Revisiting Fundamentals of the European Gas Market: The Role of Supply Substitution

Abstract

The 2022 European energy crisis exposed the central role of supply substitution in natural gas markets, as disruptions to Russian pipeline flows were accompanied by a sharp surge in liquefied natural gas (LNG) imports. This paper revisits the fundamentals of the European natural gas market by explicitly distinguishing between pipeline gas and LNG supplies. Using a Bayesian Structural Vector Autoregression (SVAR) identified through sign and elasticity restrictions, the analysis jointly identifies the contemporaneous elasticities of natural gas prices, supply by source, inventories, and euro area industrial production. The results reveal pronounced heterogeneity across supply channels: pipeline gas is highly price inelastic, whereas LNG serves as the primary adjustment channel on the supply side. Although both supply shocks affect prices, pipeline disruptions generate sharp but short-lived inflationary effects, while LNG supply shocks exert more persistent influences on price dynamics. More broadly, the dynamics of gas prices exhibit a clear horizon-dependent structure, with short-run fluctuations driven by supply shocks and inventory behavior, and medium- to long-run movements increasingly shaped by aggregate demand forces. Counterfactual scenarios of the 2022 energy crisis quantify the stabilizing role of LNG availability and demand adjustment, highlighting how supply composition critically shapes energy price dynamics in Europe.

1 Introduction

The energy crisis experienced by the European Union (EU) in 2022 reignited the debate on the role of fossil fuels in modern economies and, in particular, on the vulnerabilities associated with energy dependence. Natural gas, which accounts for roughly 25% of total energy consumption in the EU, is predominantly extracted outside Europe. As a result, unexpected disruptions to gas production or transportation in distant regions can have large and immediate effects on European energy prices and macroeconomic conditions. While oil markets have long been the focus of extensive empirical research, the modeling of natural gas markets has gained prominence only more recently, reflecting the growing importance of gas as a key input for industrial production, electricity generation, and household consumption. This increased relevance has also attracted rising attention from policymakers and central banks, especially in the context of inflationary pressures and energy security.

In contrast to crude oil markets, which are characterized by strong global integration and relatively low transportation frictions, natural gas markets have historically been fragmented across regions Szafranek and Rubaszek (2024). This fragmentation largely reflects transportation constraints. Whereas crude oil is predominantly shipped by sea, natural gas relies on two distinct delivery channels: pipeline gas and liquefied natural gas (LNG). Pipeline gas has traditionally been the dominant source of supply in Europe, accounting for approximately 85% of total net imports over the period 2010–2018. Since 2019, however, its relative importance has declined markedly, with its share fluctuating between 55% and 80% over the 2019–2024 period. The increasing role of LNG has fundamentally altered the structure of European gas supply, introducing a more flexible but capacity-constrained source that is increasingly linked to global market conditions.

Despite these structural differences, much of the existing empirical literature on natural gas markets adopts Structural Vector Autoregressive (SVAR) frameworks adapted from the crude oil literature, in which supply is typically modeled as a single aggregate. In the context of natural gas, this abstraction is problematic, as it ignores the heterogeneity across supply channels and the scope for substitution between pipeline gas and LNG. Aggregating supply obscures the information embedded in shifts in the composition of imports, which is crucial for understanding price formation and adjustment dynamics. Recent contributions, including Moll et al. (2023) and Hamilton (2023), emphasize that substitution across gas supply channels plays a central role in shaping market outcomes and that failing to model this margin explicitly can lead to a systematic underestimation of supply-side forces. The 2022 energy crisis provides a clear illustration: the sharp reduction in Russian pipeline deliveries was partially offset by a surge in LNG imports. When supply is treated as a single aggregate, this reallocation appears as only a modest change in total imports, masking the profound adjustment occurring beneath the surface.

