Universitat d'Alacant
Universidad de Alicante

MASTER'S THESIS

Heterogeneity in Geopolitical Risk Shocks, Energy Prices and Macro-Financial Transmission in the Euro Area

Author: Jaime Á. Rodríguez Almendros

Supervisor: Iván Payá Sastre

Master in Economics with Data Science
2025 – 2026

Abstract

This thesis examines whether the geographical origin of geopolitical risk affects its transmission through energy and macro-financial variables in Germany, France, Italy, and Spain. Sixteen quarterly Bayesian vector autoregressive (BVAR) systems are estimated for shocks originating from Russia, China, the Middle East and North Africa, and the Western Bloc. The results show substantial heterogeneity across origins and recipient economies. Natural gas price growth falls on impact in all systems, whereas the short-run response of Brent prices varies depending on the origin of the shock. Inflationary effects are frequent, particularly following Russian-origin shocks, while industrial-production and financial-stress responses differ across countries and horizons. Russian-origin shocks generate the most widespread responses, although this pattern is not equally robust across alternative model specifications. Overall, treating geopolitical risk as a single global shock can conceal economically relevant differences in its transmission across euro-area economies.

Keywords: Geopolitical risk, BVAR, energy prices, macro-financial transmission, euro area

JEL Classification: C11, C32, E44, F51, Q43

Contents

List of Figures

1 INTRODUCTION

Geopolitical risk refers to the risk associated with wars, terrorism, political tensions, and other adverse geopolitical events that can affect economic activity and financial conditions. In recent years, it has become a central concern for macro-financial stability in the euro area. Since 2022, successive episodes of geopolitical tension such as Russia's invasion of Ukraine, renewed tensions in the Middle East, and growing friction around Taiwan have coincided with disruptions to energy markets, inflationary pressures, and changes in financial conditions. The quantitative measurement of geopolitical risk has been substantially advanced by Caldara & Iacoviello (2022), who construct a newspaper-based Geopolitical Risk (GPR) index designed to capture fluctuations in geopolitical tensions over time. More recent extensions develop bilateral GPR indices that disaggregate geopolitical risk according to the geographical origin of the underlying event (Alonso-Alvarez et al. 2025, 2026).

The existing literature has extensively examined the economic consequences of aggregate geopolitical risk, but the geographical dimension of these shocks remains considerably less explored. These bilateral GPR indices make it possible to move beyond the assumption that geopolitical risk is a homogeneous global shock. This distinction is particularly relevant for the euro area, where exposure to geopolitical events differs according to trade relationships, energy dependence, financial linkages, and the nature of the underlying event. This raises the central research question of this thesis: does the geographical origin of geopolitical risk affect how geopolitical shocks are transmitted through energy prices and macro-financial variables across Germany, France, Italy, and Spain?

Several interconnected mechanisms suggest that this transmission may differ across geographical origins and recipient economies. Energy production and transportation are geographically concentrated and can be disrupted by conflicts, sanctions, or political instability. Geopolitical events may therefore alter expectations about energy supply, demand, and risk premia, producing different responses in oil and natural gas markets. Changes in energy prices can subsequently affect production costs, consumer prices, and economic activity. At the same time, greater uncertainty may influence consumption, investment, trade, and financial conditions. The relative importance of these channels is unlikely to be identical across geopolitical origins. In particular, the direct exposure of European economies to Russian energy supply provides a transmission mechanism that differs from those associated with China, the Middle East and North Africa (MENA), or the Western Bloc.

To investigate this question, the analysis considers Germany, France, Italy, and Spain and estimates a separate quarterly Bayesian vector autoregressive (BVAR) model for each country-origin combination, producing 16 systems in total. Each system traces the response of Brent crude oil prices, European natural gas prices, HICP inflation, industrial production, and the Country-Level Index of Financial Stress (CLIFS) to an origin-specific GPR shock. Estimating the systems separately makes it possible to compare how the same recipient economy responds to shocks from different origins and how shocks from the same origin are transmitted across different economies. The empirical specification also accounts for the exceptional volatility observed during the COVID-19 period, and additional exercises assess the robustness of the main findings. A complementary statistical analysis evaluates whether statistically detectable responses are unusually concentrated in any particular geographical origin.

The results indicate that the geographical origin of geopolitical risk matters for its transmission. Natural gas price growth falls on impact across all 16 country-origin combinations. However, since the model uses quarterly observations, the impact response captures contemporaneous adjustment only at the quarterly frequency and cannot identify the sequence of price movements within the quarter or the immediate reaction to a particular geopolitical event. Brent responses display considerably greater origin-specific heterogeneity: they are generally positive following Chinese and Russian shocks and negative following MENA shocks. Inflationary effects are frequent, particularly after Russian-origin shocks, whereas industrial-production responses differ substantially across countries. Financial-stress responses also vary across horizons, highlighting the difference between short-run market reactions and subsequent adjustment.

Russian-origin shocks generate the most widespread responses across the variables and economies considered. This pattern is consistent with the importance of trade and energy linkages between Russia and European economies, although the models do not separately identify the mechanisms responsible for it. Moreover, this result is not equally robust across alternative model specifications and should therefore be interpreted cautiously. More broadly, the findings show that treating geopolitical risk as a single aggregate shock can conceal economically relevant differences across geographical origins, transmission variables, and recipient economies.

The rest of the thesis is structured as follows. Section 2 reviews the literature on the measurement of geopolitical risk, its macro-financial and energy-market transmission, and the econometric methods relevant to the analysis. Section 3 describes the data sources, seasonal adjustment procedures, transformations, and stationarity properties of the variables. Section 4 presents the reduced-form VAR framework, the Bayesian specification, the identification strategy, the prior structure, the treatment of COVID-19 volatility, and the model-validation exercises. Section 5 presents the empirical results and evaluates the heterogeneity of geopolitical risk transmission across origins and countries. Section 6 concludes by discussing the main findings, limitations, and implications for future research. Appendix 7 provides the complete set of structural impulse responses, additional robustness analyses, and supplementary diagnostic and validation results.

2 LITERATURE REVIEW

2.1 Measurement of Geopolitical Risk

The empirical study of geopolitical risk has been shaped by the development of quantitative indicators that transform geopolitical events into observable time series. Caldara & Iacoviello (2022) established the foundational framework for this literature by constructing a newspaper-based Geopolitical Risk (GPR) index based on the share of newspaper articles containing terms related to geopolitical tensions, threats, wars, and terrorism. The index is designed to capture the intensity of geopolitical risk as reflected in contemporaneous news coverage, providing a systematic time-series measure that can be incorporated into macroeconomic and financial models. Using this measure, the authors show that higher geopolitical risk is associated with lower investment and employment, as well as greater downside risks to economic activity.

A limitation of aggregate GPR measures is that they do not identify where the underlying geopolitical tension originates. This distinction has become increasingly relevant as the literature has shown that geopolitical risk is not necessarily perceived uniformly across countries. Bondarenko et al. (2024), using Russia as a case study, construct indicators from local newspaper sources and show that shocks identified from Russian news have significant adverse effects on the Russian economy, whereas shocks measured using English-language newspapers do not produce the same effects. Their results indicate that the economic relevance of geopolitical risk depends partly on the perspective from which it is measured.

The geographical dimension is developed further by Alonso-Alvarez et al. (2025), who extend the GPR framework by distinguishing between the country from whose newspapers geopolitical risk is measured and the geographical source to which the underlying geopolitical event is attributed. This bilateral structure makes it possible to separate, within the same recipient economy, geopolitical risk associated with different countries or regions. Importantly, their VAR evidence shows that the geographical origin of a GPR shock can affect both the intensity and the sign of its estimated effect on GDP. Alonso-Alvarez et al. (2026) further develop this framework into a database covering 34 countries and four common origins of geopolitical risk: Russia, China, the Middle East and North Africa (MENA), and the Western Bloc. This common classification allows shocks originating in different regions to be compared using the same measurement methodology across countries.

2.2 Macroeconomic and Financial Transmission of Geopolitical Risk

The evidence based on aggregate GPR indicators generally points to adverse real effects following increases in geopolitical risk. Caldara & Iacoviello (2022) find that higher geopolitical risk is followed by reductions in investment and employment, with the negative effects concentrated particularly among firms and industries more exposed to geopolitical developments. This establishes an important transmission from geopolitical uncertainty to real economic activity.

For the euro area, Bondarenko et al. (2026) show that the measurement of geopolitical risk from a European perspective is important for estimating its macroeconomic consequences. Using an index constructed from local European news sources and a Bayesian VAR, they find that geopolitical risk shocks have recessionary and inflationary effects in the euro area. These effects are understated when geopolitical risk is instead measured using an index based predominantly on an Anglosphere perspective. Therefore, their results connect the measurement of geopolitical risk directly with the magnitude of its estimated macroeconomic transmission.

Geopolitical risk can also propagate through financial markets and financial intermediaries. Dieckelmann et al. (2025) construct an exposure-weighted geopolitical risk indicator for euro-area banks and find that a one-standard-deviation increase in geopolitical risk is associated with an increase of approximately 34 basis points in bank CDS spreads and a decline of around 6% in bank stock prices. This provides direct evidence that geopolitical tensions can affect financial conditions through the perceived risk and valuation of financial institutions. Complementary evidence is provided by Niepmann & Shen (2025), who use U.S. supervisory data to show that banks reduce direct cross-border lending to countries experiencing higher geopolitical risk and that increases in geopolitical risk abroad can also reduce domestic lending, particularly among banks operating foreign affiliates. These findings indicate that the effects of geopolitical risk can extend beyond asset prices and propagate through the supply of credit.

Taken together, these studies document that geopolitical risk can affect real activity, inflation, and financial conditions through several channels. More recent contributions also show that locally constructed and bilateral measures can provide information that is not captured by aggregate GPR indices. However, much of the existing literature either relies on broad aggregate measures or focuses on specific parts of the transmission process. This leaves scope for analysing macro-financial responses jointly while distinguishing the geographical origin of the geopolitical shock.

2.3 Energy Markets and Geographical Heterogeneity

Energy markets provide a particularly relevant channel through which geopolitical risk can affect macroeconomic conditions. Abid et al. (2023) compare the response of five commodity classes to geopolitical risk using a Markov-Switching framework and find that energy is the most responsive commodity category. This result supports giving energy prices a central role in the analysis of geopolitical-risk transmission.

The relationship between geopolitical risk and energy prices is nevertheless not uniform. Ferrari Minesso et al. (2023) show that the response of Brent oil prices depends on the geographical origin of the geopolitical shock. Their results indicate that geopolitical tensions often put downward pressure on oil prices through weaker expected global demand, whereas shocks associated with particular countries, including major actors in global energy and geopolitical markets, can instead increase oil prices when concerns about future supply disruptions dominate. Consequently, the sign of the energy-price response cannot be inferred from an increase in geopolitical risk alone, since it depends on the relative strength of demand and supply channels. This demand-supply ambiguity provides a useful framework for interpreting variation in the response of energy prices to geopolitical shocks.

This distinction is important for empirical interpretation because an observed energy-price response does not by itself identify the underlying transmission mechanism. Similar price movements may reflect different combinations of expected supply disruptions, changes in global demand, or broader uncertainty effects. Structural interpretations therefore require caution when these channels are not separately identified.

Recent evidence further demonstrates why this distinction matters for macroeconomic transmission. Verduzco-Bustos & Zanetti (2026) identify geopolitical oil price shocks associated with major geopolitical tensions and find that their cross-border effects on output and inflation are stronger in oil-importing economies and in countries with greater dependence on foreign energy supply. This result is particularly relevant for European economies, where external energy dependence can amplify the transmission of geopolitical disturbances.

The literature therefore points to two related dimensions of heterogeneity. First, the estimated effects of geopolitical risk can differ when risk is measured from different national perspectives or disaggregated by geographical origin (Bondarenko et al. 2024, Alonso-Alvarez et al. 2025). Second, energy-price responses can vary with the nature and geographical origin of the geopolitical disturbance (Ferrari Minesso et al. 2023). However, within the literature reviewed here, evidence jointly examining the energy-price and macro-financial responses to geopolitical risk from several geographical origins across individual euro-area economies remains limited.

This thesis addresses this gap by estimating comparable systems for Germany, France, Italy, and Spain for four geopolitical origins (Russia, China, MENA, and the Western Bloc). Therefore, the main goal is not to establish for the first time that geographical heterogeneity exists, since previous studies already provide evidence in this direction, but to examine whether this heterogeneity is also present across energy-price and macro-financial responses and whether it differs across euro-area economies.

2.4 VAR and BVAR Approaches to Geopolitical Risk Transmission

Vector autoregressive models provide a widely used empirical framework for estimating the dynamic consequences of geopolitical shocks. Since Sims (1980), VAR models have been widely used to analyse the joint dynamics of macroeconomic variables without imposing a fully specified structural model. In the geopolitical-risk literature, Caldara & Iacoviello (2022) estimate a Bayesian VAR and use posterior impulse responses to analyse the macroeconomic consequences of GPR shocks, while Alonso-Alvarez et al. (2025) employ standard VAR models to assess how the geographical origin of GPR shocks affects the magnitude and sign of their macroeconomic effects.

Bayesian estimation becomes particularly useful when the number of VAR parameters is large relative to the available sample. The Minnesota-prior approach associated with Litterman (1986) introduces shrinkage towards parsimonious prior dynamics, reducing estimation uncertainty while preserving interactions across variables. Bańbura et al. (2010) show that Bayesian shrinkage allows relatively large VAR systems to retain useful forecasting and structural properties when unrestricted estimation would otherwise suffer from parameter proliferation.

A closely related application is provided by Bondarenko et al. (2026), who estimate a monthly Bayesian VAR for the euro area using a conjugate Normal-Inverse-Wishart prior with Minnesota-type shrinkage. They relax the standard random-walk prior for the GPR indicator by setting its own first-lag prior mean to 0.8 and apply a truncated COVID-19 volatility correction to limit the influence of pandemic outliers. Geopolitical shocks are identified recursively with the euro-area GPR indicator ordered first, and posterior median impulse responses are reported together with 68% credible intervals.

The combination of Minnesota-type shrinkage, a GPR-specific prior mean, recursive GPR-first identification, and an explicit treatment of COVID-19 volatility makes their framework a particularly close methodological reference for the present analysis. However, this thesis differs by using quarterly bilateral GPR measures and estimating separate systems for each destination country and geographical origin rather than a single euro-area geopolitical risk shock.

A further methodological issue concerns the identification of geopolitical shocks. Reduced-form VAR innovations are generally contemporaneously correlated, so structural interpretation requires additional identifying restrictions. Recursive identification provides a transparent solution but makes the estimated impulse responses conditional on the assumed contemporaneous ordering (Kilian & Lütkepohl 2017). In applications where GPR is ordered first, the identifying assumption implies that macroeconomic and energy-market innovations do not affect the GPR indicator within the same observation period. This assumption is particularly relevant at quarterly frequency, where geopolitical developments and market responses may occur within the same quarter.

The treatment of persistent macroeconomic and energy variables also matters for VAR specification. Transforming persistent series into growth rates facilitates stationary modelling and focuses the analysis on short- and medium-run dynamics, but differencing may discard information about long-run relationships when variables are cointegrated (Engle & Granger 1987, Johansen 1991). This creates a trade-off between stationary short-run specifications and models that explicitly preserve long-run equilibrium relationships.

Building on this literature, the present thesis applies a Bayesian VAR framework to 16 country-origin systems, combining Minnesota-type shrinkage with recursive identification and an explicit treatment of the exceptional volatility associated with the COVID-19 period. The latter is motivated by Lenza & Primiceri (2022), who show that the sequence of extreme observations generated by the pandemic can materially affect VAR estimation if it is not explicitly accommodated. The empirical design therefore combines recent advances in the geographical measurement of geopolitical risk with the VAR and BVAR methods used to study its macroeconomic transmission. Relative to closely related euro-area applications, the analysis places particular emphasis on heterogeneity across both geopolitical origins and recipient economies. Finally, the analysis complements the country-origin Bayesian VAR results with an exact permutation test designed to assess whether detectable responses are unusually concentrated in particular geopolitical origins.

3 DATA

This section describes the data used to estimate the Bayesian VAR systems: the bilateral geopolitical risk index, energy prices, and the three macro-financial variables (HICP, industrial production, and the Country-Level Index of Financial Stress (CLIFS)), together with the transformations applied to each series before estimation.

The analysis focuses on Germany, France, Italy, and Spain, four large euro-area economies with contrasting exposure to the energy transmission channel under study. Germany has historically exhibited substantial dependence on Russian energy, France has a distinct energy mix, and Italy and Spain combine Mediterranean exposure with separate bilateral GPR coverage provided by Alonso-Alvarez et al. (2026). This country-level disaggregation extends the euro-area aggregate approach of Bondarenko et al. (2026) by examining whether transmission varies across recipient economies as well as across geographical origins.

3.1 Bilateral Geopolitical Risk Index

The geopolitical risk index used in this thesis is the bilateral extension of the original Caldara & Iacoviello (2022) Geopolitical Risk (GPR) index. The bilateral index is constructed by Alonso-Alvarez et al. (2025) through a two-stage filter applied to a wide range of newspaper articles. In the first stage, an article is counted only if it contains at least one term from a lexical dictionary composed exclusively of conflict-, threat-, and war-related vocabulary (war, coup, terrorism, military mobilisation, ultimatum). This dictionary contains no price-, market-, or energy-related terms, a feature documented in detail in Appendix A.2 Queries of Alonso-Alvarez et al. (2026). In the second stage, each article already admitted by the lexical filter is classified by the geographic region it refers to, using a proprietary regional classification filter. This second stage is a geographic partition only, i.e., it plays no role in filtering the economic content of the article.

The original working paper, Alonso-Alvarez et al. (2025), covers only six narrator countries (the United States, China, Russia, the United Kingdom, Germany, and France) and does not include Spain or Italy. The companion paper, Alonso-Alvarez et al. (2026), extends coverage to 34 countries, including Spain (from 1996) and Italy (from 1997). Therefore, both papers must be cited jointly for this thesis: Alonso-Alvarez et al. (2025) for the underlying bilateral methodology and the German and French series, and Alonso-Alvarez et al. (2026) for the extended database that additionally provides the Spanish and Italian series.

This thesis focuses on four geopolitical risk origins: Russia, China, the Middle East and North Africa (MENA), and the Western Bloc. The Western Bloc is defined by Alonso-Alvarez et al. (2026) (Section 2.5) as Europe, the United States, Canada, Australia, and New Zealand, with the receiving country's own territory always excluded from its own focal region. As a consequence, when the origin variable is constructed for a given recipient country, the Western Bloc aggregate for that country already excludes that country itself, so that, for instance, the Western Bloc origin index for Germany reflects geopolitical tension attributed to France, Spain, Italy, the United States, and the rest of the bloc, but not to Germany.

3.2 Energy Prices

Brent crude oil and European natural gas prices are obtained from the World Bank Pink Sheet and are expressed in nominal US dollars. Retaining the original quotation currency provides a direct measure of movements in international commodity prices without combining them with fluctuations in the exchange rate. Brent and European natural gas prices are common across countries and therefore enter identically in the four national systems.

3.3 HICP, Industrial Production, and Financial Stress

The Harmonised Index of Consumer Prices (HICP), all-items headline, is obtained from Eurostat. Industrial production (IPI) is also obtained from Eurostat. The Country-Level Index of Financial Stress (CLIFS) is obtained from the European Central Bank and follows the methodology of Duprey et al. (2017), who construct a country-level composite indicator of systemic financial stress from equity, sovereign-bond, and foreign-exchange markets.

Industrial production index (IPI) is used as the measure of real economic activity rather than GDP because it is more directly connected to production processes exposed to energy-price fluctuations and provides a relatively sensitive indicator of short-run changes in economic activity.

3.3.1 Seasonal Adjustment of HICP

Eurostat's HICP series is not seasonally adjusted by design, as confirmed directly in the dataset's official metadata (Eurostat 2026). An unadjusted seasonal pattern of this kind introduces a systematic and calendar-driven component into the quarterly log-differences used in the (B)VAR, which is unrelated to any economic shock and can distort the estimated dynamics. This is particularly severe for Spain and Italy, whose raw quarterly seasonal spread is approximately 2.5 times larger than that of Germany and France.

