral trade corridors and the Maritime Trade Exposure Network

The aggregate evidence indicates a persistent trade contraction and a negative medium-run trade–production differential following a transitory Oil-GPR innovation. We examine its geographical incidence by combining bilateral trade data with a Maritime Trade Exposure Network. The network assigns model-implied weights to plausible routes and measures fixed corridor exposure rather than observed vessel movements or shipping shares.[9]

We combine monthly bilateral trade data with a Maritime Trade Exposure Network to study the geographical incidence of the aggregate disturbance. The network assigns model-implied weights to plausible routes and measures fixed corridor exposure, not observed vessel movements or shipping shares.

[9]Oil production declines, the real oil price initially rises, and inventories tend to fall, consistent with an adverse supply-type disturbance. The CPB comparison preserves the direction and timing of the trade–production differential, but its magnitude and precision are measurement sensitive. Alternative propagation methods and event-based identification separately support the trade contraction. These aggregate exercises do not identify the individual contributions of freight, insurance, delays, rerouting, supplier substitution, or other cross-border margins.

4.1 Bilateral trade data and sample coverage

The source database covers 44 economies over 2000–2025: 1,892 directed pairs and approximately 590,000 potential corridor–month observations. IMF *Direction of Trade Statistics* (DOTS) distinguish exporter-reported FOB exports from importer-reported CIF imports for each corridor ij.

The final 37-country universe spans January 2001–December 2024. Continuity, coverage, and connectivity screening retains 1,271 of 1,332 potential directed pairs. This sampling frame is further restricted by outcome, lag, and common-support requirements in each regression; Appendix D.1 documents selection.

Average-response projections deflate FOB exports by the common WTPI, without recovering product-specific bilateral quantities. Section 5 distinguishes backward-looking twelve-month changes from forward outcomes relative to the pre-disturbance month.

CIF and FOB reports are not matched shipments. Their ratio captures a broad valuation wedge, potentially reflecting freight, insurance, reporting, timing, valuation conventions, and composition, rather than freight rates alone.

Six controls—exporter and importer inflation, real effective exchange-rate growth, and oil-production growth—enter with a one-month lag. Industrial production is excluded because limited coverage would substantially reduce the sample; importer oil dependence enters separately as a structural exposure.

4.2 Maritime gateways and probabilistic route exposure

Each corridor connects multiple national gateways on a stylised graph of 57 nodes and 99 navigable links representing ocean basins, coastal interfaces, and strategic passages. Fixed normalized gateway weights summarize maritime access, not observed port traffic. Figure 3 displays the graph and aggregate route usage from the construction below.

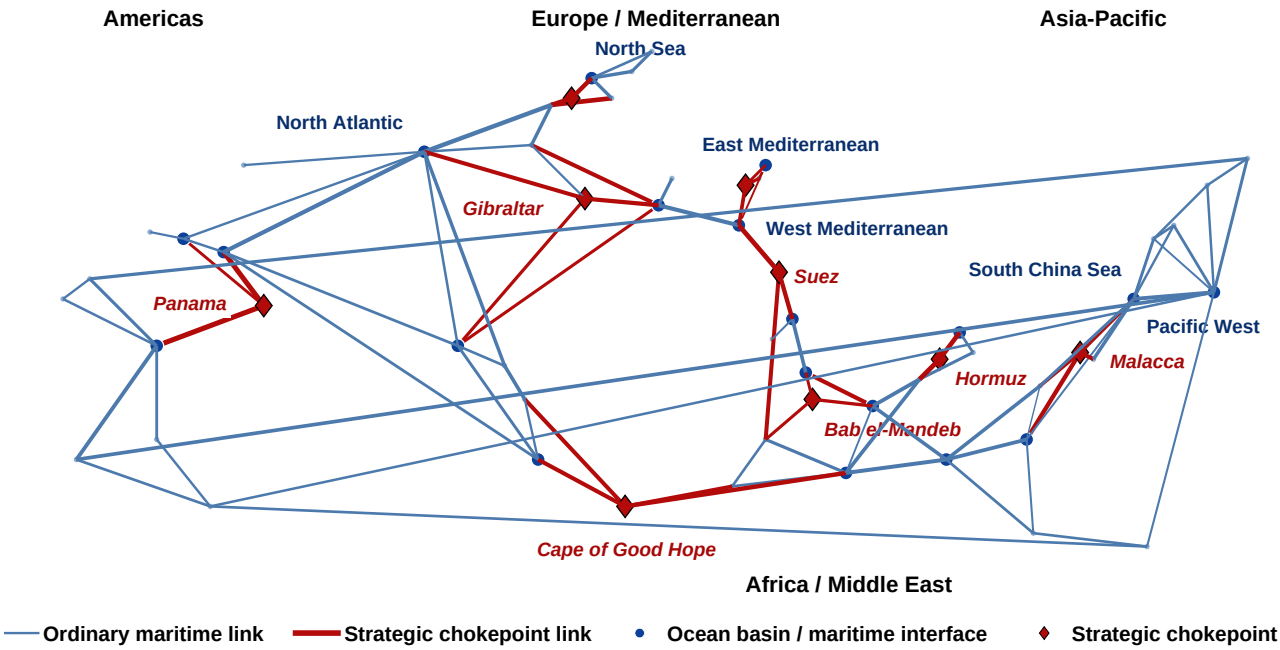
Figure 3 highlights shared connections, passage configurations, and alternative paths. Corridors can overlap on common segments, while similar distances can conceal different constrained passages or several passages in succession. Exposures consequently allow co-exposure rather than mutually exclusive route classifications. Link widths summarize model-implied usage, not cargo, vessel counts, or closure losses.

Alternative gateway paths permit partial chokepoint dependence and generate the continuous exposures interacted with Oil GPR. They represent geographical possibilities, not realized rerouting or sufficient operational capacity during a disruption.

An iterative shortest-path search generates up to five paths per gateway pair, penalizing previously used links by 1.35. Retention requires recorded search cost no greater than 1.5 times the minimum candidate cost. These penalized costs are neither physical distances nor observed monetary costs.

A multinomial-logit rule with fixed cost sensitivity 0.85 supplies conditional route weights.

Figure 3: Stylised maritime network and probabilistic route usage



Notes: The figure displays the stylised maritime graph used to construct bilateral route-exposure measures. Nodes represent ocean basins, coastal interfaces, and strategic passages. Blue links denote ordinary maritime segments; red links denote strategic passage segments. Link width reflects aggregate model-implied route usage across corridors in the 44-country route-construction universe, combining the weights of retained routes that traverse each segment. Colors identify segment types, whereas widths summarize their use within the routing construction; neither represents an estimated disruption probability. The network is a representation of plausible maritime connections, not a map of observed vessel trajectories, cargo volumes, or operational shipping capacity.

Combining these with gateway weights, aggregating identical paths, and normalizing yields corridor weights summing to one. These are construction parameters, not estimated commercial choices; Appendix D.2 details the algorithm.

Let $\mathcal{R}_{ij}$ denote the retained routes for corridor (i, j) and let π_{ijr} denote the final weight of route r. Exposure to passage c is

$$E_{ij}^{c} = \sum_{r \in \mathcal{R}_{ij}} \pi_{ijr} \mathbf{1}\{c \in r\}, \qquad \sum_{r \in \mathcal{R}_{ij}} \pi_{ijr} = 1.$$

The measure lies between zero and one and records the share of model-implied route weight traversing the passage. It allows partial exposure when feasible paths use different passages, unlike a binary indicator based on a single shortest route.

The ten passages and routing corridors are Suez, Hormuz, Bab el-Mandeb, Malacca, Panama, Bosphorus–Dardanelles, Black Sea access, the Cape of Good Hope, Gibraltar, and Dover. Their

fixed exposures use neither contemporaneous shocks nor monthly trade, but remain conditional on gateways, costs, and routing rules; fixed inputs need not predate every disturbance.

4.3 Composite exposure and exporter dependence

We summarize corridor exposure using three indices with distinct definitions. Maritime Vulnerability combines the ten passage exposures using fixed salience weights:

$$MV_{ij}^{raw} = \sum_{c \in \mathcal{C}} \nu_c^{MV} E_{ij}^c, \qquad MV_{ij} = \frac{MV_{ij}^{raw}}{\max_{k\ell} MV_{k\ell}^{raw}},$$

where the maximum is taken over the route-construction sample. The normalized index lies between zero and one. The salience weights are calibrated rather than estimated from trade outcomes.

Route Risk assigns a unit weight to each of the ten passages:

$$RR_{ij} = \sum_{c \in \mathcal{C}} E_{ij}^c.$$

It measures the expected number of strategic passages traversed by the modelled routes and may exceed one. It does not directly measure route concentration, disruption probabilities, or rerouting capacity.

Oil Chokepoint Exposure focuses on Hormuz and Bab el-Mandeb:

$$OC_{ij} = \min\left\{E_{ij}^{Hormuz} + E_{ij}^{Bab\text{-}el\text{-}Mandeb},\, 1\right\}.$$

This capped sum lies between zero and one but is not a route-level probability of crossing at least one passage. Appendix D.2 reports all definitions and salience weights.

Exporter dependence averages export-weighted exposure to Hormuz, Bab el-Mandeb, and Suez using positive exports in the stable network over 2001–2024. It captures broader exporter reliance interacting with corridor exposure, but its fixed trade weights are not predetermined before every shock.

4.4 Independent geographical validation

Independent NGA World Port Index gateways and the `searoute` graph provide a geographical cross-check without benchmark route weights or passage assignments. Across 1,271 corridors, Spearman correlations are 0.748 for Bab el-Mandeb, 0.772 for Suez, and 0.783 for the external oil indicator; correspondence remains positive after country fixed effects.

Validation is measure-specific: Maritime Vulnerability closely tracks the external expected pas-

sage count, whereas Route Risk relates more moderately to detours. Both Hormuz constructions classify the same 125 exposed pairs, but country effects absorb all external variation, validating classification rather than intensity beyond country identities.

This is geographical validation, not verification of shipping shares or causal identification. Appendix D.6 reports diagnostics, limitations, and a separate aggregate PortWatch comparison. Section 5 relates these fixed attributes to bilateral responses.

5 Bilateral trade responses and maritime transmission

We distinguish trade-weighted average real-export responses, relative maritime-exposure responses, and CIF–FOB wedge contrasts. Aggregate contraction can coexist with relative corridor improvements, and valuation changes need not equal freight-cost changes.

The panel inherits, rather than independently identifies, the aggregate SVAR disturbance. Average-response projections exploit time-series variation; Observed Oil GPR and AR innovations diagnose predictable persistence, while structural interpretation rests on the SVAR assumptions. Exposure regressions instead absorb exact-date effects and identify relative corridor gradients.

The clearest adverse differential concerns Route Risk at high exporter dependence; it survives bilateral dynamics and country-specific time controls. Other exposures yield less systematic amplification.

Country universe and estimation samples. The 37-country, 1,271-corridor frame described in Section 4.1 is selected from 1,332 potential pairs using continuity, availability, exposure coverage, and connectivity, never regression significance; Appendix D.1 reports sensitivity.

Outcome availability, lags, common support, and mirror-flow requirements produce exercise-specific samples, held fixed across LP horizons. Export and wedge projections exclude windows overlapping March 2020–December 2021, not all post-pandemic observations. Static regressions, average and heterogeneous LPs, and decompositions therefore differ in samples and outcomes.

Further details on the bilateral average-response estimation and the treatment of generated-shock uncertainty are provided in Appendices E and G.

5.1 Average dynamic response of bilateral real exports

We first estimate the trade-weighted average export response without maritime-exposure or exporter-dependence interactions, before examining its geographic allocation.

Real bilateral exports and cumulative outcomes. Let $X^N_{ij,t}$ denote nominal FOB exports from exporter i to importer j in month t. We construct WTPI-deflated bilateral exports as

$$X^R_{ij,t} = \frac{X^N_{ij,t}}{WTPI_t/100}, \qquad x^R_{ij,t} = \log X^R_{ij,t}.$$

WTPI adjusts nominal exports for common world trade prices, not product-specific bilateral price changes.

For each projection horizon $h = 0, \ldots, 12$, the dependent variable is the cumulative change in real exports relative to the month preceding the disturbance:

$$y_{ij,t+h} = 100 \left(x^R_{ij,t+h} - x^R_{ij,t-1} \right).$$

Horizon zero captures the change from $t-1$ to t; subsequent horizons measure net cumulative log-level departures from the same pre-shock month, not total export losses over intervening months.

Panel local projections. For each treatment r of the Oil-GPR signal, we estimate

$$\begin{aligned} y_{ij,t+h} = \beta^{(r)}_{0,h} s^{(r)}_t + \sum_{\ell=1}^{q_r} \psi^{(r)}_{\ell,h} s^{(r)}_{t-\ell} + \sum_{\ell=1}^{p} \rho^{(r)}_{\ell,h} \Delta x^R_{ij,t-\ell} \\ + \mathbf{X}'_{ij,t-1} \gamma^{(r)}_h + \alpha^{(h,r)}_{ij} + \sum_{m=2}^{12} \delta^{(r)}_{m,h} I_{m,t} \\ + \vartheta^{(r)}_{1,h} Trend_t + \vartheta^{(r)}_{2,h} Trend^2_t + \varepsilon^{(h,r)}_{ij,t+h}, \qquad h = 0, \ldots, 12. \end{aligned} \tag{4}$$

Here, $s^{(r)}_t$ is the standardized Oil-GPR measure and $I_{m,t}$ denotes calendar month m. In practice, we set $p = 12$, as in the aggregate VAR. We also set $q_r = p$ except in the diagnostic specification using Observed Oil GPR without shock lags, where $q_r = 0$.

Twelve export-growth lags capture corridor dynamics; $\mathbf{X}_{ij,t-1}$ contains the lagged country controls. Pair effects absorb fixed bilateral attributes, calendar-month indicators seasonality, and quadratic trends low-frequency movements. Date effects would absorb the common impulse, so $\beta^{(r)}_{0,h}$ uses conditional time-series variation, unlike the relative exposure projections below.

Weighted least squares uses fixed corridor weights proportional to mean 2000–2019 nominal exports, giving larger relationships greater influence without monthly reweighting. The reference window need not precede every shock. Conditional covariance matrices are two-way clustered by pair and date.

Alternative treatments of Oil GPR. The structural benchmark is the recursively identified SVAR Oil-GPR innovation. AR(12) and BIC-selected innovations remove predictability from the index's own history; BIC selects AR(1). These innovation measures use different information and identifying assumptions.

Observed Oil GPR enters either with twelve own lags or without them. The latter combines unexpected and predictable persistence and is diagnostic, not structural. All measures are centered and standardized over January 2000–December 2024.

Observed-index regressions describe the prevailing risk environment; AR residuals isolate own-history innovations, and the SVAR supplies the model-identified disturbance. Their comparison tests sensitivity to persistence without treating the impulses as economically identical.

Bias correction and uncertainty. We apply the split-panel jackknife before hierarchical cross-horizon smoothing, using the cluster-robust covariance of the corrected estimates. For the structural benchmark, we report both inference conditional on the median shock history and inference integrating 200 posterior BVAR histories through the complete estimation procedure.[10]

Common estimation sample. All five conditional specifications and horizons use 272,864 observations, 1,261 pairs, 37 countries, and 229 origin months from January 2002 to December 2023. Twelve-lag requirements determine the start; twelve-month forward outcomes in a panel ending December 2024 determine the endpoint.

Excluding windows intersecting March 2020–December 2021 retains both pre- and post-pandemic origins. Ten stable pairs fail the complete support requirements for outcomes, controls, histories, and horizons.

Average bilateral export responses. Figure 4 shows increasing contraction through months six to eight and only partial recovery thereafter across all five conditional specifications. Panel (a) shows attenuation after conditioning on observed-index lags; Panel (b) shows similar delayed SVAR and AR responses. The common timing is therefore not specific to one treatment of persistence.

We summarize the period of maximum contraction using the draw-by-draw average of cumulative responses over horizons six, seven, and eight:

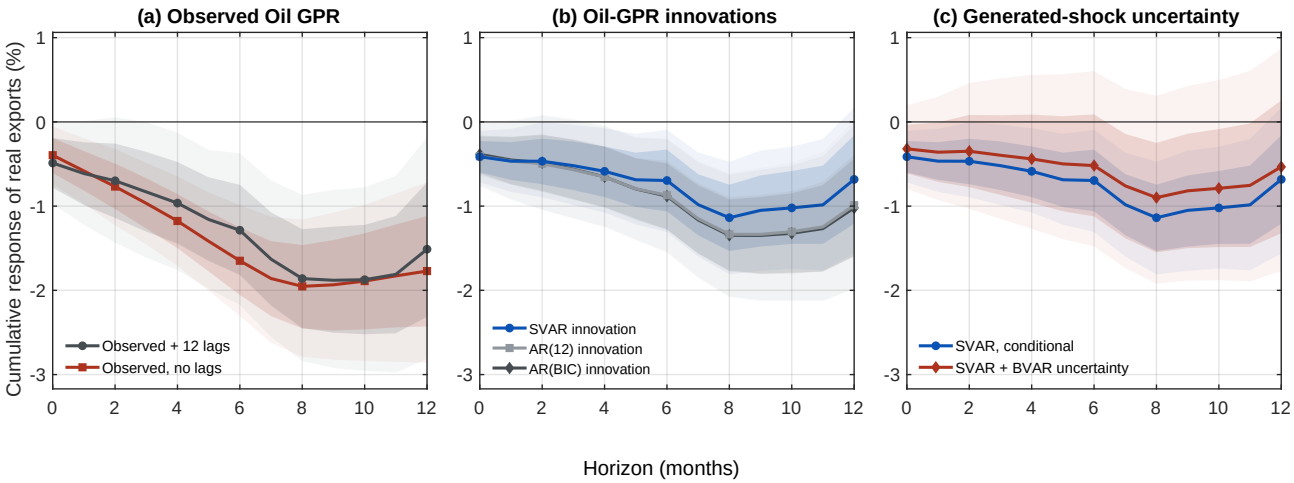
$$\overline{\beta}^{(b)}_{0,6:8} = \frac{1}{3}\sum_{h=6}^{8} \beta^{(b)}_{0,h}.$$

Draw-by-draw calculation preserves cross-horizon dependence; the statistic averages cumulative departures, not monthly export losses.

In Table 1, the conditional SVAR months-six-to-eight mean is -0.939 log-percentage points, its 90% interval is $[-1.549, -0.326]$, and its negative-sign probability is 0.995: approximately a

[10]The jackknife combines full- and half-sample estimates to address leading finite-T bias; its influence functions retain covariance across estimators and horizons. A second-difference prior regularizes curvature without directly shrinking the response level or linear slope. Integrated inference combines conditional panel uncertainty with variation across recovered shock histories. Details are provided in Appendix E.1.

Figure 4: Average response of bilateral real exports across alternative Oil-GPR specifications



Notes: The figure reports posterior mean cumulative responses of WTPI-deflated bilateral FOB exports to a one-standard-deviation increase in the corresponding Oil-GPR measure. The outcome at horizon h is $100[\log X^R_{ij,t+h} - \log X^R_{ij,t-1}]$. Dark and light shaded areas denote 68% and 90% posterior credible intervals, respectively. Panel (a) compares Observed Oil GPR conditional on twelve monthly lags of the index with the specification that omits those lags. Panel (b) compares the benchmark SVAR Oil-GPR innovation with innovations from AR(12) and BIC-selected autoregressive models of Observed Oil GPR; the BIC selects one lag. Panel (c) compares inference conditional on the posterior-median SVAR shock history with inference integrating over 200 paired posterior BVAR shock histories. All shocks are standardized over January 2000–December 2024. All specifications use the same 272,864 observations, 1,261 directed pairs, 37 countries, and 229 origin months spanning January 2002–December 2023, after excluding observations whose response windows overlap March 2020–December 2021. All regressions include twelve lags of real-export growth, lagged exporter and importer controls, directed-pair fixed effects, eleven calendar-month indicators, and linear and quadratic trends. Twelve shock lags are included in every specification except "Observed, no lags." Fixed 2000–2019 corridor weights, two-way clustering by directed pair and exact date, SPJ correction, and hierarchical Bayesian smoothing are used throughout.

0.9% export shortfall in this medium-run window.

AR(12) and AR(BIC) means are -1.118 and -1.131; both 90% intervals exclude zero and both negative-sign probabilities are 0.996. The contraction is thus reproduced outside the multivariate SVAR construction.

Observed Oil GPR yields -1.592 with twelve lags (90% interval $[-2.504, -0.659]$), versus -1.820 without them ($[-2.572, -1.103]$). Conditioning on history attenuates but does not remove the association, consistent with persistence contributing to its unlagged magnitude.

Panel (c) integrates BVAR uncertainty: the SVAR mean changes from -0.939 to -0.726, its 90% interval widens to $[-1.696, 0.418]$, and its negative-sign probability falls from 0.995 to 0.873. The posterior favors contraction but no longer excludes zero at 90%; fixing one history overstates precision without reversing the mean's sign or delayed timing.

All conditional innovations yield an economically and statistically meaningful delayed contraction, larger for Observed Oil GPR. We report the conditional structural path alongside the integrated distribution, which provides the more conservative uncertainty assessment.

Table 1: Average cumulative real-export response over months six to eight

Impulse measure	Posterior mean	90% credible interval	$\Pr(\overline{\beta}_{0,6:8} < 0)$
SVAR Oil-GPR innovation	**-0.939**	**[-1.549, -0.326]**	**0.995**
Observed Oil GPR + 12 lags	-1.592	[-2.504, -0.659]	0.998
Oil-GPR AR(12) innovation	-1.118	[-1.804, -0.437]	0.996
Oil-GPR AR(BIC) innovation	-1.131	[-1.812, -0.446]	0.996
Observed Oil GPR, no lags	-1.820	[-2.572, -1.103]	1.000
SVAR innovation, BVAR uncertainty integrated	**-0.726**	**[-1.696, 0.418]**	**0.873**

Notes: The table reports the posterior distribution of $\overline{\beta}_{0,6:8} = (\beta_{0,6} + \beta_{0,7} + \beta_{0,8})/3$, computed draw by draw so that posterior dependence across horizons is preserved. Posterior means and credible intervals are expressed in log-percentage points. The two bold rows report the benchmark SVAR Oil-GPR innovation, first conditional on the posterior-median shock history and then integrating uncertainty across 200 paired posterior BVAR shock histories. The remaining rows assess sensitivity to the persistence and predictable component of Observed Oil GPR. AR innovations are residuals from AR(12) and BIC-selected autoregressive models; the BIC selects one lag. All shock measures are standardized over January 2000–December 2024. All specifications use the same 272,864 observations, 1,261 directed pairs, 37 countries, and 229 origin dates spanning January 2002–December 2023, excluding response windows that overlap March 2020–December 2021. The regressions include the controls, fixed effects, dynamics, fixed corridor weights, and trends described in the notes to Figure 4. Reported probabilities are posterior sign probabilities and should not be interpreted as frequentist p-values.

The construction of corridor exposures, composite indices, and exporter chokepoint dependence is documented in Appendix D.2. Their descriptive properties and independent-route validation are reported in Appendices D.3 and D.6.

5.2 Continuous maritime exposure and exporter chokepoint dependence

We allow corridor exposure to interact with exporter chokepoint dependence: the same passage-dependent geography may be more consequential for an exporter whose wider trade also relies on strategic passages.

As an initial diagnostic, we relate the current disturbance to the backward-looking twelve-month change in bilateral exports. For shock specification s, the composite-exposure model is

$$\begin{aligned}\Delta_{12}\log Trade_{ijt} = {} & \beta_1^{(s)}\left(S_t^{(s)}\widetilde{E}_{ij}\right) + \beta_2^{(s)}\left(S_t^{(s)}\widetilde{D}_i\right) \\ & + \beta_3^{(s)}\left(S_t^{(s)}\widetilde{E}_{ij}\widetilde{D}_i\right) + \rho^{(s)}\left(S_t^{(s)}\widetilde{OID}_j\right) \\ & + \sum_{\ell=1}^{q_s}\boldsymbol{\pi}_\ell^{(s)\prime}\mathbf{Z}_{ij,t-\ell}^{(s)} + \mathbf{X}_{ij,t-1}'\gamma^{(s)} + \alpha_{ij} + \delta_t + \varepsilon_{ijt}^{(s)},\end{aligned} \tag{5}$$

where $\Delta_{12}\log Trade_{ijt} = \log Trade_{ijt} - \log Trade_{ij,t-12}$. The centered bilateral exposure $\widetilde{E}_{ij}$ is measured alternatively by the Maritime Vulnerability Index, Route Risk, or Oil Chokepoint Exposure. The variable $\widetilde{D}_i$ is centered exporter chokepoint dependence, $\widetilde{OID}_j$ is centered importer crude-oil net-import dependence, and $\mathbf{X}_{ij,t-1}$ contains lagged exporter and importer controls.

Directed-pair fixed effects, α_{ij}, absorb time-invariant bilateral characteristics, while exact-date fixed effects, δ_t, absorb the common shock component and all other aggregate monthly disturbances.

To determine whether the composite indices mask differences across strategic passages, we also estimate a chokepoint-specific model:

$$\begin{aligned}\Delta_{12}\log Trade_{ijt} = &\sum_{c\in\mathcal{C}}\beta_{1c}^{(s)}\left(S_t^{(s)}\widetilde{E}_{ij}^{c}\right)+\beta_2^{(s)}\left(S_t^{(s)}\widetilde{D}_i\right)\\ &+\sum_{c\in\mathcal{C}}\beta_{3c}^{(s)}\left(S_t^{(s)}\widetilde{E}_{ij}^{c}\widetilde{D}_i\right)+\rho^{(s)}\left(S_t^{(s)}\widetilde{OID}_j\right)\\ &+\sum_{\ell=1}^{q_s}\boldsymbol{\varpi}_\ell^{(s)\prime}\mathbf{W}_{ij,t-\ell}^{(s)}+\mathbf{X}_{ij,t-1}'\gamma^{(s)}+\alpha_{ij}+\delta_t+\varepsilon_{ijt}^{(s)},\end{aligned} \tag{6}$$

where $\mathcal{C}=\{\text{Hormuz},\text{Bab el-Mandeb},\text{Suez}\}$. All exposures and triple interactions enter jointly, so coefficients describe partial passage associations rather than separate regressions.

For the composite model, the lagged-interaction vector is

$$\mathbf{Z}_{ij,t-\ell}^{(s)}=\left[S_{t-\ell}^{(s)}\widetilde{E}_{ij},S_{t-\ell}^{(s)}\widetilde{D}_i,S_{t-\ell}^{(s)}\widetilde{E}_{ij}\widetilde{D}_i,S_{t-\ell}^{(s)}\widetilde{OID}_j\right]'.$$

In the chokepoint-specific model, the corresponding vector contains the exposure and triple interactions for each of the three passages, together with the exporter-dependence and importer-dependence interactions. We set $q_s=12$ for Observed Oil GPR with lagged interactions and $q_s=0$ for the remaining specifications. Thus, in the lag-augmented specification, the reported coefficients are contemporaneous effects conditional on the previous twelve months of shock interactions; they are not cumulative sums of current and lagged coefficients.

Because exact-date fixed effects absorb the common component of the shock, both models identify relative responses across corridors rather than the average effect of the disturbance on world trade. At a given level of exporter dependence, the relevant exposure gradients are

$$\frac{\partial ME^{rel,(s)}(E,D)}{\partial E}=\beta_1^{(s)}+\beta_3^{(s)}\widetilde{D},\qquad \frac{\partial ME_c^{rel,(s)}(E^c,D)}{\partial E^c}=\beta_{1c}^{(s)}+\beta_{3c}^{(s)}\widetilde{D}.$$

The triple interactions should therefore be interpreted jointly with their corresponding lower-order exposure interactions.

We use the five standardized measures introduced above: SVAR, observed index with twelve interaction lags, AR(12), AR(BIC), and observed index without lags. Their comparison diagnoses persistence rather than equating the identifying information.

Panel A of Table 2 shows negative Route Risk contrasts across all specifications: -0.0038 for SVAR, -0.0108 for Observed Oil GPR without lags, and -0.0057 to -0.0073 for lag-adjusted and prewhitened measures. Only the no-lag observed contrast is significant at 10%; it

Table 2: Static high-minus-low exposure contrasts: composite measures and individual chokepoints

Measure	SVAR innovation	Observed + 12 lags	AR(12) innovation	AR(BIC) innovation	Observed, no lags
Panel A: Composite maritime-exposure measures					
Maritime Vulnerability	-0.0017 (0.0041)	-0.0021 (0.0040)	-0.0016 (0.0043)	-0.0017 (0.0042)	-0.0030 (0.0061)
Route Risk	-0.0038 (0.0051)	-0.0073 (0.0049)	-0.0057 (0.0049)	-0.0060 (0.0048)	-0.0108* (0.0055)
Oil Chokepoint Exposure	0.0028 (0.0027)	0.0031 (0.0035)	0.0024 (0.0030)	0.0024 (0.0030)	0.0046 (0.0050)
Panel B: Chokepoint-specific exposures entered jointly					
Hormuz	0.0009 (0.0008)	0.0009 (0.0010)	0.0007 (0.0008)	0.0008 (0.0009)	0.0019* (0.0010)
Bab el-Mandeb	-0.0130* (0.0070)	-0.0182* (0.0099)	-0.0141* (0.0081)	-0.0139* (0.0079)	-0.0170* (0.0102)
Suez	0.0107 (0.0080)	0.0151 (0.0116)	0.0117 (0.0095)	0.0109 (0.0092)	0.0101 (0.0115)

Notes: Panel A reports high-minus-low contrasts from three separately estimated composite-exposure models. Panel B reports passage-specific contrasts from a single model in which exposures to Hormuz, Bab el-Mandeb, and Suez and their interactions with exporter dependence are entered jointly. For measure or passage m, the contrast is $\Delta ME_m^{(s)}(D_{75}) = (E_{H,m} - E_{L,m})[\widehat{\beta}_{1m}^{(s)} + \widehat{\beta}_{3m}^{(s)}(D_{75} - D_{50})]$, where $E_{H,m}$ and $E_{L,m}$ are mean exposure among corridors in the upper and lower quartiles of the relevant pair-level exposure distribution. The contrast is evaluated at the 75th percentile of exporter chokepoint dependence; D_{50} is its pair-level median. A negative estimate indicates that the backward-looking twelve-month change in bilateral exports is lower along highly exposed corridors. The dependent variable is $\Delta_{12} \log Trade_{ijt}$. The columns respectively use the benchmark SVAR Oil-GPR innovation; Observed Oil GPR with twelve monthly lags of all shock interactions; AR(12) and BIC-selected AR innovations; and Observed Oil GPR without lagged interactions. Every shock is standardized to one standard deviation. All specifications use the same observations and include the shock interaction with centered importer oil-import dependence, lagged exporter and importer controls, directed-pair fixed effects, and exact-date fixed effects. Standard errors in parentheses use the delta method and the full covariance matrix of the relevant lower-order and triple-interaction coefficients, with two-way clustering by directed pair and exact date. Significance markers refer to the contrast itself. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

combines predictable and unexpected variation. Maritime Vulnerability is uniformly negative but imprecise, and Oil Chokepoint Exposure uniformly positive but insignificant.