This paper models the European natural gas market by explicitly distinguishing between pipeline gas and LNG supplies. This approach contributes to the literature in two main ways. First, by jointly identifying the short-run elasticities of both supply sources, the model avoids the implicit assumption, common in previous studies, that pipeline gas and LNG interact with prices, activity, and storage in an identical manner. This distinction is strongly motivated by a growing body of work documenting substantial differences between the two supply channels in terms of contractual rigidity, cost structures, physical and logistical constraints, and short-run flexibility. Pipeline gas is typically governed by long-term contracts and fixed transport infrastructure, while LNG operates as a more flexible, globally arbitraged margin, with cargoes that can be redirected across regions in response to price differentials. As emphasized by Ritz (2019) and Sharples (2023), ignoring this heterogeneity risks understating the role of supply-side dynamics and obscuring large compositional adjustments such as those observed during the 2022 crisis.

Second, the paper contributes methodologically by modeling the European natural gas market within a Bayesian Structural Vector Autoregressive framework identified through sign and elasticity restrictions following Baumeister and Hamilton (2015) and Baumeister and Hamilton (2019). To the best of our knowledge, this identification strategy has not previously been applied to the European natural gas market. An important advantage of this approach is that it allows for the incorporation of economically meaningful priors on contemporaneous elasticities, which remain asymptotically informative under set identification. This feature is particularly relevant in energy markets, where short-run elasticities are difficult to estimate precisely from the data alone. The SVAR framework further provides a coherent structure for disentangling supply-driven and demand-driven sources of natural gas price fluctuations.

The empirical results highlight several novel findings. Supply shocks play a role that is quantitatively comparable to demand shocks in driving short-run natural gas price dynamics, contrary to much of the existing literature. Moreover, explicitly distinguishing between supply sources reveals substantial heterogeneity in their effects: pipeline supply disruptions generate sharp but relatively transitory price increases, whereas LNG supply shocks are more persistent and shape medium-run price dynamics, reflecting the role of LNG as the marginal balancing source in Europe. Storage and precautionary demand emerge as important amplification mechanisms, particularly during periods of heightened supply uncertainty, underscoring the importance of expectations and inventory behavior for price formation.

The use of this Bayesian SVAR methodology is becoming increasingly common in energy market applications. For example, Rubaszek et al. (2021) and Farag (2024) apply similar frameworks to the U.S. natural gas market and study its spillovers to the global economy through LNG exports. Building on this literature, the present study further exploits the identified structural shocks to conduct counterfactual analyses and to assess the role of supply substitution during the European energy crisis.

Finally, the paper exploits the structural VAR framework to conduct a set of counterfactual analyses aimed at quantifying the role of supply substitution and demand adjustment during the 2022 energy crisis. By manipulating the identified structural shocks, the counterfactual scenarios assess how European natural gas prices would have evolved in the absence of LNG inflows, under no demand reduction, or under a frictionless substitution of pipeline gas with LNG. These exercises provide a transparent and policy-relevant quantification of the stabilizing role played by LNG infrastructure and conservation measures, while highlighting the limits of supply flexibility in the presence of binding capacity constraints.

The remainder of the paper is organized as follows. Section 2 reviews the related literature and situates the paper within existing empirical and theoretical work on natural gas markets and energy price dynamics. Section 3 presents the econometric methodology based on the Bayesian SVAR framework with sign restrictions developed by Baumeister and Hamilton (2015). Section 4 introduces the empirical model and the data, detailing variable construction and identifying assumptions. Section 5 discusses the empirical results and several extensions, including historical decompositions and counterfactual scenarios. Section 6 concludes.

2 Literature Review

This chapter relates the present study to three strands of literature. First, it reviews the use of Structural Vector Autoregressive (SVAR) models in the crude oil market, which provides the conceptual and methodological foundation for most empirical work on energy price shocks. Second, it surveys the growing literature applying SVARs to natural gas markets, with primary emphasis on Europe, complemented by evidence from the United States and, to a lesser extent, Asia. Finally, it reviews the literature on contractual and infrastructural differences between pipeline gas and liquefied natural gas (LNG), which motivates the explicit separation of supply sources in the empirical model.