Publicly available seasonally adjusted HICP series do not exist at the national level for all four economies in this sample. The European Central Bank publishes a seasonally adjusted HICP series, but only for the euro-area aggregate, not disaggregated by member state (European Central Bank 2026). The OECD publishes country-level seasonally adjusted CPI series for a restricted set of countries (Australia, Canada, France, Germany, Japan, and the United States), explicitly excluding Spain and Italy from this coverage (OECD 2024).

Given this gap, seasonal adjustment was applied directly to the monthly HICP series for each of the four countries using the `RJDemetra` interface to the JDemetra+ engine, before aggregating the data to a quarterly frequency. JDemetra+ is the seasonal adjustment software officially recommended to the members of the European Statistical System and the European System of Central Banks since February 2015, and implements two adjustment methods: X-13ARIMA-SEATS (US Census Bureau) and TRAMO-SEATS (Salgado 2019), the latter being Eurostat's default choice within the software.

Both methods were applied to all four countries and compared empirically before selecting a final specification. TRAMO-SEATS showed a consistent and moderate advantage over X-13ARIMA-SEATS: residual-seasonality tests (the QS test and the F-test on seasonal dummies applied to the adjusted series) were equal or better for TRAMO-SEATS in all four countries, and the average reduction in the quarterly seasonal spread was higher (94.95% versus 92.22%). The AIC comparison was evenly split between the two methods, with each method preferred in two countries, while the Bayesian Information Criterion (BIC) was not directly comparable between the two implementations because the reported values were on inconsistent scales. Based on this evidence and consistent with Eurostat's own default choice, TRAMO-SEATS was adopted as the final seasonal adjustment method. The complete diagnostic comparison between both methods, including the X-13ARIMA-SEATS results, is reported in Appendix 7.1.

Table 1 reports the residual-seasonality diagnostics and the reduction in the quarterly seasonal spread achieved by TRAMO-SEATS for each country. High p-values in the residual-seasonality tests (close to 1) indicate that no seasonality remains detectable in the adjusted series.

Table 1: HICP seasonal adjustment quality diagnostics (TRAMO-SEATS), by country

Country	QS test (p)	F-test, seas. dummies (p)	Residual seasonality (p)	Spread reduction
Germany	1.000	1.000	1.000	88.4%
France	0.862	1.000	1.000	99.8%
Italy	1.000	1.000	1.000	94.9%
Spain	1.000	1.000	1.000	96.7%

All four countries show residual-seasonality p-values at or very close to 1, providing no evidence of a detectable residual seasonal pattern after adjustment. The reduction in the quarterly seasonal spread ranges from 88.4% in Germany to 99.8% in France, with Italy and Spain recording reductions of 94.9% and 96.7%, respectively. A minor residual trading-day peak was detected for Italy and Spain under the X-13ARIMA-SEATS specification (Appendix 7.1); this is not present in the diagnostics of the TRAMO-SEATS specification ultimately adopted.

3.4 Transformations and Stationarity

The variables are transformed to reduce skewness, achieve stationarity where appropriate, and ensure that the estimated impulse responses have a meaningful interpretation. Brent crude-oil prices, European natural gas prices, HICP, and industrial production enter the baseline BVAR in quarterly log-differences:

$$\Delta \log(X_t) = \log(X_t) - \log(X_{t-1}). \tag{1}$$

The corresponding impulse responses therefore represent changes in quarterly price growth, inflation, or industrial-production growth rather than changes in the levels of these variables. In the descriptive statistics, these log-differences are multiplied by 100 to facilitate interpretation as approximate quarterly percentage changes. This rescaling is used only for presentation; the BVAR is estimated using the unscaled log-differences. CLIFS is retained in levels because the stationarity tests provide support for this specification and because its level has a direct interpretation as the degree of systemic financial stress.

The bilateral GPR indices enter the baseline model as

$$\log(1 + GPR_t). \tag{2}$$

Adding one permits the transformation of zero observations, while the logarithmic specification reduces the pronounced right skewness of the indices and expresses their variation on a proportional rather than an absolute scale. The transformed index is retained in levels so that the identified innovation represents an unexpected increase in geopolitical risk. Differencing the index would instead identify an innovation to the quarterly change in geopolitical risk and would therefore answer a different empirical question.

Stationarity is evaluated using the Augmented Dickey–Fuller (ADF) and Kwiatkowski–Phillips–Schmidt–Shin (KPSS) tests. The results support the use of log-differences for Brent, Gas, HICP, and IPI and the use of levels for CLIFS. The evidence for the bilateral GPR indices is less clear. Conventional unit-root tests can have limited power when structural breaks are ignored (Perron 1989). This issue is particularly relevant for geopolitical-risk indices, which exhibit abrupt shifts around major geopolitical events.

To examine this issue, structural breaks in the transformed GPR indices are detected using the PELT algorithm, after which ADF tests are applied separately within the resulting segments. Every segment rejects the unit-root null in 8 of the 16 country-origin series. In the remaining cases, the failures to reject are concentrated mainly in short final segments beginning around the major geopolitical events of 2022, for which the statistical power of the tests is necessarily limited. Consequently, the results do not establish uniform piecewise stationarity across all bilateral GPR series. They nevertheless indicate that much of the apparent persistence is associated with discrete shifts and exceptional geopolitical episodes rather than with a smooth stochastic trend.

The baseline specification therefore retains $\log(1 + \text{GPR})$ in levels, while recognising the uncertainty surrounding its time-series properties. The sensitivity of the results to the treatment of the persistent energy and macroeconomic variables is assessed through an alternative levels specification, discussed in Section 5.3.

3.5 Descriptive Statistics

The purpose of these descriptive statistics is to characterise the scale, dispersion, and range of the variables used in the empirical analysis. They provide context for interpreting the subsequent impulse responses, identify the exceptional volatility present in some series, and document differences in the distributions of the bilateral GPR indices across geographical origins.

Table 2 summarises the transformed energy and macro-financial variables, whereas Table 3 reports the corresponding statistics for the transformed bilateral GPR indices. Both tables are based on the final aligned sample, which runs from 1997Q2 to 2025Q4 and contains 115 quarterly observations for each country. To make their magnitudes easier to interpret, log-differences are multiplied by 100 in Table 2 and can therefore be read as approximate quarterly percentage changes. CLIFS enters in levels, while the bilateral GPR indices are transformed using $\log(1 + \text{GPR})$. Brent and European natural gas prices are common across countries and are consequently reported only once.

Table 2: Descriptive statistics of transformed energy and macro-financial variables

	Germany	France	Italy	Spain
Panel A: Mean and standard deviation				
$100 \times \Delta \log$(HICP)	0.4799 (0.5687)	0.4350 (0.4182)	0.5092 (0.6014)	0.5801 (0.5744)
$100 \times \Delta \log$(IPI)	0.1679 (3.0375)	0.0674 (3.0943)	-0.1509 (3.5969)	-0.0011 (3.5185)
CLIFS	0.1141 (0.0850)	0.1177 (0.0795)	0.0947 (0.0771)	0.1138 (0.0853)
$100 \times \Delta \log$(Brent)	0.9574 (15.6142)			
$100 \times \Delta \log$(Gas)	1.1081 (21.5351)			
Panel B: Minimum and maximum				
$100 \times \Delta \log$(HICP)	[-0.4588, 2.9985]	[-0.3925, 1.9877]	[-0.3247, 4.7953]	[-0.6365, 2.7561]
$100 \times \Delta \log$(IPI)	[-20.0378, 13.4084]	[-21.1938, 19.6799]	[-18.6119, 25.5731]	[-23.0005, 22.4863]
CLIFS	[0.0135, 0.4644]	[0.0312, 0.4422]	[0.0187, 0.4468]	[0.0200, 0.4717]
$100 \times \Delta \log$(Brent)	[-72.6752, 31.4809]			
$100 \times \Delta \log$(Gas)	[-78.4979, 65.5655]			

Energy-price growth is substantially more volatile than HICP inflation and industrial-production growth. Natural gas displays both a higher standard deviation and a wider range than Brent, reflecting the particularly large fluctuations observed in the European gas market. Quarterly HICP inflation averages between approximately 0.44% and 0.58% across the four countries and exhibits relatively limited dispersion. Industrial-production growth is centred close to zero but is considerably more volatile, with large negative and positive observations consistent with the exceptional contraction and subsequent rebound surrounding the COVID-19 period. Average CLIFS levels and their variability are broadly comparable across the four economies.

Table 3: Descriptive statistics of transformed bilateral GPR indices

Recipient country	Russia	China	MENA	Western Bloc
Panel A: Mean and standard deviation				
Germany	3.7785 (1.2525)	3.9310 (1.2838)	4.4997 (0.5660)	4.4756 (0.6624)
France	4.0332 (1.1802)	4.3809 (0.9242)	4.4148 (0.7330)	4.4558 (0.5804)
Italy	3.9631 (1.0827)	4.1610 (1.0843)	4.4903 (0.5640)	4.5349 (0.5207)
Spain	3.9909 (1.0304)	4.2863 (0.9171)	4.4892 (0.5405)	4.5617 (0.4203)
Panel B: Minimum and maximum				
Germany	[1.8491, 7.4053]	[0.0000, 6.4223]	[3.1579, 6.0313]	[3.4741, 6.0683]
France	[2.4593, 7.4599]	[1.9050, 6.1096]	[2.7387, 6.2897]	[3.5887, 5.7079]
Italy	[2.0765, 7.3193]	[0.0000, 6.4964]	[3.2079, 5.9556]	[3.3002, 5.7193]
Spain	[2.0276, 7.3715]	[1.7103, 6.1876]	[3.5106, 6.2454]	[3.7876, 6.0740]

The transformed bilateral GPR indices display substantial heterogeneity across geographical origins. Russia- and China-origin indices generally exhibit the greatest variability, while Russian GPR reaches the highest maximum in each of the four countries. By contrast, the MENA and Western Bloc indices tend to have higher sample means but lower dispersion. The zero minima recorded for the Chinese index in Germany and Italy correspond to observations in which the original GPR index equals zero; these observations are preserved by the log(1+GPR) transformation. Overall, the differences in the distributions of the bilateral indices further motivate analysing geopolitical-risk transmission separately by geographical origin.

4 MODEL STRUCTURE

4.1 BVAR Specification

The empirical analysis is based on a six-variable BVAR, estimated separately for 16 country-origin systems (four destination economies: Germany, France, Italy, and Spain, each paired with four geopolitical risk origins: Russia, China, MENA, and the Western Bloc). The vector of endogenous variables is

$$Y_{i,t} = \begin{bmatrix} \log(1 + \mathrm{GPR}_{\mathrm{origin},i,t}) \\ \Delta \log(\mathrm{Brent}_t) \\ \Delta \log(\mathrm{Gas}_t) \\ \Delta \log(\mathrm{HICP}_{i,t}) \\ \Delta \log(\mathrm{IPI}_{i,t}) \\ \mathrm{CLIFS}_{i,t} \end{bmatrix}$$

All systems are estimated at quarterly frequency with two lags ($p = 2$). The dataset contains 115 observations per system after accounting for missing values, leaving an effective sample size of 113 observations for estimation after constructing the two lags. This applies identically across all 16 country-origin systems.

4.2 Levels versus Differences

Brent, Gas, HICP, and IPI enter the baseline BVAR in log-differences rather than in log-levels (Section 3.4), following the stationarity evidence discussed in the data section. This differs from the levels specification used in closely related studies (Caldara & Iacoviello 2022, Bondarenko et al. 2026).

Because differencing may discard long-run information when variables are cointegrated (Engle & Granger 1987, Johansen 1991), Johansen tests were applied to the four corresponding log-level variables. Both the trace and maximum-eigenvalue tests identify one cointegrating relationship at the 5% level in Germany, France, Italy, and Spain. Thus, while differencing provides a stationary specification, it does not explicitly model the long-run equilibrium relationship among these variables.

The sensitivity of the results to this choice is assessed by re-estimating all 16 BVAR systems with Brent, Gas, HICP, and IPI in log-levels. Across 240 variable-horizon comparisons, 74.2% of posterior median responses retain the same sign and 87 responses change their 68% significance status. Gas is particularly sensitive, with sign agreement in 28 of 48 comparisons. The geographical heterogeneity result is also not preserved: the exact permutation-test p-value for Russian GPR shocks increases from 0.0273 in the baseline to 0.7852 in the levels specification.

The differenced specification is retained as the baseline because it follows the stationary transformations adopted for the main analysis and provides a consistent specification across all country-origin systems. However, the robustness exercise shows that the treatment of persistent variables is a substantive modelling choice. Full results are reported in Appendix 7.5.

4.3 Treatment of COVID-19 Volatility

The COVID-19 pandemic generated exceptionally large observations in macroeconomic and financial series, particularly during 2020. Estimating a VAR under the assumption of constant residual variance may cause these observations to exert a disproportionate influence on the estimated dynamics of the system. Motivated by the temporary-volatility treatment of Lenza & Primiceri (2022), and its subsequent application in the geopolitical-risk BVAR literature by Bondarenko et al. (2026), the baseline specification therefore allows for a temporary increase in residual volatility during the pandemic period.

Specifically, the reduced-form innovations are assumed to satisfy

$$u_t \sim N\left(0, s_t^2 \Sigma\right),$$

where Σ denotes the baseline covariance matrix and s_t is a common time-varying volatility factor. Outside the COVID-19 episode, $s_t = 1$. During the pandemic, the volatility path is parameterised by

$$\theta = (s_0, s_1, s_2, \rho),$$

where s_0, s_1, and s_2 capture the exceptional volatility in 2020Q1, 2020Q2, and 2020Q3, respectively, while ρ governs the subsequent return towards the baseline variance. The resulting path is

$$s_t = \begin{cases} s_0, & t = 2020\text{Q1}, \\ s_1, & t = 2020\text{Q2}, \\ s_2, & t = 2020\text{Q3}, \\ 1 + (s_2 - 1)\rho^j, & t = 2020\text{Q4}, \dots, 2021\text{Q3}, \\ 1, & t \geq 2021\text{Q4}, \end{cases}$$

where $j = 1, \dots, 4$ indexes the quarters following 2020Q3.

The volatility parameters are estimated separately for each country-origin system in a first stage by maximising the concentrated Gaussian likelihood. Denoting the resulting estimates by $\hat{\theta}$, the corresponding volatility path $\hat{s}_t$ is subsequently treated as fixed when estimating the VAR and BVAR. Conditional on this path, estimation can be written as a weighted system in which both the dependent variables and regressors are scaled by the estimated volatility factor:

$$Y_t^* = \frac{Y_t}{\hat{s}_t}, \qquad X_t^* = \frac{X_t}{\hat{s}_t}.$$

Observations associated with exceptionally high pandemic volatility therefore receive less effective weight when estimating the underlying VAR dynamics, while observations outside the pandemic period are left unchanged.

This constitutes a two-stage treatment of pandemic volatility rather than a fully joint Bayesian estimation. In particular, posterior uncertainty in the BVAR coefficients and covariance matrix is evaluated conditional on $\hat{\theta}$, so that first-stage uncertainty surrounding the COVID-volatility parameters is not propagated into the posterior impulse responses.

4.4 Structural Identification: Why GPR Is Ordered First

Structural impulse responses are identified recursively through a Cholesky decomposition of the reduced-form covariance matrix. The baseline ordering is

$$[\log(1 + \text{GPR}),\ \Delta\log(\text{Brent}),\ \Delta\log(\text{Gas}),\ \Delta\log(\text{HICP}),\ \Delta\log(\text{IPI}),\ \text{CLIFS}].$$

The central identifying restriction is that GPR is ordered first. This implies that geopolitical risk does not respond contemporaneously to shocks in Brent, Gas, HICP, IPI, or CLIFS, whereas all five variables are allowed to respond contemporaneously to a GPR shock.

This restriction is not directly testable: the reduced-form covariance matrix of a VAR can be orthogonalised under alternative recursive orderings (Sims 1980). The GPR-first ordering is therefore treated as an identifying assumption whose plausibility rests on economic reasoning, the construction of the GPR measure, and empirical robustness checks.

For HICP, IPI, and CLIFS, reverse contemporaneous causality appears less plausible than the opposite direction. Geopolitical events and the associated news coverage can affect economic activity, prices, and financial stress within the quarter, whereas movements in these macro-financial variables are less likely to generate geopolitical events contemporaneously. The interpretation of the GPR index as a noisy proxy for an underlying geopolitical shock is also conceptually related to the internal-instrument result of Plagborg-Møller & Wolf (2021), who show that when an observed proxy satisfies the required instrumental-variable conditions, ordering it first in a recursive VAR can recover relative structural impulse responses even under non-invertibility and independent measurement error. This result does not establish that the GPR index satisfies these validity conditions, but it provides a useful methodological connection to specifications in which an observed proxy for the shock of interest is placed first in the system.

The energy variables, Brent and Gas, require more detailed consideration because contemporaneous feedback is more plausible in energy markets. Geopolitical developments and energy-price adjustments may occur within the same quarter, making the direction of the contemporaneous relationship more difficult to isolate, particularly for Russia-related GPR shocks. Temporal aggregation reinforces this concern. Granger-causality relationships are generally not invariant to temporal aggregation (Marcellino 1999), and aggregation can induce contemporaneous correlation even when it is absent at the underlying higher frequency (Breitung & Swanson 2002). The quarterly specification therefore cannot by itself rule out within-quarter feedback between geopolitical risk and energy prices.

Within the energy block, Brent is ordered before Gas. This choice reflects the strong historical integration between European natural gas and crude oil prices documented by Asche et al. (2013). More specifically, Lin & Li (2015) find that European gas prices are cointegrated with Brent and identify price spillovers running from crude oil to natural gas markets, without evidence of the reverse price-spillover relationship.

A further concern is mechanical contamination of the GPR measure by energy-market reporting. The construction of the bilateral GPR index mitigates this possibility because articles are selected using conflict-, threat-, and war-related vocabulary rather than price-, market-, or energy-related terms (Alonso-Alvarez et al. 2026). Energy-price movements alone are therefore insufficient for an article to enter the GPR count, although the construction cannot exclude articles discussing energy-market developments in a geopolitical context.

Lead-lag correlations between GPR and energy prices are used to provide descriptive evidence on their temporal relationships. In addition, the robustness analysis reverses the relative ordering of Brent and Gas while keeping GPR first and considers a stronger energy-first specification in which both energy variables are ordered before GPR. The corresponding diagnostics and results are reported in Appendix 7.3.

4.5 Reduced-Form VAR Framework

The reduced-form VAR provides the dynamic framework used in the subsequent BVAR specification. The VAR and BVAR share the same variables, lag structure, COVID-19 volatility adjustment, and recursive identification scheme. The classical estimation described below is therefore used to define the corresponding reduced-form quantities and to provide a frequentist benchmark for model validation, rather than as a separate empirical specification in the main results.

In reduced form, the model is written as

$$Y_t = c + A_1 Y_{t-1} + A_2 Y_{t-2} + u_t, \qquad u_t \sim N(0, s_t^2 \Sigma).$$

As described in Section 4.3, the COVID-volatility parameters are first estimated by concentrated maximum likelihood and the resulting path $\hat{s}_t$ is subsequently treated as fixed. Conditional on this path, the reduced-form VAR can be estimated by weighted least squares by scaling the dependent variables and regressors as

$$Y_t^* = \frac{Y_t}{\hat{s}_t}, \qquad X_t^* = \frac{X_t}{\hat{s}_t},$$

where $X_t = [1, Y_{t-1}', Y_{t-2}']'$. In stacked matrix form,

$$Y^* = X^* B + U^*,$$

with

$$\hat{B} = (X^{*\prime} X^*)^{-1} X^{*\prime} Y^*, \qquad \hat{\Sigma} = \frac{\hat{U}^{*\prime} \hat{U}^*}{T_{\text{eff}}},$$

where $T_{\text{eff}} = T - p = 113$ in each country-origin system. The weighting reduces the influence of observations associated with exceptional pandemic volatility while leaving observations outside the COVID-19 volatility episode unchanged.

Structural impulse responses are obtained from the reduced-form dynamics through the Wold recursion

$$\Psi_0 = I, \qquad \Psi_j = \sum_{i=1}^{\min(j,p)} A_i \Psi_{j-i},$$

combined with the recursive identification described in Section 4.4. For the classical VAR, the reduced-form covariance matrix is factorised as

$$\hat{\Sigma} = QQ',$$

where Q is the lower-triangular Cholesky factor consistent with the recursive ordering in Section 4.4. The structural impulse-response matrices are then

$$\Theta_j = \Psi_j Q.$$

The reported GPR impulse responses correspond to a one-standard-deviation innovation in the orthogonalised GPR shock, represented by the first column of Θ_j under the recursive ordering.