In Panel B, Bab el-Mandeb is negative and significant at 10% under all five measures: -0.0130 to -0.0182, approximately 1.3–1.8 log-percentage points in magnitude. At high exporter dependence, its corridors have weaker exports. Hormuz and Suez contrasts are positive; only no-lag Observed Oil GPR makes Hormuz significant, and Suez is always imprecise. Joint entry and differing support preclude interpreting these partial gradients as causal effects of closing a passage.

Table 3 shows negative Bab el-Mandeb lower-order interactions and positive significant triples: exporter dependence attenuates, but at D_{75} does not overturn, the negative gradient. Hormuz has positive lower-order and negative triple interactions, so its positive gradient diminishes with dependence. Neither Suez component is precise. These offsets motivate fitted gradients rather than isolated triple coefficients.

The static diagnostics identify Route Risk as the most stable negative composite pattern and Bab el-Mandeb as the clearest passage-specific association. Their limited precision and backward-looking outcome leave timing and persistence to the forward projections below.

Additional evidence on generated-shock uncertainty in the dynamic exposure responses is presented in Appendix G.2.

Table 3: Static interaction coefficients: composite maritime measures and individual chokepoints

Measure	Coefficient	SVAR innovation	Observed + 12 lags	AR(12) innovation	AR(BIC) innovation	Observed, no lags
Panel A: Composite maritime-exposure models						
Maritime Vulnerability	β_1	-0.0042 (0.0085)	-0.0052 (0.0084)	-0.0041 (0.0089)	-0.0042 (0.0088)	-0.0081 (0.0129)
	β_2	-0.0194 (0.0172)	-0.0118 (0.0212)	-0.0100 (0.0195)	-0.0100 (0.0188)	-0.0193 (0.0296)
	β_3	0.0453 (0.0663)	0.0590 (0.0835)	0.0476 (0.0690)	0.0501 (0.0681)	0.1239 (0.0928)
Route Risk	β_1	-0.0020 (0.0024)	-0.0036 (0.0023)	-0.0028 (0.0023)	-0.0030 (0.0023)	-0.0054** (0.0026)
	β_2	-0.0155 (0.0263)	0.0010 (0.0328)	0.0003 (0.0261)	-0.0001 (0.0260)	0.0081 (0.0298)
	β_3	0.0178 (0.0129)	0.0154 (0.0196)	0.0122 (0.0156)	0.0146 (0.0150)	0.0271 (0.0203)
Oil Chokepoint Exposure	β_1	0.0046 (0.0040)	0.0056 (0.0055)	0.0043 (0.0045)	0.0042 (0.0045)	0.0063 (0.0078)
	β_2	-0.0036 (0.0342)	0.0334 (0.0398)	0.0223 (0.0343)	0.0184 (0.0347)	-0.0375 (0.0512)
	β_3	-0.0187 (0.0567)	-0.0555 (0.0740)	-0.0390 (0.0587)	-0.0325 (0.0582)	0.0511 (0.0712)
Panel B: Chokepoint-specific model (passages entered jointly)						
Hormuz	β_{1c}	0.0179 (0.0139)	0.0202 (0.0189)	0.0160 (0.0151)	0.0181 (0.0156)	0.0384** (0.0184)
	β_{3c}	-0.1446 (0.0928)	-0.2286** (0.0984)	-0.1747* (0.0906)	-0.1797** (0.0903)	-0.2292** (0.1017)
Bab el-Mandeb	β_{1c}	-0.0307** (0.0156)	-0.0421* (0.0222)	-0.0327* (0.0182)	-0.0324* (0.0176)	-0.0405* (0.0227)
	β_{3c}	0.1907*** (0.0728)	0.2125** (0.0988)	0.1650* (0.0875)	0.1712** (0.0848)	0.2641*** (0.0946)
Suez	β_{1c}	0.0224 (0.0170)	0.0310 (0.0247)	0.0239 (0.0202)	0.0223 (0.0197)	0.0209 (0.0245)
	β_{3c}	-0.0453 (0.0583)	-0.0218 (0.0908)	-0.0168 (0.0719)	-0.0128 (0.0698)	-0.0253 (0.0888)
Exporter dependence	β_2	-0.0020 (0.0330)	0.0333 (0.0280)	0.0235 (0.0300)	0.0219 (0.0307)	0.0147 (0.0431)

Notes: Panel A reports $\widehat{\beta}_1$, $\widehat{\beta}_2$, and $\widehat{\beta}_3$ from the three composite-exposure specifications, which are estimated separately. Panel B reports coefficients from the joint chokepoint-specific specification. In that model, $\widehat{\beta}_{1c}$ is the lower-order shock–exposure interaction for passage c, $\widehat{\beta}_{3c}$ is its interaction with centered exporter chokepoint dependence, and $\widehat{\beta}_2$ is the common shock–exporter-dependence coefficient. The importer oil-dependence interaction and the coefficients on lagged interactions and controls are not reported. The five columns use the same shock definitions as in Table 2. In the second column of estimates, the contemporaneous coefficients are conditional on twelve monthly lags of all corresponding shock interactions. All models use the common estimation sample, directed-pair and exact-date fixed effects, and two-way clustering by directed pair and exact date. Standard errors are in parentheses. Significance markers refer to the individual coefficients. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

5.3 Dynamic heterogeneous effects across maritime corridors

Forward local projections trace relative export adjustment through twelve months, allowing exposure effects to vary with exporter dependence.

For each exposure measure and horizon $h = 0, \ldots, 12$, we estimate

$$\begin{aligned} 100\left(\log Trade_{ij,t+h} - \log Trade_{ij,t-1}\right) = {} & \beta_{1h}^{(s)}\left(S_t^{(s)}\widetilde{E}_{ij}\right) + \beta_{2h}^{(s)}\left(S_t^{(s)}\widetilde{D}_i\right) \\ & + \beta_{3h}^{(s)}\left(S_t^{(s)}\widetilde{E}_{ij}\widetilde{D}_i\right) + \rho_h^{(s)}\left(S_t^{(s)}\widetilde{OID}_j\right) \\ & + \sum_{\ell=1}^{q_s} \boldsymbol{\pi}_{\ell h}^{(s)\prime}\mathbf{Z}_{ij,t-\ell}^{(s)} + \mathbf{X}_{ij,t-1}'\Gamma_h \\ & + \alpha_{ij}^{(h)} + \delta_t^{(h)} + u_{ij,t+h}^{(h)}, \end{aligned} \tag{7}$$

where the lagged-interaction vector is

$$\mathbf{Z}_{ij,t-\ell}^{(s)} = \left[S_{t-\ell}^{(s)}\widetilde{E}_{ij}, S_{t-\ell}^{(s)}\widetilde{D}_i, S_{t-\ell}^{(s)}\widetilde{E}_{ij}\widetilde{D}_i, S_{t-\ell}^{(s)}\widetilde{OID}_j\right]'.$$

The benchmark sets $S_t^{(s)}$ equal to the SVAR Oil-GPR innovation. For the comparison based on the Observed Oil GPR index with explicit interaction dynamics, $q_s = 12$; the innovation specifications set $q_s = 0$. Pair fixed effects absorb time-invariant corridor characteristics and exact-date fixed effects absorb the common component of the shock. The model therefore identifies differential responses across corridors rather than the average response of bilateral real exports, which is estimated separately.

Because $\beta_{1h}^{(s)}$ and $\beta_{3h}^{(s)}$ jointly determine the exposure gradient, we report a high-minus-low exposure contrast. Let E_H and E_L denote mean exposure in the upper and lower quartiles of the pair-level distribution. Evaluated at the 75th percentile of exporter chokepoint dependence, the contrast is

$$\begin{aligned} \Delta ME_h(D_{75}) &= ME_h(E_H, D_{75}) - ME_h(E_L, D_{75}) \\ &= (E_H - E_L)\left[\beta_{1h}^{(s)} + \beta_{3h}^{(s)}(D_{75} - D_{50})\right]. \end{aligned} \tag{8}$$

A negative value indicates that cumulative exports respond more adversely along highly exposed corridors than along weakly exposed corridors. A positive value represents a relative response and does not necessarily imply that exports increase in absolute terms.

The complete coefficient vector is corrected using the split-panel jackknife before being smoothed jointly across horizons using the hierarchical Bayesian local-projection model. Estimation uses a common sample across all horizons, and observations whose response windows overlap March 2020–December 2021 are excluded.

Figure 5 shows that Route Risk provides the clearest evidence of adverse maritime transmission. Following the benchmark SVAR Oil-GPR innovation, the high-minus-low contrast averages -1.115 percentage points over months 1–6, with a 90% credible interval of $[-1.608, -0.611]$. The response remains negative over months 6–8 (-0.467) and over months 1–12 (-0.610). Hence, conditional on high exporter chokepoint dependence, cumulative exports contract more strongly along high-risk routes, and this relative contraction persists beyond the initial months after the shock.[11]

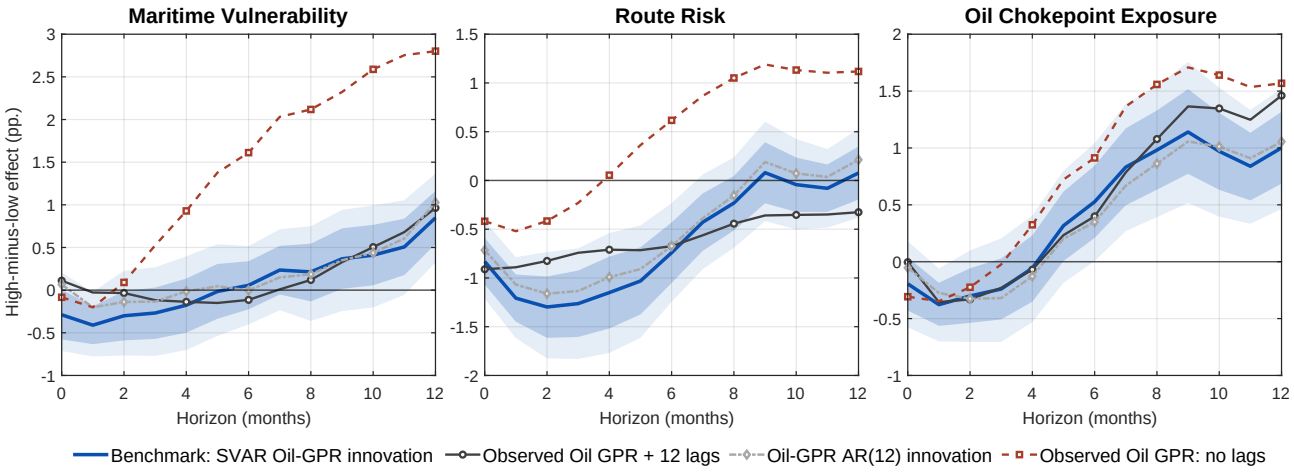
Observed Oil GPR with twelve interaction lags yields Route Risk contrasts of -0.760, -0.560, and -0.579 over the same windows; AR(12) yields -0.989, -0.405, and -0.497. Only the unlagged observed index has a markedly different medium-run profile. This convergence suggests that persistence and omitted interaction dynamics contribute to specification differences, rather than establishing a distinct mechanism.

Maritime Vulnerability remains imprecise around zero under SVAR and adjusted alternatives. Oil Chokepoint Exposure is near zero over months 1–6 and positive over 6–8 (approximately 0.780 under SVAR). This relative reversal could reflect rerouting, trade reallocation, or restoration, but establishes neither those mechanisms nor absolute export growth.

Transmission is selective: the adverse conditional differential concerns Route Risk, the expected

[11] A complementary network-permutation placebo assesses the geographical specificity of this result. Across 9,999 permutations, the months 1–6 Route Risk contrast has a one-sided permutation p-value of 0.048 and a two-sided p-value of 0.061. The evidence is weaker at longer horizons; see Appendix F.1, Figure 23.

Figure 5: Dynamic maritime-exposure contrasts following Oil-GPR disturbances



Notes: The figure reports posterior mean high-minus-low maritime-exposure contrasts at the 75th percentile of exporter chokepoint dependence. The solid blue line is the benchmark one-standard-deviation SVAR Oil-GPR innovation; the dark and light blue areas are its 68% and 90% posterior credible intervals. Gray lines report the Observed Oil GPR index with twelve lags of all shock interactions, the Observed Oil GPR index without interaction lags, and the standardized Oil-GPR AR(12) innovation. Markers and line patterns distinguish the three comparison specifications. The panels consider Maritime Vulnerability, Route Risk, and Oil Chokepoint Exposure. Estimates use 299,161 observations for 1,271 directed country pairs and include pair and exact-date fixed effects, lagged exporter and importer controls, and two-way pair/date clustering. The coefficient vector is split-panel-jackknife corrected before hierarchical Bayesian smoothing. All horizons use the same sample, excluding response windows overlapping March 2020–December 2021. Gray uncertainty bands are omitted for readability and are reported in the accompanying results table. The horizontal line denotes zero.

passage count, whereas Maritime Vulnerability is imprecise and Oil Chokepoint Exposure exhibits a later relative reversal.[12]

5.4 State dependence in global supply-chain pressures

State-dependent LPs add lagged standardized GSCPI and the complete hierarchy of identified exposure/dependence interactions, including three- and four-way shock terms. The shock–importer-oil-dependence interaction remains included; date-only terms are absorbed by exact-date effects.

Lagging GSCPI avoids defining the state by its immediate shock response, but does not identify a separate causal effect of logistics pressure.

For a given logistics state g and exporter dependence D, let $\mathcal{C}_h(g, D)$ denote the posterior high-minus-low bilateral-exposure contrast at horizon h. We define the GSCPI amplification effect

[12] As a robustness check, we add twelve monthly lags of bilateral export growth to control for trade persistence. The benchmark Route Risk contrast over months 1–6 remains nearly unchanged, moving from -1.115 to -1.054 percentage points, with a 90% credible interval that remains strictly negative. The short-run result also holds for the AR-based Oil-GPR innovations. By contrast, the delayed positive Oil Chokepoint Exposure contrast weakens and its credible interval includes zero. See Appendix F, Tables 12 and 13.

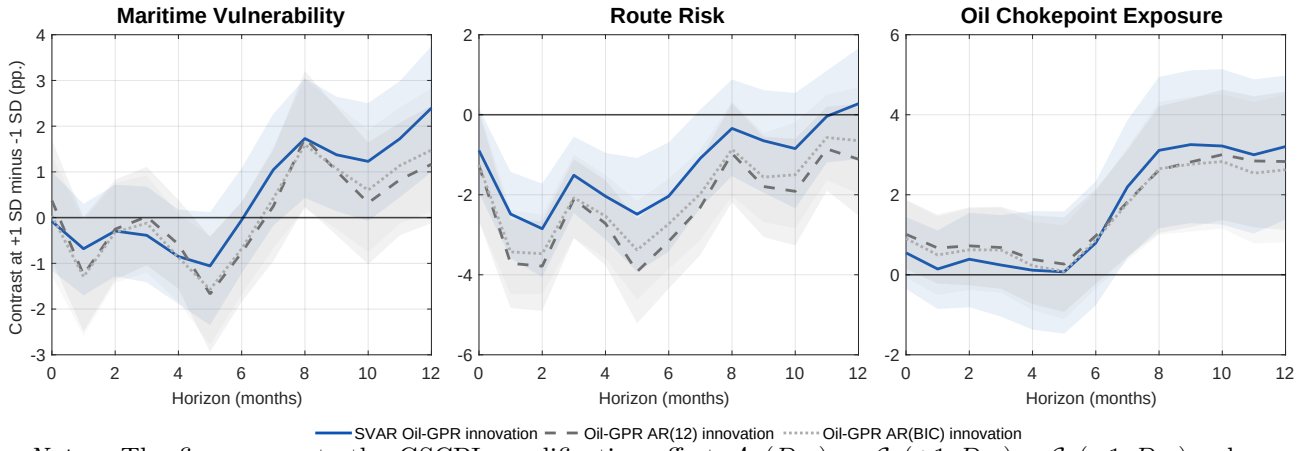
as

$$\mathcal{A}_h(D_{75}) = \mathcal{C}_h(+1, D_{75}) - \mathcal{C}_h(-1, D_{75}),$$

The states are one standard deviation above and below mean lagged GSCPI, at D_{75}. Negative amplification means a more adverse relative response under pressure; positive means more favorable, not necessarily absolute growth.

Amplification combines shock–GSCPI–exposure and four-way shock–GSCPI–exposure–dependence terms; the four-way sign alone is insufficient.

Figure 6: Global supply-chain pressures and the maritime transmission of Oil-GPR innovations



Notes: The figure reports the GSCPI amplification effect $\mathcal{A}_h(D_{75}) = \mathcal{C}_h(+1, D_{75}) - \mathcal{C}_h(-1, D_{75})$, where $\mathcal{C}_h(g, D_{75})$ is the high-minus-low bilateral-exposure contrast evaluated at logistics state g and at the 75th percentile of exporter chokepoint dependence. The solid blue line denotes the benchmark SVAR Oil-GPR innovation. The dashed and dotted gray lines denote innovations obtained from AR(12) and BIC-selected AR models of Observed Oil GPR, respectively. The BIC-selected model is AR(1). All shock measures are normalized to one standard deviation over their respective normalization samples. Shaded regions are 90% posterior credible intervals. Estimates are obtained from split-panel-jackknife hierarchical Bayesian local projections with directed-pair and exact-date fixed effects, two-way clustering by pair and date, six lagged exporter and importer macroeconomic controls, and the shock interaction with importer oil dependence. The estimation sample contains 299,161 observations, 1,271 directed country pairs, 37 countries, and 253 monthly dates between January 2001 and December 2024. Response windows overlapping March 2020–December 2021 are excluded. A negative value indicates that elevated supply-chain pressure makes the relative response of highly exposed corridors more adverse. A positive value indicates a more favorable relative response and does not necessarily imply an absolute increase in exports.

Figure 6 shows adverse SVAR Route Risk amplification of -2.23 log-percentage points over months 1–6 (90% interval $[-3.21, -1.29]$), -1.16 over 6–8 ($[-2.25, -0.05]$), and -1.34 over 1–12 ($[-2.29, -0.41]$). High pre-shock pressure therefore makes the exposure differential more adverse.

AR(12) and AR(BIC) amplification over months 1–6 is -3.24 and -2.94, with 90% intervals $[-4.11, -2.32]$ and $[-3.79, -2.08]$. Both retain precise negative amplification over 6–8 and 1–12, reproducing the pattern using own-history innovations rather than the multivariate identification.

SVAR Oil Chokepoint Exposure amplification is imprecise over months 1–6 (0.29), then positive over 6–8 (2.03; 90% interval $[0.32, 3.67]$) and 1–12 (1.65; $[0.24, 2.98]$). AR(12) and AR(BIC) yield similar 6–8 estimates of 1.80 and 1.77. These indicate delayed relative adjustment, not identified rerouting, substitution, or reallocation.

Maritime Vulnerability amplification changes from negative short-run to positive intermediate means under SVAR and AR innovations, but its 90% intervals do not establish a systematic medium-run effect.

Table 11 shows that unlagged Observed Oil GPR produces positive months-6–8 amplification of 4.03 for Maritime Vulnerability and 3.64 for Oil Chokepoint Exposure, but no credible Route Risk amplification.

Adding twelve interaction lags yields short-run Route Risk amplification of -1.33 (negative-sign probability 98.98%), while Oil Chokepoint Exposure remains positive over 6–8 (1.97) and 1–12 (1.69). Maritime Vulnerability falls from 4.03 to 2.75 over 6–8 but remains precise.

Conditioning on interaction dynamics thus moves the observed-index profile toward the innovations, suggesting persistence contributes to their differences.

State dependence is non-uniform: pressure amplifies adverse short-run Route Risk transmission and accompanies delayed relative improvement in Oil Chokepoint Exposure. The first is robust across innovation measures; the second is not proof of rerouting. These estimates condition on one SVAR history and omit both full shock-posterior propagation and the twelve trade-growth lags tested separately.

Further results for the CIF–FOB valuation wedge under generated-shock uncertainty are reported in Appendix G.3.

5.5 Maritime exposure and the CIF–FOB valuation wedge

We examine relative CIF–FOB valuation adjustment alongside export heterogeneity, using static contrasts and forward LPs rather than identifying physical freight costs directly.

Measuring the valuation wedge. For exports from country i to country j, define

$$W_{ijt} = \frac{CIF_{j\leftarrow i,t}}{FOB_{i\rightarrow j,t}},$$

where the numerator is the value of imports reported by country j on a cost-insurance-freight basis and the denominator is the corresponding export value reported by country i on a free-on-board basis. The logarithmic valuation wedge is

$$w_{ijt}^{raw} = \log W_{ijt} = \log CIF_{j\leftarrow i,t} - \log FOB_{i\rightarrow j,t}.$$

Positive w_{ijt}^{raw} means CIF exceeds FOB. As discussed in Section 4.1, unmatched mirror reports capture freight, insurance, handling, timing, reporting, valuation, and composition; they are not pure transport-cost observations.

For the static regressions, the logarithmic wedge is winsorized at the 1st and 99th percentiles of its distribution in the stable bilateral network:

$$w_{ijt} = \text{Winsor}_{1,99}\left(w_{ijt}^{raw}\right).$$

Only observations with strictly positive and finite CIF and FOB values are retained. Missing and zero trade values are not imputed. The static outcome is $100w_{ijt}$, so the estimates are expressed in log-percentage points.

Static exposure contrasts. For shock specification s, we estimate

$$\begin{aligned} 100w_{ijt} = S_t^{(s)}\left[\beta_1^{(s)}\widetilde{E}_{ij} + \beta_2^{(s)}\widetilde{D}_i + \beta_3^{(s)}\widetilde{E}_{ij}\widetilde{D}_i + \rho^{(s)}\widetilde{OID}_j\right] \\ + \sum_{\ell=1}^{q_s} \boldsymbol{\pi}_\ell^{(s)\prime}\mathbf{Z}_{ij,t-\ell}^{(s)} + \mathbf{X}_{ij,t-1}'\gamma^{(s)} + \alpha_{ij} + \delta_t + \varepsilon_{ijt}^{(s)}. \end{aligned} \tag{9}$$

The centered variables $\widetilde{E}_{ij}$, $\widetilde{D}_i$, and $\widetilde{OID}_j$ and six lagged country controls retain their benchmark definitions; exposure alternates between Maritime Vulnerability, Route Risk, and Oil Chokepoint Exposure.

For the specification using Observed Oil GPR with twelve interaction lags, $q_s = 12$ and

$$\mathbf{Z}_{ij,t-\ell}^{(s)} = \left[S_{t-\ell}^{(s)}\widetilde{E}_{ij},\quad S_{t-\ell}^{(s)}\widetilde{D}_i,\quad S_{t-\ell}^{(s)}\widetilde{E}_{ij}\widetilde{D}_i,\quad S_{t-\ell}^{(s)}\widetilde{OID}_j\right]'.$$

Other specifications set $q_s = 0$. Lag-adjusted coefficients condition on twelve previous shock interactions; they are not sums of distributed-lag coefficients.

Pair effects absorb persistent reporting and valuation differences; date effects absorb common wedge movements and the shock main effect, leaving relative corridor responses.

At a given level of exporter dependence, the exposure gradient is

$$\frac{\partial ME^{rel,(s)}(E,D)}{\partial E} = \beta_1^{(s)} + \beta_3^{(s)}\widetilde{D}.$$

The triple interaction must therefore be interpreted jointly with the lower-order shock–exposure interaction.

Using the benchmark upper/lower-quartile means E_H, E_L, the contrast at D_{75} is

$$\Delta ME^{(s)}(D_{75}) = (E_H - E_L)\left[\beta_1^{(s)} + \beta_3^{(s)}(D_{75} - D_{50})\right], \tag{10}$$

where D_{50} is the centering median. A positive contrast denotes a more positive relative wedge response, not necessarily absolute widening.

Table 4: Static maritime-exposure contrasts in the bilateral CIF–FOB valuation wedge

Specification	Maritime Vulnerability	Route Risk	Oil Chokepoint Exposure
SVAR innovation	0.6915* (0.4131)	0.4395* (0.2555)	0.7412* (0.4497)
Observed + 12 lags	0.7522 (0.5070)	0.3931 (0.2966)	0.7115 (0.5757)
AR(12) innovation	0.5787 (0.4152)	0.3034 (0.2454)	0.5451 (0.4757)
AR(BIC) innovation	0.6186 (0.4169)	0.3656 (0.2474)	0.5777 (0.4794)
Observed, no lags	1.2947** (0.5814)	0.6774* (0.3928)	1.3733** (0.6667)

Notes: High-minus-low maritime-exposure contrasts evaluated at the 75th percentile of exporter chokepoint dependence. High and low exposure are the mean values in the upper and lower corridor quartiles. The outcome is the log CIF–FOB wedge, winsorized at its 1st and 99th percentiles and multiplied by 100; units are log-percentage points. Only positive, non-missing CIF and FOB values are used. Shock measures are standardized over January 2000–December 2024. The common sample covers January 2001–December 2024, with 343,200 observations across 1,271 pairs and 37 countries. Specifications include pair and exact-date fixed effects, lagged macroeconomic controls, and the shock interaction with importer oil dependence. The Observed Oil GPR +12 lags contrast conditions on lagged shock interactions and is not a cumulative effect. Parentheses report delta-method standard errors using the full coefficient covariance matrix, two-way clustered by pair and date. The wedge is not a direct freight-rate measure. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

All static contrasts in Table 4 are positive. Conditional SVAR contrasts at high exporter dependence are 0.692 for Maritime Vulnerability, 0.440 for Route Risk, and 0.741 for Oil Chokepoint Exposure, each significant at 10%.

AR(12) and AR(BIC) contrasts range from 0.303 to 0.619 but are individually imprecise, as are the positive observed-index estimates with twelve interaction lags.

Without interaction lags, Observed Oil GPR yields larger contrasts: 1.295, 0.677, and 1.373, significant at 5%, 10%, and 5%, respectively. Because this specification combines predictable and unexpected variation, its magnitude is not evidence of a stronger causal effect; persistence influences the static association.

Dynamic exposure contrasts. To trace timing and persistence, for $h = 0, \ldots, 12$ we estimate

$$\begin{aligned} 100\left(w_{ij,t+h}^{raw} - w_{ij,t-1}^{raw}\right) = {} & \beta_{1,h}^{(s)}\left(S_t^{(s)}\widetilde{E}_{ij}\right) + \beta_{2,h}^{(s)}\left(S_t^{(s)}\widetilde{D}_i\right) \\ & + \beta_{3,h}^{(s)}\left(S_t^{(s)}\widetilde{E}_{ij}\widetilde{D}_i\right) + \rho_h^{(s)}\left(S_t^{(s)}\widetilde{OID}_j\right) \\ & + \sum_{\ell=1}^{q_s} \boldsymbol{\pi}_{\ell,h}^{(s)\prime}\mathbf{Z}_{ij,t-\ell}^{(s)} + \mathbf{X}_{ij,t-1}'\gamma_h^{(s)} + \alpha_{ij}^{(h)} + \delta_t^{(h)} + \varepsilon_{ij,t+h}^{(s,h)}. \end{aligned} \tag{11}$$

The outcome is the cumulative raw-log-wedge change from $t-1$ to $t+h$. Trimming observations outside its 1st and 99th percentiles before constructing CIF and FOB outcomes preserves their accounting identity; the components are not separately winsorized.

Joint estimation on the same sample gives the draw-by-draw identity

$$\mathcal{C}_h^w = \mathcal{C}_h^{CIF} - \mathcal{C}_h^{FOB}.$$

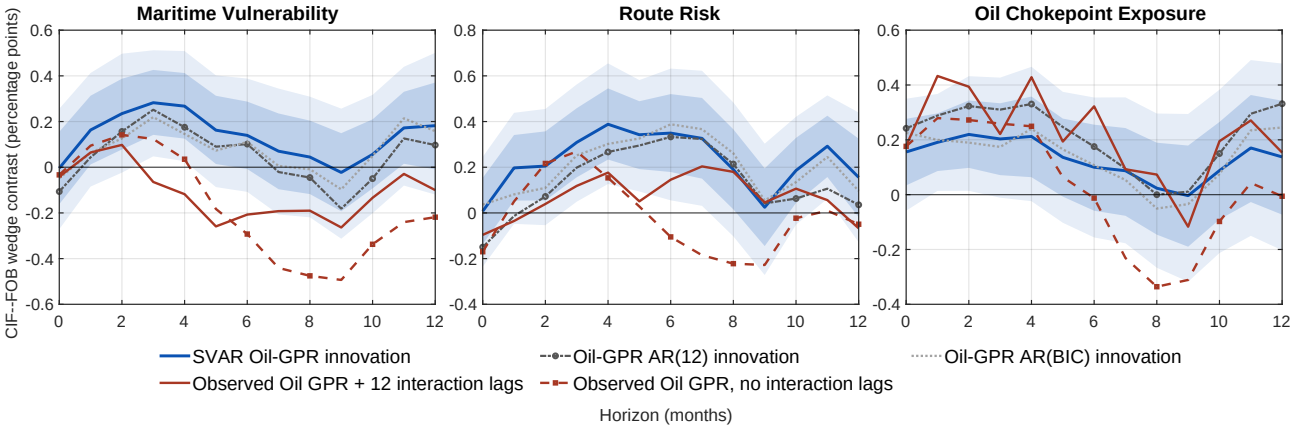
This procedure retains the posterior covariance between the two valuation components.

At each horizon, the high-minus-low wedge contrast is

$$\Delta ME_h^{(s)}(D_{75}) = (E_H - E_L)\left[\beta_{1,h}^{(s)} + \beta_{3,h}^{(s)}(D_{75} - D_{50})\right].$$

The horizon-specific estimates are corrected using the split-panel jackknife and combined through hierarchical Bayesian local-projection smoothing.

Figure 7: Dynamic maritime-exposure contrasts in the bilateral CIF–FOB valuation wedge



Notes: High-minus-low maritime-exposure contrasts in the cumulative CIF–FOB wedge response, evaluated at the 75th percentile of exporter chokepoint dependence. Blue denotes the SVAR Oil-GPR innovation, with 68% and 90% credible bands; gray denotes AR(12) and AR(BIC) innovations; brick-red denotes Observed Oil GPR with and without twelve lagged shock interactions. All measures are standardized over January 2000–December 2024. The common sample covers January 2001–December 2024, with 278,330 observations across 1,266 pairs and 37 countries, excluding COVID-overlapping response windows. CIF and FOB responses are jointly estimated after trimming the raw log wedge at its 1st and 99th percentiles; their difference is computed draw by draw. Units are log-percentage points. SVAR intervals condition on the benchmark shock history. Positive contrasts indicate relative widening along highly exposed corridors, not necessarily higher freight costs.