The modern empirical analysis of energy price shocks originates in the crude oil literature. Seminal work by Hamilton (1996) documented a strong negative relationship between oil price increases and U.S. macroeconomic activity, giving rise to an extensive literature on the identification and transmission of oil price shocks. A key methodological breakthrough was introduced by Kilian (2009), who proposed a Structural VAR framework distinguishing between oil supply shocks, global aggregate demand shocks, and oil-specific demand shocks related to precautionary or speculative motives. This approach emphasized that oil price increases can have markedly different macroeconomic effects depending on their underlying source. Subsequent contributions by Kilian and Murphy (2012) and Kilian and Murphy (2014) highlighted limitations of purely agnostic sign restrictions and argued for incorporating economically meaningful restrictions on short-run elasticities and inventory behavior. Building on these insights, Baumeister and Hamilton (2015) and Baumeister and Hamilton (2019) developed a Bayesian SVAR framework under set identification, showing that informative priors on structural elasticities can discipline the identification problem without imposing exact zero restrictions. Applied to the global oil market, this methodology revealed that demand-driven shocks account for the bulk of oil price fluctuations, while supply shocks play a more limited role.

This literature provides three key lessons that are directly relevant for natural gas markets. First, separating supply- and demand-driven shocks is crucial for understanding price dynamics. Second, short-run price elasticities of energy supply and demand are typically very small, implying that even moderate shocks can generate large price movements. Third, careful identification—often requiring priors, sign restrictions, or additional observables such as inventories—is essential to avoid misleading structural interpretations.

In contrast to crude oil, natural gas markets have historically been regionally fragmented due to transportation constraints and infrastructure specificity. In Europe, gas supply has long relied on pipeline imports under rigid long-term contracts, with LNG playing a limited role until recently, and early empirical work focused on the linkage between gas and oil prices under oil-indexed contracts. As shown by Szafranek and Rubaszek (2024), this relationship weakened markedly from the mid-2000s onward, reflecting market liberalization and the transition to hub-based pricing. More recent studies adopt SVAR frameworks to analyze European gas price dynamics, typically identifying supply shocks, demand shocks related to economic activity or weather, and storage-driven precautionary demand. While these studies document relatively inelastic short-run demand and historically constrained supply responses, the 2022 energy crisis has underscored the importance of supply-side disruptions and substitution mechanisms. In particular, Güntner et al. (2024) show that the collapse in Russian pipeline deliveries represented a dominant negative supply shock, partially offset by LNG inflows, and that precautionary storage demand amplified price pressures. A central implication of this literature is that aggregating pipeline gas and LNG into a single supply measure can mask large compositional adjustments and understate the role of supply-side shocks, motivating models that explicitly distinguish between supply sources.

Recent European SVAR studies therefore stress the importance of explicitly distinguishing between supply sources to correctly characterize price formation, elasticities, and transmission to inflation and real activity. Moll et al. (2023); Hamilton (2023)

The U.S. natural gas market differs fundamentally from the European market. It is largely self-sufficient, price formation is domestic, and supply is dominated by market-based production rather than imports. SVAR analyses of the U.S. gas market have focused on disentangling domestic supply shocks—such as technological innovations associated with the shale gas revolution—from demand shocks driven by weather, economic activity, and storage behavior. Rubaszek et al. (2021) provide a comprehensive Bayesian SVAR analysis of the U.S. gas market, estimating short-run demand elasticities that are larger in absolute value than those typically found for oil. Their results suggest that fuel-switching in electricity generation and industrial use renders U.S. gas demand more price-responsive than in Europe. They also document a non-negligible role for storage-related demand shocks, analogous to precautionary demand in the oil market. More recent work highlights the growing international dimension of the U.S. gas market following the expansion of LNG export capacity [Citar a Farag]. As the United States has become a major LNG exporter, domestic gas prices have become increasingly exposed to global demand conditions. Empirical evidence suggests that shocks originating in Europe or Asia can now transmit to U.S. gas prices through LNG arbitrage, although the magnitude of this channel remains smaller than in more import-dependent regions.