The BVAR uses the same recursion and Cholesky identification for each posterior draw of the VAR parameters. The difference is therefore in the estimation and treatment of parameter uncertainty rather than in the structural identification of the shocks.

For the classical VAR, uncertainty is assessed using a residual bootstrap with $R = 500$ replications. Standardised residual vectors are resampled with replacement, and pseudo-data are generated using the same fixed COVID-19 volatility path. The VAR and its structural impulse responses are then re-estimated in each replication using the same recursive identification scheme. This bootstrap procedure is used only for the model-validation exercises reported in Appendix 7.8; inference in the main empirical analysis is based on the posterior distributions of the BVAR.

4.6 Bayesian VAR

With six endogenous variables and two lags, each equation of the unrestricted VAR contains 13 coefficients, including the constant, corresponding to 78 coefficients across the full system,

$$n(np + 1) = 6(6 \times 2 + 1) = 78.$$

Relative to the 113 observations available for estimation in each country-origin system, the model is relatively parameter intensive. Bayesian shrinkage provides a natural way to regularise the unrestricted coefficient estimates while retaining the full six-variable information set. Following the approach commonly adopted in macroeconomic BVAR applications (Bańbura et al. 2010), the BVAR is therefore used as the main empirical specification, while classical VAR estimation is retained only for auxiliary calculations and model validation.

4.6.1 Minnesota Prior: Mixed Specification

Let $\beta = \text{vec}(B)$ denote the vector of VAR coefficients. The coefficient prior is specified independently of the residual covariance matrix as

$$\beta \sim N(\beta_0, H).$$

The prior mean follows a mixed specification. For Brent, Gas, HICP, and IPI, which enter the baseline model in log-differences, and for CLIFS, which is retained in levels following the stationarity diagnostics, the own first-lag prior mean is set to zero. All cross-variable coefficients and higher-order own lags are also centred at zero.

For GPR, the own first-lag prior mean is instead set equal to an AR(1) coefficient estimated directly from the corresponding $\log(1 + \text{GPR})$ series. This allows the prior to reflect the persistence of GPR without imposing either a random-walk prior mean of one or a zero-persistence prior. Across the 16 country-origin systems, the estimated AR(1) prior means range from 0.498 to 0.906. This follows the same principle as Bondarenko et al. (2026), who also relax the conventional random-walk prior specifically for their geopolitical-risk indicator, while the numerical prior mean used here is estimated separately for each bilateral GPR series.

The diagonal elements of the prior covariance matrix H follow the Minnesota shrinkage structure. For coefficient (i, j, l), where i denotes the equation, j the regressor variable, and l the lag,

$$H_{i,i,l} = \left(\frac{\lambda_1}{l^{\lambda_3}} \right)^2,$$

$$H_{i,j,l} = \left(\frac{\lambda_1 \lambda_2 \hat{\sigma}_i}{l^{\lambda_3} \hat{\sigma}_j} \right)^2, \qquad j \neq i,$$

and the variance assigned to the intercept in equation i is

$$H_{i,\text{const}} = (\hat{\sigma}_i \lambda_4)^2.$$

The hyperparameters are fixed at

$$\lambda_1 = 0.2, \qquad \lambda_2 = 0.5, \qquad \lambda_3 = 1, \qquad \lambda_4 = 100.$$

Here, λ_1 controls overall tightness, λ_2 applies additional shrinkage to cross-variable lags, λ_3 determines the decay of prior variance across lag order, and λ_4 assigns a diffuse prior to the constant.

The scale parameters $\hat{\sigma}_i$ are the residual standard deviations from the corresponding unadjusted homoskedastic VAR. They are computed before applying the COVID-19 volatility weighting described in Section 4.3. This preserves the baseline Minnesota scaling independently of the temporary volatility adjustment. The scales are recalculated for each country-origin specification and, in robustness exercises that alter the variable set or transformation, for the corresponding alternative VAR.

4.6.2 Prior on Σ

The coefficient prior is independent of Σ. For the residual covariance matrix, the baseline sampler sets

$$S = 0, \qquad \alpha = 0,$$

which implies the Jeffreys-type covariance-prior kernel

$$p(\Sigma) \propto |\Sigma|^{-(N+1)/2},$$

where $N = 6$ is the number of endogenous variables. This is an improper non-informative prior rather than a proper Inverse-Wishart distribution. It introduces no fixed covariance scale into the prior, leaving the scale of Σ to be determined by the likelihood.

Sensitivity to this choice is examined in Appendix 7.4, where the baseline is compared with both a scale-aware covariance prior and an identity-matrix specification.

4.6.3 Gibbs Sampler

Posterior inference is conditional on the estimated COVID-19 volatility path $\hat{s}_t$. As described in Section 4.3, the weighted system is

$$Y_t^* = \frac{Y_t}{\hat{s}_t}, \qquad X_t^* = \frac{X_t}{\hat{s}_t}.$$

Let $\hat{\beta}^*_{\text{OLS}}$ denote the vectorised weighted-OLS coefficient estimate obtained from this transformed system. Conditional on a draw of Σ, the posterior covariance and mean of the VAR coefficients are

$$V^* = \left[H^{-1} + \Sigma^{-1} \otimes (X^{*\prime} X^*)\right]^{-1}$$

$$M^* = V^* \left[H^{-1}\beta_0 + \left\{\Sigma^{-1} \otimes (X^{*\prime} X^*)\right\} \hat{\beta}^*_{\text{OLS}}\right]$$

so that

$$\beta \mid \Sigma, Y^* \sim N(M^*, V^*)$$

Given a coefficient draw β, let

$$U^*(\beta) = Y^* - X^* B(\beta),$$

where $B(\beta)$ is the coefficient matrix associated with the vectorised draw. Under the general covariance-prior parameterisation used by the sampler,

$$\Sigma \mid \beta, Y^* \sim IW\left(S + U^*(\beta)' U^*(\beta),\ T_{\text{eff}} + \alpha\right).$$

For the baseline values $S = 0$ and $\alpha = 0$, this becomes

$$\Sigma \mid \beta, Y^* \sim IW\left(U^*(\beta)' U^*(\beta),\ T_{\text{eff}}\right),$$

with $T_{\text{eff}} = 113$.

Each Gibbs iteration therefore consists of drawing β conditional on the current Σ, followed by drawing Σ conditional on the new coefficient draw. Each of the 16 country-origin systems is estimated using 10,000 Gibbs iterations. The first 5,000 draws are discarded as burn-in, leaving 5,000 retained posterior draws, with no thinning.

For each retained pair $(\beta^{(d)}, \Sigma^{(d)})$, structural impulse responses are computed using the same Wold recursion and Cholesky ordering described for the classical VAR. The posterior median is used as the point estimate, while the baseline uncertainty bands correspond to the 16th and 84th percentiles of the posterior impulse-response distribution, yielding 68% credible intervals.

4.7 Testing Geographical Heterogeneity

Beyond the country-specific impulse responses, the analysis tests whether responses whose 68% posterior credible intervals exclude zero are unusually concentrated in GPR shocks originating from a particular geographical region. The exercise focuses on the five variables responding to the GPR shock,

$$\mathcal{V} = \{\text{Brent}, \text{Gas}, \text{HICP}, \text{IPI}, \text{CLIFS}\},$$

and on the impact and one-year horizons, $h \in \{0, 4\}$.

For country i, origin o, response variable v, and horizon h, define

$$I_{i,o,v,h} = \begin{cases} 1, & \text{if the 68\% credible interval excludes zero,} \\ 0, & \text{otherwise.} \end{cases}$$

The total number of statistically significant responses associated with origin o is then

$$C_o = \sum_{i=1}^{4} \sum_{v \in \mathcal{V}} \sum_{h \in \{0,4\}} I_{i,o,v,h}.$$

Since there are four countries, five response variables, and two horizons, each origin can therefore generate at most 40 significant responses.

Statistical significance of the observed concentration is evaluated using an exact randomisation test based on within-country exchangeability of the four GPR origins. Under the null hypothesis that the geographical origin label is exchangeable within each country, one of the four origin-specific systems is selected independently for each of the four countries. This generates all

$$4^4 = 256$$

possible country-origin combinations. For each combination, the corresponding total number of significant responses is calculated, producing the exact null distribution of the count statistic.

For an observed count C_o, the one-sided exact p-value is

$$p_o = \frac{\#\left\{C^{(b)} \geq C_o\right\}}{256},$$

where $C^{(b)}$ denotes the count obtained under permutation b. A small p-value therefore indicates that the concentration of significant responses associated with a given GPR origin is unusually large relative to the distribution generated by reallocating origins across countries. Because all 256 combinations are enumerated, inference does not rely on asymptotic approximations or Monte Carlo sampling error.

The resulting p-values are origin-specific and are interpreted conditional on the within-country exchangeability assumption underlying the randomisation scheme. They are not adjusted for the simultaneous evaluation of the four geographical origins and therefore do not constitute a family-wise global test of geographical heterogeneity. The randomisation scheme also treats the origin allocation across countries independently and therefore does not explicitly preserve all forms of cross-country dependence.

4.8 Model Verification and Validation

The implementation of the estimation procedure is evaluated using a synthetic six-variable VAR(2) with dimensions comparable to those of the empirical application. Because the data-generating process is known by construction, the exercise allows us to assess whether the estimation procedure recovers the main features of the model and produces inference consistent with the known true parameters and impulse responses. The validation focuses on three aspects: parameter recovery under the Minnesota and covariance priors, sensitivity of covariance estimation to the scale imposed by the prior, and interval coverage for structural impulse responses. The COVID-19 volatility adjustment is switched off so that the exercise isolates the VAR, prior, Gibbs-sampling, Cholesky-identification, and bootstrap procedures.

The results show that the baseline specification recovers the main features of the synthetic data while displaying the shrinkage implied by the Minnesota prior. They also show that imposing a common absolute covariance scale can distort covariance estimation when innovation variances differ substantially across variables. Finally, a Monte Carlo exercise indicates that both the classical bootstrap intervals and the posterior BVAR intervals achieve coverage reasonably close to their nominal level. Overall, these exercises provide no indication of a material implementation failure. Full details of the synthetic design, parameter-recovery results, covariance-prior check, and interval-coverage experiment are reported in Appendix 7.8.

5 RESULTS

5.1 Baseline BVAR Results

The baseline BVAR analysis focuses on the impulse responses of five variables to a GPR shock across 16 country-origin systems. The discussion focuses on the impact, one-year, and two-year horizons, corresponding to $h = 0$, $h = 4$, and $h = 8$. Table 4 summarises the sign of the posterior median across the 16 systems and, in parentheses, the number of responses whose 68% credible interval excludes zero. The complete set of impulse-response plots for all 16 country-origin systems is reported in Appendix 7.7.

Table 4: Summary of baseline BVAR impulse responses

	Impact ($h = 0$)		One year ($h = 4$)		Two years ($h = 8$)	
Response	Positive	Negative	Positive	Negative	Positive	Negative
Brent	9 (5)	7 (3)	0 (0)	16 (9)	0 (0)	16 (9)
Gas	0 (0)	16 (9)	3 (0)	13 (3)	5 (0)	11 (1)
HICP	11 (6)	5 (1)	11 (3)	5 (0)	11 (3)	5 (0)
IPI	5 (2)	11 (2)	8 (3)	8 (2)	8 (3)	8 (3)
CLIFS	5 (1)	11 (3)	1 (0)	15 (8)	1 (0)	15 (8)

Notes: Entries report the number of posterior-median responses with the indicated sign across the 16 country-origin systems. Numbers in parentheses indicate how many of those responses have a 68% posterior credible interval that excludes zero.

5.1.1 Energy-Price Responses

Table 5 identifies the energy-price responses that are credibly different from zero based on their 68% posterior credible intervals. Because Brent and Gas enter the model in log-differences, the impulse responses should be interpreted as responses of quarterly price growth rather than as responses of price levels.

Table 5: Energy-price responses credibly different from zero

Response	GPR origin	Impact	One year	Two years
Brent	China	DE+, IT+, ES+	DE-, IT-	DE-, IT-
	MENA	DE-, IT-, ES-	FR-, DE-	FR-, DE-
	Russia	FR+, ES+	FR-, DE-, IT-, ES-	FR-, DE-, IT-, ES-
	Western	–	DE-	DE-
Gas	China	ES-	–	–
	MENA	–	–	–
	Russia	FR-, DE-, IT-, ES-	DE-, IT-, ES-	DE-
	Western	FR-, DE-, IT-, ES-	–	–

Notes: The table reports responses whose 68% posterior credible interval excludes zero. A + ($-$) indicates a positive (negative) posterior-median response. FR, DE, IT, and ES denote France, Germany, Italy, and Spain, respectively.

The most uniform short-run result concerns natural gas prices. The posterior median impact response of Gas is negative in all 16 country-origin systems, with nine responses credibly different from zero. This evidence is strongly concentrated in shocks originating from Russia and the Western Bloc, for which the impact response is negative and credibly different from zero in all four countries. By contrast, only Spain displays a credible impact response to Chinese GPR shocks, while none is found for MENA shocks.

The negative Gas response becomes substantially less persistent as the horizon increases. After one year, the posterior median remains negative in 13 of the 16 systems, but only 3 responses are credibly different from zero, all following Russian GPR shocks. After two years, 11 posterior medians remain negative, while only the Germany-Russia response remains credibly different from zero. Therefore, the principal Gas result is the striking uniformity of its negative short-run response rather than a persistent effect at longer horizons.

A simple energy-supply disruption mechanism does not appear sufficient to explain the negative Gas responses. In particular, shocks originating from Russia, for which supply concerns might be expected to be especially relevant to the European gas market, generate negative and credibly different from zero impact responses in all four countries and also display the greatest persistence at longer horizons. This result highlights an important distinction between a geopolitical-risk shock and a pure physical energy-supply shock. The bilateral GPR index captures increases in geopolitical tensions associated with Russia, but such episodes need not coincide contemporaneously with an unexpected reduction in gas supply.

One economically plausible explanation is that geopolitical tensions weaken expectations of economic activity and future energy demand, putting downward pressure on gas-price growth and potentially offsetting the upward pressure associated with supply concerns.

The quarterly frequency of the analysis is also important when interpreting this result. Because both GPR and gas prices are observed at quarterly frequency, the estimated impact response captures the net contemporaneous-quarter effect and cannot identify the sequence of market movements within the quarter. Energy markets can experience sharp and partially offsetting movements over shorter horizons, which may translate into a weak or negative response when aggregated to the quarterly frequency. The present BVAR also does not separately identify supply, demand, or risk-premium channels. These mechanisms should therefore be regarded as economically plausible interpretations of the observed response rather than as separately identified causal effects.

The Gas results contrast with those for Brent. The impact response of Brent varies systematically with the geographical origin of the geopolitical shock. Posterior medians are positive in all four systems following Chinese and Russian shocks, negative in all four MENA systems, and negative in three of the four Western systems. Among the credible impact responses, Chinese shocks generate positive responses in Germany, Italy, and Spain, Russian shocks generate positive responses in France and Spain, while MENA shocks generate negative responses in Germany, Italy, and Spain.

This geographical pattern can be interpreted through the interaction of two competing forces in global oil markets. On the one hand, geopolitical tensions can increase concerns about future supply disruptions, trade restrictions, and the availability of oil, raising the geopolitical risk premium embedded in Brent prices. On the other hand, geopolitical uncertainty can weaken expectations of global economic activity and future oil demand, putting downward pressure on prices (Ferrari Minesso et al. 2023). The sign of the impact response therefore depends on which of these forces dominates for a particular geopolitical origin.

For Russian-origin shocks, the positive impact response is consistent with the supply-risk channel being relatively important in the short run. Geopolitical tensions involving Russia can generate concerns about disruptions to oil exports, sanctions, or restrictions on international energy trade. Since Brent is a globally traded crude-oil benchmark, these concerns can be rapidly incorporated into prices through a higher supply-risk premium.

The positive impact response following Chinese shocks is less directly related to oil production. China is a major source of global oil demand, so a pure activity channel could instead imply lower expected demand and a negative Brent response. The positive response found here suggests that, on impact, other components of geopolitical risk — such as concerns about international trade, transportation, supply chains, or future energy availability — may dominate the demand channel. In this sense, the result is consistent with an immediate risk-premium effect rather than with a simple oil-demand interpretation.

By contrast, the negative impact response following MENA shocks shows why geopolitical risk originating in an energy-producing region should not be interpreted mechanically as an adverse oil-supply shock. The MENA index captures a broad set of geopolitical events and does not identify unexpected reductions in regional oil production. If a geopolitical episode raises uncertainty and weakens expected global activity without producing a sufficiently large contemporaneous supply disruption, the associated decline in expected oil demand may dominate the supply-risk effect, producing a negative response of Brent-price growth. Quarterly aggregation may reinforce this pattern if an initial short-lived increase in oil prices is subsequently reversed within the same quarter.

The predominantly negative posterior medians following Western-origin shocks are also consistent with a relatively stronger demand and uncertainty channel, as such shocks need not involve a direct disruption to global crude-oil supply. However, none of the impact responses is credibly different from zero, so this interpretation should be treated cautiously.

An important implication of these findings is that Brent crude and gas do not necessarily move in the same direction in response to a given geopolitical shock. This is particularly evident for Russia, where Brent-price growth increases on impact while Gas-price growth falls. This is not necessarily contradictory. A Russian GPR shock is not a common structural energy-supply shock imposed simultaneously on both markets. Brent is a globally traded crude-oil benchmark, whereas the European gas market has different pricing, supply, substitution, and adjustment mechanisms. The same geopolitical episode can therefore raise the global oil supply-risk premium while the net quarterly response of European gas prices reflects weaker expected demand or other gas-market-specific adjustment mechanisms.

Finally, the origin-specific impact pattern in Brent disappears at longer horizons. At both one and two years, the posterior median response is negative in all 16 systems, with nine responses credibly different from zero at each horizon. This convergence towards negative responses is consistent with the immediate supply-risk component being relatively short-lived, while weaker expected activity, lower oil demand, and broader uncertainty become relatively more important once the initial geopolitical reaction dissipates. As above, this interpretation should not be understood as a structural decomposition of the underlying channels, which the BVAR does not identify separately.

5.1.2 Macro-Financial Responses

Table 6 reports the macro-financial responses that are credibly different from zero. HICP and IPI enter the model in log-differences and therefore represent responses in quarterly inflation and industrial-production growth, respectively, while CLIFS enters in levels.

Table 6: Macro-financial responses credibly different from zero

Response	GPR origin	Impact	One year	Two years
HICP	China	FR+, DE+	DE+	DE+
	MENA	ES-	–	–
	Russia	FR+, DE+, IT+, ES+	DE+	DE+
	Western	–	DE+	DE+
IPI	China	FR+, DE+	DE-	DE-
	MENA	DE-	ES+	ES+
	Russia	–	DE-, ES+	DE-, ES+
	Western	IT-	ES+	DE-, ES+
CLIFS	China	DE-	FR-, DE-	FR-, DE-
	MENA	–	FR-, IT-, ES-	FR-, IT-, ES-
	Russia	FR-, ES-	ES-	ES-
	Western	FR+	FR-, ES-	FR-, ES-

Notes: The table reports responses whose 68% posterior credible interval excludes zero. A + (−) indicates a positive (negative) posterior-median response. FR, DE, IT, and ES denote France, Germany, Italy, and Spain, respectively.

HICP responses are predominantly positive. The posterior median impact response is positive in 11 of the 16 systems, and six of the seven impact responses that are credibly different from zero are positive. The clearest pattern concerns Russian GPR shocks: HICP inflation increases on impact in all four countries, with all four responses credibly different from zero. Chinese shocks also produce credible positive impact responses in France and Germany. The only credible negative impact response is observed in Spain following a

MENA shock.

The inflationary response becomes more geographically concentrated at longer horizons. Although 11 of the 16 posterior medians remain positive after both one and two years, only three responses are credibly different from zero at each horizon. All three occur in Germany, following shocks originating from China, Russia, and the Western Bloc. The evidence therefore points to a relatively broad positive short-run inflation response to Russian geopolitical risk, but much more limited persistence across countries.