Figure 7 shows conditional SVAR Route Risk wedge contrasts of 0.298, 0.289, and 0.247 log-percentage points over months 1–6, 6–8, and 1–12, with credible intervals excluding zero at the reported levels. The positive differential emerges early and persists through the year.

AR(BIC) Route Risk contrasts are 0.243, 0.338, and 0.218, all credibly positive; AR(12) also gives a positive profile, including 0.290 over 6–8. Thus, the conditional pattern is not specific to the SVAR history.

Maritime Vulnerability is positive over 1–6 under SVAR (0.208), but smaller and imprecise later and over the full window, as with the AR profiles.

SVAR Oil Chokepoint Exposure is positive but imprecise in all windows. AR(12) yields positive 1–6 and 1–12 responses, while AR(BIC) is concentrated in the short window: widening is not uniform across measures.

Observed-index paths are sensitive to dynamics: twelve interaction lags generally yield imprecision, although short-run Oil Chokepoint Exposure remains positive. Without lags, intermediate contrasts turn negative, particularly Maritime Vulnerability (-0.403 over 6–8), and less strongly Route Risk and Oil Chokepoint Exposure.

This sensitivity supports prioritizing the structural innovation and using observed-index regressions to diagnose persistence across horizons.

Separating CIF and FOB responses. For Route Risk over 1–6, the CIF contrast is -0.835 and FOB is -1.133, yielding 0.298. Widening therefore reflects a more adverse exporter-side contrast, not a positive importer-side contrast.

The configuration persists over 6–8 and 1–12; over the full window, CIF is -0.519 and FOB -0.766, yielding 0.247.

For Maritime Vulnerability over months 1–6, the CIF and FOB contrasts are -0.368 and -0.576, yielding a positive wedge contrast of 0.208. The component responses become considerably less precise at longer horizons.

Both SVAR Oil Chokepoint Exposure component contrasts are positive, with an imprecise difference; Route Risk provides the clearest component-level evidence.

These are relative contrasts, not absolute declines in either group's valuations. The identity explains wedge adjustment but does not isolate freight, insurance, handling, rerouting, timing, reporting, or composition.

Economic interpretation. Congestion, rerouting, slow steaming, fleet reallocation, insurance, handling, and delays can all alter border-to-border valuations; the data cannot separate their contributions.

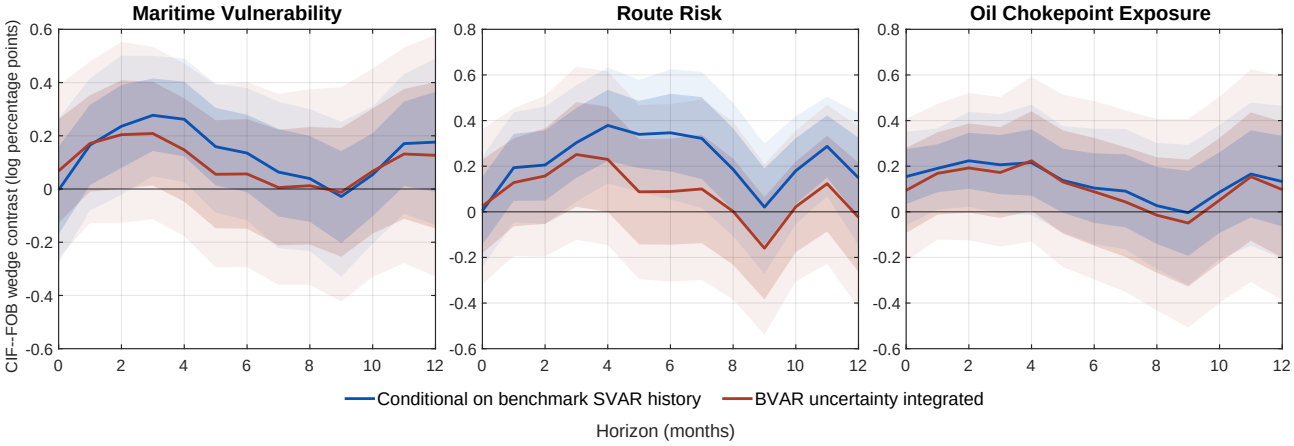
Static exposure contrasts are positive under the structural innovation; larger unlagged observed-index estimates attenuate after addressing persistence.

The conditional dynamic evidence is selective: Route Risk widening is reproduced by AR innovations and reflects relatively more adverse FOB than CIF adjustment.

These exercises support geographically differentiated valuation adjustment, not a pure freight-cost effect, causal mediation of export losses, or exclusion of reporting, timing, and composition alternatives.

Figure 21 integrates shock uncertainty. The Route Risk months-1–6 wedge contrast falls from

Figure 8: CIF–FOB valuation-wedge responses and BVAR shock uncertainty



Notes: High-minus-low maritime-exposure contrasts in the CIF–FOB wedge following a one-standard-deviation SVAR Oil-GPR innovation, evaluated at the 75th percentile of exporter chokepoint dependence. Blue lines condition on the benchmark shock history; brick-red lines integrate uncertainty over 1,000 posterior BVAR histories. Dark and light bands denote 68% and 90% credible intervals. The wedge is computed as the CIF-minus-FOB contrast in each draw. The sample covers January 2001–December 2024, with 278,330 observations across 1,266 pairs and 37 countries, excluding COVID-overlapping response windows. Positive values indicate greater relative widening along highly exposed corridors, not necessarily higher freight costs.

0.29 to 0.16; its 90% interval widens from $[0.10, 0.48]$ to $[-0.15, 0.46]$, and its positive-sign probability is approximately 80%. Some short-horizon pointwise 68% intervals exclude zero. Maritime Vulnerability and Oil Chokepoint Exposure are likewise positive but imprecise short-run patterns. The mechanism evidence is therefore suggestive and less precise than the export response after first-stage uncertainty is incorporated.

The conditional and uncertainty-integrated country-date fixed-effect exercises are documented in Appendices G.4 and G.5. Appendix F.1 reports the network-permutation placebo.

5.6 Controlling for country-specific time-varying shocks

Country-specific supply, demand, exchange-rate, or policy shocks may confound exposure contrasts. We re-estimate dynamic LPs with exporter- and importer-date effects to absorb arbitrary monthly conditions common to an exporter's destinations and an importer's suppliers.

For each horizon $h = 0, \ldots, 12$, we estimate

$$100\,(\log X_{ij,t+h} - \log X_{ij,t-1}) = \beta_{1,h}\left(S_t \widetilde{E}_{ij}\right) + \beta_{3,h}\left(S_t \widetilde{E}_{ij} \widetilde{D}_i\right) + \alpha_{ij,h} + \lambda^X_{it,h} + \mu^M_{jt,h} + \varepsilon_{ij,t+h}. \tag{12}$$

Here S_t is the SVAR innovation and exposures retain their centered benchmark definitions. Pair effects are $\alpha_{ij,h}$; exporter- and importer-date effects $\lambda^X_{it,h}, \mu^M_{jt,h}$ absorb country conditions, common date effects, $S_t \widetilde{D}_i$, the importer-oil interaction, and all country controls, whose coefficients

are consequently unidentified.

Remaining bilateral differences within the exporter-date/importer-date structure identify the exposure response. This restricts country-level confounding but cannot eliminate every bilateral time-varying disturbance.

At exporter dependence D_q, the exposure gradient is

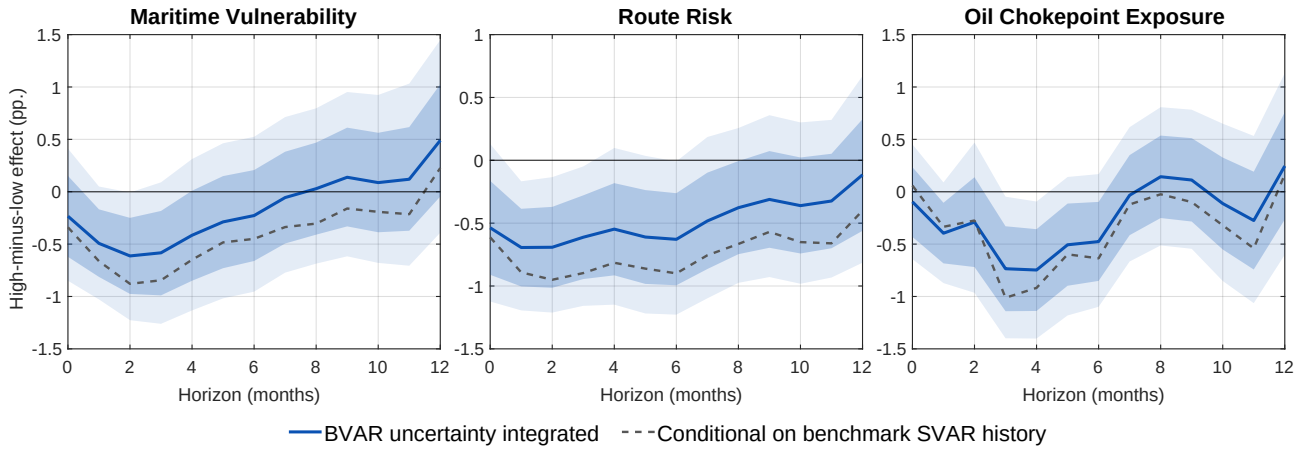
$$\frac{\partial^2 ME_h^{rel}(E,D)}{\partial S_t \, \partial E_{ij}} = \beta_{1,h} + \beta_{3,h}(D_q - D_{50}),$$

and the corresponding high-minus-low contrast is

$$\mathcal{C}_h(D_q) = (E_H - E_L)\left[\beta_{1,h} + \beta_{3,h}(D_q - D_{50})\right].$$

The benchmark upper/lower-quartile exposure means are unchanged; Figure 9 evaluates the contrast at D_{75}.

Figure 9: Dynamic maritime-exposure contrasts with exporter-by-date and importer-by-date fixed effects



Notes: The figure reports posterior distributions of the dynamic high-minus-low maritime-exposure contrast, $\mathcal{C}_h(D_{75})$, following a one-standard-deviation SVAR Oil-GPR innovation. The contrast is evaluated at the 75th percentile of exporter chokepoint dependence. The dependent variable at horizon h is $100[\log X_{ij,t+h} - \log X_{ij,t-1}]$. Estimates use directed-pair, exporter-by-exact-date, and importer-by-exact-date fixed effects; the latter two absorb the common shock, its lower-order interaction with exporter dependence, its interaction with importer oil dependence, and all country-specific macroeconomic controls. Local projections are corrected using the split-panel jackknife and smoothed using the hierarchical Bayesian procedure. Inference integrates both panel-estimation uncertainty and variation across 200 posterior BVAR shock histories. The estimation sample contains 299,160 observations, 1,271 directed pairs, 37 countries, and 253 monthly dates from January 2001 to December 2024. Response windows overlapping March 2020–December 2021 are excluded. Uncertainty bands are 68% and 90% posterior credible intervals. A negative value indicates a less favorable relative export response for highly exposed corridors; it is not an estimate of the absolute response of those corridors.

This check focuses on the SVAR innovation, with conditional and integrated inference. Figure 9

shows an integrated months-1–6 Route Risk contrast of -0.631 (90% interval $[-1.155, -0.075]$; negative-sign probability 96.9%), versus -0.886 conditionally ($[-1.321, -0.412]$). The short-run differential thus survives country-date effects and first-stage uncertainty together.

Integrated Route Risk means remain negative over 6–8 (-0.496) and 1–12 (-0.480), with negative-sign probabilities 90.8% and 91.9%. Their 68%, but not 90%, intervals exclude zero, making the first six months the most precisely supported window.

Integrated short-run Maritime Vulnerability is -0.437 (88.4% negative) and Oil Chokepoint Exposure -0.525 (94.0% negative; 90% interval $[-1.068, 0.035]$). Neither has a precise medium-run contrast under first-stage uncertainty.

Shock recovery contributes approximately 21% of short-run Route Risk variance, versus 43% for Maritime Vulnerability and 38% for Oil Chokepoint Exposure. Route Risk therefore provides the strongest forward-response robustness to both country-specific monthly conditions and generated-shock uncertainty, unlike a backward-looking static check.

5.7 Maritime disruption episodes and alternative-route capacity

Event studies test whether corridors with plausible alternatives adjust differently after identifiable passage disruptions, distinguishing exposure from geographical substitutability.

Bai et al. (2026) distinguish potential and effective capacity and emphasize fleet reallocation, utilization, and congestion beyond exogenous iceberg costs. We instead test ex ante corridor geography, without observing vessel movements or operational spare capacity.

Episodes and alternative-route measures. We report all four episodes irrespective of significance: Hormuz tanker attacks in May–June 2019, the March 2021 Suez blockage, Panama restrictions from August 2023, and the Red Sea crisis from December 2023. Each has predetermined passage exposure.

For episode e, E^e_{ij} is affected-passage exposure and A^e_{ij} alternative-route capacity: Cape plausibility for Suez/Red Sea and Suez-or-Cape access for Panama. These describe feasible geography, not realized rerouting, vessel availability, or operational spare capacity.

Hormuz uses one minus exporter Hormuz dependence as a diversification proxy, not physical pipeline, port, terminal, or bypass capacity.

Event-study specification. We estimate the following specification separately for each episode:

$$
\begin{aligned}
100\Delta_{12}\log X^{real}_{ijt} = \sum_{\substack{k=-K_e^- \\ k\neq -1}}^{K_e^+} \mathbf{1}\{\tau_{et}=k\}\Big[\beta^e_k \widetilde{E}^e_{ij} + \eta^e_k \widetilde{A}^e_{ij} + \theta^e_k \widetilde{E}^e_{ij}\widetilde{A}^e_{ij} \\
+ \lambda^e_k \widetilde{D}_i + \rho^e_k \widetilde{OID}_j\Big] + \mathbf{X}'_{ij,t-1}\gamma + \alpha_{ij} + \delta_t + \varepsilon^e_{ijt}.
\end{aligned} \tag{13}
$$

Here, τ_{et} denotes months relative to episode onset and $k = -1$ is the omitted reference month. The outcome is the twelve-month log change in WTPI-deflated bilateral exports, multiplied by 100. Exposure and alternative-route capacity are centered at their pair-level medians.

Pair effects absorb persistent corridor attributes and date effects absorb common movements, including event-time main effects. Event interactions with exporter and importer dependence and lagged country controls address further heterogeneity. Standard errors are clustered by directed pair.

Panama includes event-time Red Sea-exposure interactions because its post-window overlaps that crisis. Hormuz ends in February 2020 (month nine) to avoid Covid; other long windows end at month twelve.

The outcome traces backward-looking annual-growth differences relative to the omitted pre-event month, not cumulative forward level responses. Date effects cannot remove concurrent shocks correlated with exposure or capacity, particularly during broader disruptions.

Exposure contrasts and the capacity differential. At event month k, the high-minus-low exposure contrast evaluated at alternative capacity A is

$$\mathcal{C}_k^e(A) = (E_H^e - E_L^e)\left[\beta_k^e + \theta_k^e\,(A - A_{50}^e)\right], \tag{14}$$

where E_H^e and E_L^e are the mean exposure levels in the upper and lower tails of the event-specific distribution. Exporter dependence and importer oil dependence are held fixed in this comparison.

Our principal statistic compares the exposure contrast at the 75th and 25th percentiles of alternative-route capacity:

$$\begin{aligned}\Delta\mathcal{C}_k^e &= \mathcal{C}_k^e(A_{75}^e) - \mathcal{C}_k^e(A_{25}^e)\\ &= (E_H^e - E_L^e)\,\theta_k^e\,(A_{75}^e - A_{25}^e)\,.\end{aligned} \tag{15}$$

Positive values denote a more favorable high-exposure response with greater alternative capacity, not absolute export growth or observed shipment rerouting.

This difference of exposure contrasts is not a direct high/low-capacity comparison. Its scale depends on episode-specific exposure and capacity gaps, so it is not a common structural elasticity.

Episode-specific results. Table 5 shows no uniform capacity advantage. Suez and Red Sea intervals include zero in every window, establishing imprecision rather than irrelevance of alternatives.

Hormuz estimates are large and uncertain. Significant pre-event dynamics in the high-capacity profile undermine causality despite non-rejection for the capacity differential: these test different

Table 5: Maritime disruption episodes and alternative-route capacity

Event	Months 0–2	Months 0–6	Long window	Pre-trend p-value
Hormuz tanker attacks	86.30 [−11.13, 183.73]	83.39 [−5.77, 172.55]	89.44 [−9.17, 188.06]	0.510
Suez Canal blockage	−21.94 [−56.37, 12.48]	−23.64 [−55.56, 8.28]	−14.76 [−47.43, 17.92]	0.787
Panama Canal restrictions	**29.31** [**4.32**, **54.30**]	**20.59** [**−3.86**, **45.04**]	**15.68** [**−10.01**, **41.36**]	**0.513**
Red Sea crisis	−18.73 [−49.61, 12.15]	−8.59 [−36.98, 19.81]	1.13 [−27.35, 29.60]	0.546
Long-window endpoint			Month 12; month 9 for Hormuz	

Notes: Average differences in the high-minus-low exposure contrast between the 75th and 25th percentiles of predetermined alternative-route capacity. Positive values indicate more favorable relative export adjustment with greater capacity. Brackets report 90% confidence intervals. The outcome is twelve-month growth in WTPI-deflated exports, expressed in log-percentage points. Samples cover March 2020–March 2022 (Suez), January 2022–December 2024 (Red Sea), February 2022–August 2024 (Panama), and May 2018–February 2020 (Hormuz). Specifications include pair and exact-date fixed effects, event-time interactions with exporter and importer dependence, and lagged controls; standard errors are clustered by pair. Panama additionally controls for event-time interactions with Red Sea exposure. Pre-trend p-values test joint equality to zero of pre-event capacity differentials. Long windows end at month 12, except Hormuz (month 9). Hormuz estimates are exploratory because of pre-event dynamics in the high-capacity exposure profile (p=0.003). Only the Panama short-window differential has a 90% interval excluding zero.

hypotheses. The exercise remains exploratory.

Panama's months-0–2 differential is 29.31 log-percentage points (90% interval $[4.32, 54.30]$), with pre-trend $p = 0.513$. It is the clearest short-run relative capacity advantage.

This large, imprecise annual log-growth differential is not a 29.31% export-level increase. It falls to 20.59 over 0–6 and 15.68 over 0–12, with both intervals including zero: persistence is not precisely established.

One significant window among multiple episode-window comparisons and a non-rejecting pre-trend test cannot establish parallel counterfactual trends or exclude correlated shocks. The Panama finding is suggestive, not definitive identification of rerouting gains.

Identification is also uneven: Suez/Red Sea exposure and Cape alternatives have limited common support; Hormuz variation is concentrated in a smaller network, especially Iran. Panama offers broader, more stable variation, while the other exercises are underpowered or weakly identified checks.

Implications for maritime adjustment. The evidence is consistent with, but does not test, the operational mechanisms of Bai et al. (2026): congestion, rerouting, and common-fleet reallocation can transmit disruptions beyond affected routes. We distinguish predetermined route availability from those realized adjustments.

Bilateral trade cannot separate rerouting, fleet reallocation, congestion, delays, or composition from other costs. The estimates are reduced-form export-adjustment differences, not individual mechanism estimates.

Panama suggests short-run benefits from alternatives, not systematic protection across episodes. Vessel availability, congestion, and voyage duration can prevent feasible geography from becoming usable capacity.

Maritime resilience is episode specific: exposure locates potential incidence, whereas substitution possibilities yield an advantage only when the disruption and wider network constraints permit adjustment.

6 Conclusion

Oil-specific geopolitical disturbances have persistent trade consequences and a distinct maritime incidence. Our aggregate SVAR yields a transitory Oil-GPR response, a persistent trade contraction, a negative medium-run trade–production differential, and modest trade-price adjustment. Bilateral projections locate the clearest adverse short-run differential in Route Risk among exporters dependent on strategic passages.

This differential survives additional trade dynamics and, under exporter- and importer-date fixed effects, integrated structural-shock uncertainty. Conditional estimates indicate stronger adverse transmission when supply chains are strained. Other exposure measures yield less uniform responses.[13]

CIF–FOB decompositions suggest more adverse FOB than CIF adjustment on high-risk corridors, although the wedge becomes imprecise under integrated shock uncertainty. Disruption episodes likewise show that geographical alternatives can facilitate adjustment without guaranteeing usable capacity.[14]

The findings motivate models combining energy markets, trade, shipping networks, and constrained adjustment, without decomposing the aggregate differential into maritime causes or evaluating particular resilience policies. More granular operational data, further state-dependent uncertainty analysis, and structural modelling could identify the relevant margins and quantify policy gains.

References

Anderson, J.E., van Wincoop, E., 2003. Gravity with gravitas: A solution to the border puzzle. American Economic Review 93, 170–192.

Bai, X., Fernández-Villaverde, J., Li, Y., Marto, R., Zanetti, F., 2026. Shipping to America. Working Paper 35808. National Bureau of Economic Research.

Baumeister, C., Hamilton, J.D., 2019a. Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. American Economic Review 109, 1873–1910.

[13] Average bilateral exports weaken most clearly around months six to eight, but integrating shock uncertainty reduces precision. Longer-window exposure contrasts are also less precise, and the delayed positive Oil Chokepoint Exposure contrast is sensitive to bilateral dynamics.

[14] The wedge can reflect freight, insurance, reporting, timing, valuation, and composition. Panama provides the clearest short-run capacity differential; the other events are imprecise or face identifying concerns. Neither exercise establishes realized rerouting or a unique causal mechanism.

Baumeister, C., Hamilton, J.D., 2019b. Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. American Economic Review 109, 1873–1910.

Baumeister, C., Korobilis, D., Lee, T.K., 2022. Energy markets and global economic conditions. Review of Economics and Statistics 104, 828–844.

Brignone, D., Gambetti, L., Ricci, M., 2026. Geopolitical risk shocks: When size matters. European Economic Review 188, 105392.

Caldara, D., Iacoviello, M., 2022. Measuring geopolitical risk. American Economic Review 112, 1194–1225.

Chen, S.S., Hsu, C.C., 2013. Oil price volatility, economic growth and trade. Energy Economics 36, 70–78.

Chow, G.C., Lin, A.L., 1971. Best linear unbiased interpolation, distribution, and extrapolation of time series by related series. The Review of Economics and Statistics 53, 372–375. doi:10.2307/1928739.

Feyrer, J., 2011. Distance, trade, and income: The 1967 to 1975 closing of the suez canal as a natural experiment. American Economic Journal: Applied Economics 3, 189–213.

Hamilton, J.D., 1983. Oil and the macroeconomy since world war ii. Journal of Political Economy 91, 228–248.

Head, K., Mayer, T., 2014. Gravity equations: Workhorse, toolkit, and cookbook. Handbook of International Economics 4, 131–195.

Hummels, D., Lugovskyy, V., 2006. Are matched partner trade statistics a usable measure of transportation costs? Review of International Economics 14, 69–86. doi:10.1111/j.1467-9396.2006.00561.x.

Hummels, D.L., Schaur, G., 2013. Time as a trade barrier. American Economic Review 103, 2935–2959. doi:10.1257/aer.103.7.2935.

Iacoviello, M., Tong, J., 2026. The AI-GPR Index: Measuring Geopolitical Risk using Artificial Intelligence. Working Paper. Federal Reserve Board of Governors.

Känzig, D.R., Raghavan, R., 2026. Supply Chain Shocks and the Macroeconomy: Evidence from Global Shipping Disruptions. Working Paper. Northwestern University.

Kilian, L., 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. American Economic Review 99, 1053–1069.

Kilian, L., Murphy, D.P., 2014. The role of inventories and speculative trading in the global market for crude oil. Journal of Applied Econometrics 29, 454–478.

Kilian, L., Rebucci, A., Spatafora, N., 2009. Oil shocks and external balances. Journal of International Economics 77, 181–194. doi:10.1016/j.jinteco.2009.01.001.

Känzig, D.R., 2021. The macroeconomic effects of oil supply news: Evidence from opec announcements. American Economic Review 111, 1092–1125.

Li, D., Plagborg-Møller, M., Wolf, C.K., 2024. Local projections vs. vars: Lessons from thousands of dgps. Journal of Econometrics 244, 105722. doi:10.1016/j.jeconom.2024.105722.

Pinchetti, M., 2026. Geopolitical risk and inflation: The role of energy markets. Journal of International Economics 164, 104336.

Pratson, L.F., 2023. Assessing impacts to maritime shipping from marine chokepoint closures. Communications in Transportation Research 3, 100083. doi:10.1016/j.commtr.2022.100083.

UNCTAD, 2020. Review of maritime transport 2020. United Nations Conference on Trade and Development.

U.S. Energy Information Administration, 2017. World maritime chokepoints. EIA Background Report.

U.S. Energy Information Administration, 2026. World oil transit chokepoints. URL: https://www.eia.gov/international/content/analysis/special_topics/World_Oil_Transit_Chokepoints/. updated March 3, 2026.

Verschuur, J., Koks, E.E., Hall, J.W., 2022. Ports' criticality in international trade and global supply-chains. Nature Communications 13, 4351. doi:10.1038/s41467-022-32070-0.

Verschuur, J., Lumma, J., Hall, J.W., 2025. Systemic impacts of disruptions at maritime chokepoints. Nature Communications 16, 10421. doi:10.1038/s41467-025-65403-w.

Zanetti, F., Verduzco-Bustos, G., 2026. The Effects of Geopolitical Oil Price Shocks. Technical Report.

A Construction and validation of world trade indicators

This appendix describes the construction of the monthly World Trade Volume Index and World Trade Price Index used in the aggregate SVAR. The main text summarizes the economic role of these variables; the appendix provides the detailed construction, validation, and coverage procedures.

A.1 World Trade Volume Index

The World Trade Volume Index is based on the annual world merchandise-trade volume index published by the World Trade Organization. Let Q_t^{WTO} denote the annual WTO benchmark. Because the empirical analysis is conducted at monthly frequency, we disaggregate the annual benchmark using the Chow–Lin procedure (Chow and Lin, 1971).

The Chow–Lin procedure preserves the annual WTO benchmark exactly while allowing intra-annual variation to be determined by monthly indicators. Candidate indicator sets combine the OECD+6 world industrial-production measure of Baumeister and Hamilton (2019b), the real price of oil, U.S. real imports, world trade shares, and an aggregate export indicator constructed from IMF Direction of Trade Statistics data. For every candidate specification, the monthly series is constrained to aggregate to the annual WTO benchmark.

Let z_m denote a vector of monthly indicators. The temporally disaggregated monthly trade series is written generically as

$$WTRADE_m = CL\left(Q_t^{WTO}, z_m\right), \tag{16}$$

where $CL(\cdot)$ denotes the Chow–Lin temporal-disaggregation operator.

The resulting monthly series is normalized to have an average of 100 in 2015. The benchmark SVAR uses the logarithm of the seasonally adjusted series after quadratic detrending.

A.2 Validation against the CPB World Trade Monitor

The CPB World Trade Monitor is not an input into temporal disaggregation. It is used to validate and select the monthly WTO-based construction over the overlapping sample, and is also used directly as an alternative trade measure in the separate robustness exercise in Appendix C.4. Consequently, CPB provides an alternative data construction, but is not an untouched out-of-sample validation series. The candidate Chow–Lin specifications are compared with the CPB series over their common period using correlations in levels and growth rates, root mean squared errors, volatility ratios, directional accuracy, and turning-point accuracy.

The preferred specification is selected using a composite validation score:

$$\begin{aligned} Score = {} & 0.20R^2 + 0.20Corr^{level} + 0.25Corr^{growth} \\ & + 0.15S^{RMSE,level} + 0.10S^{RMSE,growth} + 0.05S^{vol} + 0.05S^{turn}. \end{aligned} \tag{17}$$

The validation exercise confirms that the selected monthly series reproduces the principal movements of the CPB World Trade Monitor over their common sample while preserving the longer annual WTO benchmark. The CPB-based augmented SVAR is discussed in Appendix C.4.

A.3 Alternative monthly trade measure

As a robustness check, we construct an alternative world trade measure directly from monthly IMF trade data. Let $X_{i,m}$ denote nominal exports reported by country i in month m. The world aggregate is

$$X_m^W = \sum_{i=1}^{N_m} X_{i,m}, \tag{18}$$

where N_m denotes the number of reporting countries.

Because country coverage changes over time, we adjust aggregate exports for the number of reporting countries:

$$\log X_m^W = \alpha + \beta \log N_m + \epsilon_m. \tag{19}$$

The adjusted monthly trade measure is constructed from the residual component:

$$\widetilde{X}_m^W = \exp\left(\epsilon_m + \overline{\log X_m^W}\right). \tag{20}$$

This direct monthly measure is not used as the benchmark because it lacks the annual volume-benchmarking property of the WTO-based index. Coverage adjustment alone does not convert nominal exports into a volume index. Any SVAR comparison using this series must document its deflator, effective sample, transformations, and normalization before being interpreted as a trade-volume robustness exercise.

A.4 World Trade Price Index

Country-level export and import price indices are obtained from the International Monetary Fund (IMF) *International Financial Statistics* (IFS). Let $P_{i,t}^X$ and $P_{i,t}^M$ denote the export and import price indices of country i in month t, and let $X_{i,t}$ and $M_{i,t}$ denote its nominal exports and imports. Because national price indices may use different base years, the aggregation is performed on monthly log changes rather than on index levels.

Define $\mathcal{I}_t^X$ as the set of countries for which $\Delta \log P_{i,t}^X$ is available and usable in month t. The global export-price growth rate is the trade-weighted average of the national export-price changes:

$$g_t^X = \sum_{i \in \mathcal{I}_t^X} w_{i,t}^X \Delta \log P_{i,t}^X, \qquad w_{i,t}^X = \frac{X_{i,t}}{\sum_{j \in \mathcal{I}_t^X} X_{j,t}}. \tag{21}$$

The weights therefore sum to one among the countries reporting a usable export-price change in that month.

Similarly, let $\mathcal{I}_t^M$ denote the set of countries with an available and usable import-price change. The global import-price growth rate is

$$g_t^M = \sum_{i \in \mathcal{I}_t^M} w_{i,t}^M \Delta \log P_{i,t}^M, \qquad w_{i,t}^M = \frac{M_{i,t}}{\sum_{j \in \mathcal{I}_t^M} M_{j,t}}. \tag{22}$$

Because g_t^X and g_t^M are weighted log changes, the corresponding global export and import price indices are chain-linked using the exponential transformation:

$$P_t^X = P_{t-1}^X \exp(g_t^X), \qquad P_t^M = P_{t-1}^M \exp(g_t^M). \tag{23}$$

This formulation avoids the approximation that would arise from applying $1+g_t$ to a log growth rate.