A large body of research shows that pipeline gas and LNG represent fundamentally different supply margins. Pipeline gas is typically governed by long-term, route-specific contracts with take-or-pay and destination clauses, which imply low short-run price elasticity and limited flexibility in response to shocks Molnar (2022). LNG supply, by contrast, has become increasingly flexible as contract durations shorten, destination restrictions weaken, and spot market trading expands, allowing cargoes to be redirected across regions in response to price differentials Sharples (2023). This flexibility gives LNG a strategic role beyond its volume share, as it provides a credible outside option that constrains the market power of pipeline exporters Ritz (2019). Evidence from the 2022 European energy crisis illustrates these mechanisms clearly: rising LNG imports partially offset the collapse in Russian pipeline flows, though only at exceptionally high spot prices, underscoring both the substitutability and the distinct pricing dynamics of the two supply channels.

Taken together, this literature implies that pipeline gas and LNG differ markedly in contractual rigidity, cost structures, elasticity, and adjustment speed. Ignoring these differences risks mischaracterizing both the sources of natural gas price fluctuations and their macroeconomic consequences. This insight provides a central motivation for the empirical strategy adopted in this chapter, which explicitly distinguishes between pipeline and LNG supplies within a Bayesian SVAR framework.

3 Methodology

The model can be rewritten in matrix form as a typical structural VAR (SVAR).

$$\mathbf{A}y_t = \mathbf{B}x_{t-1} + u_t \tag{1}$$

$$u_t \sim N(0_{5x1}, \mathbf{D}) \tag{2}$$

where $y_t = (y_{1t}, y_{2t}..., y_{nt})'$ is the $n \times 1$ vector of endogenous variables, $\mathbf{A}$ is the $n \times n$ matrix that contains all the contemporaneous elasticities across the endogenous variables, $x_{t-1} = (y_{t-1}, ...y_{t-m})'$ is the $k \times 1$ vector of lagged regressor, $\mathbf{B}$ is the $n \times k$ autoregressive matrix of coefficients and $\mathbf{D} = diag(d_{11}, ..., d_{nn})$ is the variance of the structural innovations.

But it also has a reduced form, which can be interpreted as a conventional VAR:

$$y_t = \Phi x_{t-1} + \varepsilon_t \tag{3}$$
$$\varepsilon_t \sim N(0_{5x1}, \Omega) \tag{4}$$
$$\Phi = A^{-1}B \tag{5}$$
$$\Omega = A^{-1}D(A^{-1})' \tag{6}$$

The model is estimated using Bayesian methods and sign-restrictions. Under set-identified models, the prior of these elasticities does not disappear asymptotically as Baumeister and Hamilton (2015) (henceforth BH) demonstrated. Later on this methodology was enhanced and applied for the global crude oil market in Baumeister and Hamilton (2019) and for the American natural gas market in Rubaszek et al. (2021) and Farag (2024). The use of BH methodology brings the flexibility to center all the priors of the elements in $\mathbf{A}$ at reasonable values given by the literature. Hence, centering the location and scale of the prior for the structural elasticities in reasonable values already acknowledged makes the model to incorporate valuable information that is not available in the data. In this section I provide a brief description of the methodology.

3.1 Prior

The starting point for setting the prior is decomposing $p(A, B, D)$ as:

$$p(\mathbf{A}, \mathbf{B}, \mathbf{D}) = p(\mathbf{A}) \times P(\mathbf{D}|\mathbf{A}) \times p(\mathbf{B}|\mathbf{A}, \mathbf{B}) \tag{7}$$

The prior for the diagonal variance-covariance matrix $\mathbf{D}$ consists in setting that the inverse of its elements d_{ii}^{-1} are jointly independent and follow a Gamma distribution with $\kappa_i, \tau_i(\mathbf{A})$ as parameters.