The positive HICP responses, particularly following Russian GPR shocks, are consistent with geopolitical risk operating partly as an adverse supply disturbance, raising production costs through supply shortages, trade disruptions, sanctions, and higher imported input costs (Bondarenko et al. 2026). This interpretation is particularly plausible for Russian shocks, given the potential exposure of European production to energy and trade disruptions. Chinese shocks may operate through a related but broader imported-input and trade channel: geopolitical tensions involving China can increase uncertainty surrounding international production networks and imported intermediate goods, putting upward pressure on production costs even in the absence of a direct energy-supply disruption.

The negative Spain–MENA impact response constitutes an exception to this predominantly inflationary pattern. One possible explanation is that weaker demand has a stronger effect than higher supply costs in this case. This is also consistent with the negative Brent response found for MENA shocks, although the BVAR does not establish a causal energy-price mediation mechanism. Since this is the only credible negative HICP impact response, it should not be overinterpreted as a systematic MENA effect.

At longer horizons, the concentration of credible positive responses in Germany suggests greater persistence of the inflationary transmission in that economy. Differences in industrial structure, reliance on imported intermediate inputs, and exposure to external energy and trade conditions may contribute to this persistence, although the present model does not identify which of these mechanisms is responsible.

In particular, Bondarenko et al. (2026) find that shortages account for an important share of the inflationary effects of geopolitical risk in the euro area. However, the present BVAR does not separately identify these transmission channels. Moreover, the mixed responses of Brent and Gas suggest that the inflationary effect cannot be attributed mechanically to a single energy-price channel.

Industrial production exhibits considerably greater heterogeneity. At impact, 11 of the 16 posterior medians are negative, while the four credible responses are evenly divided between positive and negative effects. At both one and two years, positive and negative posterior medians are split evenly across the 16 systems. A notable cross-country contrast nevertheless emerges. Germany exhibits negative medium-run IPI responses across all four GPR origins, with credible declines following Chinese and Russian shocks at one year and an additional Western response after two years. Spain, by contrast, displays positive and credible medium-run responses following MENA, Russian, and Western shocks. The estimates therefore provide little evidence of a common real-activity response across the four economies.

The heterogeneous IPI responses contrast with the predominantly contractionary effects of geopolitical risk documented in the broader literature (Caldara & Iacoviello 2022, Bondarenko et al. 2026). Higher geopolitical risk can weaken industrial activity through greater uncertainty, delayed investment, trade and supply-chain disruptions, and weaker external demand. These mechanisms are consistent with the negative responses observed for Germany following MENA shocks on impact and Chinese and Russian shocks at longer horizons, as well as with the negative Italian response to Western geopolitical risk.

The positive impact responses to Chinese shocks in France and Germany require a more cautious interpretation. Since IPI enters in quarterly log-differences, the estimated impact response refers to industrial production growth within the same quarter as the GPR innovation, rather than to an instantaneous increase in the level of industrial output after the geopolitical event. If the geopolitical shock occurs late in the quarter, a substantial part of observed production may reflect decisions and orders made before the shock. Quarterly aggregation can therefore distort the sequence, within a single quarter, between geopolitical events and industrial activity. These positive impact responses should consequently not be interpreted as evidence that geopolitical risk immediately stimulates aggregate industrial production.

Sectoral reallocation could in principle also contribute to positive industrial responses if geopolitical tensions increase production in particular industries, including defence-related sectors. However, aggregate IPI does not allow this mechanism to be identified, so it should be regarded only as a possible interpretation rather than the main explanation of the impact result.

At medium horizons, a notable cross-country contrast emerges. Germany displays negative responses across GPR origins, which is consistent with the greater sensitivity of a highly manufacturing- and export-oriented economy to trade disruption, weaker external demand, and uncertainty. Spain instead exhibits positive and credible responses following MENA, Russian, and Western shocks. Part of this difference may reflect the predominantly negative medium-run responses of Brent and Gas, which imply lower energy-price growth relative to the no-shock path and may reduce input-cost pressures.

Importantly, a positive IPI response at one or two years refers to higher quarterly industrial-production growth relative to the baseline at that horizon; it does not necessarily imply that the level of industrial production is permanently above the no-shock path. Such positive responses may therefore partly reflect recovery or mean reversion following earlier weakness. Assessing the cumulative effect on the level of industrial production would require accumulating the responses of $\Delta \log(IPI)$.

The balance between these forces may differ across countries because of differences in industrial structure, trade exposure, and energy dependence. Since the BVAR does not separately identify their contribution, the absence of a common IPI response is itself an economically relevant result.

Finally, CLIFS responses become predominantly negative beyond the impact horizon. At both one and two years, 15 of the 16 posterior medians are negative and eight responses are credibly different from zero, all with a negative sign. Since higher CLIFS values represent greater systemic financial stress, the estimated responses indicate that the identified GPR shocks do not generate a persistent increase in financial stress.

The impact responses are more mixed. The positive response for France following a Western shock is consistent with the usual idea that greater geopolitical uncertainty can increase financial stress. However, the negative impact responses found for Germany after Chinese shocks and for France and Spain after Russian shocks are more difficult to interpret as immediate reactions to geopolitical events.

One reason is the quarterly frequency of the model. Financial markets can react to geopolitical news within hours or days, while the GPR index and CLIFS are analysed at quarterly frequency. Therefore, the impact response captures the net effect within the whole quarter, not the immediate market reaction. For example, financial stress could increase just after a geopolitical event and then fall again before the end of the quarter. In that case, the quarterly response could appear weak or even negative. For this reason, the negative impact responses should not be interpreted as evidence that geopolitical shocks immediately reduce financial stress.

The negative responses after one and two years are easier to explain. If the financial effects of geopolitical events are mainly short-lived, markets may gradually adjust once the initial uncertainty decreases. This may lead CLIFS to fall relative to its no-shock path as financial stress normalises over time. However, this does not mean that geopolitical risk improves financial conditions. It only means that financial stress is estimated to be lower than its no-shock path at those horizons. Since CLIFS combines information from equity, sovereign-bond, and foreign-exchange markets, different parts of the financial system may also recover at different speeds (Duprey et al. 2017). The BVAR does not identify the exact mechanism behind these negative responses, and the robustness analysis also shows that CLIFS is more sensitive to changes in the model specification than HICP or IPI.

5.2 Geographical Heterogeneity

The baseline results indicate that the transmission of geopolitical risk differs not only across destination countries but also across the geographical origin of the shock. To assess whether responses whose 68% posterior credible intervals exclude zero are disproportionately concentrated in a particular GPR origin, Table 7 reports the number of such responses across the five response variables, four countries, and the impact and one-year horizons. The corresponding exact permutation test evaluates whether the observed concentration could arise under exchangeability of GPR origins within each country.

Table 7: Concentration of detectable responses by GPR origin

GPR origin	Responses excluding zero	Exact p-value
Russia	23	0.0273
China	15	0.5430
MENA	11	0.9063
Western	11	0.9063

Notes: Counts refer to responses whose 68% posterior credible interval excludes zero across five response variables, four countries, and horizons $h = 0$ and $h = 4$, giving a maximum of 40 responses per origin. Exact upper-tail p-values are obtained from the permutation test described in the methodology.

The results show a clear concentration of detectable responses following Russian GPR shocks. Of the 40 country-variable-horizon responses associated with each origin, 23 exclude zero for Russia, compared with 15 for China and 11 for both MENA and the Western Bloc. The permutation distribution has a mean count of 15 responses per origin. The exact permutation test assigns the Russian count a p-value of 0.0273, whereas none of the other origins displays an unusually high concentration of detectable responses.

This result provides evidence that the geographical origin of geopolitical risk matters for its macro-financial transmission. In particular, Russian GPR shocks generate a broader set of detectable responses across the four euro-area economies than would be expected under exchangeability of shock origins. This is consistent with several patterns documented in the baseline results, including the uniform negative impact response of Gas and the broad positive impact response of HICP following Russian shocks.

The result should nevertheless be interpreted as evidence about the concentration of detectable responses rather than their economic magnitude. The permutation test does not imply that Russian shocks always generate larger responses, nor that the remaining GPR origins are economically irrelevant. Indeed, the preceding results show important origin-specific effects for other variables, particularly the heterogeneous impact responses of Brent. Rather, the test indicates that responses whose posterior credible intervals exclude zero are more systematically concentrated in shocks associated with Russia.

Heterogeneity is also apparent across destination countries. The direction, persistence, and detectability of the responses can differ substantially across economies even when the geographical origin of the GPR shock is held fixed. The clearest example is industrial production. Germany exhibits negative medium-run IPI responses across all four GPR origins. The 68% credible interval lies entirely below zero following Chinese and Russian shocks at one year and following Western shocks after two years. By contrast, Spain displays positive and credible medium-run responses following MENA, Russian, and Western shocks. HICP provides another example: although Russian GPR shocks generate positive and credible impact responses in all four countries, all credible HICP responses beyond the impact horizon are concentrated in Germany. These differences suggest that exposure to a given geopolitical shock depends importantly on the characteristics of the receiving economy.

At the same time, cross-country heterogeneity is not uniform across variables. Gas provides the clearest counterexample: its posterior median impact response is negative in all 16 country-origin systems, indicating a short-run pattern that is shared across receiving economies despite differences in shock origin. Other variables, particularly IPI and the persistence of HICP responses, display much stronger cross-country differentiation. Taken together, the results therefore indicate that geographical heterogeneity operates along both dimensions considered in the analysis: the origin of geopolitical risk and the economy receiving the shock.

5.3 Robustness and Specification Sensitivity

The robustness analysis examines the main modelling choices that could potentially drive the baseline results. The first set of exercises assesses sensitivity to system dimensionality, the prior on the reduced-form covariance matrix, and the treatment of persistent variables. These exercises allow a direct comparison with the baseline using common sign and posterior-detectability metrics and are summarised in Table 8. A separate set of checks focuses specifically on the negative short-run response of Gas and examines its sensitivity to structural identification and sample composition. Full results are reported in the Appendix 7.

Table 8: Summary of robustness and specification checks

Robustness check	IRF comparisons	Sign agreement	Detectability changes
Reduced system	144	94.4%	5
Scale-aware Σ prior	240	100.0%	2
Levels specification	240	74.2%	87

Notes: Sign agreement refers to the proportion of posterior-median responses with the same sign as in the baseline specification. Detectability changes count cases in which the 68% posterior credible interval excludes zero in one specification but not in the other. Full results are reported in Appendices 7.2, 7.4, and 7.5.

The reduced-system exercise assesses whether the macro-financial responses depend on the dimensionality of the six-variable baseline or on the inclusion of the two energy-price variables. To this end, Brent and Gas are removed and the BVAR is re-estimated using only GPR, HICP, IPI, and CLIFS. The resulting responses display 94.4% sign agreement with the corresponding baseline estimates, with only 5 detectability changes across 144 comparisons. Four of these changes concern CLIFS, while none affects IPI. The macro-financial conclusions therefore do not appear to depend mechanically on the inclusion of the energy-price variables, although the greater sensitivity of CLIFS is consistent with the caution applied to its baseline interpretation. Full results are reported in Appendix 7.2.

The covariance-prior exercise examines whether posterior inference is sensitive to the treatment of the reduced-form covariance matrix, particularly given the substantial differences in scale across the variables included in the BVAR. Replacing the non-informative baseline prior with a scale-aware prior based on the diagonal of the homoskedastic VAR covariance matrix preserves the sign of all 240 responses and changes detectability in only 2 cases, while the credible intervals are, on average, only 0.8% wider than in the baseline specification. By contrast, the $S = I$ specification produces larger changes in both signs and posterior uncertainty. Since a common identity scale does not reflect the different empirical variances of the variables, the $S = I$ specification is used only as a sensitivity check and not as an equally plausible alternative prior. The results therefore suggest that the baseline conclusions are not materially driven by the covariance prior when an empirically scaled alternative is used. Full results are reported in Appendix 7.4.

The levels specification tests whether the baseline results depend on transforming the persistent variables into log-differences. The differenced specification is used in the baseline because the analysis focuses on the short and medium run transmission of GPR shocks, while Brent, Gas, HICP, and IPI are transformed to achieve stationarity and to provide a direct interpretation in terms of quarterly changes. At the same time, the cointegration evidence among these persistent variables makes the levels specification a relevant robustness check. Across 240 comparable responses, sign agreement falls to 74.2%, with 87 changes in whether the 68% credible interval excludes zero. The sensitivity is particularly marked for Gas, for which only 28 of 48 comparisons preserve the baseline sign. The origin-heterogeneity result is also affected: under the levels specification, the exact permutation p-value for the concentration of detectable Russian responses rises from 0.0273 in the baseline to 0.7852. Hence, the formal evidence that detectable transmission is disproportionately concentrated in Russian GPR shocks is not robust to the levels specification. Therefore, the differences between the two specifications represent an important source of model sensitivity, particularly for Gas and geographical heterogeneity. Cointegration diagnostics and the full levels-versus-differences comparison are reported in Appendix 7.5.

A separate set of robustness exercises focuses on the negative short-run response of Gas. Given that the posterior median impact response is negative in all 16 baseline systems, these checks assess whether this finding is an artefact of the recursive ordering imposed on GPR, Brent, and Gas or is disproportionately driven by the exceptional European energy-market conditions surrounding 2022.

Table 9 reports the number of country-origin systems with a negative posterior-median Gas response under each robustness exercise. Under the energy-first ordering, the impact response at $h = 0$ is zero by construction. Because Gas is ordered before GPR in the recursive Cholesky identification, Gas cannot respond contemporaneously to a GPR shock. The comparison therefore starts at $h = 1$, the first horizon at which GPR can affect Gas.

Table 9: Robustness of the negative Gas response

Robustness check	Horizon	Systems with negative Gas response
Brent–Gas inversion	$h = 0$	16/16
Energy before GPR	$h = 1$	16/16
	$h = 2$	16/16
	$h = 4$	15/16
	$h = 8$	11/16
Excluding 2021Q4–2022Q4	$h = 0$	15/16
Sample ending 2022Q2	$h = 0$	16/16

Reversing the ordering of Brent and Gas leaves the negative impact response unchanged in all 16 systems. When the energy variables are ordered before GPR, the negative Gas response appears at the first horizon after the shock and is present in all 16 systems at $h = 1$ and $h = 2$. It remains negative in 15 of 16 systems after one year and in 11 of 16 after two years. The short-run sign is also highly stable when 2021Q4–2022Q4 is excluded and when the estimation sample ends in 2022Q2. These results indicate that the negative Gas response is not driven solely by the baseline recursive ordering or by the most extreme phase of the European energy crisis. Its magnitude and persistence, however, are more sensitive to the specification than its short-run sign. Full results are reported in Appendix 7.3.

Overall, the robustness exercises distinguish between conclusions that are relatively stable and those that depend more strongly on specification choices. The negative short-run Gas response survives several alternative identification and sample exercises, while the macro-financial responses are largely preserved in the reduced system and under a scale-aware covariance prior. By contrast, the levels specification produces substantial changes in signs, detectability, and the geographical concentration of responses. Therefore, the baseline results should be interpreted as robust to several modelling choices, but not as invariant to the treatment of persistence and variable transformation.

6 CONCLUSIONS

This thesis examines whether the geographical origin of geopolitical risk affects its transmission through energy prices and macro-financial variables across Germany, France, Italy, and Spain. To answer this question, 16 quarterly BVAR systems are estimated by combining four recipient economies with four geopolitical origins: Russia, China, the Middle East and North Africa (MENA), and the Western Bloc. Each system traces the responses of Brent crude oil prices, European natural gas prices, HICP inflation, industrial production, and financial stress to an origin-specific GPR shock.

The central conclusion is that geopolitical risk should not be treated as a homogeneous global shock. Its transmission depends on the geographical origin of the underlying event, the market or macroeconomic variable considered, and the economy receiving the shock. This heterogeneity is particularly visible in energy markets. Natural gas price growth falls on impact in all 16 baseline systems, although the response becomes less persistent over time. Brent instead responds positively on impact to Chinese and Russian shocks and negatively to MENA shocks, before becoming negative across all systems at longer horizons. The contrast between the two markets shows that a GPR shock cannot be interpreted mechanically as a physical energy-supply shock.

The macro-financial results reinforce this conclusion. HICP inflation responds predominantly positively, particularly following Russian-origin shocks, whereas industrial-production responses differ markedly across countries, with predominantly negative medium-run responses in Germany and several positive responses in Spain. Financial-stress responses also vary across horizons: CLIFS does not exhibit a common positive impact response and becomes predominantly negative after one and two years. This longer-run decline should be interpreted as normalisation relative to the no-shock path rather than as evidence that geopolitical risk improves financial conditions.

Russian-origin shocks generate the broadest set of detectable responses in the baseline specification, and the permutation test identifies their concentration as unusual under exchangeability of shock origins. However, this result concerns the prevalence rather than the magnitude of the responses and disappears when the persistent energy and macroeconomic variables are modelled in levels. Other findings are more stable: the main macro-financial patterns are largely preserved in the reduced system and under an empirically scaled covariance prior, while the negative short-run Gas response survives alternative recursive orderings and sample restrictions. Thus, the broader conclusion that geographical origin matters is supported across several modelling choices, whereas the stronger Russia-specific result depends on the baseline treatment of persistent variables.

Overall, the results show the value of disaggregating geopolitical risk by both origin and recipient economy. An aggregate GPR index can conceal differences in the direction, persistence, and detectability of the responses generated by events originating in different regions. At the same time, a common geopolitical origin does not produce identical effects across euro-area economies. Differences in energy exposure, industrial structure, trade relationships, and financial adjustment may contribute to this cross-country variation, although the present models do not separately identify these mechanisms.

The main limitation of the analysis is its quarterly frequency. Energy and financial markets react to geopolitical news within hours or days, while quarterly observations capture only the net adjustment occurring within the period. The models therefore cannot recover the sequence of market movements within a quarter or distinguish an initial reaction from a subsequent reversal. This limitation is also relevant for the recursive identification strategy, which imposes a contemporaneous ordering that cannot be observed directly at the quarterly frequency. In addition, the two-stage treatment of COVID-19 volatility does not propagate uncertainty about the estimated volatility path into the posterior impulse responses, while the baseline differenced specification does not explicitly model possible long-run cointegrating relationships.

Future research could address these limitations by combining bilateral GPR measures with daily or weekly energy and financial data. A higher-frequency analysis would make it possible to study immediate market reactions around major geopolitical events and determine whether the quarterly responses reflect an initial movement, a subsequent reversal, or the net result of several offsetting adjustments. Monthly macroeconomic data could provide a complementary link between these immediate market responses and their later transmission to inflation and economic activity. Extending the analysis to additional euro-area economies and incorporating observable measures of trade, industrial, and energy exposure would also help explain why shocks from the same geopolitical origin are transmitted differently across recipient countries.

7 APPENDIX

7.1 HICP Seasonal Adjustment – Full Diagnostic

Table 10 reports the complete diagnostic comparison between X-13ARIMA-SEATS and TRAMO-SEATS underlying the choice of seasonal adjustment method described in Section 3.3. The Q statistic is an X-11-specific diagnostic and is not produced by TRAMO-SEATS, which performs a model-based ARIMA/SEATS decomposition rather than relying on empirical filters. The Bayesian Information Criterion (BIC) is reported for completeness but is not directly comparable between the two implementations: both report the corrected-for-length BIC of the shared RegARIMA pre-adjustment step, but on markedly different scales (a consistent offset of approximately 4.4–5.2× across all four countries), and was therefore excluded from the method selection.