Before aggregation, the export and import components are placed on a common scale. The World Trade Price Index is then defined as

$$WTPI_t = \alpha_t P_t^X + (1 - \alpha_t) P_t^M, \qquad \alpha_t = \frac{X_t^W}{X_t^W + M_t^W}, \tag{24}$$

where X_t^W and M_t^W are total world nominal exports and imports. The time-varying coefficient α_t gives the export and import components weights proportional to their respective shares of total recorded world trade. The resulting index is normalized to average 100 in 2015.

We impose two data-quality restrictions to prevent the index from being driven by sparse or highly concentrated reporting. Export- and import-price coverage are defined as

$$Coverage_t^X = \frac{\sum_{i \in \mathcal{I}_t^X} X_{i,t}}{X_t^W}, \qquad Coverage_t^M = \frac{\sum_{i \in \mathcal{I}_t^M} M_{i,t}}{M_t^W}. \tag{25}$$

A monthly observation is retained only when both coverage ratios exceed 30%. In addition, we require

$$\max_{i \in \mathcal{I}_t^X} w_{i,t}^X \leq 0.70, \qquad \max_{i \in \mathcal{I}_t^M} w_{i,t}^M \leq 0.70. \tag{26}$$

Thus, a month is excluded when a single country represents more than 70% of the available export or import trade weight. These coverage and concentration criteria limit changes in the global index arising solely from variation in the composition of reporting countries.

The construction consequently combines country-level price information in a common, benchmark-consistent monthly index while avoiding aggregation across incomparable national index levels. It should nevertheless be interpreted as a trade-weighted measure of price movements among reporting countries, subject to the stated coverage restrictions, rather than as a direct price index for every transaction in world trade.

B Benchmark Bayesian SVAR: data, prior, and identification

This appendix provides the technical details underlying the benchmark monthly Bayesian SVAR reported in Section 3. The main text describes the economic motivation, identification logic, and impulse responses. This appendix documents the sample alignment, variable transformations, recursive identification, Bayesian prior, and posterior simulation procedure.

B.1 Data construction and sample alignment

The benchmark input calendar spans January 1984–December 2024 at monthly frequency, as in the main text. The effective regression sample additionally reflects the lag requirements described below. Let

$$\mathcal{T} = \{t_0, t_0 + 1, \ldots, T\}$$

denote the common sequence of monthly dates. Each series is aligned to this calendar and padded with missing values when it is not available. Estimation uses observations with the required consecutive monthly history; missing dates are not removed in a way that makes nonconsecutive months adjacent. The first usable regression date additionally reflects the twelve lagged observations required by the VAR.

Oil-specific geopolitical risk is measured by the Oil GPR index of Iacoviello and Tong (2026). The Oil GPR index is constructed through a two-stage large-language-model classification of international newspaper articles. The first stage identifies articles containing geopolitical risk information. The second stage selects articles referring to actual or prospective disruptions to oil or energy supply, including conflict, sanctions, threats to energy infrastructure, and risks affecting energy transport. The monthly Oil GPR series is standardized over the benchmark estimation sample.

World crude-oil production, WOP_t, is obtained from the U.S. Energy Information Administration's *Monthly Energy Review*. The real price of oil, $RRAC_t$, is the refiners' acquisition cost of imported crude oil deflated by the U.S. consumer price index for all urban consumers. World oil inventories, $WSTOCKS_t$, are constructed following Kilian and Murphy (2014). Total OECD petroleum inventories are multiplied by the ratio of U.S. crude-oil inventories to total U.S. petroleum inventories, which produces a proxy for the worldwide crude-oil inventory stock.

The benchmark World Trade Volume Index, $WTRADE_t$, is a new long-run monthly measure constructed from annual WTO merchandise-trade volume benchmarks and monthly indicators. The World Trade Price Index, $WTPI_t$, aggregates country-level export and import price indices from the IMF *International Financial Statistics* using time-varying nominal trade weights. The construction, validation, country coverage, and alternative trade indicators are described in Appendix A.

World industrial production, WIP_t, is not included in the six-variable benchmark SVAR. It is introduced in the augmented specification used to estimate the trade–production differential in Section 3.4. The activity index is the OECD+6 industrial-production measure of Baumeister and Hamilton (2019b).

B.2 Seasonal adjustment and transformations

World trade volume and world oil inventories display pronounced seasonal variation and are seasonally adjusted using X-13 before estimation. Let X_t denote a strictly positive series. The logarithmic transformation is

$$x_t = \log X_t.$$

World oil production, seasonally adjusted world trade volume, the world trade price index, and seasonally adjusted world inventories are detrended using a quadratic deterministic trend:

$$x_t = a_0 + a_1 t + a_2 t^2 + \widetilde{x}_t, \tag{27}$$

where $\widetilde{x}_t$ denotes the cyclical component. The real oil price enters in logarithmic levels.

The benchmark SVAR uses

$$GPR_t^{oil}, \quad \log\widetilde{WOP}_t, \quad \log\widetilde{WTRADE}_t, \quad \log\widetilde{WTPI}_t, \quad \log RRAC_t, \quad \log\widetilde{WSTOCKS}_t.$$

The use of cyclical log components removes deterministic low-frequency movements while preserving medium-run fluctuations relevant for oil-market and trade dynamics. When impulse responses of log or log-cycle variables are multiplied by 100, they can be interpreted approximately as percentage deviations from their pre-shock paths.

B.3 Benchmark SVAR

The benchmark monthly SVAR is

$$y_t = A_0 + \sum_{\ell=1}^{12} A_\ell y_{t-\ell} + u_t, \qquad u_t \sim \mathcal{N}(0, \Sigma_u), \tag{28}$$

where

$$y_t = \begin{bmatrix} GPR_t^{oil} \\ \log\widetilde{WOP}_t \\ \log\widetilde{WTRADE}_t \\ \log\widetilde{WTPI}_t \\ \log RRAC_t \\ \log\widetilde{WSTOCKS}_t \end{bmatrix}. \tag{29}$$

Here A_0 is a vector of intercepts. The benchmark has twelve monthly lags, allowing for delayed adjustment over an annual cycle. Alternative lag lengths and transformations are examined in Appendix C.

B.4 Recursive identification

Structural innovations satisfy

$$u_t = B\varepsilon_t, \qquad \mathbb{E}(\varepsilon_t\varepsilon_t') = I, \qquad \Sigma_u = BB', \tag{30}$$

where B is lower triangular. Oil GPR is ordered first. Thus $u_{1t} = B_{11}\varepsilon_{1t}$: the first structural innovation is *proportional*, rather than generally numerically identical, to the reduced-form residual in the Oil GPR equation. Identification rules out within-month feedback from the other variables to that innovation, while allowing it to affect all remaining variables contemporaneously. Identifying this first shock does not require choosing a contemporaneous ordering among the remaining variables. The restriction is not a general guarantee against anticipation or omitted common shocks.

For each retained posterior draw, impulse responses are normalized so that the impact response of the standardized Oil GPR index equals one:

$$IRF_{GPR^{oil}}(0) = 1. \tag{31}$$

The normalization fixes the impact increase in the standardized index at one, or one standard deviation of its observed benchmark-sample level. This is not necessarily a one-standard-deviation structural disturbance under a unit-variance shock normalization. Responses are traced over horizons from zero to 36 months.

B.5 Bayesian estimation and posterior simulation

The benchmark system is estimated using a Bayesian block-Minnesota prior implemented through dummy observations. The prior shrinks the system toward parsimonious univariate dynamics but allows stronger interactions within economically related groups of variables than across groups.

The overall tightness parameter is

$$\lambda = 0.10.$$

Relative tightness is

$$\phi_{\text{within}} = 0.50, \qquad \phi_{\text{between}} = 0.25.$$

The lag-decay parameter is

$$d = 1.50.$$

The prior standard deviation associated with the intercept is 10. All own-lag prior means are set to zero rather than to one. This is a prior specification; it does not assert that the log real oil price has been detrended or establish its stationarity.

Posterior inference is based on 45,000 Gibbs iterations. The first 5,000 draws are discarded as burn-in. Every twentieth subsequent draw is retained, yielding 2,000 posterior draws. The recursive identification and the shock normalization are applied separately to every retained

draw. The reported median impulse response is the median across these draws. The 68% and 90% credible intervals correspond to the 16th–84th and 5th–95th posterior percentiles, respectively.

B.6 Augmented trade–production specification

The six-variable benchmark contains oil-market variables, Oil GPR, and world trade. To estimate the differential in trade–production response, we supplement the set of benchmark information with the global industrial-production index of Baumeister and Hamilton (2019b). World industrial production is placed after world oil production and before world trade. The augmented system retains the same monthly sample, lag length, prior structure, recursive identification of the Oil GPR shock, and draw-by-draw shock normalization.

For the posterior draw b and the horizon h, the trade–production response differential is

$$\mathcal{D}_h^{(b)} = IRF_{WTRADE,h}^{(b)} - IRF_{WIP,h}^{(b)}. \tag{32}$$

The posterior distribution of $\mathcal{D}_h$ is calculated directly from the draw-specific difference, preserving the posterior covariance between the trade and production responses. A negative differential indicates that world trade contracts more than world industrial production following the oil-specific geopolitical shock.

C Robustness of the aggregate SVAR responses

The robustness exercises are presented in the following order: alternative transformations, lag-length sensitivity, pre-pandemic estimation, measurement of the trade–production differential, alternative propagation methods, projections using observed Oil GPR, and the event-based measure. These exercises address different aspects of the aggregate results. In particular, a contraction in trade in a six-variable system does not, by itself, establish a trade–production differential.

C.1 Alternative transformations

We examine whether the responses to an oil-specific geopolitical risk shock depend on the transformation of the variables or on the dynamic specification of the Bayesian SVAR. In addition to the cyclical benchmark specification described in Appendix B.2, we consider models estimated using logarithmic levels and monthly log differences.

Figure 10 compares the impulse responses obtained from two alternative specifications, one using variables in log levels and the other using growth rates. Both specifications yield a rapid decline in Oil GPR following the initial innovation, a persistent fall in world trade, and broadly similar responses of world oil production and inventories. The decline in world trade is somewhat more pronounced at medium horizons in the levels specification. The responses

of the world trade price also differ in magnitude, although their credible intervals are wide. Overall, the comparison suggests that the main qualitative features of the responses are not driven by the choice between these two transformations, while the magnitude and persistence of some responses remain sensitive to the specification.

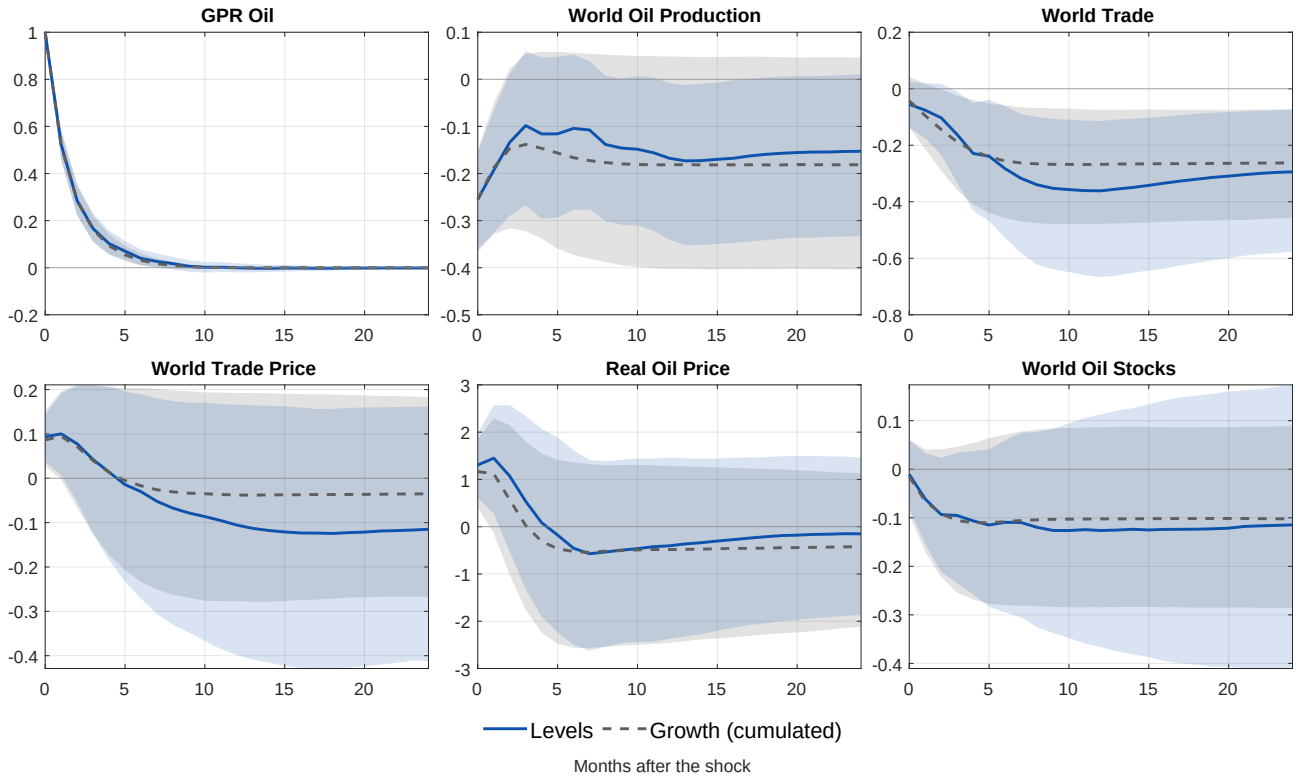


Figure 10: Sensitivity of impulse responses to data transformations

Notes: The figure reports posterior median impulse responses to an Oil-GPR innovation under two alternative specifications. The solid blue lines correspond to the specification in log levels, and the dashed grey lines to the growth-rate specification. Shaded areas indicate 90% posterior credible intervals. In the growth-rate specification, responses are cumulated at each posterior draw to obtain responses in levels. Oil GPR is standardized and the real oil price is included in log levels in both specifications. The impact response of Oil GPR is normalized to one in each posterior draw. Both models include 12 lags, use the same recursive identification and are estimated over the same effective regression sample, selected jointly after differencing and lag requirements. The common input calendar spans January 1984–December 2024. Own-lag prior means follow the specification: random-walk means for the log-level variables and zero means for the growth-rate specification. The horizontal axis measures months after the shock.

C.2 Sensitivity to the number of lags

Figure 11 compares the benchmark impulse responses with those obtained using 3, 6, and 9 lags. All specifications are estimated over a common effective sample and retain the benchmark data transformations, identification, and prior settings.

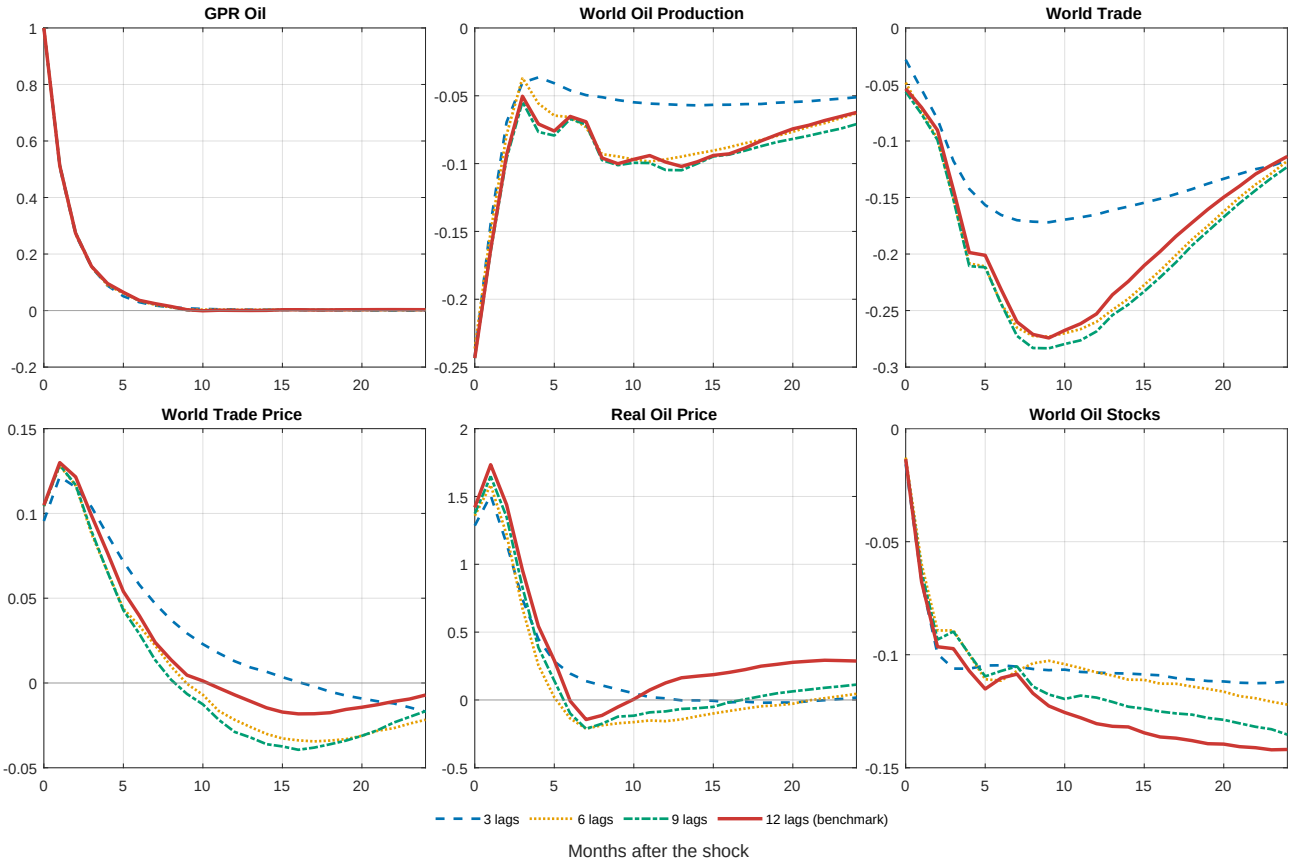


Figure 11: Sensitivity of impulse responses to the number of lags

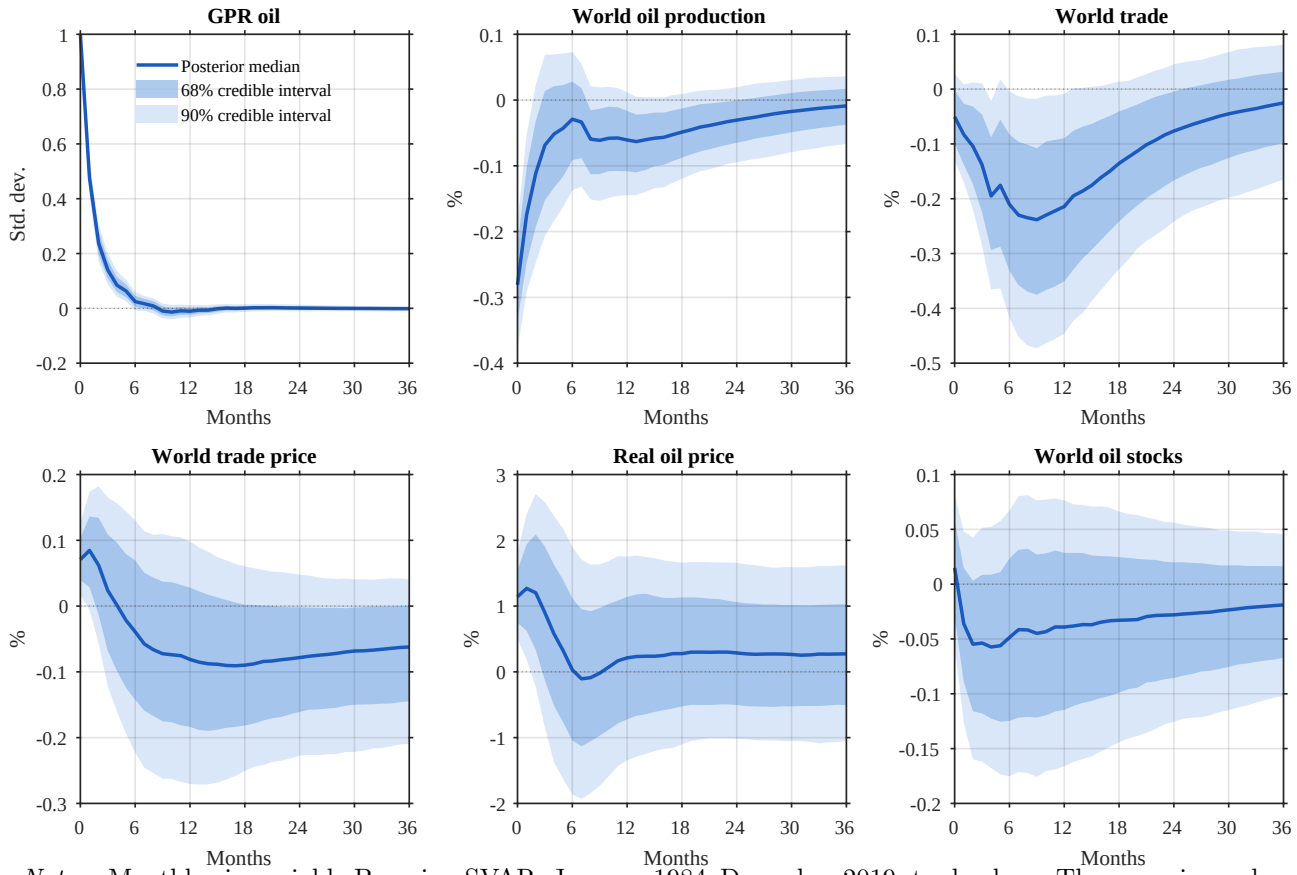
Notes: The figure reports posterior median impulse responses to an Oil-GPR innovation for specifications with 3, 6, 9, and 12 lags. The 12-lag specification corresponds to the benchmark. All models use the same effective regression sample after the common twelve-lag presample allowance. The input calendar spans January 1984–December 2024. All models use the same data transformations as the benchmark. They also use the same recursive identification and prior settings. The impact response of Oil GPR is normalized to one in each posterior draw. The horizontal axis measures months after the shock. Credible intervals are omitted for clarity.

C.3 Pre-pandemic sample

The six-variable benchmark is re-estimated using the January 1984–December 2019 input calendar, with twelve presample observations required to construct the regression lags. Seasonal adjustment, quadratic detrending, and Oil GPR standardization are recomputed using data through December 2019. The twelve lags, recursive ordering, prior settings, and impact normalization are otherwise retained. The exercise changes both the estimation period and the inputs used for transformation; it is not a formal test of equality between the two periods. Because Oil GPR is standardized within each sample, a unit impact need not represent the same increase in the raw index across samples.

Figure 12 documents the qualitative pre-pandemic response profile described in the main text. The median trade contraction persists, while the shorter sample yields less precise responses. Inventory movements are interpreted cautiously and do not, on their own, identify precautionary demand or a physical supply disruption.

Figure 12: Aggregate responses in the pre-pandemic sample



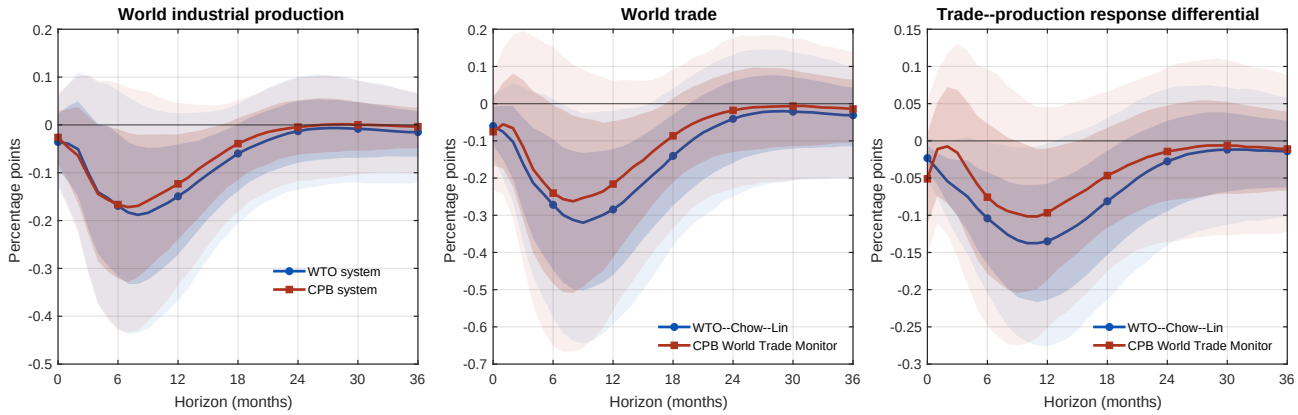
Notes: Monthly six-variable Bayesian SVAR, January 1984–December 2019; twelve lags. The recursive order is Oil GPR, world oil production, world trade volume, the World Trade Price Index, the real oil price, and oil inventories. Transformations and index standardization are recomputed within the pre-pandemic sample. At each draw, the impact response of standardized Oil GPR is set to one. Lines report posterior medians; dark and light bands report central 68% and 90% credible intervals. Log-variable responses are multiplied by 100. The displayed horizon is 0–36 months.

C.4 CPB measurement robustness for the trade–production differential

Because global industrial production is among the candidate indicators used in the temporal disaggregation of the annual WTO series, we assess whether the qualitative differential persists with a separately produced trade series. We replace our monthly WTO-based quantity index with the seasonally adjusted world-trade volume index from the CPB World Trade Monitor and re-estimate both systems on the same effective regression sample using the common January 2000–December 2024 input calendar and identical lag requirements. All remaining data, transformations, priors, lag length, recursive ordering, and shock normalization are held fixed.

Figure 13 shows that the direction and timing of the result do not depend on the benchmark monthly trade construction. With the CPB index, the trade response reaches a trough of approximately −0.25 percentage point around months 7–9, compared with approximately −0.17 percentage point for industrial production. The resulting differential reaches approximately −0.10 percentage point before gradually returning toward zero. The CPB-based contraction and

Figure 13: Trade–production responses using alternative world-trade measures



Notes: The figure compares posterior responses from two augmented BVARs estimated on a common effective regression sample using the January 2000–December 2024 input calendar and identical lag requirements. Blue lines use the benchmark world-trade quantity index constructed from annual WTO volume benchmarks and monthly indicators. Brick-red lines use the seasonally adjusted CPB World Trade Monitor index. Both systems contain Oil GPR, world oil production, world industrial production, world trade, the World Trade Price Index, the real price of oil, and oil inventories, in that recursive order. Each includes 12 monthly lags and uses the transformations described in the text. Responses are normalized to a one-standard-deviation increase in Oil GPR on impact and are expressed in percentage points. The trade–production response differential is calculated at each posterior draw. Dark and light shaded regions denote 68% and 90% posterior credible intervals.

differential are smaller than their WTO-based counterparts on the same sample, and the CPB 90% credible interval generally includes zero, although the posterior median remains negative over most of the relevant horizon and the 68% interval provides some medium-run evidence. Thus, the qualitative ranking persists with a separately produced monthly trade series, while its magnitude and posterior precision remain sensitive to measurement. Because CPB also informs selection of the WTO-based construction, this exercise is an alternative-measure robustness check rather than wholly independent validation of the selection procedure.

C.5 Local-projection comparisons under recursive identification

Let s_t denote the fixed historical SVAR Oil-GPR innovation used in the matched LP exercise. For each variable j and horizon $h = 0, \ldots, 36$,

$$y_{j,t+h} = a_{j,h} + \theta_{j,h} s_t + \sum_{\ell=1}^{12} \Gamma_{j,h,\ell} y_{t-\ell} + e_{j,t+h}. \tag{33}$$

The vector y_t has the same six transformed variables as the benchmark. An observation is usable only if the shock, lagged controls, and outcome at the specified calendar horizon are observed. For local projections estimated separately at each horizon, the effective sample may vary with the horizon because of endpoint losses. The hierarchical Bayesian local projections instead use a common sample across the horizons included in the estimation, from month 0 through month 60. The shock scaling follows the historical series used in the corresponding

regression.

The hierarchical Bayesian local projections use the saved posterior median historical BVAR innovation, rescaled using the benchmark impact-normalization factor. Each regression includes a constant and twelve lags of all six benchmark variables. Lag coefficients receive Minnesota-type shrinkage, with horizon-specific tightness selected by marginal likelihood over a prespecified grid. The current-shock coefficient receives a diffuse prior.

The response paths are estimated over horizons 0–60 months on a common sample, while only horizons 0–36 are reported. The additional horizons provide a buffer against terminal-boundary effects. For each response variable, a hierarchical prior shrinks second differences of the response path toward zero:

$$D_2\boldsymbol{\theta}_j \mid \tau_j^2 \sim \mathcal{N}(0, \tau_j^2 I), \qquad \tau_j \sim \text{half-Cauchy}(0, 1), \tag{34}$$

where the prior is applied after response-specific rescaling. The level and linear slope of the response path remain unpenalized. The prior is not centered on the BVAR response.

Cross-horizon dependence is approximated using a Bartlett score-HAC correlation matrix with bandwidth 60, rescaled by the horizon-specific BLP posterior standard deviations. This covariance matrix enters the second-stage Gaussian smoothing model. Posterior simulation retains 5,000 response paths per variable after 10,000 burn-in iterations, keeping every fourth draw from 30,000 iterations. Reported credible intervals are pointwise. Inference conditions on the estimated historical shock series and the selected first-stage shrinkage parameters.

Figure 14 compares the Bayesian SVAR responses with those obtained from alternative local-projection estimators. The hierarchical Bayesian local projections use the estimated historical BVAR innovation as a fixed regressor. The frequentist local projections instead identify the innovation recursively from the system data, with Oil GPR ordered first. Thus, the comparison does not use the same estimated shock history for every method and does not independently test the recursive identifying restriction.