$$p(\mathbf{D} \mid \mathbf{A}) = \prod_{i=1}^{n} p(d_{ii} \mid \mathbf{A}) \tag{8}$$
$$d_{ii}^{-1} \mid \mathbf{A} \sim \Gamma(\kappa_i, \tau_i(\mathbf{A})), \tag{9}$$

Similarly, rows of matrix $\mathbf{B}$ are assumed to be jointly independent and to follow a multivariate normal distribution with parameters m_i, M_i.

$$p(\mathbf{B} \mid \mathbf{A}, \mathbf{D}) = \prod_{i=1}^{n} p(b_i \mid \mathbf{D}, \mathbf{A}) \tag{10}$$
$$b_i \mid \mathbf{A}, \mathbf{D} \sim N(m_i,\, d_{ii} M_i), \tag{11}$$

The prior for $\mathbf{A}$ can be freely chosen to reflect any knowledge about the structural relations between the endogenous variables. In the case of this study it will be the joint distribution of several truncated Student-t.

3.2 Likelihood

Innovation of the structural model is assumed to be normal.

$$p(Y_T \mid \mathbf{A}, \mathbf{D}, \mathbf{B}) = (2\pi)^{-Tn/2} \,|\det(\mathbf{A})|^T \det(\mathbf{D})^{-T/2} \tag{12}$$

$$\times \exp\left\{ -\tfrac{1}{2} \sum_{t=1}^{T} (\mathbf{A}y_t - \mathbf{B}x_{t-1})' \mathbf{D}^{-1} (\mathbf{A}y_t - \mathbf{B}x_{t-1}) \right\}, \tag{13}$$

3.3 Posterior

To proceed with the posterior distribution, we need a similar decomposition as before:

$$p(\mathbf{A}, \mathbf{B}, \mathbf{D} \mid Y_t) = p(\mathbf{A} \mid Y_t) \times P(\mathbf{D}|\mathbf{A} \mid Y_t) \times p(\mathbf{B}|\mathbf{A}, \mathbf{B} \mid Y_t) \tag{14}$$

Where the conditional posterior of the elements of the diagonal variance covariance $P(\mathbf{D}|\mathbf{A} \mid Y_t)$ is expressed as the joint distribution of all the elements in the diagonal:

$$p(\mathbf{D} \mid \mathbf{A}Y_t) = \prod_{i=1}^{n} p(d_{ii} \mid \mathbf{A}Y_t) \tag{15}$$

$$d_{ii}^{-1} \mid \mathbf{A} \sim \Gamma(\kappa_i^*, \tau_i^*(\mathbf{A})), \tag{16}$$

These parameters $\kappa_i^*, \tau_i^*(\mathbf{A})$ that fully characterize this posterior are a function of the data and of $mathbfA$ as it follows.

$$\kappa_i^* = \kappa_i + T/2 \tag{17}$$

$$\tau_i^*(\mathbf{A}) = \tau_i(\mathbf{A}) + \xi_i^*(\mathbf{A}) \tag{18}$$

And the term ξ_i^* refers to the sum of squared residuals from regressing $\widetilde{Y}_i(\mathbf{A})$ on $\widetilde{X}_i$ defined as:

$$\xi_i^*(\mathbf{A}) = \left(\widetilde{Y}_i(\mathbf{A})'\widetilde{Y}_i(\mathbf{A})\right) - \left(\widetilde{Y}_i(\mathbf{A})'\widetilde{X}_i\right)\left(\widetilde{X}_i'\widetilde{X}_i\right)^{-1}\left(\widetilde{X}_i'\widetilde{Y}_i(\mathbf{A})\right) \tag{19}$$

$$\widetilde{Y}_i = \begin{bmatrix} y_1'a_i & \cdots & y_T'a_i & m_i'P_i \end{bmatrix}'_{(T+k)\times 1} \tag{20}$$

$$\widetilde{X}_i = \begin{bmatrix} x_0' & \cdots & x_{T-1}' & P_i \end{bmatrix}'_{(T+k)\times k} \tag{21}$$

where $M_i = P_i'P_i$ and a_i are the rows of the matrix $\mathbf{A}$.