Table 10: Full seasonal adjustment diagnostic comparison, X-13ARIMA-SEATS vs. TRAMO-SEATS

	AIC		BIC (not comparable)		QS test (p)		F-test, seas. dummies (p)	
Country	X-13	TRAMO-S.	X-13	TRAMO-S.	X-13	TRAMO-S.	X-13	TRAMO-S.
Germany	74.93	129.50	−2.525	−11.53	1.000	1.000	0.993	1.000
France	32.67	25.46	−2.698	−11.86	0.479	0.862	1.000	1.000
Italy	108.40	107.50	−2.401	−11.50	0.543	1.000	0.864	1.000
Spain	206.00	209.10	−2.146	−11.23	0.976	1.000	0.681	1.000

Table 11: Q statistic and quarterly seasonal spread reduction, X-13ARIMA-SEATS vs. TRAMO-SEATS

Country	Q (X-13 only)	Resid. seasonality, X-13 (p)	Spread reduction, X-13	Spread reduction, TRAMO-S.
Germany	0.434	1.000	83.0%	88.4%
France	0.382	1.000	93.4%	99.8%
Italy	0.375	0.998	95.3%	94.9%
Spain	0.546	1.000	97.2%	96.7%

Averaging across the four countries, TRAMO-SEATS achieves a higher mean spread reduction (94.95%) than X-13ARIMA-SEATS (92.22%), and matches or exceeds X-13ARIMA-SEATS on the QS test and the F-test on seasonal dummies in all four countries. AIC is tied between the two methods (two countries each). These results, together with TRAMO-SEATS being Eurostat's default method within JDemetra+, motivated its adoption as the final specification.

7.2 Reduced System (Excluding Energy Prices)

To assess whether the dimensionality of the baseline system materially affects posterior inference, a reduced four-variable BVAR was estimated after removing Brent and Gas:

$$[\log(1 + \text{GPR}),\ \Delta \log(\text{HICP}),\ \Delta \log(\text{IPI}),\ \text{CLIFS}]$$

The exercise changes the number of endogenous variables while retaining the same lag length, frozen COVID-19 volatility path, Minnesota hyperparameters, covariance prior, and recursive identification of the GPR shock. Prior scales are recalculated from the corresponding unadjusted homoskedastic four-variable VAR, consistently with the baseline prior construction. Reducing the system from six to four variables lowers the number of coefficients per equation from 13 to 9 and raises the ratio of usable observations to coefficients from approximately 8.7 to 12.6.

The comparison is conducted directly between the six-variable and four-variable BVAR for the three common response variables, HICP, IPI, and CLIFS, at the impact, one-year, and two-year horizons. Across the resulting 144 comparisons, the posterior median responses retain the same sign in 94.4% of cases. Only 5 of 144 comparisons change significance status according to the 68% credible interval.

The reduction in posterior uncertainty is also quantitatively small. The mean ratio between the width of the 68% credible band in the four-variable system and that in the six-variable system is 0.989, with a median of 0.985 and an interquartile range of 0.961–1.010. Thus, excluding the energy variables narrows the credible bands by only around 1–1.5% on average. Four of the five changes in significance status concern CLIFS, while IPI exhibits no changes.

Overall, the reduced-system exercise provides little evidence that the dimensionality of the six-variable specification materially weakens posterior inference. Therefore, the full system is retained because it incorporates the energy-price transmission channel central to the analysis while imposing only a negligible cost in terms of posterior uncertainty.

7.3 Structural Identification and Gas-Response Robustness

This section evaluates the empirical sensitivity of the identification strategy discussed in Section 4.4: that GPR is ordered first in the Cholesky factorisation. Alternative recursive orderings cannot establish the validity of the baseline identifying assumption, but they can assess whether particular impulse-response patterns depend mechanically on the contemporaneous restrictions imposed by the baseline ordering.

Three checks were carried out. The first examines the lead-lag correlation profile between GPR and energy prices directly to test whether energy prices systematically precede GPR. The second reverses the relative ordering of Brent and Gas while keeping GPR first, testing whether the negative Gas response is mechanically generated by the baseline Brent-before-Gas restriction. The third places both Brent and Gas before GPR while preserving Brent before Gas, to assess the sensitivity of the Gas response to the baseline GPR-first assumption.

Additional diagnostics reported at the end of this section assess the robustness of the negative Gas response to posterior uncertainty and to the exceptional European energy-market observations surrounding 2022.

7.3.1 Lead-Lag Profile: GPR versus Energy Prices

For each of the 16 country-origin systems, the Pearson correlation between $\log(1 + \mathrm{GPR}_{\mathrm{origin}})$, shifted by k quarters, and quarter-on-quarter energy price growth (Brent or Gas) is computed for $k = -4$ to $k = 4$. Positive k corresponds to GPR leading energy prices (past GPR against current energy growth) and negative k corresponds to energy prices leading GPR.

Figure 1 reports the simple average of this correlation across all 16 systems, for each energy variable and lag.

The correlation is weaker on average when energy prices lead GPR ($k < 0$) than when GPR leads energy prices ($k > 0$). Brent's most negative correlation occurs at $k = 1$ (one quarter after the GPR shock), while Gas's most negative correlation occurs at $k = 3$ (three quarters after the GPR shock). This asymmetry supports the identification assumption defended in Section 4.4: energy-price movements do not systematically precede changes in the bilateral GPR indices. However, since the lead-lag exercise is descriptive, it should be interpreted as evidence consistent with the GPR-first ordering rather than as a formal test of structural exogeneity.

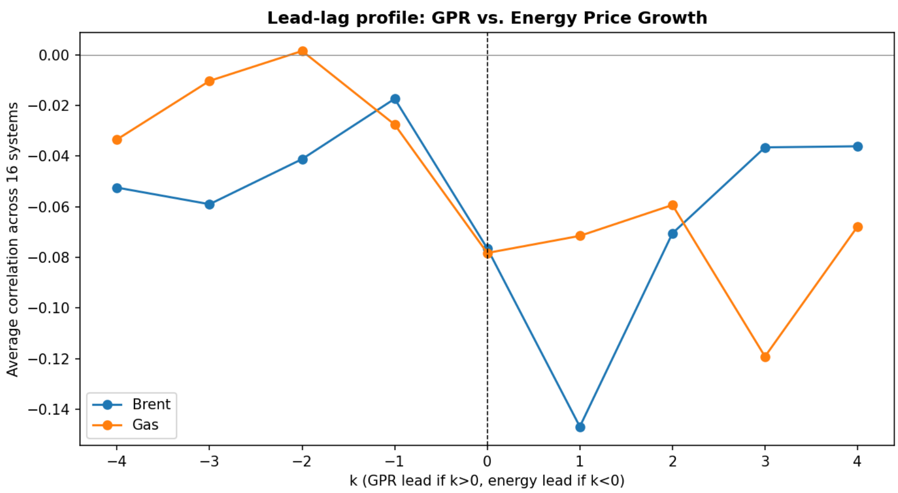

Figure 1: Lead-lag profile: average correlation between GPR and energy price growth, across all 16 systems

7.3.2 Brent/Gas Inversion

The baseline ordering places Brent before Gas, reflecting the historical linkage between European natural gas and oil prices documented in the literature. Previous evidence finds strong integration between European gas prices and Brent, including price spillovers running from crude oil towards natural gas markets.

$$[\log(1+\text{GPR}),\ \Delta\log(\text{Brent}),\ \Delta\log(\text{Gas}),\ \Delta\log(\text{HICP}),\ \Delta\log(\text{IPI}),\ \text{CLIFS}]$$

The alternative specification reverses this relative ordering to assess whether the resulting recursive restriction materially affects the estimated Gas response.

$$[\log(1+\text{GPR}),\ \Delta\log(\text{Gas}),\ \Delta\log(\text{Brent}),\ \Delta\log(\text{HICP}),\ \Delta\log(\text{IPI}),\ \text{CLIFS}]$$

GPR remains ordered first in both specifications, so the impact response of Gas to a GPR shock is directly comparable. Table 12 reports the posterior median impact responses for all 16 country-origin systems.

No system exhibits a sign change when the Brent/Gas ordering is reversed (0 of 16), and the estimated magnitudes remain very similar across the two specifications. This indicates that the negative impact response of Gas observed in the baseline is not mechanically generated by placing Brent before Gas in the Cholesky factorisation.

Table 12: Gas impact response to a GPR shock: baseline vs. inverted Brent/Gas ordering

Country	Origin	Baseline	Gas before Brent	Sign flip
Germany	Russia	−0.02308	−0.02310	No
Germany	China	−0.01052	−0.01058	No
Germany	MENA	−0.00975	−0.01033	No
Germany	Western	−0.03240	−0.03276	No
France	Russia	−0.02596	−0.02597	No
France	China	−0.00976	−0.01028	No
France	MENA	−0.00006	−0.00042	No
France	Western	−0.02577	−0.02625	No
Italy	Russia	−0.03100	−0.03135	No
Italy	China	−0.01508	−0.01553	No
Italy	MENA	−0.01385	−0.01449	No
Italy	Western	−0.02386	−0.02434	No
Spain	Russia	−0.03386	−0.03409	No
Spain	China	−0.01699	−0.01754	No
Spain	MENA	−0.00600	−0.00656	No
Spain	Western	−0.03823	−0.03827	No

7.3.3 Energy Before GPR

The third robustness exercise reverses the central contemporaneous restriction underlying the baseline identification. Instead of ordering GPR before energy prices, the alternative specification is

$$[\Delta \log(\text{Brent}),\ \Delta \log(\text{Gas}),\ \log(1 + \text{GPR}),\ \Delta \log(\text{HICP}),\ \Delta \log(\text{IPI}),\ \text{CLIFS}]$$

This places both energy variables before GPR while preserving the Brent-before-Gas ordering justified in Section 4.4. Under this recursive structure, the contemporaneous response of Gas to a GPR shock is mechanically restricted to zero. Therefore, the impact response under the alternative ordering should not be interpreted as empirical evidence of the absence of an immediate effect.

The relevant comparison begins at $h = 1$. The median Gas response is negative in all 16 country-origin systems at both $h = 1$ and $h = 2$, remains negative in 15 of 16 systems at $h = 4$, and is still negative in 11 of 16 systems at $h = 8$. Thus, even when the contemporaneous response is eliminated by construction, the negative Gas response re-emerges immediately through the dynamic propagation of the VAR.

This result indicates that the negative Gas dynamics are not solely an artefact of the baseline GPR-first contemporaneous restriction. The precise short-run magnitude remains sensitive to the recursive ordering, but the negative response itself persists strongly once the model evolves beyond the impact period. More generally, assessing the sensitivity of GPR impulse responses to alternative recursive orderings follows the identification robustness approach employed by Caldara & Iacoviello (2022).

Table 13: Gas response to a GPR shock: GPR-first vs. energy-first ordering

Country	Origin	Ordering	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$
Germany	Russia	GPR-first	−0.02308	−0.00571	−0.00439	−0.00802	−0.00954	−0.00882	−0.00783	−0.00707	−0.00648
Germany	Russia	energy-first	0.00000	−0.00459	−0.00877	−0.00991	−0.00952	−0.00866	−0.00780	−0.00703	−0.00639
Germany	China	GPR-first	−0.01052	0.00384	0.00506	−0.00292	−0.00641	−0.00593	−0.00495	−0.00428	−0.00376
Germany	China	energy-first	0.00000	−0.00437	−0.00413	−0.00548	−0.00607	−0.00558	−0.00487	−0.00422	−0.00369
Germany	MENA	GPR-first	−0.00975	−0.01425	−0.01112	−0.00582	−0.00294	−0.00208	−0.00174	−0.00139	−0.00107
Germany	MENA	energy-first	0.00000	−0.00532	−0.00619	−0.00510	−0.00366	−0.00271	−0.00207	−0.00165	−0.00128
Germany	Western	GPR-first	−0.03240	−0.01634	−0.00825	−0.00630	−0.00541	−0.00483	−0.00437	−0.00396	−0.00356
Germany	Western	energy-first	0.00000	−0.00727	−0.00854	−0.00758	−0.00612	−0.00538	−0.00477	−0.00427	−0.00381
France	Russia	GPR-first	−0.02596	0.00006	0.00111	−0.00422	−0.00626	−0.00576	−0.00495	−0.00435	−0.00381
France	Russia	energy-first	0.00000	−0.00115	−0.00587	−0.00736	−0.00686	−0.00583	−0.00489	−0.00423	−0.00360
France	China	GPR-first	−0.00976	−0.00312	−0.00019	−0.00056	−0.00063	−0.00027	0.00004	0.00018	0.00029
France	China	energy-first	0.00000	−0.00195	−0.00156	−0.00148	−0.00114	−0.00064	−0.00028	0.00001	0.00018
France	MENA	GPR-first	−0.00006	−0.01009	−0.00868	−0.00435	−0.00194	−0.00092	−0.00054	−0.00035	−0.00022
France	MENA	energy-first	0.00000	−0.00535	−0.00492	−0.00358	−0.00235	−0.00151	−0.00109	−0.00083	−0.00058
France	Western	GPR-first	−0.02577	−0.00864	−0.00436	−0.00117	0.00045	0.00130	0.00159	0.00157	0.00140
France	Western	energy-first	0.00000	−0.00255	−0.00562	−0.00298	−0.00084	0.00038	0.00087	0.00097	0.00099
Italy	Russia	GPR-first	−0.03100	−0.01027	−0.00660	−0.01000	−0.00920	−0.00707	−0.00525	−0.00394	−0.00294
Italy	Russia	energy-first	0.00000	−0.00439	−0.00976	−0.01188	−0.00994	−0.00753	−0.00555	−0.00410	−0.00303
Italy	China	GPR-first	−0.01508	−0.00113	0.00205	−0.00349	−0.00557	−0.00510	−0.00416	−0.00337	−0.00276
Italy	China	energy-first	0.00000	−0.00128	−0.00229	−0.00521	−0.00596	−0.00530	−0.00424	−0.00340	−0.00268
Italy	MENA	GPR-first	−0.01385	−0.01921	−0.01169	−0.00364	0.00062	0.00129	0.00116	0.00090	0.00068
Italy	MENA	energy-first	0.00000	−0.00754	−0.00542	−0.00290	−0.00041	0.00043	0.00053	0.00046	0.00037
Italy	Western	GPR-first	−0.02386	−0.01664	−0.01246	−0.00849	−0.00461	−0.00295	−0.00216	−0.00161	−0.00119
Italy	Western	energy-first	0.00000	−0.00581	−0.00968	−0.00871	−0.00568	−0.00391	−0.00286	−0.00210	−0.00152
Spain	Russia	GPR-first	−0.03386	−0.00821	−0.00597	−0.00908	−0.00898	−0.00711	−0.00540	−0.00410	−0.00312
Spain	Russia	energy-first	0.00000	−0.00363	−0.01006	−0.01129	−0.00962	−0.00726	−0.00553	−0.00423	−0.00324
Spain	China	GPR-first	−0.01699	−0.00158	0.00227	−0.00105	−0.00221	−0.00194	−0.00151	−0.00121	−0.00101
Spain	China	energy-first	0.00000	−0.00281	−0.00256	−0.00304	−0.00271	−0.00216	−0.00170	−0.00137	−0.00113
Spain	MENA	GPR-first	−0.00600	−0.01072	−0.00881	−0.00381	−0.00067	0.00032	0.00045	0.00041	0.00037
Spain	MENA	energy-first	0.00000	−0.00244	−0.00422	−0.00301	−0.00089	0.00012	0.00037	0.00033	0.00029
Spain	Western	GPR-first	−0.03823	−0.01668	−0.00527	−0.00055	0.00156	0.00188	0.00157	0.00120	0.00091
Spain	Western	energy-first	0.00000	−0.00166	−0.00264	−0.00112	0.00081	0.00132	0.00116	0.00087	0.00066

7.3.4 Additional Robustness of the Negative Gas Response

Additional diagnostics examine whether the negative Gas response is disproportionately driven by the exceptional European energy-market observations surrounding 2022.

The first sample-sensitivity exercise excludes 2021Q4–2022Q4 while preserving the calendar and lag structure. The Gas impact response remains negative in 15 of the 16 country-origin systems. The only sign change occurs for France–MENA, whose baseline response is already close to zero.

The second exercise re-estimates the BVARs using observations only up to 2022Q2. The Gas impact response remains negative in all 16 systems and retains the same sign as the corresponding full-sample estimate in every case. The 68% credible interval changes from excluding to including zero, or vice versa, in only 3 of the 16 systems.

Together with the alternative-ordering exercises reported above, these results indicate that the negative Gas response is not solely driven by the recursive identification or by observations associated with the 2022 energy crisis. Its magnitude and longer-horizon persistence nevertheless remain specification-sensitive.

7.4 Prior on Σ: Specification and Consequences

The treatment of the residual covariance matrix, Σ, requires particular care because the variables entering the BVAR are measured on markedly different scales. The baseline specification used in this thesis sets

$$S = 0, \qquad \alpha = 0,$$

which yields the Jeffreys-type kernel

$$p(\Sigma) \propto |\Sigma|^{-(N+1)/2}.$$

This specification introduces no fixed scale matrix into the covariance prior, leaving the scale of Σ primarily determined by the likelihood.

The choice is also broadly consistent with the diffuse treatment of the residual covariance matrix in the replication code of Caldara & Iacoviello (2022). In their main specification (`GPRBASELINE`), the covariance-prior inputs are set so that no additional fixed scale information is introduced. Their implementation is not identical to the one used here, but both approaches avoid imposing an arbitrary common unit scale on the residual variances.

To assess sensitivity to this choice, two alternative specifications are considered while keeping $\alpha = 0$ and all remaining elements of the BVAR unchanged. The first is a scale-aware specification,

$$S = \operatorname{diag}(\hat{\Sigma}_{\text{OLS}}),$$

where $\hat{\Sigma}_{\text{OLS}}$ is obtained from the corresponding unadjusted homoskedastic VAR. The same homoskedastic regression is used to construct the equation-specific scale parameters entering the Minnesota prior. This alternative therefore introduces non-zero scale information into the covariance prior while preserving the empirical relative scale of the individual equations. Because S is calibrated from the same sample, this specification is interpreted as a scale-aware sensitivity exercise rather than as an independently elicited prior.

The scale-aware specification produces results that are almost indistinguishable from the baseline. Across the 240 country-origin-variable-horizon comparisons, the posterior median response retains the same sign in every case. Only 2 of 240 responses change their 68% significance status. The width of the 68% credible intervals increases by only 0.8% on average, with a median ratio of 1.007 relative to the baseline. The two significance changes concern the impact response of IPI to Chinese GPR in France and the one-year CLIFS response to MENA GPR in Spain; in both cases the posterior median is effectively unchanged.

The geographical heterogeneity result is similarly stable. Under the baseline, the numbers of significant responses at the impact and one-year horizons are 23 for Russia, 15 for China, 11 for MENA, and 11 for the Western bloc. Under the scale-aware specification the corresponding counts are 23, 14, 10, and 11. The exact permutation test continues to identify Russia as the only origin displaying a statistically unusual concentration of significant responses, with a p-value of 0.0234 compared with 0.0273 in the baseline.

As an additional stress test, the model is also estimated under

$$S = I_N, \qquad \alpha = 0.$$

Unlike the scale-aware specification, the identity matrix assigns the same absolute prior scale to every residual variance. This is a particularly strong perturbation in the present application because the variables are expressed in substantially different units and exhibit very different residual variances. The specification is therefore not invariant to rescaling of the endogenous variables and should be interpreted as a deliberately stringent prior-sensitivity exercise.

The consequences of this stress test are substantially larger. Posterior median signs agree with the baseline in 87.9% of the 240 comparisons, while 66 responses change their 68% significance status. The median credible-band width increases by 39.1%, and the mean width ratio rises to 3.284, reflecting several particularly large increases for variables with small residual scales. The number of significant responses falls to 7 for Russia, 4 for China, 3 for MENA, and 4 for the Western bloc. Russia remains the origin with the largest count, although its exact permutation p-value increases to 0.0781.

Table 14 summarises the overall posterior sensitivity, while Table 15 reports the corresponding geographical heterogeneity results.

Table 14: Sensitivity of posterior inference to the covariance-prior scale matrix

Specification	Sign agree.	Sig. changes	Mean width	Median width
Scaled S	100.0%	2/240	1.008	1.007
$S = I_N$	87.9%	66/240	3.284	1.391

Notes: Scaled S denotes $S = \mathrm{diag}(\hat{\Sigma}_{\mathrm{OLS}})$. Sign agreement refers to the sign of the posterior median relative to the baseline $S = 0$ specification. Significance is defined by whether the 68% credible interval excludes zero. Mean and median width report the ratio of the credible-band width relative to the corresponding baseline interval.