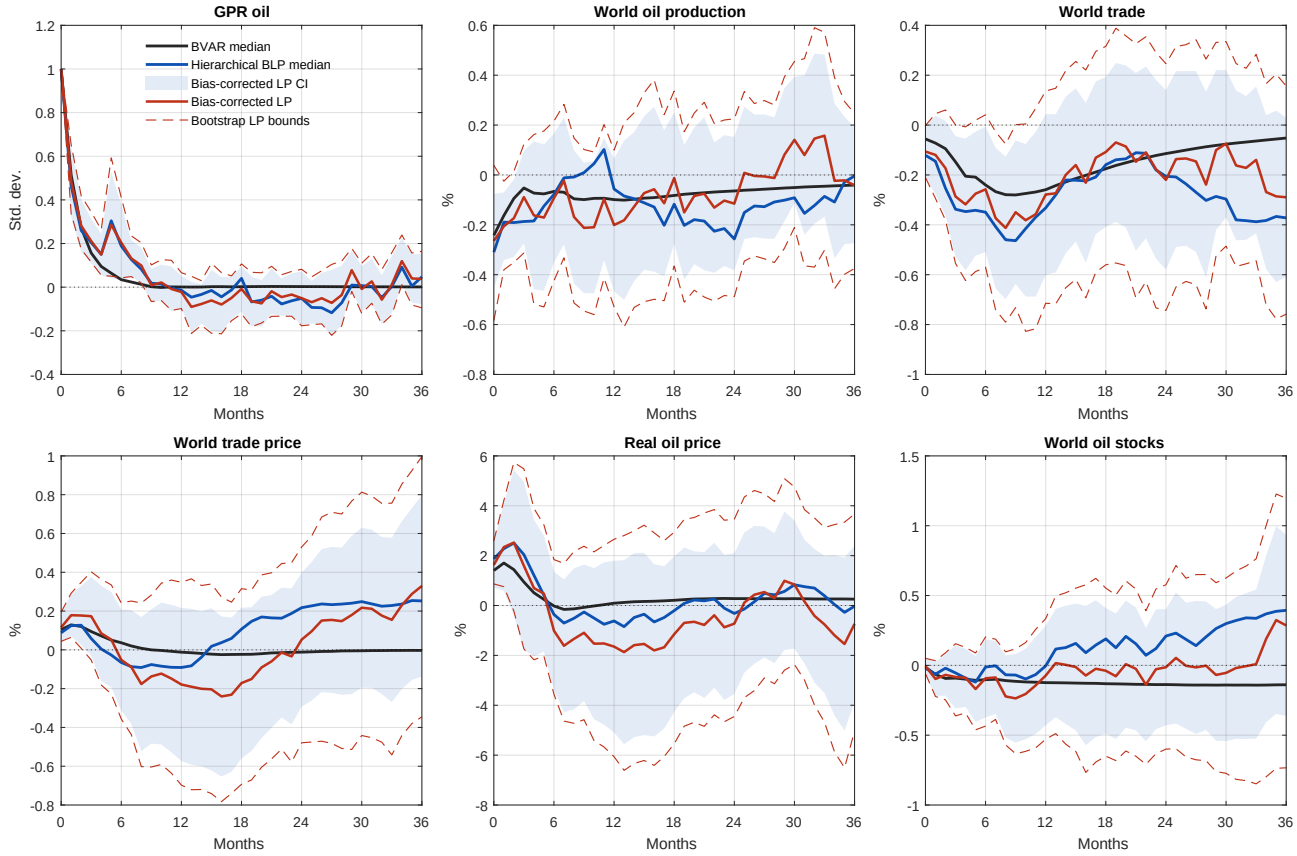
The hierarchical Bayesian intervals condition on the estimated historical shock series and the selected first-stage shrinkage parameters. They therefore do not incorporate uncertainty from estimating the historical shock.

C.6 Local projections using observed Oil GPR

Replace s_t in Equation (33) with the standardized observed Oil GPR index, maintaining the matched response variables and lagged controls. A unit increase is an increase of one standard deviation of the observed index, not a one-standard-deviation identified structural shock. Observed Oil GPR combines predictable and unexpected variation, so these coefficients describe conditional associations. They need not equal the responses to an SVAR innovation, even when their signs are similar.

Figure 15 reports this comparison. Its purpose is to assess whether the negative trade profile is confined to the particular innovation series, without treating the observed index as an

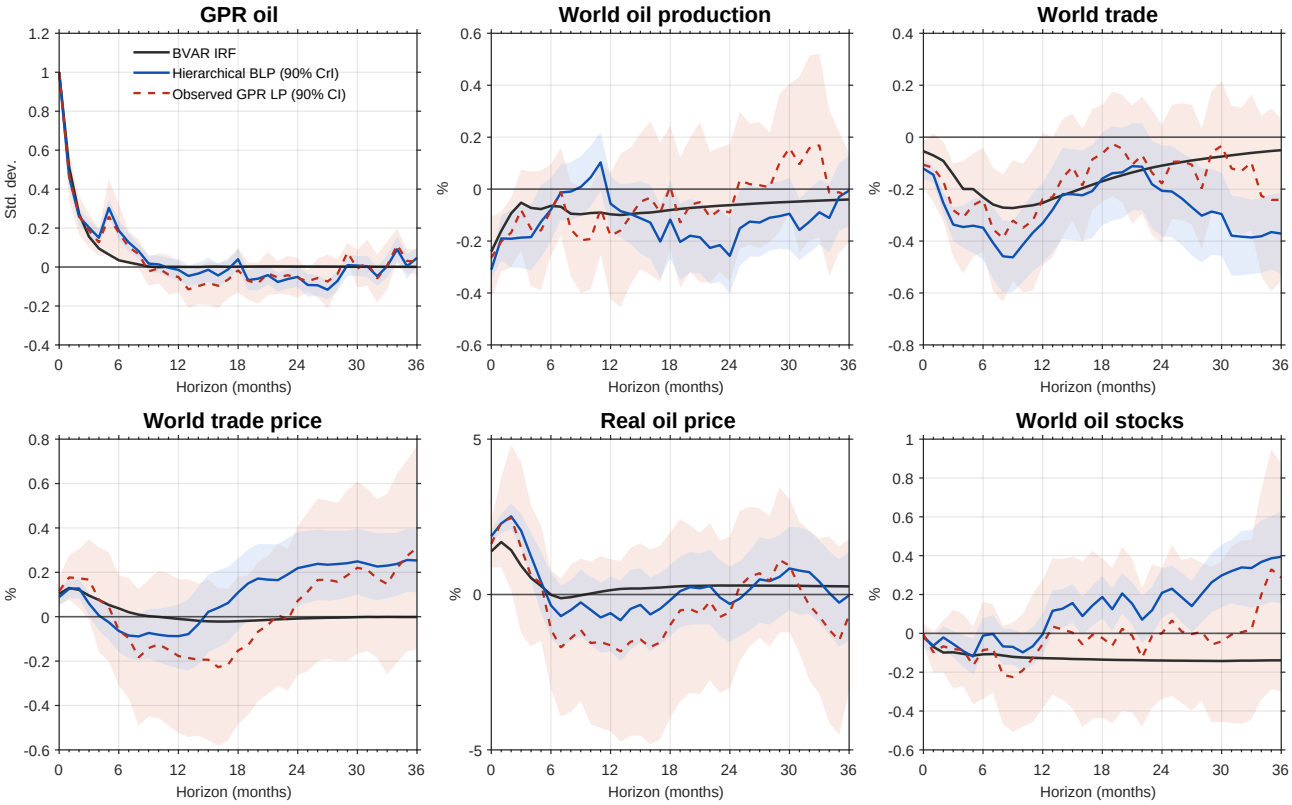
Figure 14: Alternative propagation estimates under recursive identification



Notes: The black line reports the benchmark Bayesian SVAR posterior median, the blue line reports the hierarchical Bayesian local-projection posterior median, and the red line reports the bias-corrected frequentist local-projection estimate. The shaded region displays the 95% confidence interval returned by the bias-corrected LP routine. Dashed red lines display 95% bootstrap confidence bounds based on 1,000 replications using VAR residual block resampling, with block length $\lceil 5.03T^{1/4} \rceil$, where T is the number of observations supplied to the routine. The frequentist LP uses twelve lags and identifies the innovation recursively with Oil GPR ordered first. It is estimated directly from the system data rather than using the saved BVAR historical shock series. The hierarchical BLP instead uses the fixed estimated historical BVAR innovation and applies shrinkage to lagged controls and hierarchical smoothing across horizons. Its uncertainty conditions on that estimated shock history. Only median curves are displayed for the BVAR and hierarchical BLP; the shaded confidence region and dashed bounds refer to the frequentist LP. Responses are reported over horizons 0–36 months. Responses of log and log-cycle variables are multiplied by 100.

alternative structurally identified shock.

Figure 15: Structural-shock responses and projections on observed Oil GPR



Notes: The black line reports the benchmark Bayesian SVAR posterior median. The blue line and blue shaded region report the hierarchical Bayesian local-projection posterior median and pointwise 90% credible interval, conditional on the saved historical BVAR innovation. The hierarchical response path is estimated through month 60 on a common sample and displayed through month 36. The dashed red line and corresponding shaded region report OLS local projections using the contemporaneous standardized observed Oil GPR index, with pointwise 90% normal-approximation confidence intervals. These regressions include a constant and twelve lags of all six system variables. Standard errors use a Bartlett HAC covariance estimator with bandwidth equal to the projection horizon and a degrees-of-freedom correction. The direct-LP estimation sample varies with the horizon. The direct-LP regressor is standardized using its sample mean and standard deviation. The BVAR and hierarchical BLP responses are rescaled to the same observed-index unit. This common scale does not make the observed-index coefficient a structural response. Observed-index projections describe conditional associations; the hierarchical BLP conditions on the estimated structural shock history. Log-variable responses are multiplied by 100.

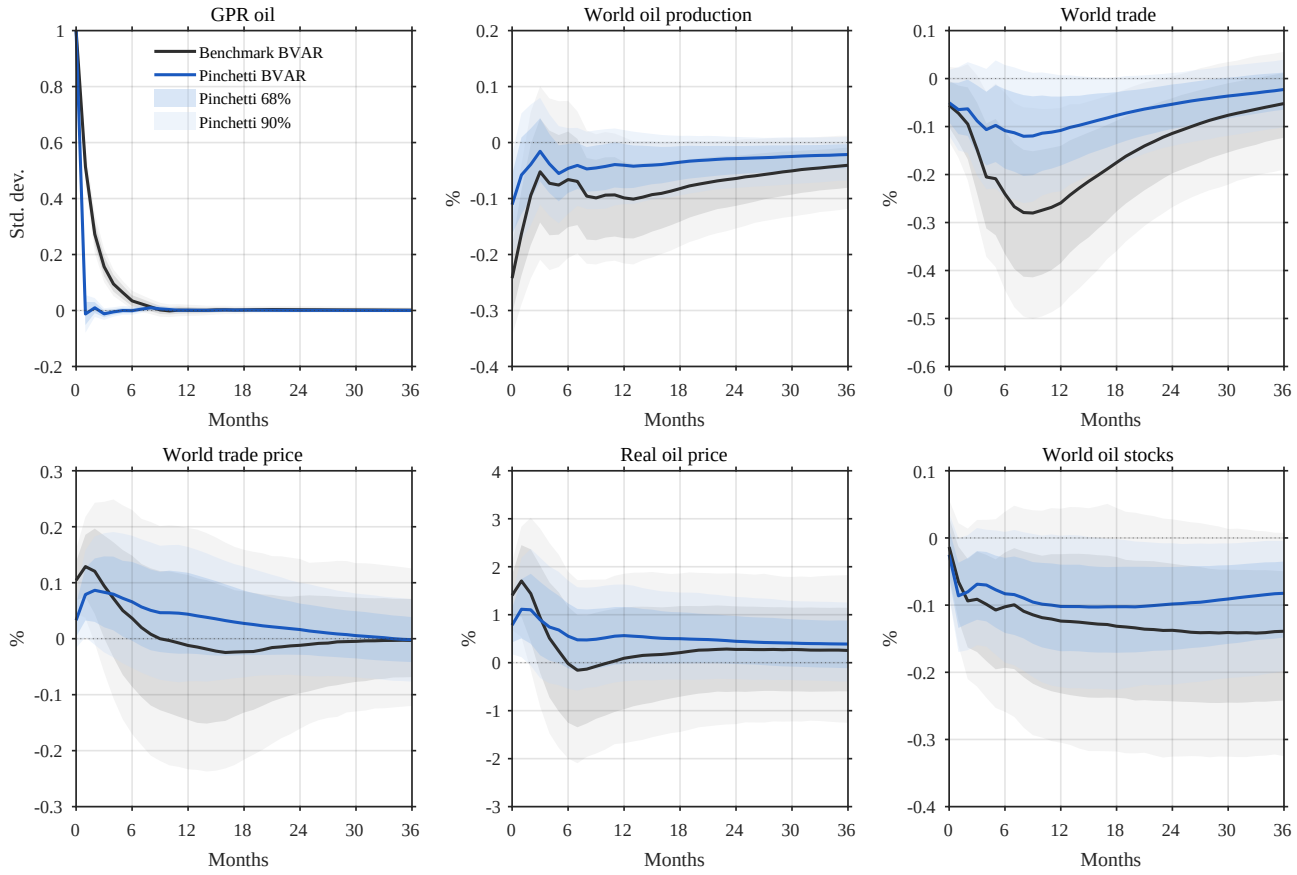
C.7 Event-based measure in a Bayesian VARX

The alternative measure follows the event-based logic of Pinchetti (2026). Sharp geopolitical-risk events are classified using their short-window WTI-price comovement and aggregated to the monthly frequency. The resulting series enters the Bayesian system as an exogenous regressor. This construction is distinct from both observed monthly Oil GPR and the recursively identified benchmark innovation.

An event-based variable entering a VARX is not automatically an external instrument. A proxy-SVAR would instead use its covariance with reduced-form innovations under explicit relevance and exclusion conditions. Accordingly, the present exercise changes the risk measure and its model treatment; it does not establish equivalence of the three shock concepts.

Figure 16 documents the event-based response profile. The comparison concerns the qualitative contraction in world trade. Unless both trade and industrial production are included and their response difference is evaluated jointly, it does not independently re-identify the trade–production differential.

Figure 16: Benchmark and event-based oil-geopolitical responses



Notes: The black line reports the benchmark Bayesian SVAR posterior median. The blue line reports the posterior median response from the alternative Bayesian VARX, in which the event-based oil-geopolitical series enters exogenously. The event series is constructed following the classification approach of Pinchetti (2026), using geopolitical-risk events and short-window WTI price movements. Gray and blue shaded regions display the corresponding benchmark and alternative uncertainty bands; darker and lighter shading denotes 68% and 90% posterior credible intervals, respectively. Both responses are scaled to a one-unit impact increase in standardized Oil GPR. This common impact normalization facilitates comparison but does not make the underlying disturbances identical. Responses are reported over 36 months, in standard-deviation units for Oil GPR and in percent for the remaining variables.

Taken together, these exercises support the initial decline in oil production, the contraction in world trade, and the short-run increase in the real price of oil. The local-projection estimates display greater horizon-to-horizon variation, while the medium- and long-run responses of trade prices and inventories are more sensitive to the estimation method. This pattern is consistent

with the bias–variance trade-off documented by Li et al. (2024), who show that local projections can exhibit substantially higher variance than VAR-based estimators at intermediate and long horizons. The event-based specification also produces a persistent trade contraction, although its median response is smaller than in the benchmark. These comparisons document qualitative similarities rather than formal equality across response functions. The structural-shock projections retain the benchmark identification, whereas projections on observed Oil GPR have a conditional, reduced-form interpretation.

C.8 Scope of the robustness evidence

Table 6: Purpose and interpretation of the aggregate robustness exercises

Exercise	What changes	Interpretation
Alternative transformations	Log levels versus growth rates, with specification-specific prior means	Assesses sensitivity to transformations and associated prior conventions.
Alternative lag lengths	3, 6, 9, and 12 lags on a common effective sample	Assesses sensitivity of median response paths to lag length.
Pre-pandemic SVAR	Sample and within-sample transformations	Checks dependence on pandemic-era observations; not a formal stability test.
CPB augmented SVAR	Trade measure on a common calendar	Directly evaluates measurement sensitivity of the jointly computed differential.
Local-projection comparison	Propagation estimator; shock history also differs for frequentist LP	Compares propagation estimates under recursive identification, not an independent identification test.
Observed-index LP	Risk regressor	Conditional association; not a structural shock response.
Event-based Bayesian VARX	Event measure and its model treatment	Alternative trade-response exercise; not proxy-SVAR identification.

Notes: Evidence for trade contraction and evidence for a trade–production differential are distinguished. All interval statements are pointwise unless simultaneous inference is explicitly implemented. Differences in whether separate intervals include zero do not constitute a test of equality between models or response paths.

The exercises support a comparison of propagation methods and risk measures, while the CPB augmented-system exercise specifically addresses trade measurement in the differential. They do not identify the contribution of individual maritime frictions. That question motivates the bilateral analysis.

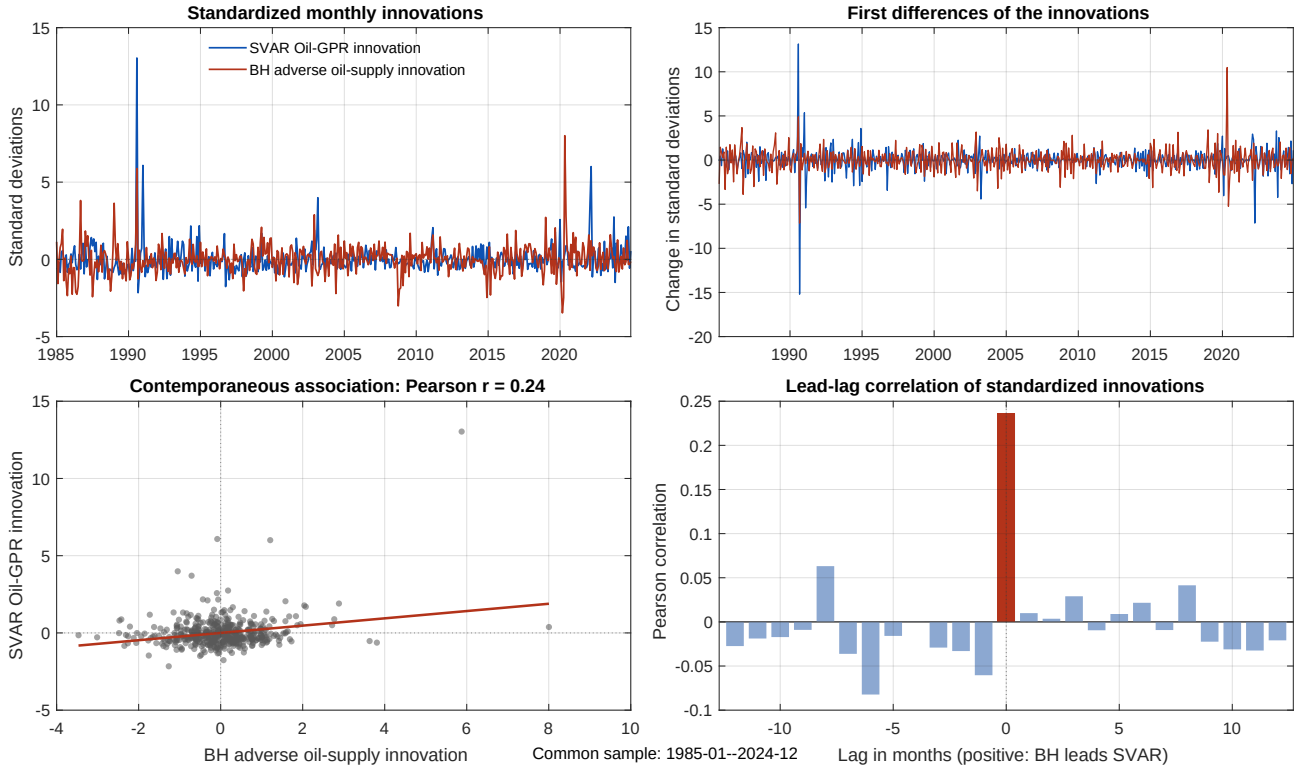
C.9 Comparison with the BH adverse oil-supply innovation

We compare the benchmark Oil-GPR innovation with the adverse oil-supply innovation associated with Baumeister and Hamilton (2019a). Positive BH realizations are oriented to represent

adverse supply innovations. Both series are compared over January 1985–December 2024. The word "levels" below denotes the innovation series themselves, not levels of oil production or geopolitical risk.

For each retained posterior SVAR shock history, we calculate Pearson and Spearman correlations with the supplied BH series, both for the innovations and for their first differences. The reported quantiles summarize variation across posterior SVAR histories. They do not incorporate posterior uncertainty in the BH shock recovery.

Figure 17: Oil-GPR and BH adverse oil-supply innovations



Notes: The upper panels display standardized monthly innovations and their first differences. The lower-left panel shows their contemporaneous association for the displayed SVAR history; its Pearson correlation is approximately 0.24. The lower-right panel reports lead–lag correlations of the standardized innovations, with positive lags indicating that BH leads the SVAR innovation. The common sample is January 1985–December 2024. The lead–lag profile is descriptive and is not a causality or predictive-content test. The contemporaneous correlation for the displayed history differs from the median across posterior histories reported in Table 7.

The association is positive across almost all evaluated histories but modest in magnitude. Median Pearson correlation is 0.228 for the innovations, compared with a median rank correlation of 0.076. The difference suggests that agreement in relative monthly rankings is much weaker than linear comovement, potentially reflecting influential large episodes; it does not establish their influence without additional diagnostics.

First-difference correlations are somewhat larger but remain modest. Differencing an innovation series changes its statistical properties and is used here only as a descriptive sensitivity check. The figure's lead–lag profile has its largest displayed correlation at zero, without establishing

Table 7: Association with the BH adverse oil-supply innovation across posterior SVAR histories

Correlation measure	Median	Central 68% interval	Central 90% interval	Share positive
Pearson: innovations	0.22805	[0.21473, 0.24088]	[0.20537, 0.24883]	1.0000
Spearman: innovations	0.075646	[0.050240, 0.10118]	[0.032834, 0.11703]	0.9992
Pearson: first differences	0.26325	[0.24468, 0.28167]	[0.23291, 0.29368]	1.0000
Spearman: first differences	0.11699	[0.084866, 0.15015]	[0.063822, 0.16905]	0.9996

Notes: Each statistic is computed separately for each posterior SVAR shock history against the same supplied BH series. The central intervals use the 16th–84th and 5th–95th percentiles of these correlations. "Share positive" is the fraction of evaluated posterior histories with positive correlation, not a frequentist significance level. First differences are changes in innovations, not alternative identified structural shocks. The comparison conditions on the BH series and does not jointly propagate uncertainty from both shock-identification procedures.

the absence of predictive relationships or a causal direction.

Overall, the comparison supports limited overlap between oil-geopolitical and adverse oil-supply innovations. It does not imply identical shocks, independence, or a decomposition of Oil-GPR innovations into supply and non-supply components.

D Bilateral trade data and the Maritime Trade Exposure Network

This appendix documents the bilateral sampling frame, the construction of fixed maritime-exposure measures, and their geographical cross-validation. It distinguishes the original 44-country route-construction universe from the 37-country universe and 1,271 stable directed corridors used in the panel analysis. Regression-specific observation counts depend on outcome availability and the requirements of each empirical exercise.

D.1 Country coverage and sample selection

Country coverage is not determined solely by raw data availability. We complement the initial availability requirement with a screening, specified before estimation, of bilateral trade continuity over the January 2001–December 2024 sample. A country pair is classified as stable when positive bilateral trade is observed in at least 90 percent of the 288 sample months, the pair itself is observed for at least 90 percent of those months, all three maritime-exposure measures are available for at least 90 percent of the period, and no inactive spell exceeds twelve consecutive months. We additionally require the resulting country network to satisfy a minimum-degree condition under which each country must retain at least ten stable trading partners.

We first restrict the panel to the final universe of 37 countries. Over January 2001–December 2024, this universe contains 1,332 potential directed country pairs and 383,616 pair–month rows. After applying the continuity, exposure-coverage, and network connectivity criteria, the stable core contains 1,271 directed pairs and 366,048 pair–month rows, corresponding to 95.42%

of the initial pairs and rows. All 37 countries remain in the stable network.

Each retained pair has a row for every month in the 288-month calendar. The presence of a row does not imply that the export value is positive or non-missing: these conditions are assessed separately when applying the continuity criteria.

The composition of the stable core is insensitive to reasonable alternative thresholds. Varying the minimum active-month requirement between 70 and 90 percent and the minimum number of stable partners between 5 and 15 retains the same 1,271 directed pairs and all 37 countries. Moreover, the resulting V9 core is identical to the former V8 stable core after restricting that core to the same final 37-country universe. These diagnostics indicate that the estimation sample forms a dense and stable trading network. Importantly, the 37-country universe is defined before estimating the trade equations; it is not selected on the basis of coefficient estimates or statistical significance.

Table 8: Bilateral Trade-Sample Stability Diagnostics

Diagnostic	Number	Share / value
Sample period	288 months	2001:M1–2024:M12
Initial countries	44	100.0%
Initial directed country pairs	1,892	100.0%
Initial pair-month observations	544,900	100.0%
Pairs active in at least 90% of months	1,743	92.1%
Pairs satisfying all continuity conditions	1,723	91.1%
Pair-month observations in the stable core	496,220	91.1%
Countries retained in the stable core	44	100.0%
Minimum number of stable partners	30	—
Mean number of stable partners	41.3	—
Median number of stable partners	43	—
Maximum number of stable partners	43	—
Pairs with complete maritime-exposure coverage	1,892	100.0%
Pairs retained under alternative thresholds	1,723	91.1%
Countries retained under alternative thresholds	44	100.0%

Notes: The initial universe contains all 1,332 potential directed bilateral relationships among the pre-defined 37 countries. A pair is classified as stable when positive trade is observed in at least 90 percent of the 288 sample months, the pair is observed in at least 90 percent of those months, the three maritime-exposure measures are available for at least 90 percent of the period, and no inactive spell exceeds twelve consecutive months. The resulting network must also form a k-core with at least ten stable partners per country. Partner counts are calculated on the corresponding undirected network. Sensitivity checks vary the active-month requirement between 70 and 90 percent and the minimum-degree requirement between 5 and 15 partners. All nine combinations retain the same 1,271 directed pairs and all 37 countries. The criteria are fixed before estimation and do not depend on estimated coefficients, standard errors, or statistical significance.

D.2 Exact construction of composite exposure indices

This subsection describes the construction of the bilateral chokepoint exposures, the three composite maritime indices, and the exporter-level chokepoint dependence measure used in the panel analysis. The bilateral geographical exposures are fixed across time for each exporter–importer pair. They are model-implied measures based on a maritime network and are not observed shares of vessels or shipments.

D.2.1 *Maritime network and gateway weights*

The benchmark maritime network is represented by an undirected graph whose edges connect maritime nodes and carry baseline network costs. The network used in the route-construction files contains 57 nodes and 99 edges. Its edge costs are routing indices rather than observed voyage distances or transport costs.

Let $\mathcal{P}_i$ and $\mathcal{P}_j$ denote the sets of maritime gateways associated with exporter i and importer j. Each gateway $a \in \mathcal{P}_i$ is assigned a country-specific weight $\omega_{ia} > 0$, normalized so that

$$\sum_{a \in \mathcal{P}_i} \omega_{ia} = 1.$$

The preliminary weight assigned to a gateway pair (a, b) is therefore

$$\omega_{ij}^{ab} = \omega_{ia}\omega_{jb}.$$

The weights are fixed inputs to the route-construction procedure. They do not represent estimated bilateral port-use shares. In the 44-country route-construction sample, the gateway-weight file contains 151 country–gateway entries, with between two and five entries per country.

D.2.2 *Candidate paths and route weights*

For each exporter-gateway–importer-gateway combination, the routing algorithm generates up to $K = 5$ candidate paths. The first path minimizes the baseline network cost. Alternative paths are generated iteratively by multiplying the cost of edges used by previously selected paths by a penalty factor of 1.35. The search stops when no additional path is found or when the next path duplicates a previously selected path.

Let d_{ijr}^{ab} denote the cost returned by the algorithm for candidate route r between gateways a and b. For alternative paths, this cost incorporates the penalties accumulated during the route search. Define the minimum candidate cost for the gateway pair as

$$d_{ij,\min}^{ab} = \min_{s \in \mathcal{R}_{ij}^{ab}} d_{ijs}^{ab}.$$

Candidate paths are retained when their returned cost satisfies

$$d_{ijr}^{ab} \leq 1.5\, d_{ij,\min}^{ab}.$$

The threshold therefore applies to the algorithm's path-cost index, including the penalties used during the search. It should not be interpreted as a bound on observed voyage distance.

Conditional on a gateway pair, the weight assigned to a retained route is

$$\pi_{ijr|ab} = \frac{\exp\left(-\lambda d_{ijr}^{ab}\right)}{\sum_{s \in \mathcal{R}_{ij}^{ab}} \exp\left(-\lambda d_{ijs}^{ab}\right)}, \qquad \lambda = 0.85,$$

where $\mathcal{R}_{ij}^{ab}$ is the set of retained candidate paths for the gateway pair. Shorter paths in terms of the returned cost receive greater weight. These weights are imposed by the routing model and are not estimated from vessel movements.

The preliminary weight of a candidate path is

$$\widetilde{\pi}_{ijr}^{ab} = \omega_{ia}\omega_{jb}\pi_{ijr|ab}.$$

Exact duplicate paths arising from different gateway combinations are collapsed and their preliminary weights are added. The resulting weights are normalized across the retained paths for country pair ij:

$$\pi_{ijr} = \frac{\sum_{a \in \mathcal{P}_i} \sum_{b \in \mathcal{P}_j} \mathbf{1}\{r \in \mathcal{R}_{ij}^{ab}\}\, \widetilde{\pi}_{ijr}^{ab}}{\sum_{s \in \mathcal{R}_{ij}} \sum_{a \in \mathcal{P}_i} \sum_{b \in \mathcal{P}_j} \mathbf{1}\{s \in \mathcal{R}_{ij}^{ab}\}\, \widetilde{\pi}_{ijs}^{ab}}, \qquad \sum_{r \in \mathcal{R}_{ij}} \pi_{ijr} = 1.$$

If two gateways correspond to the same network node, the algorithm records a zero-cost path with no traversed edge. Such a path has no passage flagged through an edge. This convention reflects the aggregation of gateways into network nodes.

The saved V5 route-exposure file covers 44 countries and all 1,892 directed country pairs, with no missing routes. After combining gateway pairs and collapsing duplicate paths, the mean number of retained routes is 33.11 per directed country pair. This number can exceed five because each country pair can involve multiple gateway pairs.

D.2.3 *Bilateral chokepoint exposure*

For each strategic passage c, model-implied bilateral exposure is defined as

$$E_{ij}^{c} = \sum_{r \in \mathcal{R}_{ij}} \pi_{ijr}\, \mathbf{1}\{c \in r\}, \qquad 0 \leq E_{ij}^{c} \leq 1,$$

where $\mathbf{1}\{c \in r\}$ indicates whether an edge of route r is flagged as using passage c. The measure integrates over the retained candidate paths and the country-specific gateway weights. It is a geographical exposure implied by the model, rather than an observed fraction of bilateral trade using the passage.

The set of passages is

$$\mathcal{C} = \{\text{Suez, Hormuz, Bab el-Mandeb, Malacca, Panama,}$$
$$\text{Bosphorus–Dardanelles, Black Sea, Cape of Good Hope, Gibraltar, Dover}\}.$$

D.2.4 *Route Risk Index*

The Route Risk Index is the unweighted sum of bilateral exposures across the ten passages:

$$RR_{ij} = \sum_{c \in \mathcal{C}} E_{ij}^{c}.$$

Each passage receives unit weight. Because a path can traverse several passages, the index can exceed one. Under the model-implied route weights, it is the expected number of the ten flagged passages traversed by a retained path. In the saved 44-country route-construction sample, its maximum is 3.95.