Now the conditional posterior of the autoregressive matrix $p(\mathbf{B}|\mathbf{A}, \mathbf{B})$ can be defined as the joint posterior of its rows:

$$p(\mathbf{B} \mid \mathbf{A}, \mathbf{D}, Y_t) = \prod_{i=1}^{n} p(b_i \mid \mathbf{D}, \mathbf{A}) \tag{22}$$

$$b_i \mid \mathbf{A}, \mathbf{D}Y_t \sim N(m_i^*,\, d_{ii}M_i^*) \tag{23}$$

where

$$m_i^* = \left(\widetilde{X}_i'\widetilde{X}_i\right)^{-1}\left(\widetilde{X}_i'\widetilde{Y}_i\right), \tag{24}$$

$$M_i^* = \left(\widetilde{X}_i'\widetilde{X}_i\right)^{-1}. \tag{25}$$

Finally, the posterior of A is:

$$p(\mathbf{A} \mid \mathbf{Y}_T) = \frac{k_T\, p(\mathbf{A}) \left[\det(\mathbf{A}\hat{\Omega}_T\mathbf{A}')\right]^{T/2}}{\prod_{i=1}^{n}\left[(2\tau_i^*/T)^{\kappa_i^*}\right]} \prod_{i=1}^{n} \left\{ \frac{|\mathbf{M}_i^*|^{1/2}}{|\mathbf{M}_i|^{1/2}} \frac{\tau_i^{\kappa_i}}{\Gamma(\kappa_i)} \Gamma(\kappa_i^*) \right\}.$$

where $\hat{\Omega}_T = T^{-1}\sum_{t=1}^{T}\hat{\varepsilon}_t\hat{\varepsilon}_t'$ is the sample variance covariance matrix estimated by maximum likelihood.

4 Empirical model and data

4.1 Data

For the construction of the baseline model, the dataset comprises five variables. The first variable, q_{pip}, measures the net imports of natural gas through pipelines into the EU on a monthly basis, obtained from the IEA's *Gas Trade Flows* database. The second variable, ip, corresponds to the industrial production index of the manufacturing sector in the EU-19, sourced from Eurostat. The third variable, p, captures the real price of natural gas traded at the Dutch TTF, which serves as the benchmark price for the European market, available from the World Bank Commodity Price Database. The fourth variable, s, represents total natural gas inventories, included given their key role in consumption smoothing as well as precautionary and speculative behavior, also sourced from the IEA's *Gas Trade Flows* database. Finally, the fifth variable, q_{LNG}, denotes net imports of liquefied natural gas (LNG) into the EU, likewise obtained from the IEA's *Gas Trade Flows* database.

$$y_t = [q_{\text{pip}}, ip, p, s, q_{\text{LNG}}] \tag{26}$$

All series are available at monthly frequency, and are expressed in logarithmic growth rates to remove unit roots in prices, industrial production, and gas imports.

4.2 Specification and assumptions

This section describes the structural model used in the paper to make the identification for the European natural gas market. The main characteristics of the contemporaneous elasticities among endogenous variables is equivalent to the one used in Baumeister and Hamilton (2019); Rubaszek et al. (2021); Farag (2024).

We begin with the structural equations of the two supply components. Both pipeline and LNG imports are assumed to respond contemporaneously only to innovations in the gas price. Consistent with previous studies, both supply types are considered to be mildly elastic; that is, the corresponding supply price elasticities γ_{qp} and β_{qp} are assumed to differ from zero. Because the production and delivery of natural gas involve adjustment costs, imports are expected to react to real activity shocks only with a delay of at least one month. Moreover, since the output of these two technologies is extracted simultaneously, they cannot be adjusted relative to one another within the same period.