Table 15: Origin heterogeneity under alternative covariance-prior scale matrices

Origin	Baseline $S = 0$	Scaled S	$S = I_N$
Russia	23 (0.0273)	23 (0.0234)	7 (0.0781)
China	15 (0.5430)	14 (0.5664)	4 (0.7188)
MENA	11 (0.9063)	10 (0.9531)	3 (1.0000)
Western	11 (0.9063)	11 (0.8359)	4 (0.7188)

Notes: Scaled S denotes $S = \mathrm{diag}(\hat{\Sigma}_{\mathrm{OLS}})$. Entries report the number of responses whose 68% credible interval excludes zero at the impact and one-year horizons. Exact permutation-test p-values are reported in parentheses. The permutation distribution is obtained by enumerating all $4^4 = 256$ possible origin assignments across the four countries.

Overall, the scale-aware sensitivity exercise provides little evidence that the substantive results depend on the zero-scale covariance prior: posterior signs are unchanged, credible-band widths are almost identical, and the geographical heterogeneity result for Russian GPR shocks is preserved. By contrast, the identity-matrix stress test materially weakens posterior precision. Given the heterogeneous measurement scales of the endogenous variables, this latter result is interpreted as evidence that imposing a common absolute unit scale can be consequential, rather than as evidence that the baseline findings are generally fragile to reasonable changes in the covariance prior.

7.5 Levels versus Differences

The baseline specification uses log-differences for Brent, Gas, HICP, and IPI, while GPR enters as log(1+GPR) and CLIFS in levels. As a robustness exercise, all 16 country-origin BVARs are re-estimated using

$$[\log(1 + \mathrm{GPR}),\ \log(\mathrm{Brent}),\ \log(\mathrm{Gas}),\ \log(\mathrm{HICP}),\ \log(\mathrm{IPI}),\ \mathrm{CLIFS}].$$

The lag length, recursive ordering, Minnesota hyperparameters, covariance prior, frozen COVID-19 volatility path, Gibbs-sampling settings, and impulse-response horizon are maintained. Prior scales are recalculated from the corresponding unadjusted homoskedastic VAR, consistently with the baseline prior construction. The own first-lag prior mean is set to one for Brent, Gas, HICP, and IPI, while GPR retains its estimated AR(1) prior mean and CLIFS retains a prior mean of zero.

7.5.1 Cointegration Diagnostics

Johansen tests are first applied to the four log-level variables that enter the baseline in first differences:

$$[\log(\text{Brent}),\ \log(\text{Gas}),\ \log(\text{HICP}),\ \log(\text{IPI})].$$

With one lag in differences, corresponding to the two-lag VAR specification, both the trace and maximum-eigenvalue tests identify a cointegration rank of one at the 5% level in Germany, France, Italy, and Spain. In every country, the null of no cointegration is rejected, while the null of rank at most one is not rejected.

Table 16: Johansen cointegration rank for the persistent-variable block

Country	Observations	Trace rank	Max-eigen rank
Germany	115	1	1
France	115	1	1
Italy	115	1	1
Spain	115	1	1

Notes: The tested variables are log Brent, log Gas, log HICP, and log IPI. Ranks are determined sequentially at the 5% level. The Johansen specification uses a constant deterministic component and one lag in differences.

As an additional diagnostic, the Johansen procedure is also applied to the 16 complete six-variable systems,

$$[\log(1+\text{GPR}),\ \log(\text{Brent}),\ \log(\text{Gas}),\ \log(\text{HICP}),\ \log(\text{IPI}),\ \text{CLIFS}].$$

The trace and maximum-eigenvalue procedures select the same rank in 9 of the 16 systems. Since this exercise combines the four persistent variables with GPR and CLIFS, which have different baseline treatments, the full-system ranks are interpreted only as a supplementary diagnostic.

Table 17: Johansen rank combinations in the full six-variable systems

Trace rank	Max-eigen rank	Number of systems
1	0	1
1	1	3
2	0	2
2	1	4
2	2	6

The rank-one result for the four-variable block indicates that the differenced baseline does not explicitly model the long-run equilibrium relationship among Brent, Gas, HICP, and IPI.

7.5.2 Direct Levels Robustness

The baseline and levels specifications are compared at the impact, one-year, and two-year horizons. Because impulse responses of log-differenced and log-level variables have different quantitative interpretations, the exercise focuses on qualitative comparisons of posterior-median signs and 68% significance status.

Across 240 country-origin-variable-horizon comparisons, 178 responses retain the same sign, corresponding to an agreement rate of 74.2%. In addition, 87 of 240 responses change significance status.

Table 18: Sign agreement between baseline and levels specifications

Response variable	Same sign	Agreement rate
Brent	41/48	85.4%
Gas	28/48	58.3%
HICP	39/48	81.3%
IPI	34/48	70.8%
CLIFS	36/48	75.0%
Overall	178/240	74.2%

Sign agreement is highest on impact, at 67 of 80 responses (83.8%), compared with 55 of 80 (68.8%) at one year and 56 of 80 (70.0%) at two years. Gas displays the greatest sensitivity to the transformation choice, retaining the same sign in 28 of 48 comparisons.

7.5.3 Origin Heterogeneity under the Levels Specification

The exact origin-heterogeneity test is repeated under the levels specification using the same criterion as in the baseline: the number of responses whose 68% credible interval excludes zero at the impact and one-year horizons. The permutation distribution is obtained by enumerating all $4^4 = 256$ possible origin assignments across the four countries.

Table 19: Origin heterogeneity: baseline versus levels specification

Origin	Baseline count	Baseline p	Levels count	Levels p
Russia	23	0.0273	16	0.7852
China	15	0.5430	14	0.9180
MENA	11	0.9063	21	0.1016
Western	11	0.9063	18	0.4453

Notes: Counts refer to significant responses at the impact and one-year horizons, with significance defined by a 68% credible interval excluding zero. Exact p-values are obtained from all $4^4 = 256$ possible origin assignments across countries.

The geographical concentration associated with Russian GPR shocks in the baseline is not preserved under the levels specification. Russia's count falls from 23 to 16 and its exact permutation-test p-value increases from 0.0273 to 0.7852. MENA records the largest count under the levels specification, with 21 responses, but its p-value of 0.1016 remains above conventional significance levels. Thus, no origin exhibits a statistically unusual concentration of significant responses at the 5% level under the levels specification.

Overall, the exercise shows that the treatment of the persistent variables is a substantive modelling choice. The differenced specification is retained as the baseline because the main objective is to study the short and medium run transmission of GPR shocks, while Brent, Gas, HICP, and IPI are transformed to achieve stationarity and to express their responses in terms of quarterly changes. The cointegration evidence refers to long-run relationships among these persistent macro-energy variables and therefore motivates the levels specification as an important robustness check. The results show that some conclusions, particularly the Gas response and the geographical heterogeneity finding, are sensitive to this modelling choice.

7.6 COVID-19 Volatility Treatment

The baseline estimation allows for a temporary increase in reduced-form innovation volatility around the COVID-19 period. Following the treatment described in Section 4.3, the reduced-form innovations satisfy

$$u_t \sim \mathcal{N}\left(0, s_t^2\Sigma\right),$$

where s_t is a scalar volatility factor common to all equations of a given VAR at date t. The specification preserves the contemporaneous covariance structure in Σ while allowing the overall magnitude of reduced-form shocks to increase temporarily during the pandemic.

7.6.1 Volatility Path

The temporary volatility path is parameterised by

$$\theta = (s_0, s_1, s_2, \rho),$$

where s_0, s_1, and s_2 determine the volatility factors in the first three quarters of 2020 and ρ governs the subsequent decay towards the baseline variance. Specifically,

$$
s_t = \begin{cases} s_0, & t = 2020\text{Q}1, \\ s_1, & t = 2020\text{Q}2, \\ s_2, & t = 2020\text{Q}3, \\ 1 + (s_2 - 1)\rho^j, & t = 2020\text{Q}4, \dots, 2021\text{Q}3, \\ 1, & \text{otherwise}, \end{cases}
$$

where $j = 1, \dots, 4$ for the four quarters from 2020Q4 to 2021Q3. The volatility factor therefore returns mechanically to one from 2021Q4 onwards.

The parameters satisfy the restrictions

$$
s_0 > 0, \qquad s_1 > 0, \qquad s_2 > 0, \qquad 0 < \rho < 1,
$$

ensuring positive volatility scales and a declining post-2020Q3 adjustment.

7.6.2 First-Stage Estimation

The volatility parameters are estimated separately for each country-origin VAR using concentrated Gaussian maximum likelihood. For a candidate value of θ, the corresponding volatility path $s_t(\theta)$ is constructed and the VAR is transformed as

$$
Y_t^* = \frac{Y_t}{s_t(\theta)}, \qquad X_t^* = \frac{X_t}{s_t(\theta)}.
$$

Conditional on θ, the transformed system is homoskedastic,

$$
Y_t^* = X_t^* B + u_t^*, \qquad u_t^* \sim \mathcal{N}(0, \Sigma).
$$

The VAR coefficients and covariance matrix can therefore be concentrated out for each candidate volatility path. The resulting concentrated Gaussian likelihood is maximised with respect to (s_0, s_1, s_2, ρ), producing the first-stage estimate

$$
\hat{\theta} = (\hat{s}_0, \hat{s}_1, \hat{s}_2, \hat{\rho}).
$$

The corresponding estimated quarterly volatility path is then constructed for the full sample.

7.6.3 Second-Stage VAR and BVAR Estimation

The estimated volatility parameters are held fixed in the second stage. Let

$$\hat{s}_t = s_t(\hat{\theta}).$$

After accounting for the two VAR lags, the effective volatility path is aligned with the estimation sample and both the dependent variables and regressors are divided by $\hat{s}_t$:

$$Y_t^* = \frac{Y_t}{\hat{s}_t}, \qquad X_t^* = \frac{X_t}{\hat{s}_t}.$$

The classical VAR is estimated by OLS on this transformed system, while the BVAR applies the Minnesota coefficient prior and covariance prior described in Section 4.6 to the same transformed observations. Consequently, both estimators use an identical COVID-19 volatility correction.

Importantly, the Minnesota residual scales used to construct the prior variance matrix are obtained from the corresponding unadjusted homoskedastic VAR before applying the COVID-19 weighting. The temporary volatility correction therefore affects the likelihood in the second-stage BVAR but does not redefine the scale calibration of the Minnesota prior.

This procedure is deliberately two-stage. The first-stage estimate $\hat{\theta}$ is treated as fixed when estimating the VAR and BVAR, so uncertainty surrounding the temporary-volatility parameters is not propagated into the posterior distribution or the classical bootstrap intervals. The approach therefore differs from a fully joint Bayesian model in which θ, B, and Σ would be sampled simultaneously.

7.6.4 Implementation Checks

The estimated volatility parameters and complete quarterly scale paths used in the production models are stored prior to estimation and subsequently reused without re-estimation in the baseline and robustness exercises. This ensures that comparisons across alternative priors, recursive orderings, and reduced systems do not inadvertently reflect changes in the estimated COVID-19 volatility path.

The estimated volatility paths were checked to ensure that the adjustment is restricted to the intended COVID-19 period and returns to unity from 2021Q4 onwards.

7.7 Full Structural IRFs – All 16 Country-Origin Systems

This section reports the complete set of structural impulse response functions for all 16 country-origin systems, with 68% credible bands (shaded areas) around the posterior medians (solid lines). The main text (Section 5) summarises the responses across countries, origins, variables, and selected horizons, while the complete impulse-response plots are reported here for completeness.

Each panel uses its own y-axis scale, reflecting the different magnitudes of the six variables. Cross-system comparisons in the main text are therefore based on posterior summaries rather than direct visual comparisons of panel scales.