D.2.5 *Oil Chokepoint Exposure Index*

The Oil Chokepoint Exposure Index combines bilateral exposure to Hormuz and Bab el-Mandeb:

$$OC_{ij} = \min\left\{E_{ij}^{Hormuz} + E_{ij}^{Bab\text{-}el\text{-}Mandeb},\, 1\right\}.$$

Both passages receive unit weight and the sum is capped at one, so that

$$0 \leq OC_{ij} \leq 1.$$

This index is a capped sum of two model-implied passage exposures. It is not calculated as the route-level probability of crossing at least one of the two passages. The corresponding variable in the V9 estimation panel is `OilChokepointRoute_Prob_V9`.

D.2.6 *Maritime Vulnerability Index*

The Maritime Vulnerability Index combines the ten bilateral passage exposures using fixed salience weights:

$$MV_{ij}^{raw} = \sum_{c \in \mathcal{C}} \nu_{c}^{MV} E_{ij}^{c}.$$

The weights are reported in Table 9.

Table 9: Chokepoint weights used in the Maritime Vulnerability Index

Strategic passage	Weight ν_c^{MV}
Hormuz	1.00
Malacca	0.90
Suez	0.80
Bab el-Mandeb	0.70
Gibraltar	0.40
Panama	0.30
Bosphorus–Dardanelles	0.25
Cape of Good Hope	0.20
Dover	0.15
Black Sea	0.15

Notes: The table reports the fixed salience weights used to aggregate model-implied bilateral chokepoint exposures. The weights are calibration parameters rather than estimated coefficients. The Route Risk Index instead assigns a unit weight to each of the ten passages.

The raw index is normalized by its maximum in the original 44-country route-construction sample:

$$MV_{ij} = \frac{MV_{ij}^{raw}}{\max_{k,\ell} MV_{k\ell}^{raw}}.$$

The saved maximum of the raw index is 2.6225, and the normalized index has a maximum of one in that construction sample. The V9 panel uses the normalized index inherited from the original network construction.

D.2.7 *Export-weighted chokepoint dependence*

The panel analysis also uses an exporter-level measure of dependence on three passages:

$$\mathcal{C}^E = \{\text{Hormuz}, \text{Bab el-Mandeb}, \text{Suez}\}.$$

For exporter i and passage $c \in \mathcal{C}^E$, dependence is the export-value-weighted average of bilateral exposure:

$$Dep_i^c = \frac{\sum_{j:(i,j)\in\mathcal{S}} \sum_{t\in\mathcal{T}} \mathbf{1}_{ijtc}\, X_{ijt} E_{ij}^c}{\sum_{j:(i,j)\in\mathcal{S}} \sum_{t\in\mathcal{T}} \mathbf{1}_{ijtc}\, X_{ijt}},$$

where X_{ijt} denotes nominal bilateral exports, $\mathcal{S}$ is the stable-core network of 1,271 directed corridors, and $\mathcal{T}$ covers January 2001–December 2024. The indicator $\mathbf{1}_{ijtc}$ equals one when exports are finite and strictly positive and the corresponding passage exposure is available. Other observations do not contribute to the weighted average.

The exporter-level composite is the simple average of the three passage-specific measures:

$$ExporterChokepointDep_i = \frac{1}{3}\left(Dep_i^{Hormuz} + Dep_i^{Bab\text{-}el\text{-}Mandeb} + Dep_i^{Suez}\right).$$

This dependence measure is calculated once using the fixed 2001–2024 weighting window and is then treated as constant over time for each exporter. Because the weighting window overlaps the analysis period, the measure should not be interpreted as being constructed exclusively from information available before every shock in the sample.

D.3 Distribution of route-specific bilateral exposure

The probabilistic routing procedure generates continuous bilateral exposure measures rather than simple indicators for whether a chokepoint lies on a unique shortest route. Table 10 summarizes their pair-level distributions in the original 44-country construction universe. These construction-sample descriptives are distinct from the validation statistics below, which use the final 1,271 stable directed corridors.

Table 10: Route-specific exposure in the original maritime-network construction

Exposure	Mean	Std. dev.	Min	P25	Median	P75	Max	Positive (%)	High (%)
Suez	0.179	0.282	0.000	0.000	0.000	0.292	1.000	46.934	15.116
Hormuz	0.169	0.316	0.000	0.000	0.000	0.237	0.938	25.687	16.490
Bab el-Mandeb	0.176	0.280	0.000	0.000	0.002	0.262	1.000	50.846	12.896
Malacca	0.178	0.256	0.000	0.000	0.005	0.364	1.000	52.114	16.702
Panama	0.234	0.310	0.000	0.000	0.019	0.432	1.000	52.167	21.089
Bosphorus–Dardanelles	0.050	0.175	0.000	0.000	0.000	0.000	0.815	8.985	4.545
Black Sea	0.020	0.069	0.000	0.000	0.000	0.000	0.365	8.985	0.000
Cape of Good Hope	0.177	0.280	0.000	0.000	0.009	0.232	1.000	52.537	17.970
Gibraltar	0.200	0.287	0.000	0.000	0.027	0.369	1.000	53.911	16.121
Dover	0.269	0.374	0.000	0.000	0.000	0.500	1.000	43.869	24.789

Notes: Notes: The table reports unweighted descriptive statistics for the 1,892 directed corridors in the original 44-country maritime-network construction. The panel analysis retains 37 countries and 1,271 stable directed corridors after applying trade continuity, coverage, and connectivity criteria. The corresponding maritime exposures are matched from the original network to the retained corridors. The statistics reported here therefore describe the construction universe rather than the final estimation sample.

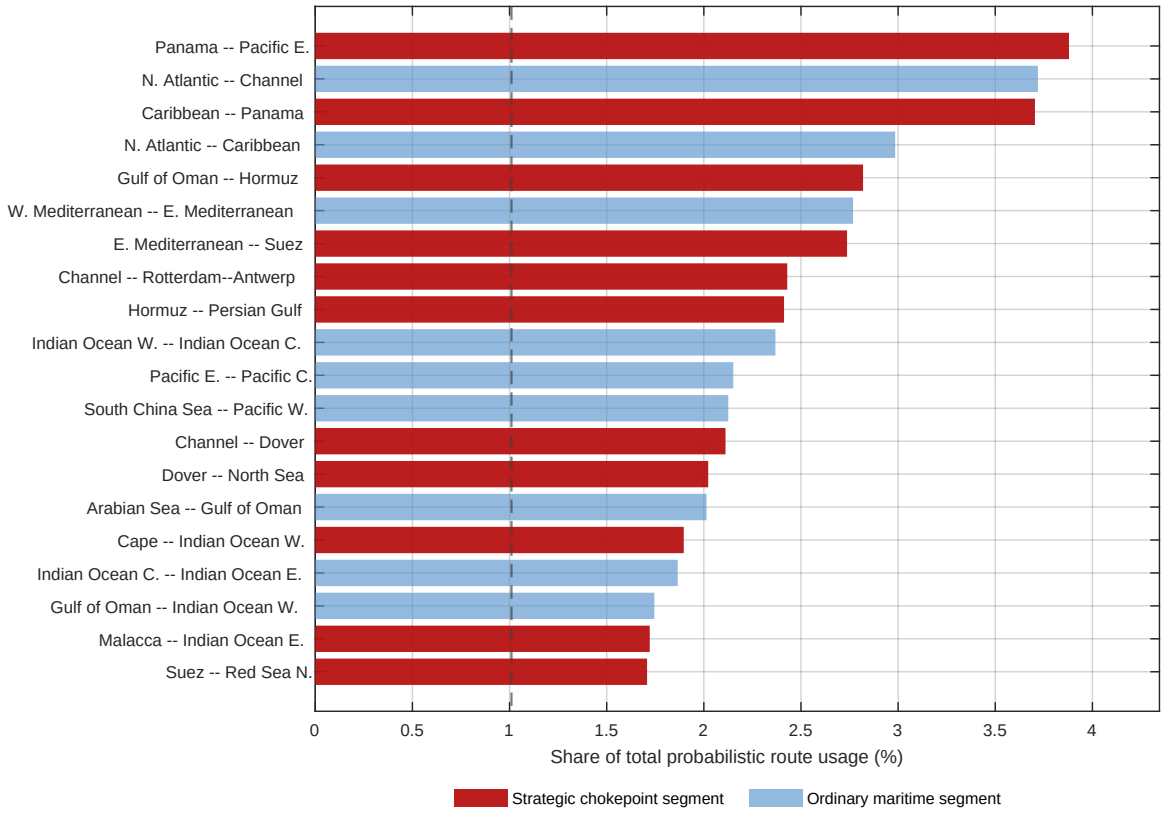
The distributions are strongly right-skewed, generally combining a substantial mass at zero with a smaller group of highly exposed corridors. At the same time, intermediate and near-unit values reveal considerable variation along the intensive margin. These patterns illustrate the cross-corridor variation generated by the probabilistic routing procedure in the original construction universe. They should not be read as descriptive statistics for the final estimation sample.

D.4 Top maritime segments by aggregate route usage

Figure 18 reports the twenty maritime segments with the highest aggregate probabilistic route usage. Strategic chokepoints such as Panama, Hormuz, Suez and Malacca naturally rank among

the most frequently traversed segments, but several ordinary links also appear because they connect major ocean basins, confirming that the routing algorithm distributes traffic across multiple key segments rather than a single shortest path.

Figure 18: Top maritime segments by aggregate route usage



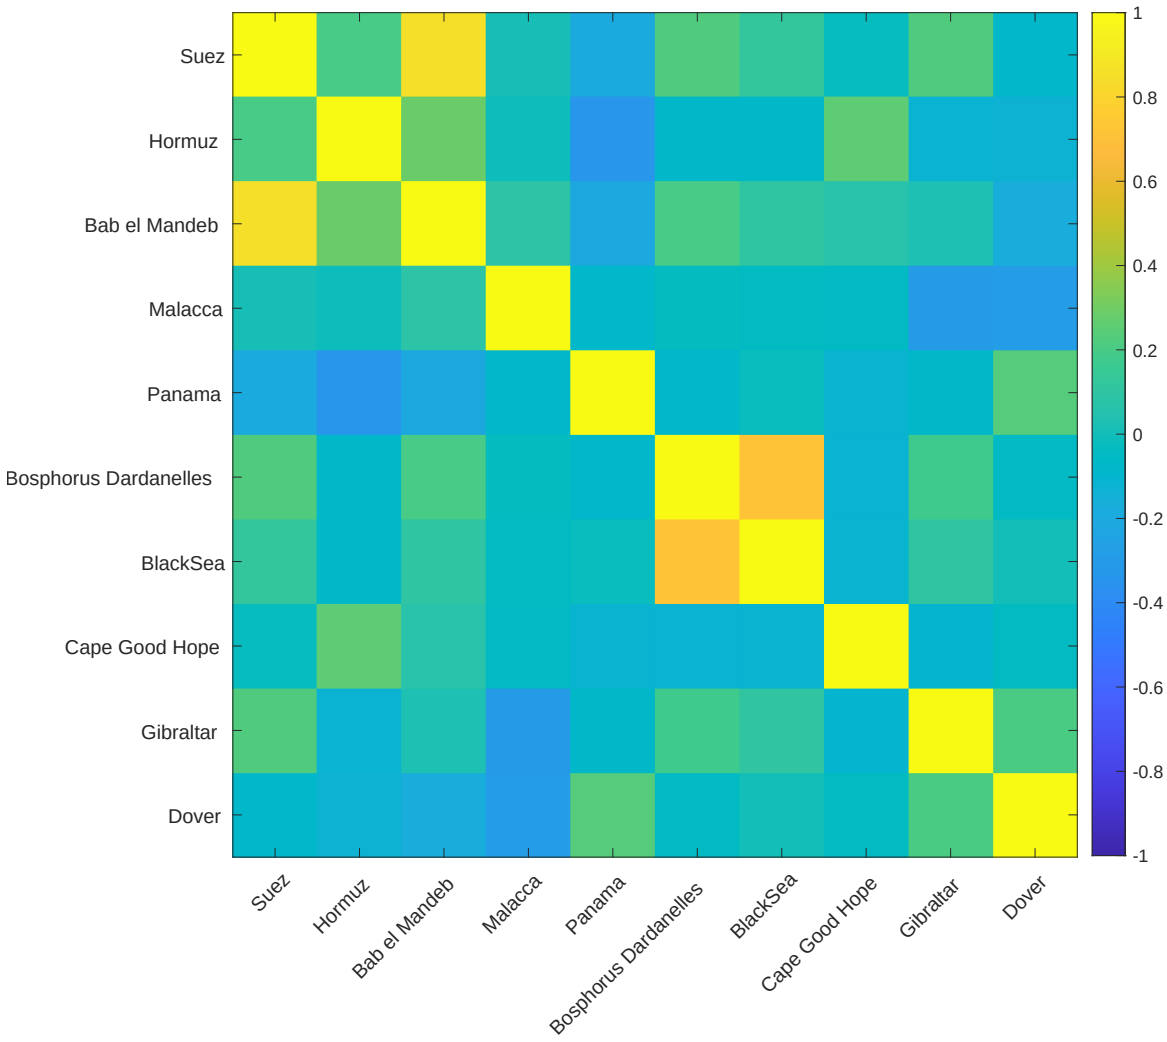
Notes:Notes: The figure reports the twenty maritime segments with the highest aggregate model-implied route usage in the original network covering 44 countries and 1,892 directed corridors. Usage is measured as each segment's share of total edge-level route-probability mass across all constructed corridors. Red bars identify strategic passage segments, while blue bars identify ordinary maritime links. These measures describe usage implied by the routing model, not observed vessel traffic. The panel estimations use the corresponding exposures matched to the 1,271 stable corridors retained among 37 countries.

D.5 Correlation structure of chokepoint exposure indices

Figure 19 reports correlations between bilateral passage-exposure measures across the 1,892 directed corridors of the original 44-country network. The matrix describes how exposures to different passages overlap across corridors. Geographically connected passages tend to exhibit stronger co-exposure, while other passage pairs capture more distinct dimensions of maritime geography. These construction-sample correlations do not directly describe the variation remaining in the final estimation sample after fixed effects are removed.

First, co-exposure is systematically stronger for geographically and functionally linked passages. The Bosphorus–Dardanelles–Black Sea block exhibits high positive correlations, reflecting the fact that any corridor connecting to Black Sea ports must necessarily transit the Bosphorus and the Dardanelles. Similarly, routes that traverse Suez often have non-negligible exposure to

Figure 19: Correlation matrix of bilateral chokepoint exposure indices



Notes: The figure reports pairwise correlations between model-implied bilateral exposures to Suez, Hormuz, Bab el-Mandeb, Malacca, Panama, Bosphorus–Dardanelles, Black Sea access, the Cape of Good Hope, Gibraltar, and Dover. Correlations are calculated across the 1,892 directed corridors in the original 44-country network, with each corridor entering once. Diagonal entries equal one. The panel estimations use the corresponding exposures matched to the 1,271 stable corridors retained among 37 countries. The matrix describes geographical co-exposure rather than observed shipping flows or conditional correlations after removing regression fixed effects.

Bab el-Mandeb and, for long-haul East Asia–Europe corridors, to Malacca. Second, exposure to energy-critical Gulf and Asian chokepoints (Hormuz, Bab el-Mandeb, Malacca) is only moderately correlated with exposure to Atlantic and European chokepoints (Panama, Gibraltar, Dover), which is consistent with distinct geographical dimensions of exposure. These correlations alone do not establish independence from maritime distance.

From an econometric standpoint, the correlation structure shows that the route-specific measures capture related but non-identical dimensions of maritime exposure. The absence of perfect collinearity preserves useful cross-sectional variation for the panel estimations, although the route-specific coefficients should be interpreted cautiously when geographically linked passages are entered jointly.

D.6 Validation using alternative maritime routes

We assess the geographical content of the Maritime Trade Exposure Network using a separate reconstruction of bilateral maritime routes. This exercise combines an external port database with a different routing graph and does not reuse the benchmark network's route probabilities, passage weights, or chokepoint assignments. It examines whether benchmark exposure measures correspond to recognizable geographical features under an alternative representation of maritime connections.

The comparison provides geographical cross-validation. It does not validate exposure against observed vessel movements or shipment-level routing decisions. Both constructions reflect the same physical geography, but differ in their gateway selection, routing representation, and aggregation rules. Agreement therefore supports the geographical interpretation of the benchmark measures without establishing that their assigned probabilities coincide with realized shipping shares.

External port data and gateway selection. Port locations are obtained from the World Port Index (WPI) published by the U.S. National Geospatial-Intelligence Agency (NGA).[15] For the final 37-country universe, we use an explicit, documented selection of up to three maritime gateways per country. The selection provides a parsimonious geographical representation of maritime access; it is not estimated from bilateral trade, observed port throughput, or the benchmark exposure measures.

Selected gateways receive equal weights within each country. Writing $\mathcal{P}_i$ for the gateway set of country i, the weight of gateway p is

$$\omega_{ip} = \frac{1}{|\mathcal{P}_i|}, \qquad p \in \mathcal{P}_i. \tag{35}$$

[15]The supplied World Port Index snapshot contains 3,807 ports. The NGA World Port Index is available through the Maritime Safety Office's public FeatureServer: https://vcps.nga.mil/nauticalpubs-feature/rest/services/WPI/World_Port_Index_Viewer/FeatureServer. Routes are computed using version 1.4.3 of the open-source Python package `searoute`: https://github.com/genthalili/searoute-py. The replication files preserve the WPI snapshot, explicit gateway list, port coordinates, aggregation weights, package versions, and construction diagnostics.

The construction uses 105 gateways across the 37 countries. For landlocked Switzerland, Rotterdam is designated as the maritime gateway, representing access through the Rhine corridor. This is a geographical modelling convention rather than an estimate of Swiss trade shares across foreign ports.

The gateway system and equal weights are fixed conventions. The WPI snapshot does not establish the historical availability or relative importance of these gateways throughout the estimation period. Accordingly, the resulting indicators are static geographical proxies rather than historical port-use shares.

Route reconstruction and passage detection. For each selected origin-port–destination-port combination, we reconstruct a shortest path on the `searoute` maritime graph. The construction covers all 1,271 stable directed country pairs and produces 10,221 port-pair routes, with no failed routes.

Distances include the maritime graph path and great-circle connectors between the selected ports and their associated graph nodes. These connectors are recorded separately for diagnostic purposes and are not used to classify passage crossings. Passage use is determined from the routing graph's passage annotations for Suez, Hormuz, Bab el-Mandeb, Panama, and the South Africa corridor. The latter is used as the Cape-related routing proxy.

Let I_{pq}^{c} indicate whether the graph path connecting gateways p and q uses passage c. The independently reconstructed transit share for directed country pair ij is

$$\widehat{P}_{ij}^{c} = \sum_{p \in \mathcal{P}_i} \sum_{q \in \mathcal{P}_j} \omega_{ip} \omega_{jq} I_{pq}^{c}. \tag{36}$$

Since gateway weights sum to one within each country, $\widehat{P}_{ij}^{c} \in [0, 1]$. It measures the weighted share of reconstructed gateway paths using passage c, conditional on the selected ports and routing graph. It is not an observed fraction of shipments or vessels.

For each port pair, let L_{pq}^{route} denote the reconstructed distance, including the port-to-graph connectors, and let L_{pq}^{direct} denote the great-circle distance between the selected ports. The relative detour is

$$\text{Detour}_{pq} = \frac{L_{pq}^{\text{route}}}{L_{pq}^{\text{direct}}} - 1. \tag{37}$$

Country-pair distances and detour ratios are aggregated using the same gateway weights:

$$\widehat{Z}_{ij} = \sum_{p \in \mathcal{P}_i} \sum_{q \in \mathcal{P}_j} \omega_{ip} \omega_{jq} Z_{pq}. \tag{38}$$

The country-pair detour is therefore the weighted mean of port-pair detour ratios, rather than the ratio of weighted mean distances. It describes distance elongation relative to a direct geographical connection, not realized rerouting, voyage duration, or transport cost.

External composite indicators. The external oil-chokepoint indicator measures the weighted share of reconstructed paths using at least one of Hormuz, Bab el-Mandeb, or Suez:

$$\widehat{P}_{ij}^{\text{oil}} = \sum_{p \in \mathcal{P}_i} \sum_{q \in \mathcal{P}_j} \omega_{ip} \omega_{jq} \mathbf{1} \left\{ I_{pq}^{\text{Hormuz}} + I_{pq}^{\text{Bab}} + I_{pq}^{\text{Suez}} > 0 \right\}. \tag{39}$$

The union is evaluated at the individual-route level before aggregation. Thus, a route crossing several of these passages contributes once to this indicator.

We also construct an expected passage count:

$$\widehat{N}_{ij}^{\text{oil}} = \widehat{P}_{ij}^{\text{Hormuz}} + \widehat{P}_{ij}^{\text{Bab}} + \widehat{P}_{ij}^{\text{Suez}}. \tag{40}$$

This is the weighted mean number of the three passages used by a reconstructed route and lies between zero and three. Unlike the union indicator, it counts each passage separately.

Benchmark Oil Chokepoint Exposure is compared with the external union indicator. Maritime Vulnerability is compared with both the expected passage count and the detour ratio, while Route Risk is compared with the detour ratio. These external indicators provide complementary geographical comparisons rather than identical reconstructions of the benchmark indices. In particular, the external oil indicator includes Suez, so agreement should be interpreted as geographical correspondence rather than equality of index definitions.

Validation design. Each of the 1,271 directed country pairs enters the comparison once. We report unconditional Pearson and Spearman correlations. We then remove exporter and importer fixed effects from the benchmark and external measures to examine their residual association. These conditional comparisons assess agreement beyond additive exporter and importer components.

For each matched pair of indicators, we estimate

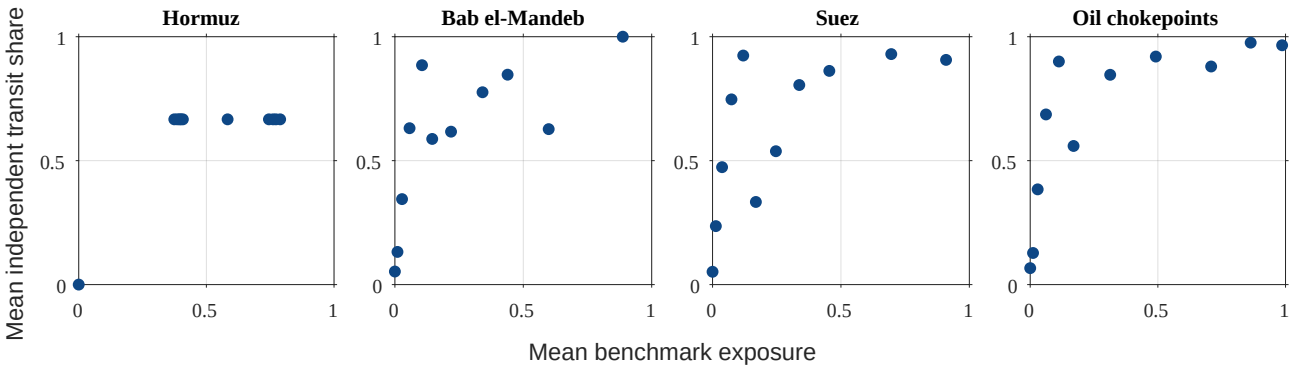
$$\widehat{Z}_{ij} = a_i + b_j + \gamma \frac{B_{ij} - \overline{B}}{\sigma_B} + \nu_{ij}, \tag{41}$$

where B_{ij} is the benchmark measure and $\widehat{Z}_{ij}$ its external counterpart. The benchmark measure is standardized across the validation sample before estimation. Consequently, γ measures the conditional association with a one-standard-deviation difference in benchmark exposure, expressed in the external indicator's units. Standard errors are two-way clustered by exporter and importer.

Conditional validation statistics are reported only when both measures retain variation after removing the country fixed effects. A measure fully absorbed by these effects cannot provide a conditional comparison.

Results. The corrected comparisons support substantial geographical correspondence for Bab el-Mandeb, Suez, and oil-chokepoint exposure. Their unconditional Spearman correlations are approximately 0.748, 0.772, and 0.783, respectively. After removing exporter and importer fixed

Figure 20: Geographical validation using independently reconstructed bilateral maritime routes



Notes: The figure presents descriptive profiles comparing benchmark exposures with external transit indicators across 1,271 directed country pairs. Corridors with zero benchmark exposure form a separate group. Positive exposures are grouped using up to ten quantile intervals, with identical exposure values kept together. Each point reports the group's mean benchmark exposure and mean external indicator. Group sizes may differ substantially. The external passage shares use explicitly selected WPI gateways, equal within-country gateway weights, and passage annotations from the `searoute` graph. The oil-chokepoint indicator is the weighted share of paths crossing at least one of Hormuz, Bab el-Mandeb, or Suez, evaluated at the route level. The profiles assess geographical correspondence and do not represent observed shipment shares or causal relationships.

effects, the corresponding partial Pearson correlations are 0.657, 0.677, and 0.627. Agreement therefore extends beyond additive country-specific components.

Maritime Vulnerability is strongly associated with the external expected passage count. Their partial Pearson correlation is 0.771. Conditional on exporter and importer fixed effects, a one-standard-deviation difference in benchmark Maritime Vulnerability is associated with approximately 0.938 additional passages in the external weighted count. This is an association between geographical indicators, not a causal effect on passage use.

Route Risk has a positive but more moderate relationship with the external detour ratio. Their unconditional Spearman correlation is 0.395, while the partial Pearson correlation is 0.190. The conditional standardized slope is 0.103, with a two-way clustered standard error of 0.038. By contrast, Maritime Vulnerability has little conditional linear association with the detour ratio: the partial Pearson correlation is 0.033, and the regression slope is not statistically distinguishable from zero ($p = 0.552$). The external comparisons therefore support specific geographical dimensions of the benchmark measures rather than a uniform relationship with every routing characteristic.

Hormuz requires a distinct interpretation. The benchmark and external constructions identify exactly the same 125 exposed corridors; the remaining 1,146 corridors are unexposed in both constructions. Their unconditional Pearson and Spearman correlations are 0.946 and 0.998, respectively.

However, the external Hormuz share takes only two substantive values: zero and 2/3. Within

the matched sample, it satisfies

$$\widehat{P}_{ij}^{\text{Hormuz}} = \frac{2}{3}\left[\mathbf{1}\{i \in \{\text{IRN}, \text{SAU}\}\} + \mathbf{1}\{j \in \{\text{IRN}, \text{SAU}\}\}\right]. \tag{42}$$

The matched sample contains no corridor for which both indicators in brackets equal one. The value 2/3 reflects the equal-weight gateway convention, not an observed trade share.

Exporter and importer fixed effects consequently absorb all variation in the external Hormuz measure. Its conditional validation statistics are therefore undefined and are not interpreted. This result supports agreement in the geographical location of exposed corridors, but does not independently validate differences in exposure intensity among those corridors.

Construction diagnostics and sensitivity. All 1,271 country pairs have complete reconstructed route coverage. Comparisons of available reverse routes reveal no discrepancies in the audited passage classifications.

Port-to-graph connectors are also examined. A diagnostic threshold of 100 nautical miles flags 303 country pairs with at least one large connector. This threshold is a screening convention, not a test of navigability or route feasibility.

Restricting the comparison to the 968 pairs without a flagged connector preserves positive conditional associations for Bab el-Mandeb, Suez, the oil-chokepoint indicator, and Maritime Vulnerability against the passage count. Their partial Pearson correlations are approximately 0.609, 0.642, 0.592, and 0.743, respectively.

This restriction changes the country composition, retaining 32 countries and excluding Brazil, Indonesia, Iran, Malaysia, and Nigeria entirely. It therefore provides a sensitivity check for connector size, rather than a comparison holding country coverage fixed. In particular, it cannot establish robustness of the Hormuz comparison with Iran retained.

Scope and limitations. The exercise supports the geographical content of the benchmark network, particularly for Suez, Bab el-Mandeb, oil-chokepoint exposure, and Maritime Vulnerability measured against passage counts. For Hormuz, its contribution is the agreement in exposed-corridor membership.

The reconstruction remains conditional on explicit gateway choices, equal aggregation weights, shortest graph paths, and the graph's passage representation. It does not identify actual port-use shares, commercial route choice, vessel-specific feasibility, congestion, or operational substitutes during a disruption. The regression standard errors do not incorporate uncertainty from alternative gateway selections, port weights, or routing assumptions.

Accordingly, this exercise provides a geographical cross-check of the benchmark exposures. It does not calibrate their probabilities to observed shipping shares or establish the causal identification of the panel estimates.

Descriptive comparison with PortWatch. A separate comparison uses IMF PortWatch cargo traffic across five passages. The benchmark network and PortWatch portfolio shares have a Pearson correlation of −0.639 and a Spearman rank correlation of −0.300. The benchmark share for Panama is 0.637, compared with 0.084 in PortWatch. These are descriptive comparisons across only five passages; the constructed network and PortWatch traffic differ in coverage and in what they measure. Their divergence is therefore informative about limitations in aggregate calibration and is not by itself a test of the accuracy of individual bilateral route probabilities.

E Bilateral panel methods and additional robustness

The bilateral average-response local projections are specified in Section 5.1. Their dependent variable is the cumulative log change in WTPI-deflated bilateral FOB exports, given in Equation (5.1). It is a measure adjusted by a common world trade-price index, not a directly observed bilateral quantity index. The specifications use directed-pair fixed effects, calendar-month indicators, trends, lagged bilateral export growth, and exporter and importer controls. Exact-date fixed effects are excluded because they would absorb the common current Oil-GPR measure. Fixed corridor weights are based on average nominal exports over 2000–2019; this window is not prior to every shock in the sample. See Equation (4) for the specification and the main-text figure notes for the sample and shock-lag choices.

The reported dynamic paths apply the split-panel jackknife correction before hierarchical smoothing, as described by Equations (43) and (44). Standard errors condition on the chosen shock history and allow for clustering by directed pair and date. For the SVAR innovation, the analysis also propagates uncertainty across paired BVAR shock histories; its variance decomposition is given in Equation (45). The sample, number of histories, and normalization must be read from the corresponding exercise and figure notes. The observed Oil GPR and AR-based specifications use different information sets and are not automatically equivalent to the SVAR shock.

E.1 Split-panel correction, hierarchical smoothing, and shock uncertainty

Fixed-effects local projections can exhibit finite-T bias. For each horizon, the split-panel-jackknife estimator is

$$\widehat{\beta}_{0,h}^{SPJ} = 2\widehat{\beta}_{0,h}^{F} - \frac{1}{2}\left(\widehat{\beta}_{0,h}^{A} + \widehat{\beta}_{0,h}^{B}\right), \tag{43}$$

where F denotes the full sample and A and B its two contiguous time blocks. The cross-horizon covariance matrix uses the same linear combination of observation-level influence functions, retaining covariance between the full- and half-sample estimators.