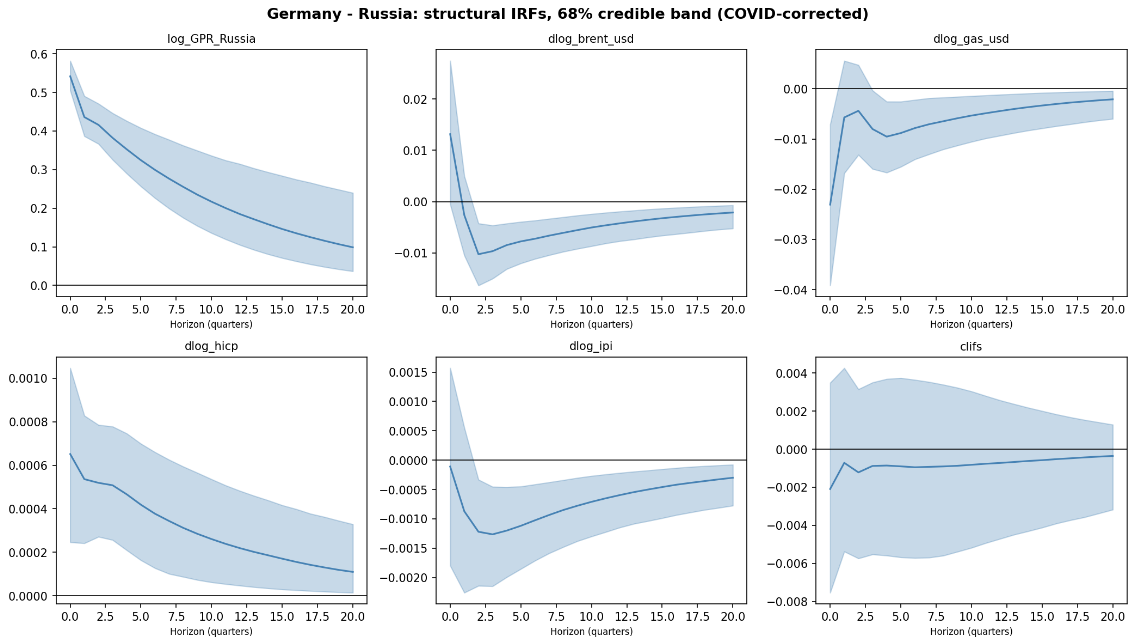

Figure 2: Structural IRFs, Germany–Russia system, 68% credible band

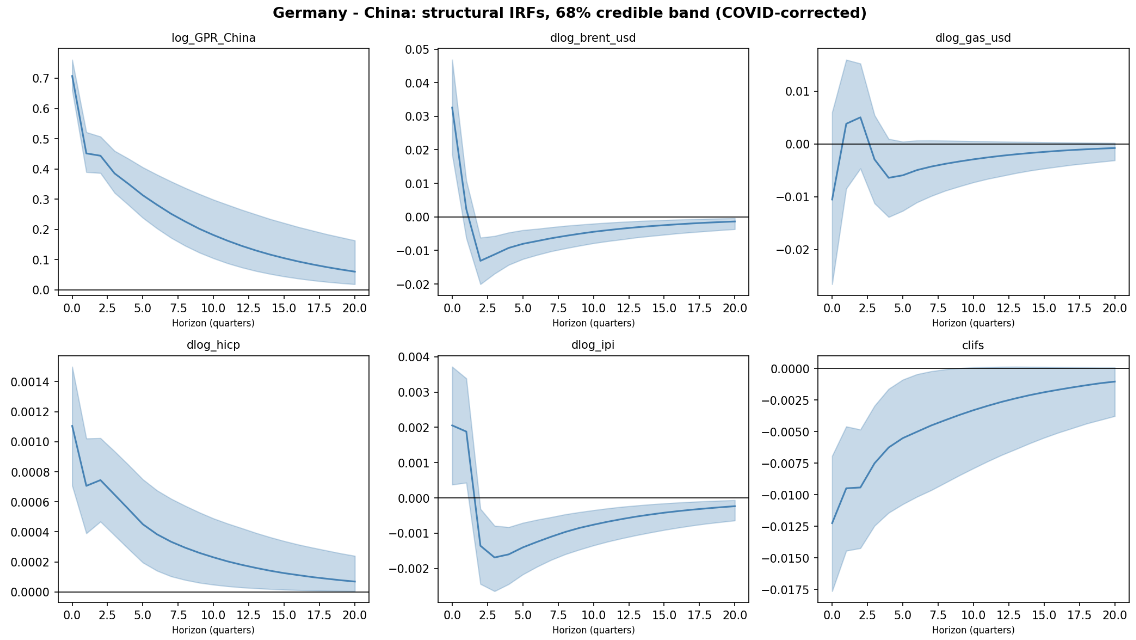

Figure 3: Structural IRFs, Germany–China system, 68% credible band

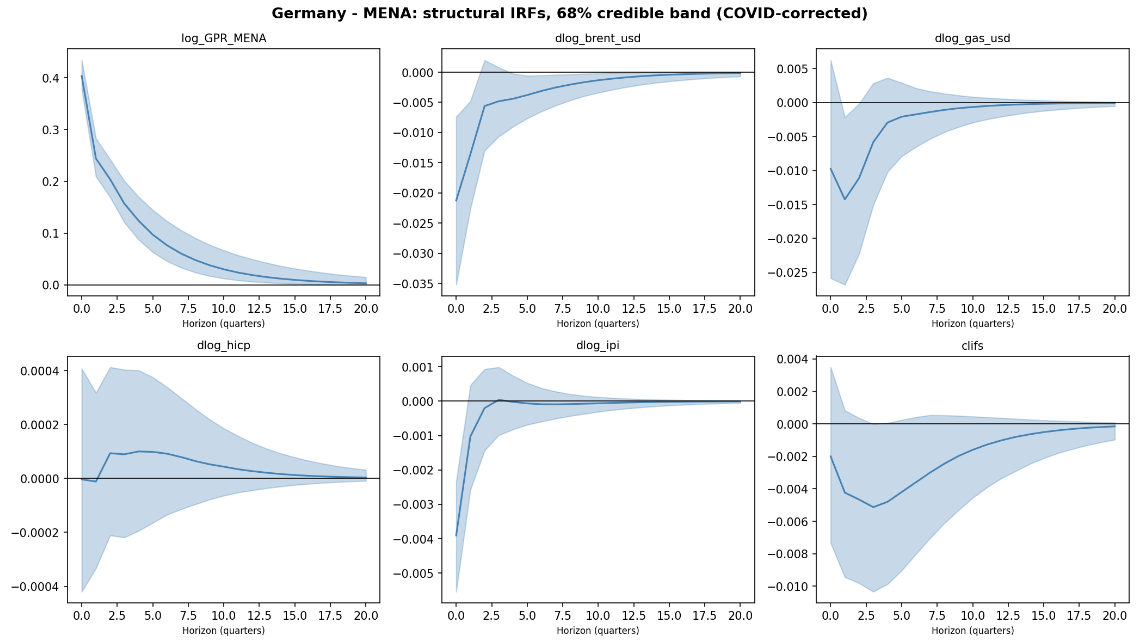

Figure 4: Structural IRFs, Germany–MENA system, 68% credible band

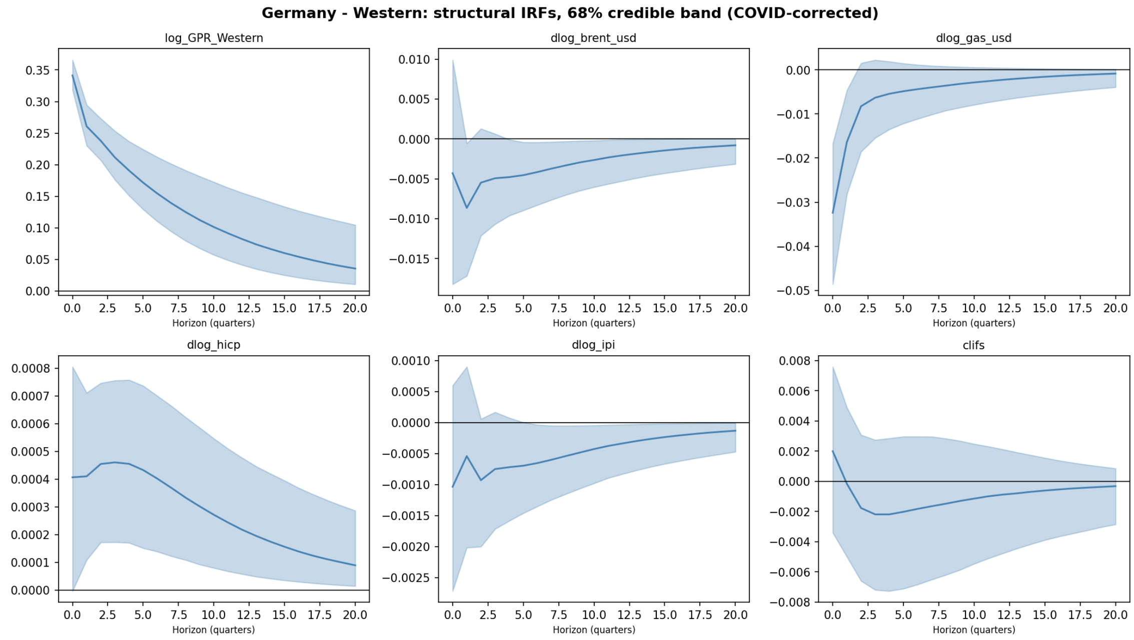

Figure 5: Structural IRFs, Germany–Western system, 68% credible band

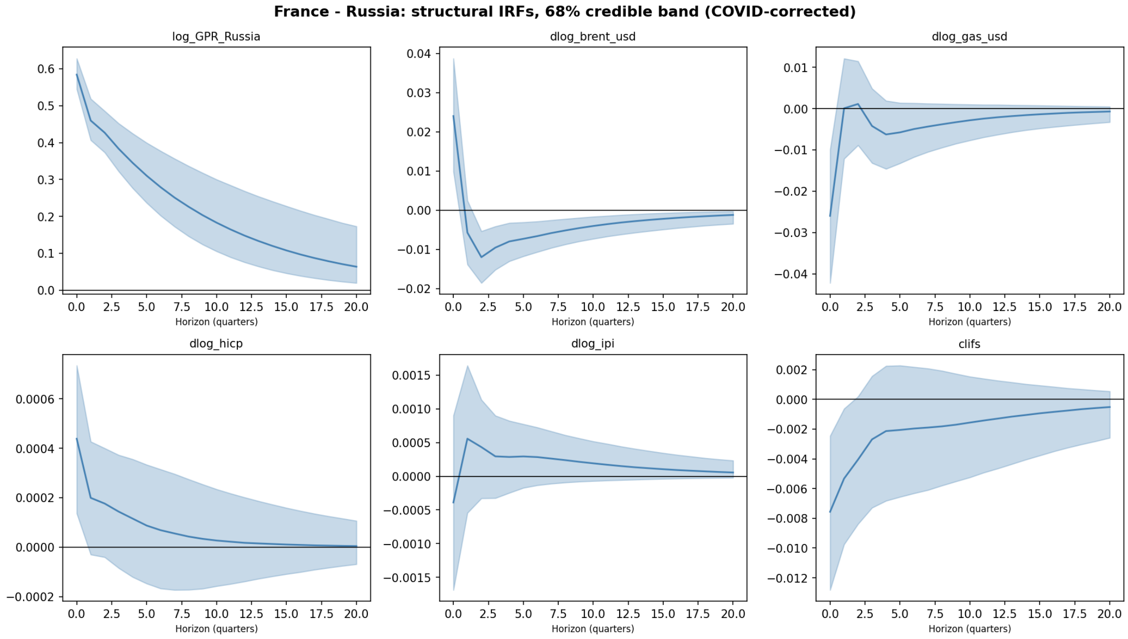

Figure 6: Structural IRFs, France–Russia system, 68% credible band

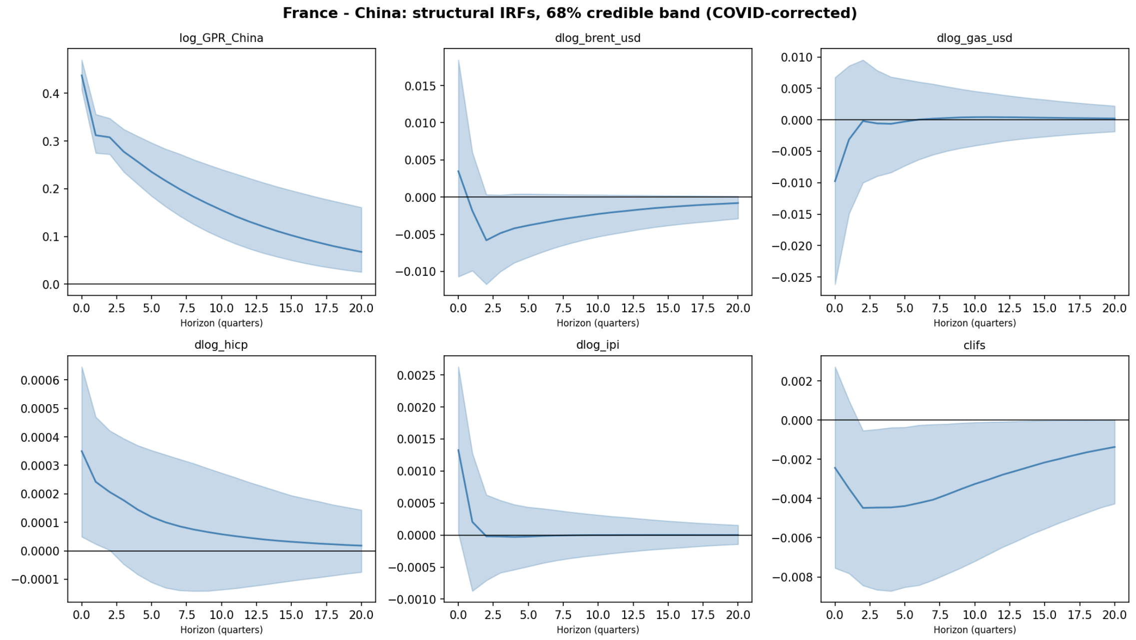

Figure 7: Structural IRFs, France–China system, 68% credible band

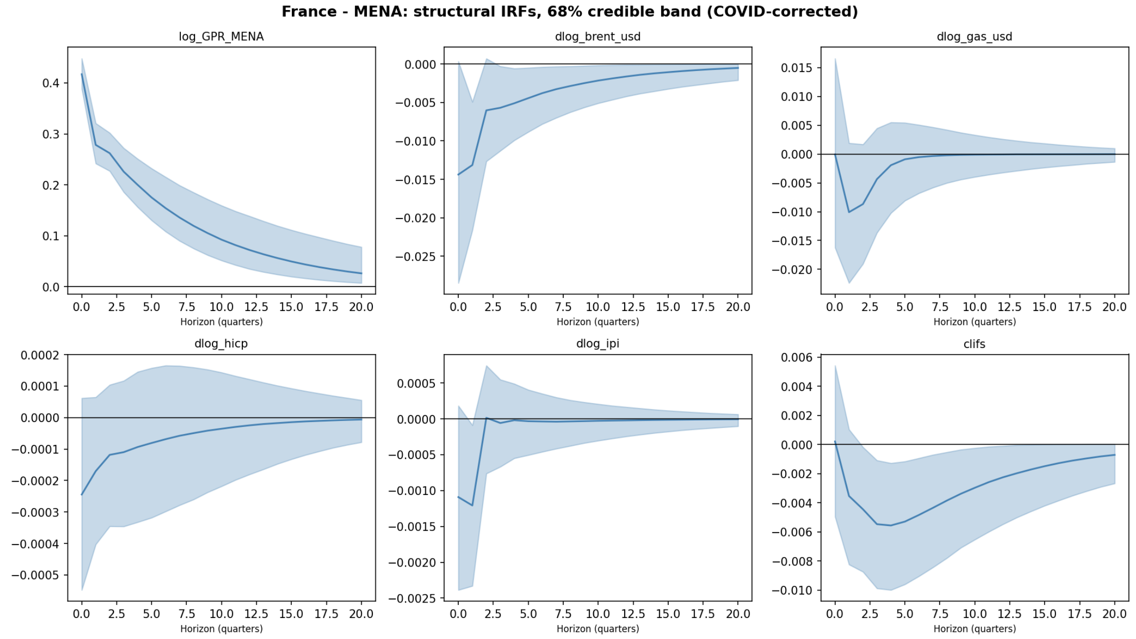

Figure 8: Structural IRFs, France–MENA system, 68% credible band

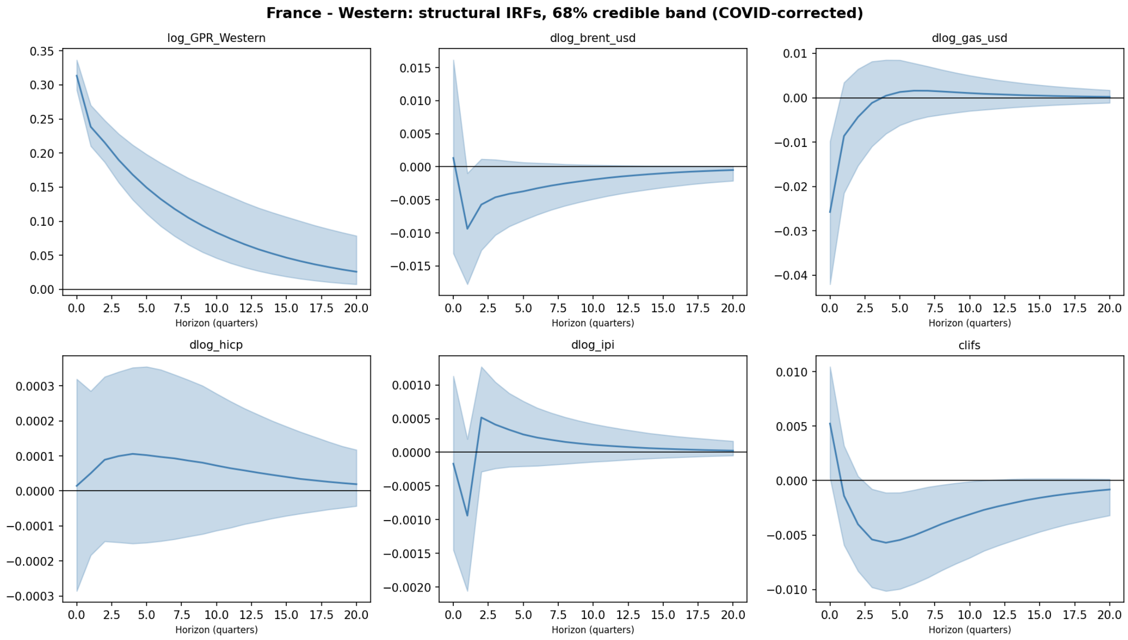

Figure 9: Structural IRFs, France–Western system, 68% credible band

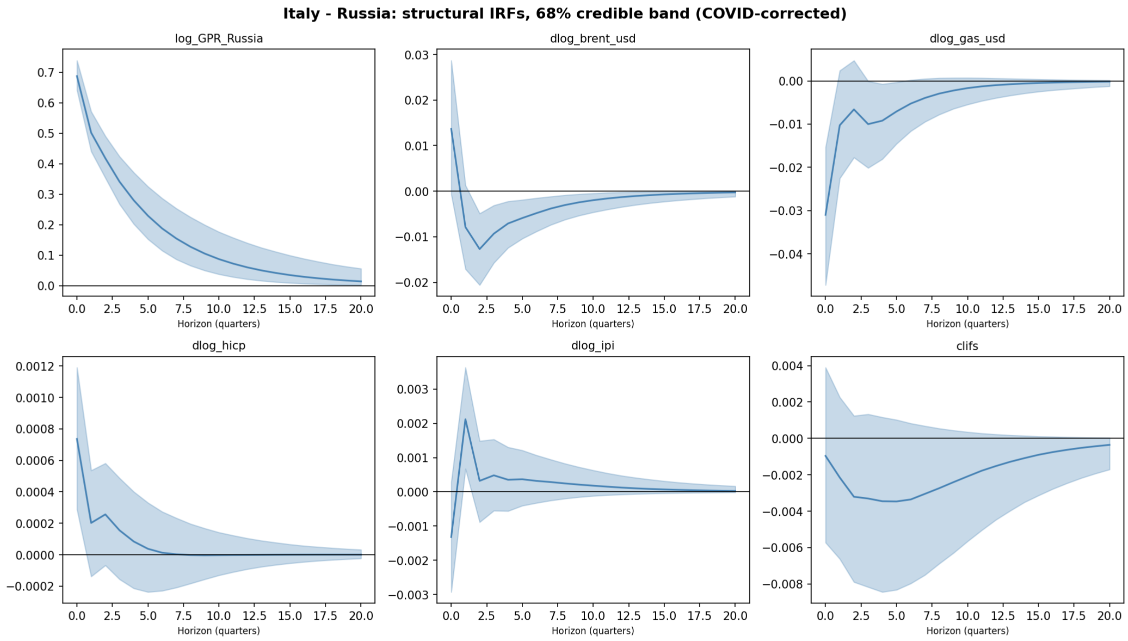

Figure 10: Structural IRFs, Italy–Russia system, 68% credible band

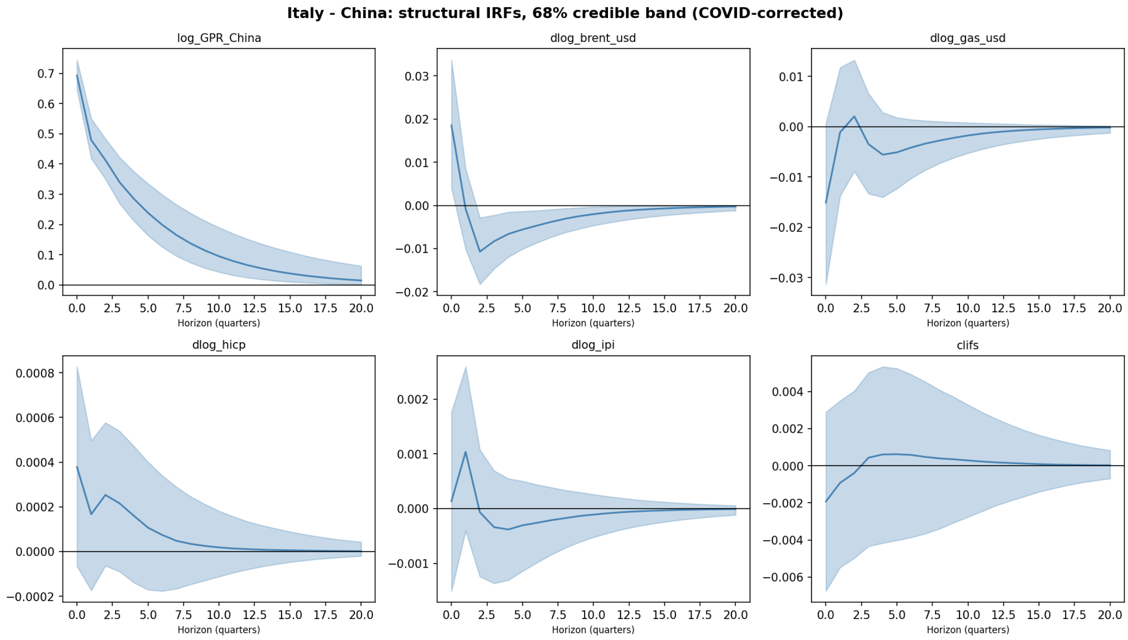

Figure 11: Structural IRFs, Italy–China system, 68% credible band

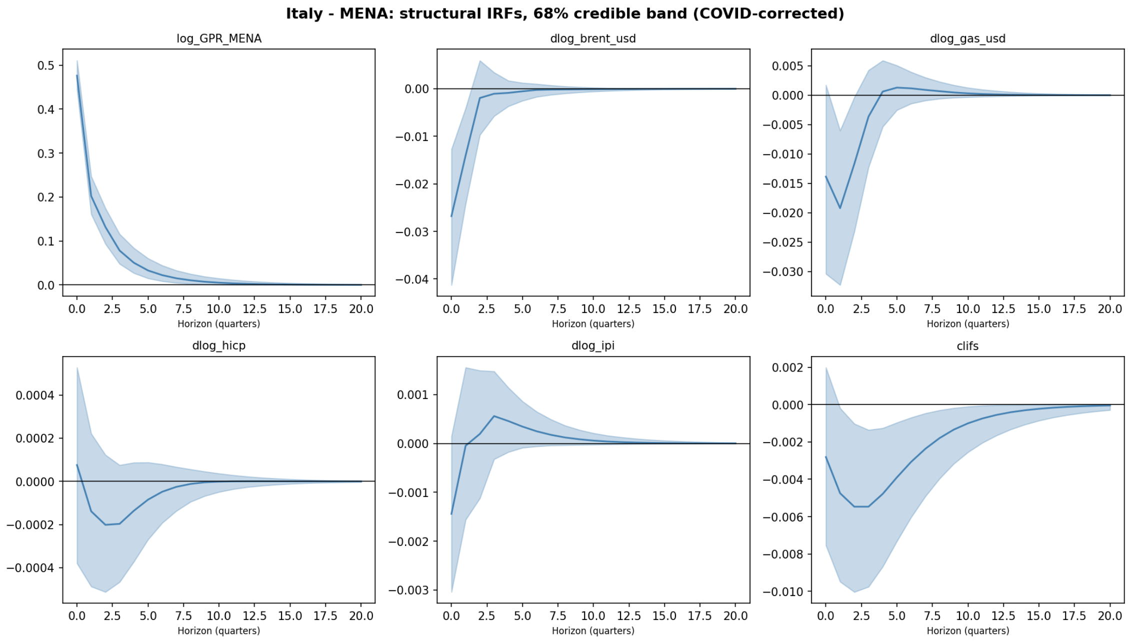

Figure 12: Structural IRFs, Italy–MENA system, 68% credible band

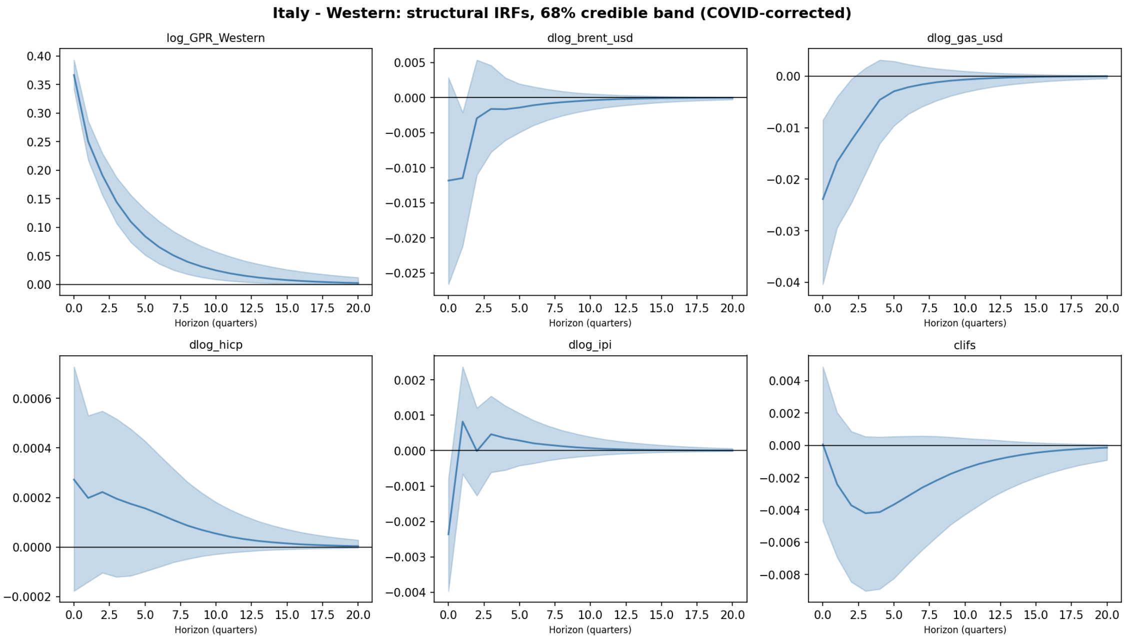

Figure 13: Structural IRFs, Italy–Western system, 68% credible band

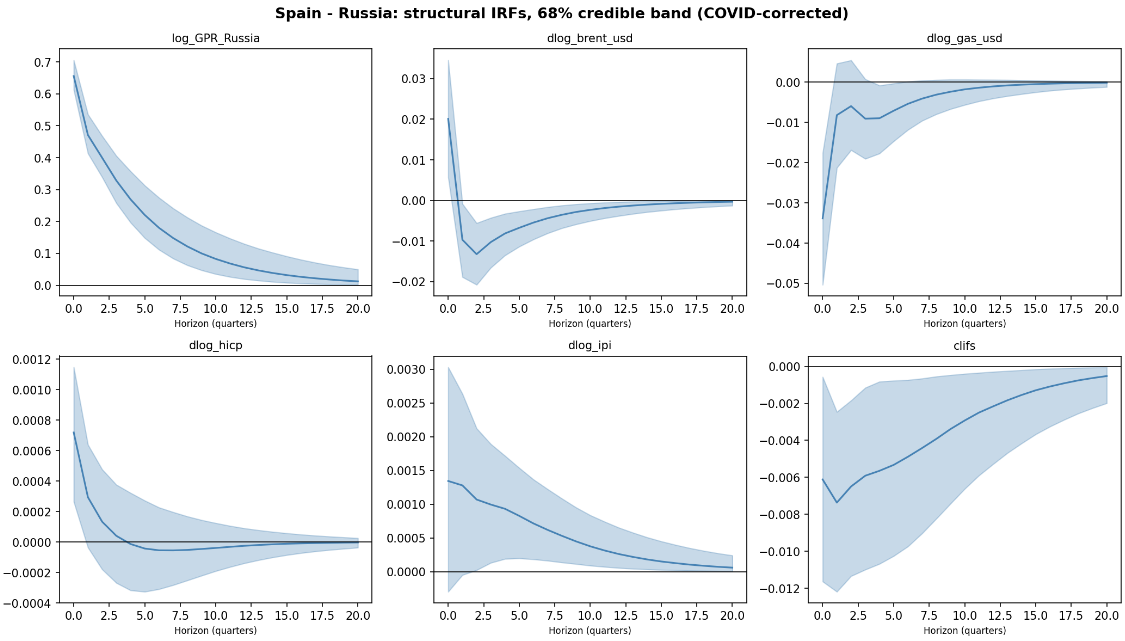

Figure 14: Structural IRFs, Spain–Russia system, 68% credible band

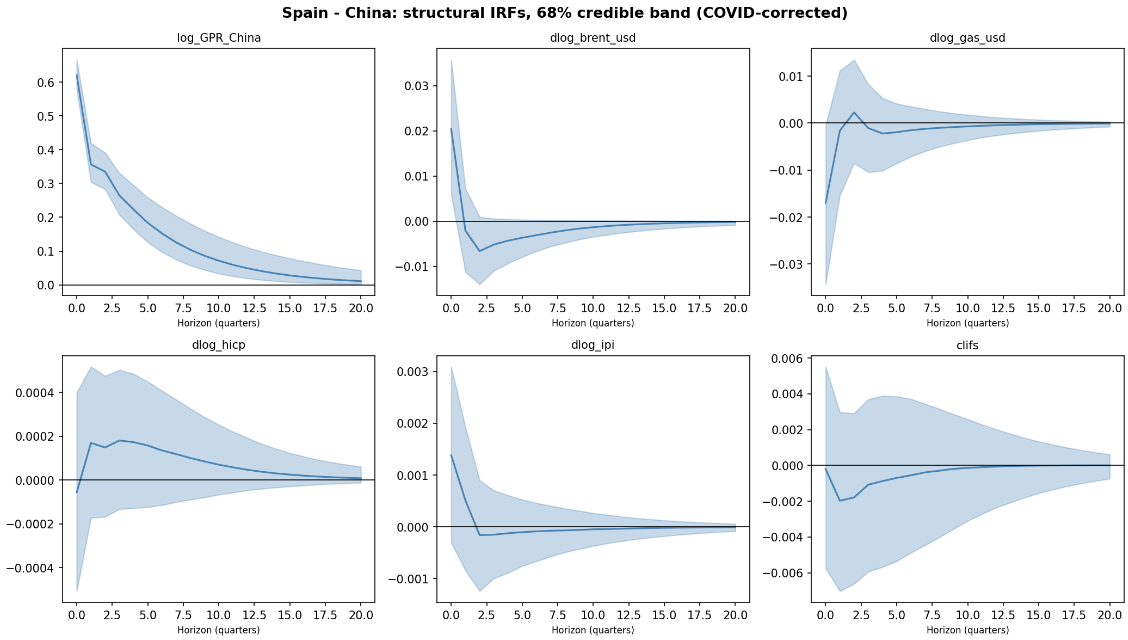

Figure 15: Structural IRFs, Spain–China system, 68% credible band

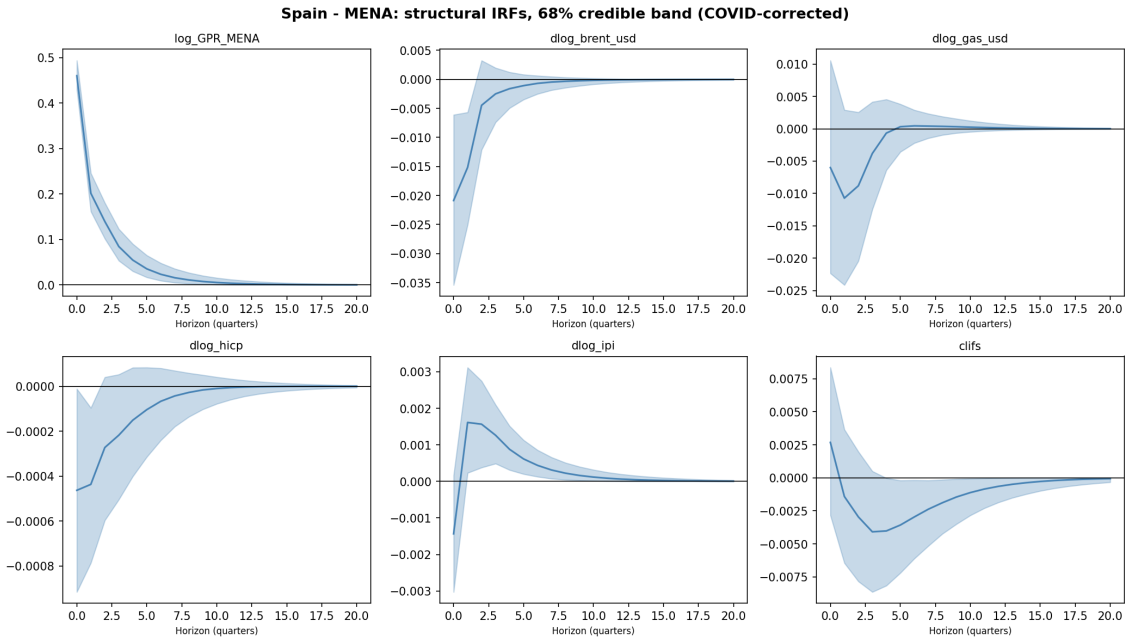

Figure 16: Structural IRFs, Spain–MENA system, 68% credible band

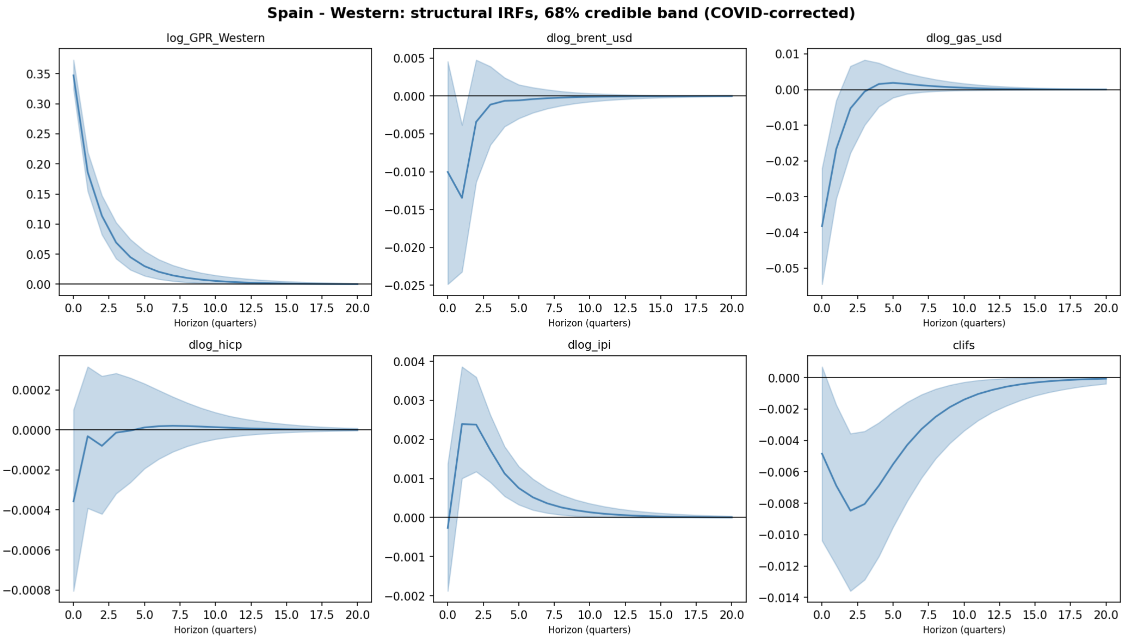

Figure 17: Structural IRFs, Spain–Western system, 68% credible band

7.8 Model Verification and Validation

This appendix reports numerical checks of the implementation used for the baseline BVAR and classical VAR. The purpose of these exercises is not to establish the validity of the empirical identifying assumptions, but to verify that the estimation, prior, covariance, structural-IRF, and interval-construction routines behave as expected when applied to data generated from a known process.

7.8.1 Synthetic Data-Generating Process

A stable six-variable VAR(2) is simulated with the same number of endogenous variables, lag order, and sample length as the empirical systems. The sample contains $T = 115$ observations, leaving $T_{\text{eff}} = 113$ observations after accounting for two lags. The largest modulus of the companion-matrix eigenvalues is

$$\rho(\mathcal{A}) = 0.9547 < 1,$$

confirming dynamic stability.

The first-lag own-variable coefficients of the simulated system are

$$\text{diag}(A_1) = (0.85, 0.20, 0.20, 0.15, 0.15, 0.20),$$

with an additional dynamic effect from GPR to Gas,

$$(A_1)_{\text{Gas,GPR}} = 0.05.$$

The second-lag coefficient matrix is set to $A_2 = 0.1 I_N$. A positive contemporaneous correlation of 0.30 is also imposed between the GPR and Gas innovations. The resulting true structural GPR-to-Gas impulse response equals

$$\Theta_{\text{Gas,GPR}}(0) = 0.0450$$

on impact and

$$\Theta_{\text{Gas,GPR}}(4) = 0.00943$$

after four quarters.

The COVID-19 volatility path is fixed at unity throughout the simulation so that the exercise evaluates the remaining components of the estimation procedure independently of the temporary-volatility correction.

7.8.2 Parameter Recovery

The true residual standard deviations used to generate the data are

$$(0.150,\ 0.100,\ 0.150,\ 0.006,\ 0.015,\ 0.020).$$

Under the baseline covariance specification, the square roots of the posterior mean variances are

$$(0.1641,\ 0.0942,\ 0.1532,\ 0.0060,\ 0.0155,\ 0.0222).$$

The maximum relative deviation from the true residual scale is 10.8%, indicating that the sampler reproduces the strongly heterogeneous innovation scales reasonably closely in a sample of the same size as the empirical application.

For the own first-lag coefficients, the posterior means are

$$(0.780,\ 0.102,\ 0.073,\ 0.064,\ 0.097,\ 0.103),$$

compared with the true values

$$(0.850,\ 0.200,\ 0.200,\ 0.150,\ 0.150,\ 0.200).$$

These differences are expected under the Minnesota prior. The prior pulls the coefficients of the differenced variables and CLIFS towards zero, while the prior mean for the first own lag of GPR is obtained from an auxiliary AR(1) regression. In this simulated sample, the estimated GPR prior mean is 0.850.

7.8.3 Covariance-Prior Scale Check

The covariance-prior implementation is additionally evaluated by comparing the baseline $S = 0$ specification with the deliberately restrictive alternative $S = I_N$.

Under $S = 0$, the ratios between posterior and true residual scales are

$$(1.09,\ 0.94,\ 1.02,\ 0.99,\ 1.03,\ 1.11),$$

which remain close to unity across all six variables.

Under $S = I_N$, the corresponding ratios become

$$(1.28,\ 1.37,\ 1.22,\ 16.30,\ 6.60,\ 5.02).$$

The distortion is particularly large for variables whose true innovation variances are small. This numerical experiment illustrates why a common absolute prior scale can strongly affect covariance inference when variables have substantially different residual magnitudes, consistent with the empirical sensitivity analysis reported in Appendix 7.4.

7.8.4 Monte Carlo Interval Coverage

A final experiment evaluates the coverage of structural impulse-response intervals. One hundred independent samples are generated from the synthetic VAR. For each replication, the classical VAR is estimated and 500 residual-bootstrap replications are used to construct a 90% interval. The BVAR is estimated using the same Gibbs sampler as in the empirical application, and the 5th and 95th percentiles of the posterior IRF distribution are used to construct a comparable 90% posterior interval.

Coverage is evaluated for the structural response of Gas to the GPR shock at the impact horizon and after four quarters.

Table 20: Monte Carlo coverage of the GPR-to-Gas structural response

	$h = 0$	$h = 4$
Classical VAR, 90% bootstrap interval	89%	85%
BVAR, 90% posterior interval	92%	90%

Notes: Results are based on 100 independently simulated samples. Coverage denotes the proportion of replications in which the corresponding interval contains the known structural impulse response. The true GPR-to-Gas responses are 0.0450 at $h = 0$ and 0.00943 at $h = 4$. The BVAR uses a 90% posterior interval in this validation exercise to facilitate comparison with the 90% classical bootstrap interval. Baseline empirical BVAR results are reported using 68% credible intervals.

The coverage rates are broadly consistent with the nominal 90% level. The classical bootstrap performs almost exactly at the nominal rate on impact and somewhat below it at the four-quarter horizon, while the BVAR achieves 92% and 90% coverage, respectively. Given the finite number of Monte Carlo replications, these results provide no evidence of a substantial distortion in the interval-construction procedures.

REFERENCES

Abid, I., Dhaoui, A., Kaabia, O. & Tarchella, S. (2023), 'Geopolitical risk on energy, agriculture, livestock, precious and industrial metals: New insights from a Markov-Switching model', *Resources Policy* **85**, 103925.
URL: *https://doi.org/10.1016/j.resourpol.2023.103925*

Alonso-Alvarez, I., Bukina, E., Diakonova, M., Khitarishvili, N., Pérez, J. J. & Piqueras, P. (2026), Geopolitical risk: A database of general and bilateral indices, Occasional Paper 2603, Banco de España.
URL: *https://doi.org/10.53479/42445*

Alonso-Alvarez, I., Diakonova, M. & Pérez, J. J. (2025), Rethinking GPR: The sources of geopolitical risk, Working Paper 2522, Banco de España.
URL: *https://doi.org/10.53479/39685*

Asche, F., Misund, B. & Sikveland, M. (2013), 'The relationship between spot and contract gas prices in Europe', *Energy Economics* **38**, 212–217.
URL: *https://doi.org/10.1016/j.eneco.2013.02.010*

Bańbura, M., Giannone, D. & Reichlin, L. (2010), 'Large Bayesian vector auto regressions', *Journal of Applied Econometrics* **25**(1), 71–92.
URL: *https://doi.org/10.1002/jae.1137*

Bondarenko, Y., Kang, N., Lewis, V., Rottner, M. & Schüler, Y. (2026), Geopolitical risk in the Euro Area: Measurement and transmission, Discussion Paper 14/2026, Deutsche Bundesbank. Also published as BIS Working Paper No. 1348.
URL: *https://doi.org/10.71734/DP-2026-14*

Bondarenko, Y., Lewis, V., Rottner, M. & Schüler, Y. (2024), 'Geopolitical risk perceptions', *Journal of International Economics* **152**, 104005.
URL: *https://doi.org/10.1016/j.jinteco.2024.104005*

Breitung, J. & Swanson, N. R. (2002), 'Temporal aggregation and spurious instantaneous causality in multiple time series models', *Journal of Time Series Analysis* **23**(6), 651–665.
URL: *https://doi.org/10.1111/1467-9892.00284*

Caldara, D. & Iacoviello, M. (2022), 'Measuring geopolitical risk', *American Economic Review* **112**(4), 1194–1225.
URL: *https://www.aeaweb.org/articles?id=10.1257/aer.20191823*

Dieckelmann, D., Larkou, C., McQuade, P., Pancaro, C. & Rößler, D. (2025), 'Geopolitical risk and Euro Area bank CDS spreads and stock prices: Evidence from a new index',

Economics Letters **254**, 112461.
URL: *https://doi.org/10.1016/j.econlet.2025.112461*

Duprey, T., Klaus, B. & Peltonen, T. (2017), 'Dating systemic financial stress episodes in the EU countries', *Journal of Financial Stability* **32**, 30–56.
URL: *https://doi.org/10.1016/j.jfs.2017.07.004*

Engle, R. F. & Granger, C. W. J. (1987), 'Co-integration and error correction: Representation, estimation, and testing', *Econometrica* **55**(2), 251–276.
URL: *https://doi.org/10.2307/1913236*

European Central Bank (2026), 'HICP inflation rate—total—index, Euro Area, monthly', ECB Data Portal.
URL: *https://data.ecb.europa.eu/data/datasets/HICP/HICP.M.U2.Y.000000.4F0.INX*

Eurostat (2026), 'Harmonised index of consumer prices (HICP) (prc_hicp)', Eurostat metadata.
URL: *https://ec.europa.eu/eurostat/cache/metadata/en/prc_hicp_esms.htm*

Ferrari Minesso, M., Lappe, M.-S. & Rößler, D. (2023), 'Geopolitical risk and oil prices', ECB Economic Bulletin, Issue 8/2023, Box 2.
URL: *https://www.ecb.europa.eu/press/economic-bulletin/focus/2024/html/ecb.ebbox202308_02~ed883ebf56.en.html*

Johansen, S. (1991), 'Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models', *Econometrica* **59**(6), 1551–1580.
URL: *https://doi.org/10.2307/2938278*

Kilian, L. & Lütkepohl, H. (2017), *Structural Vector Autoregressive Analysis*, Cambridge University Press, Cambridge.
URL: *https://doi.org/10.1017/9781108164818*

Lenza, M. & Primiceri, G. E. (2022), 'How to estimate a vector autoregression after March 2020', *Journal of Applied Econometrics* **37**(4), 688–699.
URL: *https://doi.org/10.1002/jae.2895*

Lin, B. & Li, J. (2015), 'The spillover effects across natural gas and oil markets: Based on the VEC–MGARCH framework', *Applied Energy* **155**, 229–241.
URL: *https://doi.org/10.1016/j.apenergy.2015.05.123*

Litterman, R. B. (1986), 'Forecasting with Bayesian vector autoregressions—five years of experience', *Journal of Business & Economic Statistics* **4**(1), 25–38.
URL: *https://doi.org/10.1080/07350015.1986.10509491*

Marcellino, M. (1999), 'Some consequences of temporal aggregation in empirical analysis', *Journal of Business & Economic Statistics* **17**(1), 129–136.
URL: *https://doi.org/10.1080/07350015.1999.10524802*

Niepmann, F. & Shen, L. S. (2025), Geopolitical risk and global banking, International Finance Discussion Papers 1418, Board of Governors of the Federal Reserve System.
URL: *https://doi.org/10.17016/IFDP.2025.1418*

OECD (2024), 'Consumer price indices—frequently asked questions (FAQs)', OECD Data Explainer.
URL: *https://www.oecd.org/en/data/insights/data-explainers/2024/07/consumer-price-indices-frequently-asked-questions-faqs.html*

Perron, P. (1989), 'The great crash, the oil price shock, and the unit root hypothesis', *Econometrica* **57**(6), 1361–1401.
URL: *https://doi.org/10.2307/1913712*

Plagborg-Møller, M. & Wolf, C. K. (2021), 'Local projections and VARs estimate the same impulse responses', *Econometrica* **89**(2), 955–980.
URL: *https://doi.org/10.3982/ECTA17813*

Salgado, D. (2019), Manual de usuario de JDemetra+, Technical report, Instituto Nacional de Estadística, Madrid.
URL: *https://www.ine.es/clasifi/manual_jdemetra.pdf*

Sims, C. A. (1980), 'Macroeconomics and reality', *Econometrica* **48**(1), 1–48.
URL: *https://doi.org/10.2307/1912017*

Verduzco-Bustos, G. & Zanetti, F. (2026), The effects of geopolitical oil price shocks, CESifo Working Paper 12606, CESifo, Munich.
URL: *https://doi.org/10.65864/qapdei68on*