The corrected response vector and its cluster-robust cross-horizon covariance define the working likelihood for hierarchical estimation. The prior

$$\beta_{0,h} - 2\beta_{0,h-1} + \beta_{0,h-2} \sim \mathcal{N}(0, \tau_\beta^2), \qquad h = 2, \ldots, 12, \tag{44}$$

with a weakly informative half-Cauchy prior on τ_β regularizes irregular curvature without directly shrinking the level or linear slope.

Conditional inference holds the median SVAR shock history fixed. Integrated inference repeats panel estimation, jackknife correction, covariance construction, and smoothing for 200 paired posterior BVAR shock histories. Pooling the resulting response paths combines the two uncertainty components:

$$\begin{aligned}\operatorname{Var}\left(\boldsymbol{\beta}_0 \mid \mathcal{D}\right) = \mathbb{E}_s\left[\operatorname{Var}\left(\boldsymbol{\beta}_0 \mid s, \mathcal{D}\right)\right] \\ + \operatorname{Var}_s\left[\mathbb{E}\left(\boldsymbol{\beta}_0 \mid s, \mathcal{D}\right)\right].\end{aligned} \tag{45}$$

The first term captures panel uncertainty conditional on a recovered history; the second captures variation across posterior histories.

E.2 Supply-chain pressure amplification across Oil-GPR specifications

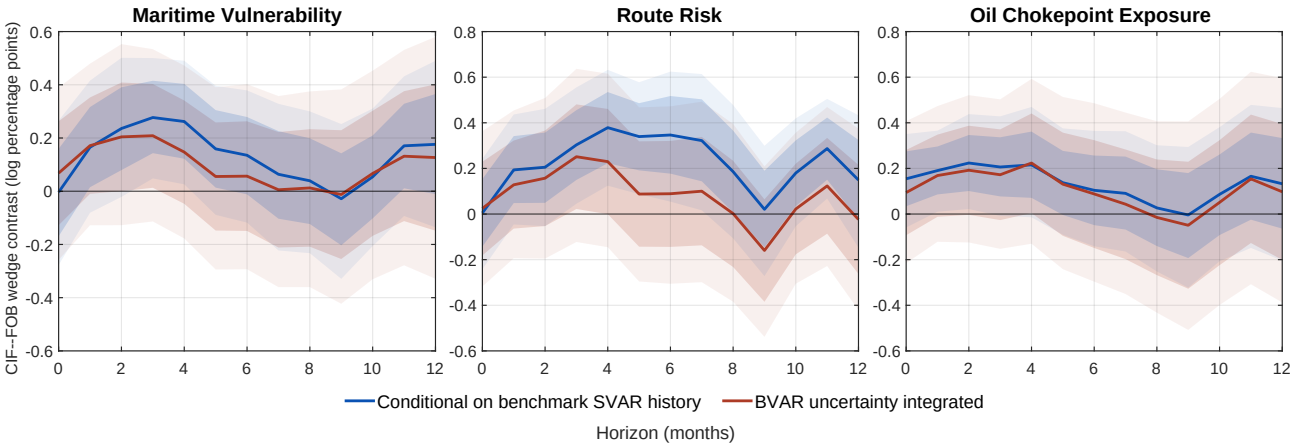
Table 11: GSCPI amplification across alternative Oil-GPR specifications

Specification	$h = 1$–6	$h = 6$–8	$h = 1$–12
Panel A: Maritime Vulnerability			
SVAR Oil-GPR innovation	−0.554 (0.537)	0.908 (0.694)	0.515 (0.545)
Observed Oil GPR + 12 interaction lags	1.040 (0.710)	2.749*** (0.743)	2.394*** (0.676)
Oil-GPR AR(12) innovation	−0.750 (0.536)	0.407 (0.669)	0.068 (0.541)
Oil-GPR AR(BIC) innovation	−0.809 (0.529)	0.444 (0.625)	0.121 (0.531)
Observed Oil GPR, no interaction lags	1.047 (0.778)	4.027*** (1.040)	3.180*** (0.858)
Panel B: Route Risk			
SVAR Oil-GPR innovation	−2.234*** (0.586)	−1.155* (0.669)	−1.341** (0.573)
Observed Oil GPR + 12 interaction lags	−1.330** (0.615)	−0.419 (0.704)	−0.408 (0.633)
Oil-GPR AR(12) innovation	−3.235*** (0.542)	−2.147*** (0.612)	−2.364*** (0.543)
Oil-GPR AR(BIC) innovation	−2.936*** (0.520)	−1.856*** (0.575)	−2.061*** (0.511)
Observed Oil GPR, no interaction lags	−0.904 (0.806)	0.572 (1.079)	0.493 (0.905)
Panel C: Oil Chokepoint Exposure			
SVAR Oil-GPR innovation	0.293 (0.731)	2.034** (1.008)	1.645* (0.830)
Observed Oil GPR + 12 interaction lags	−0.253 (0.853)	1.967** (0.898)	1.693** (0.846)
Oil-GPR AR(12) innovation	0.616 (0.568)	1.804** (0.791)	1.636** (0.676)
Oil-GPR AR(BIC) innovation	0.487 (0.576)	1.769** (0.824)	1.510** (0.706)
Observed Oil GPR, no interaction lags	0.811 (0.782)	3.641*** (1.156)	3.009*** (0.955)

Notes: The table reports posterior means of the GSCPI amplification effect $\mathcal{A}_h(D_{75})$, averaged over the indicated horizon windows. Posterior standard deviations are reported in parentheses. The contrast compares lagged GSCPI states of +1 and −1 standard deviation and is evaluated at the 75th percentile of exporter chokepoint dependence. A negative estimate means that elevated supply-chain pressure makes the high-minus-low exposure differential more adverse. The five specifications use the benchmark SVAR Oil-GPR innovation; Observed Oil GPR conditional on twelve monthly lags of all shock-dependent interactions; innovations obtained from AR(12) and BIC-selected AR models of Observed Oil GPR; and Observed Oil GPR without lagged shock interactions. In the lag-augmented specification, the reported amplification is based on the contemporaneous shock interactions conditional on the previous twelve months; it is not a cumulative distributed-lag coefficient. All specifications are estimated on the same 299,161 observations covering 1,271 directed pairs, 37 countries, and 253 monthly dates between January 2001 and December 2024. They include directed-pair and exact-date fixed effects, six lagged exporter and importer macroeconomic controls, and the shock interaction with importer oil dependence. Response windows overlapping March 2020–December 2021 are excluded. Split-panel-jackknife correction is applied before hierarchical Bayesian smoothing. Stars summarize posterior sign probabilities: *** $\Pr(\text{sign}) \geq 0.995$, ** $\Pr(\text{sign}) \geq 0.975$, and * $\Pr(\text{sign}) \geq 0.95$.

E.3 CIF–FOB valuation-wedge responses under structural-shock uncertainty

Figure 21: CIF–FOB valuation-wedge responses and BVAR shock uncertainty



Notes: The figure reports the dynamic high-minus-low maritime-exposure contrast in the CIF–FOB valuation wedge following a one-standard-deviation SVAR Oil-GPR innovation. The contrast is evaluated at the 75th percentile of exporter chokepoint dependence. The wedge is constructed draw by draw as the difference between the corresponding CIF and FOB contrasts. The blue line conditions on the benchmark SVAR shock history. The brick-red line integrates uncertainty across 1,000 posterior BVAR shock histories. Dark and light shaded areas denote 68% and 90% posterior credible intervals, respectively.
Estimates are obtained from split-panel-jackknife hierarchical Bayesian local projections with directed-pair and exact-date fixed effects. Standard errors underlying the conditional estimates are two-way clustered by directed country pair and exact date. The common dynamic sample contains 278,330 observations for 1,266 directed country pairs, 37 countries, and 253 monthly dates spanning January 2001 to December 2024. Response windows overlapping March 2020–December 2021 are excluded, removing 40,415 observations. The CIF–FOB wedge is restricted to the common-support interval defined by its first and ninety-ninth percentiles, equal to -2.390 and 2.094, respectively.
A positive estimate indicates that the CIF–FOB wedge responds more positively along highly exposed corridors than along weakly exposed corridors. Because the wedge may also reflect insurance, shipment timing, reporting practices, valuation conventions, and commodity composition, it should not be interpreted as a direct measure of freight costs.

F Robustness of the bilateral maritime-exposure results

The benchmark local projections control for lagged exporter and importer conditions but do not include the bilateral trade outcome's own dynamics. Persistent corridor-specific trade movements could therefore be correlated with the interaction between the common shock and maritime exposure. To examine this possibility, we augment each horizon-specific regression

with twelve exact monthly lags of bilateral export growth:

$$\begin{aligned} 100\left(\log Trade_{ij,t+h} - \log Trade_{ij,t-1}\right) = {} & \beta_{1,h}\left(S_t\widetilde{E}_{ij}\right) + \beta_{2,h}\left(S_t\widetilde{D}_i\right) \\ & + \beta_{3,h}\left(S_t\widetilde{E}_{ij}\widetilde{D}_i\right) + \rho_h\left(S_t\widetilde{OID}_j\right) \\ & + \sum_{\ell=1}^{12}\phi_{\ell,h}\left[100\Delta\log Trade_{ij,t-\ell}\right] + \mathbf{X}'_{ij,t-1}\gamma_h \\ & + \alpha_{ij} + \delta_t + \varepsilon_{ij,t+h}. \end{aligned} \tag{46}$$

The lagged trade-growth variables are constructed using exact calendar-month matching within each directed pair, so an unobserved month is never treated as a one-month lag. Twelve lags are used as a conservative annual dynamic specification. We do not interpret this choice as an estimate of the optimal lag order or as evidence that longer trade dynamics are necessarily absent.

The additional lag requirements reduce the estimation sample from 299,161 to 297,711 observations and from 253 to 252 shock dates. The sample begins in February 2001 rather than January 2001 and ends in December 2024 in both specifications. All 1,271 directed country pairs and all 37 countries remain represented. The pair-level exposure distribution used to construct the high-minus-low contrasts is held fixed across specifications. As in the benchmark analysis, response windows overlapping March 2020–December 2021 are excluded.

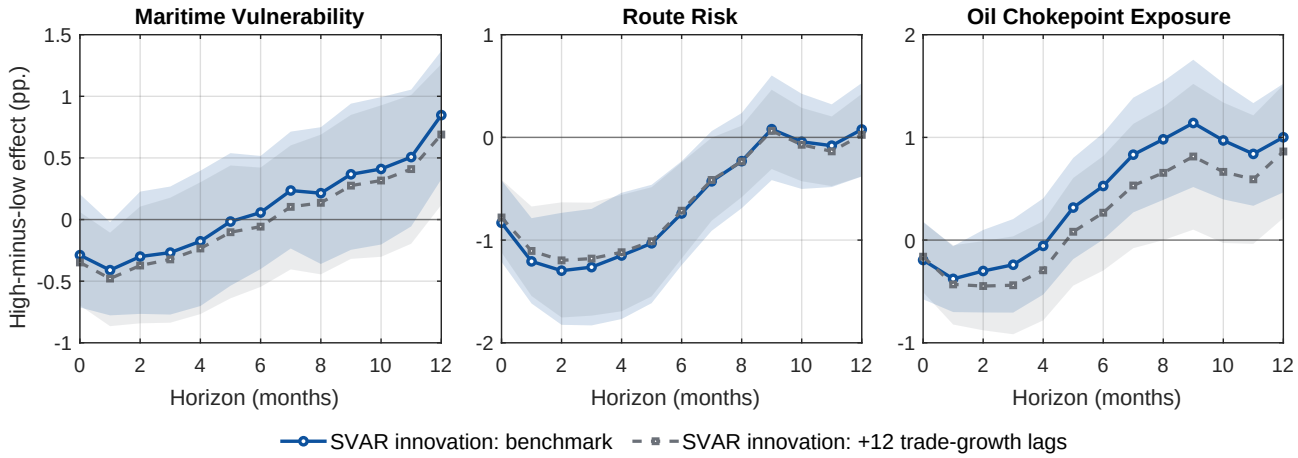
Figure 22 shows that controlling for bilateral trade persistence leaves the principal Route Risk response almost unchanged. The two posterior paths are particularly close during the first six months, when the adverse exposure differential is strongest. The inclusion of lagged bilateral trade growth also modestly reduces posterior uncertainty at several intermediate horizons. By contrast, the delayed positive response associated with Oil Chokepoint Exposure becomes smaller and less precisely estimated. Maritime Vulnerability remains close to zero over the principal summary windows.

Table 12 confirms the graphical evidence. For Route Risk, the average high-minus-low contrast over months 1–6 changes only from -1.115 to -1.054 percentage points after twelve trade-growth lags are introduced. The corresponding 90% credible interval remains strictly negative, at $[-1.543, -0.562]$, and the posterior probability of a negative effect is 99.98%.

The medium- and longer-window estimates are equally stable. Over months 6–8, the Route Risk contrast moves from -0.467 to -0.454, with a 90% credible interval of $[-0.830, -0.079]$ and a posterior probability of a negative response of 97.84%. Over months 1–12, it changes from -0.610 to -0.591; its 90% credible interval is $[-0.966, -0.217]$, and its posterior probability of being negative is 99.54%. The short-run Route Risk amplification is therefore not explained by the omission of these twelve monthly bilateral trade-growth lags.

The results for Oil Chokepoint Exposure are more sensitive. Its positive contrast over months 6–8 declines from 0.780 to 0.484 percentage points. The 90% credible interval widens from $[0.244, 1.305]$ to $[-0.110, 1.068]$, while the posterior probability of a positive effect falls from 98.74% to 91.40%. Over months 1–12, the estimate falls from 0.469 to 0.238, and the posterior

Figure 22: Dynamic maritime-exposure responses with bilateral trade-growth lags



Notes: The figure reports posterior means of the high-minus-low bilateral-exposure contrast following a one-standard-deviation SVAR Oil-GPR innovation. The contrast is evaluated at the 75th percentile of exporter chokepoint dependence. The blue solid line shows the benchmark specification, while the gray dashed line shows the specification that additionally controls for twelve exact monthly lags of bilateral export growth. Shaded areas denote 90% credible intervals. The dependent variable is the cumulative change in bilateral exports from month $t-1$ through horizon $t+h$, expressed in percentage points. Both specifications include directed-pair and exact-date fixed effects, lagged exporter and importer controls, two-way clustering by pair and date, split-panel-jackknife correction, and hierarchical Bayesian smoothing. Response windows overlapping March 2020–December 2021 are excluded. The benchmark uses 299,161 observations and 253 shock dates over January 2001–December 2024. The twelve-lag specification uses 297,711 observations and 252 shock dates over February 2001–December 2024. Both specifications retain 1,271 directed pairs and 37 countries.

probability of a positive effect declines from 96.58% to 79.52%. Thus, the delayed relative improvement remains visible in the posterior mean but is no longer credibly different from zero after controlling for bilateral trade dynamics. Maritime Vulnerability remains small and imprecisely estimated over all three windows.

The broader comparison across shock treatments, reported in Appendix Table 13, yields the same central conclusion. For the AR(12) and BIC-selected AR innovations, the months 1–6 Route Risk contrasts remain close to -1 after the trade-growth lags are added: -1.031 and -1.017, respectively. Their posterior probabilities of a negative effect exceed 99.9%. The corresponding medium- and longer-window Route Risk contrasts also remain negative and become more precisely estimated.

The results based on Observed Oil GPR are more specification dependent. When twelve lags of its interactions are already included, adding the bilateral trade-growth lags reduces the months 1–6 Route Risk contrast from -0.760 to -0.407. Its 90% credible interval nevertheless remains slightly below zero, at $[-0.849, -0.014]$. Over the months 6–8 and 1–12 windows, however, the corresponding intervals include zero. When Observed Oil GPR is used without lagged shock interactions, the positive Maritime Vulnerability response remains precisely estimated, while the positive Oil Chokepoint Exposure response becomes larger after controlling for bilateral trade persistence. These differences reinforce the importance of separating the persistent observed geopolitical-risk environment from unexpected Oil-GPR innovations.

Table 12: Robustness of maritime-exposure contrasts to bilateral trade persistence

Exposure	Window	Benchmark SVAR	+12 trade-growth lags	Change
Maritime Vulnerability	h1–h6	-0.186 (0.271)	-0.262 (0.275)	-0.076
	h6–h8	0.169 (0.276)	0.060 (0.296)	-0.108
	h1–h12	0.122 (0.243)	0.030 (0.267)	-0.092
Route Risk	h1–h6	-1.115*** (0.307)	-1.054*** (0.297)	0.061
	h6–h8	-0.467* (0.279)	-0.454** (0.230)	0.012
	h1–h12	-0.610** (0.259)	-0.591*** (0.226)	0.018
Oil Chokepoint Exposure	h1–h6	-0.022 (0.237)	-0.211 (0.254)	-0.189
	h6–h8	0.780** (0.326)	0.484 (0.355)	-0.296
	h1–h12	0.469* (0.256)	0.238 (0.299)	-0.231

Notes: The table reports posterior means of the high-minus-low bilateral-exposure contrast following a one-standard-deviation SVAR Oil-GPR innovation. Posterior standard deviations are reported in parentheses. The contrast is evaluated at the 75th percentile of exporter chokepoint dependence and averaged over the indicated horizons. "Benchmark SVAR" denotes the benchmark dynamic specification without lagged bilateral trade growth. "+12 trade-growth lags" adds twelve exact monthly lags of $100\Delta \log Trade_{ij,t-\ell}$. "Change" is the posterior mean from the twelve-lag specification minus the benchmark posterior mean. Because the two specifications are estimated separately and differ slightly in their usable samples, this change is descriptive and is not accompanied by a posterior significance test.
Both specifications retain the stable network of 1,271 directed pairs among 37 countries and use the same pair-level exposure support. They include directed-pair and exact-date fixed effects, lagged exporter and importer controls, two-way clustering by pair and date, split-panel-jackknife correction, hierarchical Bayesian smoothing, and the exclusion of response windows overlapping March 2020–December 2021. The benchmark contains 299,161 observations and 253 shock dates over January 2001–December 2024. The twelve-lag specification contains 297,711 observations and 252 shock dates over February 2001–December 2024. Stars summarize the posterior probability that the contrast has the sign of its posterior mean: *** Pr(sign) ≥ 0.995, ** Pr(sign) ≥ 0.975, and * Pr(sign) ≥ 0.95. A negative contrast indicates a more adverse relative export response along highly exposed corridors. It does not measure the absolute export response of either exposure group.

Overall, the robustness exercise strengthens the paper's principal dynamic finding. The adverse response of high-Route-Risk corridors to the benchmark SVAR Oil-GPR innovation is highly stable after controlling for twelve months of bilateral trade persistence. By contrast, the delayed positive response associated with Oil Chokepoint Exposure is attenuated and becomes substantially less precise. The exercise supports the robustness of the short-run Route Risk contrast to the inclusion of bilateral trade dynamics, while the delayed Oil Chokepoint Exposure response remains more sensitive to specification.

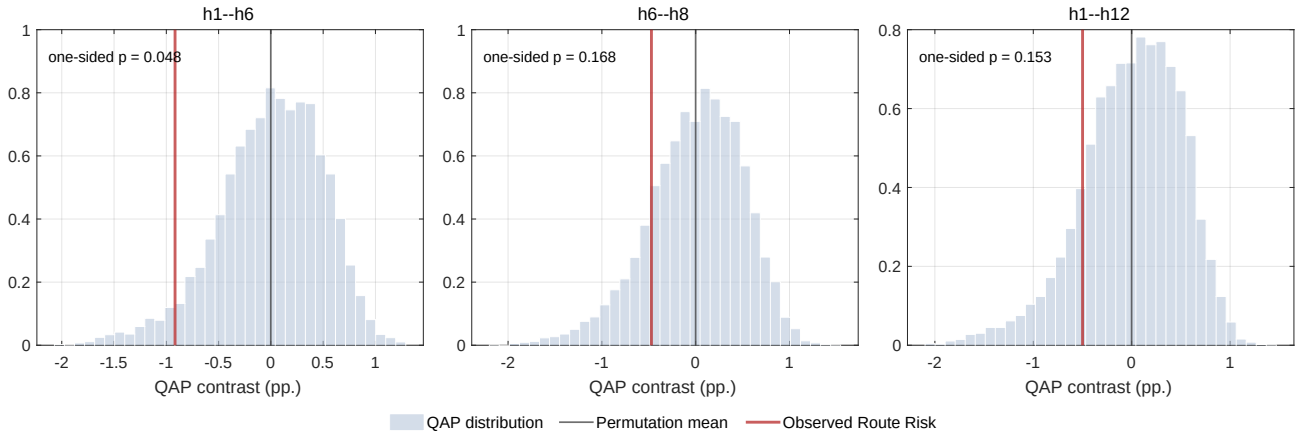
Table 13: Bilateral trade persistence across alternative Oil-GPR specifications

Specification	Exposure	Window	No trade lags	+12 trade-growth lags	Change
SVAR Oil-GPR innovation	Maritime Vulnerability	h1–h6	-0.186 (0.271)	-0.262 (0.275)	-0.076
		h6–h8	0.169 (0.276)	0.060 (0.296)	-0.108
		h1–h12	0.122 (0.243)	0.030 (0.267)	-0.092
	Route Risk	h1–h6	-1.115*** (0.307)	-1.054*** (0.297)	0.061
		h6–h8	-0.467* (0.279)	-0.454** (0.230)	0.012
		h1–h12	-0.610** (0.259)	-0.591*** (0.226)	0.018
	Oil Chokepoint Exposure	h1–h6	-0.022 (0.237)	-0.211 (0.254)	-0.189
		h6–h8	0.780** (0.326)	0.484 (0.355)	-0.296
		h1–h12	0.469* (0.256)	0.238 (0.299)	-0.231
Observed Oil GPR + 12 interaction lags	Maritime Vulnerability	h1–h6	-0.097 (0.232)	-0.277 (0.241)	-0.181
		h6–h8	0.005 (0.304)	-0.192 (0.331)	-0.197
		h1–h12	0.169 (0.256)	-0.029 (0.283)	-0.198
	Route Risk	h1–h6	-0.760** (0.284)	-0.407* (0.260)	0.353
		h6–h8	-0.560** (0.271)	-0.161 (0.209)	0.398
		h1–h12	-0.579** (0.261)	-0.185 (0.210)	0.395
	Oil Chokepoint Exposure	h1–h6	-0.059 (0.209)	-0.103 (0.238)	-0.044
		h6–h8	0.754** (0.299)	0.482 (0.367)	-0.272
		h1–h12	0.577** (0.220)	0.474 (0.305)	-0.104
Oil-GPR AR(12) innovation	Maritime Vulnerability	h1–h6	-0.074 (0.240)	-0.241 (0.247)	-0.166
		h6–h8	0.111 (0.231)	-0.027 (0.282)	-0.138
		h1–h12	0.193 (0.188)	0.043 (0.236)	-0.149
	Route Risk	h1–h6	-0.989*** (0.317)	-1.031*** (0.309)	-0.042
		h6–h8	-0.405 (0.281)	-0.527** (0.204)	-0.122
		h1–h12	-0.497** (0.253)	-0.586** (0.214)	-0.088
	Oil Chokepoint Exposure	h1–h6	-0.083 (0.175)	-0.175 (0.189)	-0.092
		h6–h8	0.627** (0.236)	0.584** (0.289)	-0.043
		h1–h12	0.422** (0.180)	0.370 (0.238)	-0.052
Oil-GPR AR(BIC) innovation	Maritime Vulnerability	h1–h6	0.041 (0.260)	-0.186 (0.253)	-0.228
		h6–h8	0.222 (0.240)	0.018 (0.298)	-0.203
		h1–h12	0.302 (0.202)	0.097 (0.242)	-0.205
	Route Risk	h1–h6	-1.023*** (0.324)	-1.017*** (0.315)	0.006
		h6–h8	-0.420 (0.280)	-0.518** (0.216)	-0.098
		h1–h12	-0.498* (0.256)	-0.573*** (0.218)	-0.075
	Oil Chokepoint Exposure	h1–h6	0.086 (0.195)	-0.020 (0.202)	-0.106
		h6–h8	0.772*** (0.239)	0.681** (0.290)	-0.091
		h1–h12	0.576*** (0.195)	0.490** (0.241)	-0.086
Observed Oil GPR, no interaction lags	Maritime Vulnerability	h1–h6	0.722*** (0.271)	0.816** (0.299)	0.094
		h6–h8	1.920*** (0.395)	2.038*** (0.414)	0.118
		h1–h12	1.579*** (0.324)	1.672*** (0.349)	0.093
	Route Risk	h1–h6	-0.022 (0.327)	-0.379 (0.355)	-0.356
		h6–h8	0.845** (0.362)	0.467 (0.366)	-0.378
		h1–h12	0.527 (0.321)	0.121 (0.337)	-0.406
	Oil Chokepoint Exposure	h1–h6	0.227 (0.193)	0.608* (0.336)	0.381
		h6–h8	1.279*** (0.327)	1.879*** (0.450)	0.600
		h1–h12	0.895*** (0.276)	1.472*** (0.384)	0.577

F.1 Network-permutation placebo

A remaining concern is that Route Risk may proxy for persistent features of the bilateral trade network rather than for the geographical transmission mechanism of interest. We address this concern using a quadratic assignment procedure (QAP). In each permutation, exporter and importer labels are simultaneously relabelled in the directed Route Risk matrix. This preserves the distribution and network structure of the measure while breaking its empirical correspondence with the actual trade corridors. The benchmark local-projection specification is then re-estimated using each of 9,999 permuted networks.

Figure 23: Network-permutation placebo for Route Risk



Notes: The figure presents QAP permutation distributions for the high-minus-low Route Risk contrast evaluated at the 75th percentile of exporter chokepoint dependence. Each distribution is based on 9,999 simultaneous permutations of exporter and importer labels in the directed Route Risk matrix, thereby preserving its topology while reallocating exposure across bilateral corridors. The vertical red line denotes the observed unsmoothed split-panel-jackknife estimate, and the black line denotes the mean of the permutation distribution. The dependent variable is the cumulative response of bilateral exports to a one-standard-deviation SVAR Oil-GPR innovation. The benchmark specifications include country-pair and exact-date fixed effects, the full set of lower-order interactions, lagged controls, and two-way clustering by country pair and date. Months whose outcome windows overlap the Covid period are excluded. Reported permutation probabilities use the finite-sample plus-one correction. The QAP probabilities concern the unsmoothed SPJ statistic and are distinct from posterior probabilities obtained from the hierarchical Bayesian local projections.

Figure 23 shows that the principal short-run Route Risk effect is unusually negative relative to effects obtained after randomly reallocating the network across country pairs. Over horizons $h = 1, \ldots, 6$, the observed SPJ estimate is -0.917, compared with a permutation mean close to zero. Its one-sided QAP p-value is 0.048, while the two-sided value is 0.061. The main short-run result is therefore significant at the 5% level against the directional alternative and at the 10% level under a two-sided test.

The corresponding QAP evidence is weaker over the $h = 6, \ldots, 8$ and $h = 1, \ldots, 12$ windows, for which the one-sided probabilities are 0.168 and 0.153, respectively. This distinction is informative: the placebo exercise provides its clearest support for the initial trade response, rather than for the persistence of that response. The hierarchical BLP estimates remain negative over all three windows, but these posterior results should not be interpreted as QAP probabilities because smoothing and network-permutation uncertainty represent different inferential objects.

As an additional falsification exercise, we replace Route Risk with exposure to Panama and Dover and compare it with the Bosphorus–Dardanelles route. Although a few individual coefficients are marginally significant before adjustment, none survives the Holm correction for multiple testing. Taken together, the results indicate that the benchmark short-run Route Risk effect is not easily reproduced by arbitrary reallocations of the maritime network or by alternative passage measures. They support the network specificity of the principal result, while stopping short of providing an independent causal identification test.

G Generated-shock uncertainty and country-date fixed effects

G.1 Static estimates

Table 14: Static maritime-exposure contrasts and SVAR shock uncertainty

Exposure measure	Conditional estimate	SVAR-integrated estimate	Pr(sign)	SVAR variance share
Panel A. Without importer oil-dependence control				
Maritime Vulnerability	−0.0018 [−0.0086, 0.0051]	−0.0017 [−0.0101, 0.0072]	0.649	0.279
Route Risk	−0.0038 [−0.0122, 0.0046]	−0.0038 [−0.0125, 0.0056]	0.758	0.162
Oil Chokepoint Exposure	0.0025 [−0.0021, 0.0071]	0.0022 [−0.0046, 0.0092]	0.703	0.368
Panel B. With importer oil-dependence control				
Maritime Vulnerability	−0.0017 [−0.0085, 0.0050]	−0.0017 [−0.0097, 0.0074]	0.623	0.280
Route Risk	−0.0038 [−0.0122, 0.0046]	−0.0038 [−0.0127, 0.0057]	0.751	0.162
Oil Chokepoint Exposure	0.0028 [−0.0017, 0.0073]	0.0025 [−0.0040, 0.0090]	0.740	0.357
Observations		343,840		
Directed country pairs		1,271		
SVAR posterior histories		200		

Notes: The table reports high-minus-low maritime-exposure contrasts evaluated at the 75th percentile of exporter chokepoint dependence following a one-standard-deviation SVAR Oil-GPR innovation. Low and high exposure correspond to the mean exposure among country pairs below the 25th percentile and above the 75th percentile, respectively. The conditional estimates treat the posterior-median SVAR Oil-GPR innovation as observed; The static contrasts are expressed in log units, not multiplied by 100; 90% confidence intervals based on pair/date two-way clustered standard errors are reported in brackets. The SVAR-integrated estimates propagate first-stage uncertainty using 200 posterior shock histories, each standardized separately over the same 300 months from January 2000 to December 2024. Their brackets report total 90% intervals combining the average conditional panel variance with the between-SVAR-draw variance. Pr(sign) is the probability that the total-uncertainty draw has the same sign as the corresponding integrated mean. The SVAR variance share is the fraction of total variance attributable to variation across posterior shock histories. All specifications include six lagged exporter and importer controls, directed-pair fixed effects, and exact-date fixed effects. Panel B additionally controls for the interaction between the shock and centered importer crude-oil dependence.

Table 14 examines whether inference on the static maritime-exposure contrasts is sensitive to the fact that the SVAR shock is a generated regressor. The conditional estimates reproduce the benchmark approach, which treats the posterior-median shock as observed, whereas the SVAR-integrated estimates re-estimate the panel model across 200 posterior shock histories and incorporate the resulting first-stage uncertainty.

The integrated posterior means remain very close to their conditional counterparts. With the importer oil-dependence control, the contrast changes from −0.0017 to −0.0017 for Maritime

Vulnerability, from −0.0038 to −0.0038 for Route Risk, and from 0.0028 to 0.0025 for Oil Chokepoint Exposure. First-stage uncertainty therefore does not materially shift the point estimates or reverse their signs.

It nevertheless contributes meaningfully to overall uncertainty. The SVAR component accounts for approximately 16% of total variance for Route Risk, 28% for Maritime Vulnerability, and 36% for Oil Chokepoint Exposure. Consequently, the SVAR-integrated intervals are wider than the intervals conditional on the posterior-median shock. None of the static contrasts is distinguishable from zero at the 10% level once both sources of uncertainty are incorporated. The posterior sign probabilities remain suggestive, however: the Route Risk contrast is negative with a probability of 75.1%, whereas the Oil Chokepoint Exposure contrast is positive with a probability of 74.0% in the specification including the importer oil-dependence control.

Finally, the near-identical results across the two panels show that the importer oil-dependence interaction has little influence on the estimated maritime contrasts. Overall, accounting for generated-shock uncertainty weakens statistical precision but does not alter the qualitative static patterns. Since the paper's main transmission results concern the timing of the response, the more consequential exercise is to propagate the same SVAR uncertainty through the local projections.

G.2 Dynamic maritime-exposure responses

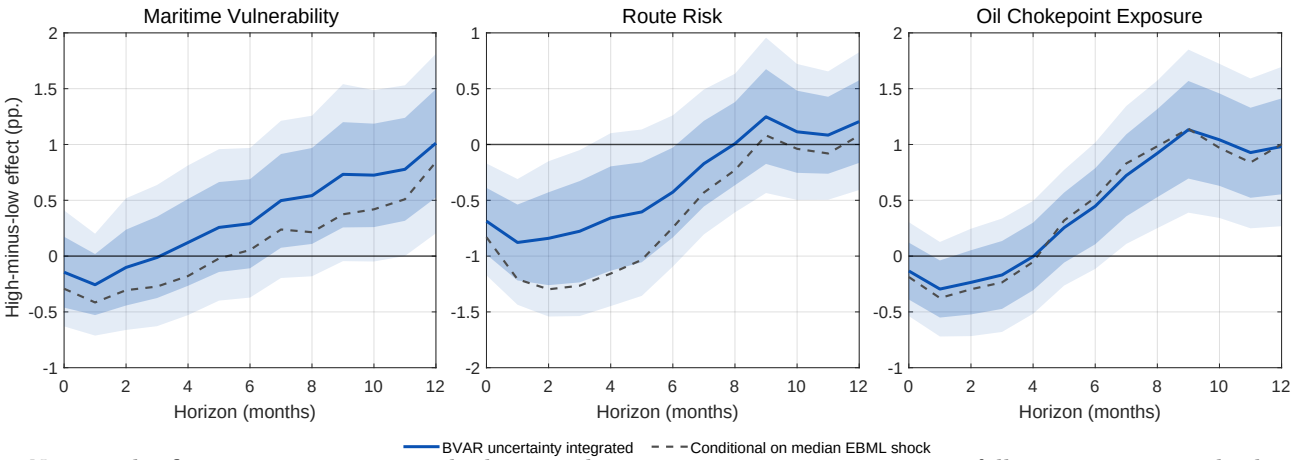
We next propagate first-stage BVAR uncertainty through the dynamic panel estimates. For each of 200 posterior BVAR draws, we reconstruct the complete history of the SVAR Oil-GPR innovation and standardize it separately to one standard deviation over the common January 2000–December 2024 window. We then re-estimate the complete local-projection system, including the split-panel jackknife correction and hierarchical smoothing of the three coefficient paths. This procedure produces 200 conditional hierarchical paths for each BVAR history, for a total of 40,000 paths per maritime-exposure measure.

Figure 24 compares the benchmark responses, which condition on the posterior-median SVAR Oil-GPR innovation, with the responses obtained after integrating over BVAR shock uncertainty. The figure reports the high-minus-low exposure contrast evaluated at the 75th percentile of exporter chokepoint dependence.

Propagating first-stage uncertainty preserves the principal short-run Route Risk contrast, while reducing its magnitude and weakening evidence at longer horizons. The clearest short-run result concerns Route Risk. Over months 1–6, the BVAR-integrated contrast is −0.70 percentage points, compared with −1.12 when conditioning on the posterior-median shock. Its 90% total interval remains entirely negative, $[-1.37, -0.03]$, and the probability of a negative response is 95.8%. Hence, even after accounting for uncertainty in the recovered structural shock, highly exposed routes experience a significantly less favorable short-run trade response.

The Route Risk effect subsequently dissipates. Over months 6–8, its integrated mean equals −0.20 percentage points, with a 90% interval of $[-0.81, 0.44]$ and a 70.8% probability of being negative. Over the full months 1–12 window, the estimate remains negative at −0.31, but its interval also includes zero. The generated-regressor correction therefore preserves the short-run

Figure 24: Dynamic maritime-exposure responses with BVAR shock uncertainty



Notes: The figure reports posterior high-minus-low maritime-exposure contrasts following a one-standard-deviation SVAR Oil-GPR innovation. The contrasts compare the mean exposure of country pairs above the 75th percentile with that of pairs below the 25th percentile and are evaluated at the 75th percentile of exporter chokepoint dependence. The dashed gray lines report the hierarchical local-projection estimates conditional on the posterior-median SVAR Oil-GPR innovation. The solid blue lines integrate uncertainty over 200 posterior BVAR shock histories. Each history is standardized separately over the same 300 months from January 2000 to December 2024 and is passed through the complete panel, split-panel-jackknife, and hierarchical-smoothing procedure. The dark and light shaded areas denote, respectively, 68% and 90% total uncertainty intervals around the BVAR-integrated responses. Total inference combines panel-estimation and first-stage shock-recovery uncertainty. The specification includes six lagged exporter and importer controls, the interaction between the shock and centered importer crude-oil dependence, directed-pair fixed effects, and exact-date fixed effects. Response windows overlapping March 2020–December 2021 are excluded. Positive values indicate a more favorable *relative* response of highly exposed corridors and do not necessarily imply an absolute increase in their exports.

contraction while weakening the evidence of persistence.

Oil Chokepoint Exposure displays the opposite timing pattern. Its short-run contrast over months 1–6 is essentially zero. A positive differential response then emerges in the medium run: over months 6–8, the integrated estimate is 0.70 percentage points, with a 90% total interval of $[0.11, 1.29]$ and a 97.2% probability of being positive. The corresponding months 1–12 estimate is 0.48, but its lower 90% bound is slightly negative, at -0.02, despite a posterior probability of a positive effect of 94.4%. The evidence therefore supports a delayed relative rebound along corridors with high oil-chokepoint exposure, but is strongest over the targeted medium-run window.

For Maritime Vulnerability, integrating over BVAR histories shifts the response upward relative to the conditional estimate. The integrated contrasts equal 0.44 over months 6–8 and 0.38 over months 1–12, with probabilities of a positive sign of approximately 86%–87%. Their 90% intervals nevertheless include zero, so these estimates remain suggestive rather than conclusive.

First-stage uncertainty is quantitatively important. Depending on the exposure measure and response window, variation across BVAR shock histories accounts for approximately 27% to 52% of total uncertainty. Accounting for the generated nature of the SVAR Oil-GPR innovation therefore affects both precision and, for some exposure measures, the posterior mean. The

short-run Route Risk contrast retains a negative 90% interval. The positive Oil Chokepoint Exposure contrast over months 6–8 is specific to this specification: it weakens when bilateral trade-growth lags or country-date fixed effects are introduced, as documented in Appendices **??** and G.4.

G.3 BVAR shock uncertainty and the bilateral CIF–FOB valuation wedge

This exercise assesses whether the exposure-dependent response of the bilateral CIF–FOB valuation wedge remains distinguishable from zero after accounting for uncertainty in the estimated BVAR shock. We estimate the CIF and FOB responses jointly, using the split-panel jackknife correction, country-pair and exact-date fixed effects, and the common sample excluding COVID overlap. The wedge response is constructed as the difference between the CIF and FOB responses in each posterior draw. This joint construction preserves their covariance when evaluating uncertainty in the wedge.

Figure 25 compares the high-minus-low exposure contrasts conditional on the central BVAR shock history with those integrating uncertainty over 200 posterior shock histories. The exposure contrasts and sample restrictions are held fixed across the two calculations.

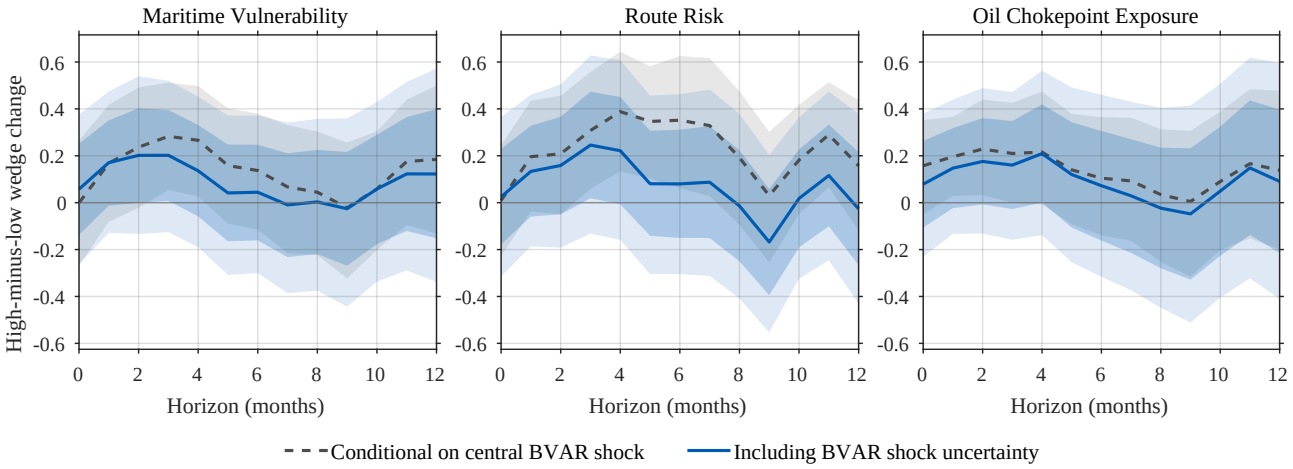


Figure 25: Exposure contrasts in the bilateral CIF–FOB valuation wedge: accounting for BVAR shock uncertainty

Notes: Each panel reports the high-minus-low exposure contrast in the response of the bilateral CIF–FOB valuation wedge, using the exposure contrast defined in the benchmark exercise. The dashed grey line reports the posterior mean conditional on the central BVAR shock history; the grey band denotes its pointwise 90% posterior interval. The solid blue line reports the posterior mean after integrating uncertainty over 200 BVAR shock histories. Darker and lighter blue bands denote pointwise 68% and 90% posterior intervals, respectively. The wedge response is calculated as the difference between jointly estimated CIF and FOB responses, draw by draw. Estimates employ the split-panel jackknife correction, country-pair and exact-date fixed effects, and the common sample excluding COVID overlap. All panels use the same vertical scale. The intervals are pointwise rather than simultaneous across horizons.

Conditional on the central shock history, the posterior mean wedge contrasts are generally positive at short horizons. The most pronounced response occurs for Route Risk, with a conditional mean reaching approximately 0.39 around month four. Integrating shock uncertainty lowers

this peak to approximately 0.24 and substantially weakens the subsequent response. Maritime Vulnerability and Oil Chokepoint Exposure also display positive posterior means during the first few months, but their responses become small around the middle of the horizon range.

Once BVAR shock uncertainty is included, the pointwise 90% intervals contain zero at every horizon for all three exposure measures. The window averages lead to the same conclusion. Over months one to six, the total posterior mean wedge contrasts are approximately 0.133 for Maritime Vulnerability, 0.153 for Route Risk, and 0.148 for Oil Chokepoint Exposure, with posterior probabilities of a positive response of approximately 79–80%. Nevertheless, all corresponding 90% posterior intervals include zero. Uncertainty in the BVAR shock accounts for approximately 42–45% of the total posterior variance of these short-horizon window averages.

The exercise therefore provides suggestive evidence of a short-run increase in the wedge for more exposed corridors, but does not establish a precisely estimated positive response after accounting for shock uncertainty. This qualification concerns the wedge: individual CIF and FOB responses can remain distinguishable from zero even when their difference does not. Moreover, the valuation wedge is not a pure measure of freight costs; insurance, differences in reporting and valuation, and changes in trade composition can also affect it.

G.4 Country-date fixed effects: conditional shock histories

The benchmark local projections control for directed-pair and exact-date fixed effects. A remaining concern is that the estimated maritime-exposure gradients may capture time-varying shocks specific to exporting or importing countries. We therefore estimate a more demanding dynamic specification containing directed-pair, exporter-by-date, and importer-by-date fixed effects:

$$\begin{aligned} 100\left(\log X_{ij,t+h} - \log X_{ij,t-1}\right) = \beta_{1,h}\left(Shock_t \times \widetilde{E}_{ij}\right) \\ + \beta_{3,h}\left(Shock_t \times \widetilde{E}_{ij} \times \widetilde{D}_i\right) + \alpha_{ij} + \lambda_{it}^{X} + \mu_{jt}^{M} + \varepsilon_{ij,t+h}, \end{aligned} \tag{47}$$

where λ_{it}^{X} and μ_{jt}^{M} absorb arbitrary monthly conditions common to all exports from country i and all imports into country j, respectively. These effects therefore control nonparametrically for country-specific supply, demand, policy, exchange-rate, and logistics shocks. The lower-order interaction $Shock_t \times \widetilde{D}_i$, whose coefficient is $\beta_{2,h}$ in the benchmark specification, is perfectly collinear with the exporter-by-date fixed effects and is consequently absorbed rather than estimated. The same is true of the common shock, the importer oil-dependence interaction, and the lagged exporter- and importer-specific macroeconomic controls.

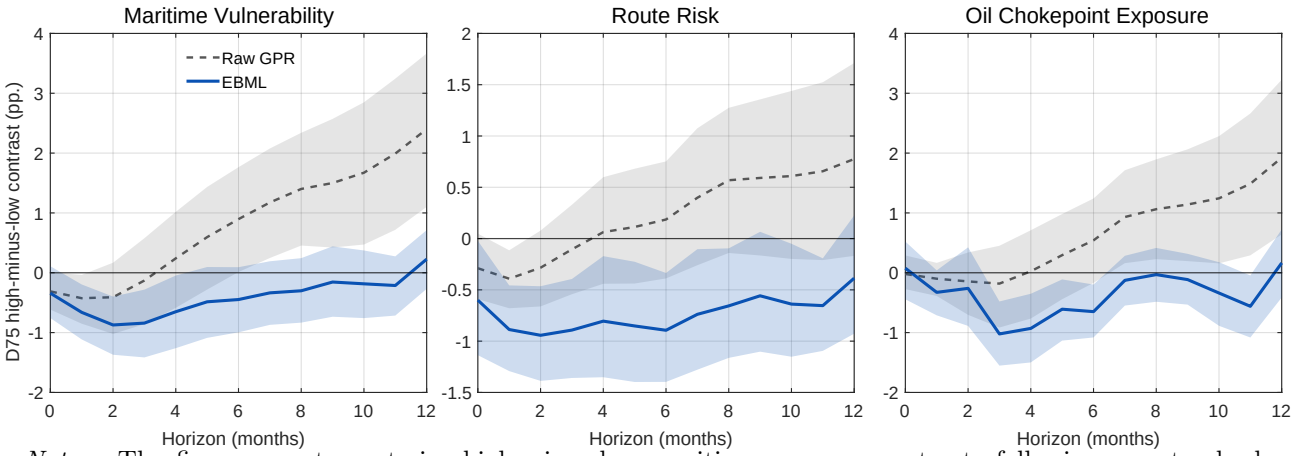
Figure 26 reports the high-minus-low exposure contrast evaluated at the 75th percentile of exporter chokepoint dependence:

$$\mathcal{C}_h(D_{75}) = \left(E^H - E^L\right)\left[\beta_{1,h} + \beta_{3,h}\left(D_{75} - D_{50}\right)\right]. \tag{48}$$

This remains a relative effect: a negative value means that highly exposed corridors perform less favorably than weakly exposed corridors after the shock, conditional on exporter- and

importer-specific monthly conditions.

Figure 26: Dynamic maritime-exposure contrasts with exporter-by-date and importer-by-date fixed effects



Notes: The figure reports posterior high-minus-low maritime-exposure contrasts following one-standard-deviation standardized Observed Oil GPR and the posterior-median SVAR Oil-GPR innovation, respectively. Contrasts are evaluated at the 75th percentile of exporter chokepoint dependence. All specifications include directed-pair, exporter-by-date, and importer-by-date fixed effects. The interaction between the common shock and exporter chokepoint dependence is absorbed by the exporter-by-date effects and is not separately estimated. Estimation uses the stable 37-country network and excludes observations whose response windows overlap March 2020–December 2021. The coefficient paths are split-panel-jackknife corrected and smoothed jointly across horizons using the hierarchical Bayesian local-projection procedure. Inference uses two-way clustering by directed pair and exact date and is conditional on the posterior-median SVAR Oil-GPR innovation. Shaded areas denote 90% credible intervals. Positive values indicate a more favorable *relative* response of highly exposed corridors and do not necessarily imply an absolute increase in exports.

The strongest result concerns Route Risk. Under the SVAR Oil-GPR innovation, the contrast is negative throughout most of the projection horizon and remains precisely estimated after arbitrary exporter- and importer-specific monthly variation has been absorbed. Its posterior mean is -0.88 percentage points over months 1–6, with a 90% credible interval of $[-1.32, -0.40]$ and a 99.7% posterior probability of being negative. The effect remains negative over months 6–8, at -0.76 percentage points with an interval of $[-1.27, -0.21]$, and over months 1–12, at -0.74 with an interval of $[-1.18, -0.26]$. The route-risk result is therefore not explained by contemporaneous domestic conditions affecting all trade relationships of a given exporter or importer.

The broader exposure measures also indicate an adverse short-run response to the structural shock. For Maritime Vulnerability, the SVAR contrast averages -0.66 percentage points over months 1–6, with a 90% credible interval of $[-1.15, -0.19]$, but becomes less precisely estimated thereafter. For Oil Chokepoint Exposure, the corresponding short-run estimate is -0.63, with an interval of $[-1.06, -0.21]$. Its average over months 1–12 remains negative at -0.40, with an interval of $[-0.78, -0.04]$.

The latter result qualifies the delayed positive oil-chokepoint contrast obtained in the benchmark specification. Once exporter- and importer-specific monthly shocks are absorbed, the medium-run rebound is no longer observed. The evidence for a delayed relative recovery along oil-

chokepoint-exposed corridors should therefore be interpreted more cautiously, as it appears sensitive to the treatment of country-specific time variation.

The Observed Oil GPR results display a different profile. Maritime Vulnerability and Oil Chokepoint Exposure become positive at medium horizons, whereas the Route Risk estimates are comparatively weak and imprecise. Overall, this robustness exercise reinforces the distinction between the raw geopolitical risk index and the structural SVAR Oil-GPR innovation. Most importantly, the negative and persistent Route Risk response to the SVAR Oil-GPR innovation survives the most demanding country-time fixed-effect structure and provides further evidence of geographically differentiated trade responses.

G.5 Country-date fixed effects: generated-shock uncertainty

A demanding concern is that the baseline dynamic responses may partly reflect time-varying conditions specific to exporting or importing countries. We therefore re-estimate the maritime local projections with directed-pair, exporter-by-date, and importer-by-date fixed effects. For each exposure measure, the specification is

$$100\big(\log X_{ij,t+h} - \log X_{ij,t-1}\big) = \beta_{1,h}\left(Shock_t \widetilde{E}_{ij}\right) + \beta_{3,h}\left(Shock_t \widetilde{E}_{ij} \widetilde{D}_i\right) + \alpha_{ij} + \lambda_{it}^{X} + \mu_{jt}^{M} + \varepsilon_{ij,t+h}. \tag{49}$$

The exporter-by-date effects, λ_{it}^{X}, absorb arbitrary time-varying conditions common to all exports from country i, while the importer-by-date effects, μ_{jt}^{M}, absorb arbitrary conditions common to all imports entering country j. These fixed effects also absorb the common shock effect, the interaction between the shock and exporter dependence, the interaction with importer oil dependence, and the country-specific macroeconomic controls. Identification therefore comes from differences in bilateral maritime exposure across trading partners within the same exporter and importer at each date.

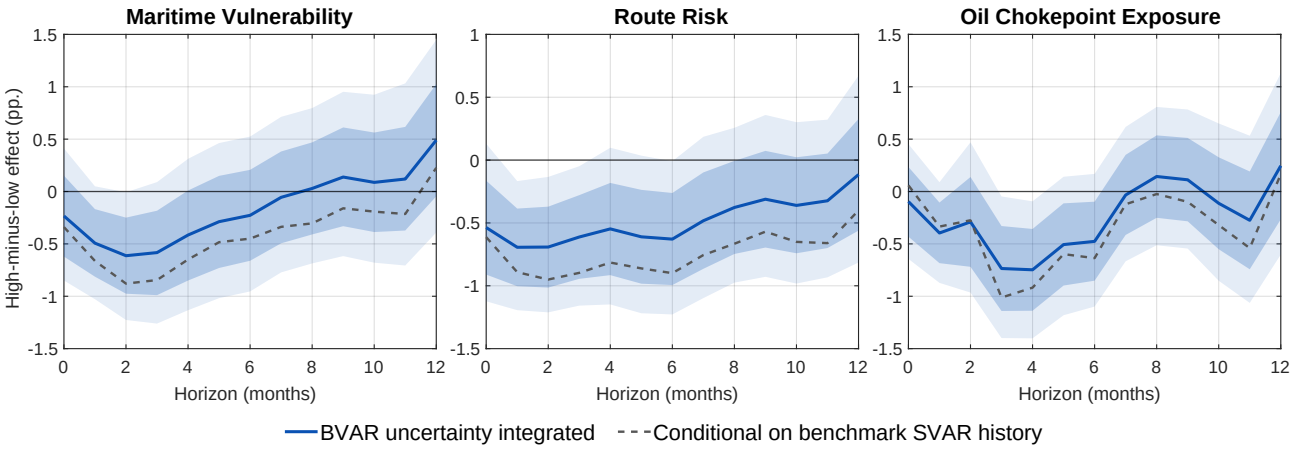
As in the benchmark analysis, the reported object is the high-minus-low exposure contrast evaluated at the 75th percentile of exporter chokepoint dependence:

$$\mathcal{C}_h(D_{75}) = (E_H - E_L)\left[\beta_{1,h} + \beta_{3,h}\left(D_{75} - D_{50}\right)\right]. \tag{50}$$

The estimation retains the split-panel-jackknife correction and hierarchical Bayesian smoothing. In addition, the uncertainty associated with recovering the structural shock is propagated through the panel estimation using 200 posterior BVAR shock histories.

Figure 27 provides particularly strong support for the short-run Route Risk result. After integrating BVAR shock uncertainty, the average high-minus-low contrast over months 1–6 is -0.63 percentage points, with a 90% credible interval of $[-1.16, -0.07]$ and a posterior probability of a negative effect of 96.9%. The response remains negative over months 6–8 and 1–12, with posterior probabilities of 90.8% and 91.9%, respectively, although the corresponding 90% intervals marginally include zero. Thus, the relative contraction of high-route-risk corridors remains after controlling for exporter- and importer-specific monthly shocks.

Figure 27: Maritime-exposure contrasts with exporter- and importer-date fixed effects



Notes: The figure reports posterior high-minus-low maritime-exposure contrasts following a one-standard-deviation structural oil-specific geopolitical shock. Contrasts are evaluated at the 75th percentile of exporter chokepoint dependence. The solid blue lines integrate uncertainty over 200 posterior BVAR shock histories; the dark and light shaded areas denote 68% and 90% total credible intervals, respectively. Dashed gray lines report estimates conditional on the posterior-median SVAR Oil-GPR innovation history. All specifications include directed-pair, exporter-by-date, and importer-by-date fixed effects. The common shock effect and its interaction with exporter dependence are absorbed by these fixed effects. Estimation uses the stable 37-country network, the No-Covid-common sample, split-panel-jackknife correction, hierarchical Bayesian smoothing, and two-way clustering by directed pair and exact date. Positive values indicate a more favorable *relative* response of highly exposed corridors and do not necessarily imply an absolute increase in exports.

The other exposure measures display a similar, although less precisely estimated, short-run pattern. For Maritime Vulnerability, the integrated contrast over months 1–6 is −0.44 percentage points, with an 88.4% posterior probability of being negative. For Oil Chokepoint Exposure, the corresponding estimate is −0.53 percentage points and the probability of a negative response reaches 94.0%; its 90% credible interval only narrowly includes zero, with an upper bound of 0.04. Both responses subsequently dissipate.

Accounting for generated-shock uncertainty generally moves the posterior means toward zero and widens the credible intervals relative to inference conditional on the median SVAR Oil-GPR innovation. For example, the Route Risk estimate over months 1–6 changes from −0.89 conditionally to −0.63 after integrating across BVAR histories. Nevertheless, its total 90% credible interval continues to exclude zero. The Route Risk result therefore survives two demanding adjustments simultaneously: arbitrary exporter- and importer-specific time variation and uncertainty generated by the first-stage recovery of the structural shock. This supports the robustness of the short-run relative export response to these two adjustments without identifying a unique causal maritime mechanism.

G.6 CIF–FOB valuation wedges with country-date fixed effects

As an additional mechanism exercise, we examine whether geopolitical oil shocks generate a larger importer–exporter valuation wedge along more exposed maritime corridors. The under-

lying log CIF–FOB valuation wedge is

$$w_{ijt} = 100\,(\log CIF_{ijt} - \log FOB_{ijt})\,,$$

and the projections use its change from $t-1$ to $t+h$, as in the main-text dynamic decomposition. The reported object is the high-minus-low exposure contrast evaluated at the 75th percentile of exporter chokepoint dependence. CIF and FOB responses are estimated jointly, and the wedge response is constructed draw by draw as

$$\mathcal{C}_h^w = \mathcal{C}_h^{CIF} - \mathcal{C}_h^{FOB}.$$

This construction preserves the posterior covariance between the CIF and FOB components.

The specification includes directed-pair, exporter-by-date, and importer-by-date fixed effects. It therefore absorbs arbitrary monthly shocks common to all exports from a given country and to all imports entering a given country. We additionally propagate uncertainty in the generated SVAR Oil-GPR innovation by re-estimating the panel model across posterior BVAR shock histories. The solid curves in Figure 28 consequently incorporate both second-stage panel uncertainty and first-stage BVAR shock uncertainty.

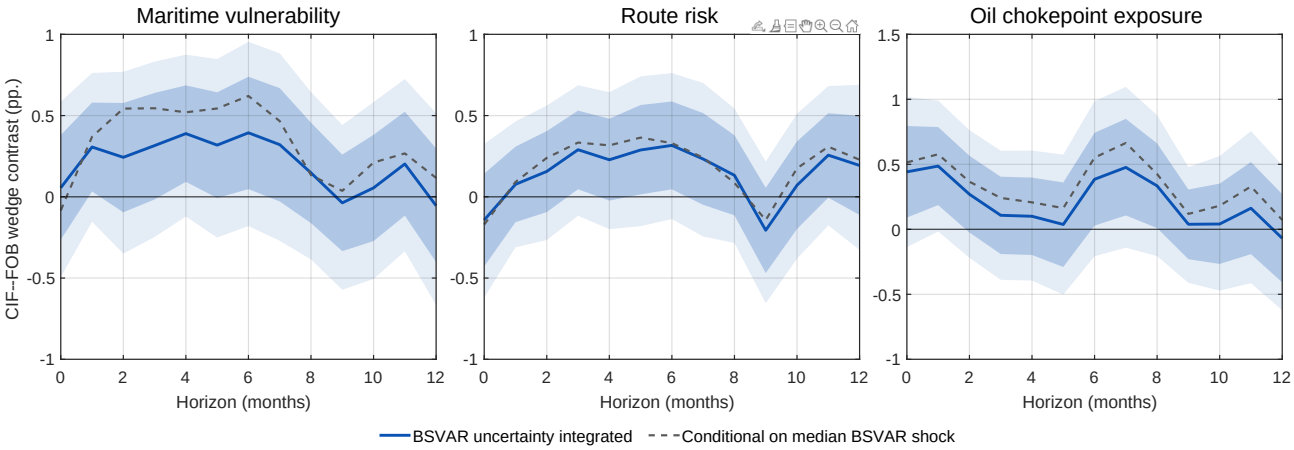


Figure 28: Dynamic CIF–FOB valuation-wedge responses under generated-shock uncertainty

Notes: The figure reports the posterior dynamic response of the high-minus-low maritime-exposure contrast in the CIF–FOB valuation wedge, evaluated at the 75th percentile of exporter chokepoint dependence. The dependent variables underlying the decomposition are expressed in 100 log points. CIF and FOB local projections are estimated jointly, and the wedge is constructed draw by draw as $\mathcal{C}_h^w = \mathcal{C}_h^{CIF} - \mathcal{C}_h^{FOB}$, thereby retaining their posterior covariance. The solid blue curves integrate uncertainty across posterior BVAR shock histories, while the dashed gray curves condition on the posterior-median SVAR Oil-GPR innovation. Dark and light shaded regions denote, respectively, 68% and 90% posterior credible intervals for the uncertainty-integrated estimates. All specifications include directed-pair, exporter-by-date, and importer-by-date fixed effects, use pair/date two-way clustering and split-panel-jackknife correction, and apply hierarchical Bayesian smoothing across horizons. The estimation uses the common no-Covid-overlap sample. A positive value indicates that the CIF–FOB valuation wedge rises more strongly for highly exposed corridors. Because this wedge may reflect freight, insurance, handling, shipment timing, valuation practices, reporting discrepancies, or changes in commodity composition, it should not be interpreted as a direct measure of physical freight rates.

Figure 28 reveals a broadly positive, but imprecisely estimated, differential response of the CIF–

FOB wedge. For Maritime Vulnerability, the posterior mean becomes positive shortly after the shock and remains predominantly above zero through the middle horizons. The Route Risk response follows a similar pattern, rising during approximately the first six months before weakening temporarily around month nine. Oil Chokepoint Exposure also displays a positive early and medium-run profile, although the estimated response is less persistent.

Accounting for uncertainty in the generated shock generally attenuates the positive conditional estimates and widens their posterior intervals. In particular, the solid uncertainty-integrated curves tend to lie below the responses obtained by conditioning on the posterior-median shock. The reported 90% credible regions do not provide clear evidence of an exposure differential distinguishable from zero across the displayed horizons. The figure therefore does not provide decisive evidence of a systematically positive valuation-wedge response once both first-stage shock uncertainty and country-specific monthly conditions are taken into account.

The exercise is nonetheless informative about the mechanism. The posterior means are predominantly positive across all three exposure measures and do not indicate a systematic decline in the wedge along more exposed corridors. This directional consistency is compatible with greater disruption-related costs or valuation discrepancies on exposed routes. However, the substantial posterior uncertainty prevents attributing the estimated trade contraction primarily to a measured increase in the CIF–FOB wedge. Accordingly, the results are best interpreted as suggestive supporting evidence for a broad maritime-disruption channel, rather than as a direct identification of a freight-cost channel.