$$q_t^{pip} = \gamma_{qp}p_t + \mathbf{b}_1'x_{t-1} + \varepsilon_t^{\text{pipeline supply}} \tag{27}$$

$$q_t^{LNG} = \beta_{qp}p_t + \mathbf{b}_5'x_{t-1} + \varepsilon_t^{\text{LNG supply}} \tag{28}$$

Real activity is broadly proxied by industrial production, which is assumed to respond contemporaneously only to the gas price. This assumption is motivated by the fact that a supply disruption affects industrial activity primarily through its impact on prices. In the short run, however, a shortfall in gas supply can be partially offset by substituting other components of the energy mix, such as biogas or alternative electricity sources. Nevertheless, any increase in the gas price is expected to exert a negative effect on industrial production.

$$ip_t = \alpha_{yp} p_t + \mathbf{b}_2' x_{t-1} + \varepsilon_t^{\text{aggregate demand}} \tag{29}$$

Demand for natural gas is separated into *flow demand* and *precautionary demand*. Flow demand captures shifts in the portion of gas that is effectively consumed, and can be mostly associated with changes in market conditions that make gas more desirable to consumers. The structural equation for gas demand allows industrial production to affect the gas price contemporaneously, with the elasticity of industrial demand denoted by $\alpha_{py} > 0$. The remaining parameter of the demand's structural equation, $\alpha_{py} < 0$, reflects the price elasticity of the demand. Since consumers are indifferent to the source of natural gas supply, any unit of pipeline gas, LNG, or stored gas is treated as identical, so aggregate demand is defined as $Q_t = q_t^{pip} + q_t^{LNG} - s_t$. In this case, the structural demand equation is given by:

$$Q_t = \alpha_{py} ip_t + \alpha_{pq} p_t + \mathbf{b}_3' x_{t-1} + \varepsilon_t^{\text{flow demand}}, \tag{30}$$

or equivalently, in inverse demand form:

$$\alpha_{qp} p_t = q_t^{pip} + q_t^{LNG} - s_t + \alpha_{py} ip_t + \mathbf{b}_3' x_{t-1} + \varepsilon_t^{\text{flow demand}}. \tag{31}$$

Precautionary demand captures changes in market participants' expectations that lead to adjustments in natural gas inventories. Such behavior may arise either because agents anticipate a potential supply disruption or because a rise in prices triggers speculative motives among inventory holders. Stocks of gas are adjusted to the current imports gas, which makes inventories to be a potential substitute for them. But inventories also react to the price innovations within the same month, because any trigger that ends in a rise of the gas price will have two effects in the stocks: On one hand, the rise in the price brings the opportunity of inventory holders to make a profit by selling their gas at a higher price than the one they faced when they bought the commodity, which implies that stocks decline given a increase of the price $\psi_2 < 0$. But on the other hand, if inventory holders foresee that the current rise of the price is preluding additional future price increments, then it is likely that they prefer to increase current inventories for selling more gas in the future, $\psi_2 > 0$. These two conflicting forces make the sign of the price elasticity ψ_2 theoretically ambiguous, and thus, like the other two supply elasticities, it is left without a sign restriction.

$$s_t = \psi_1 q_t^{pip} + \psi_2 p_t + \psi_3 q_t^{LNG} + \mathbf{b}_4' x_{t-1} + \varepsilon_t^{\text{precautionary demand}} \tag{32}$$

Table 1: Contemporaneous elasticities in $\mathbf{A}$

Parameter	Meaning	Sign restriction
γ_{qp}	Pipeline supply elasticity	$\gamma_{qp} > 0$
β_{qp}	LNG supply elasticity	$\beta_{qp} > 0$
α_{yp}	Effect of p on economic activity	$\alpha_{yp} < 0$
α_{qy}	Income elasticity of gas demand	$\alpha_{qy} > 0$
α_{qp}	Gas demand elasticity	$\alpha_{qp} < 0$
ψ_1	Effect of q^{pip} on oil inventories	none
ψ_2	Effect of p on oil inventories	none
ψ_3	Effect of q^{LNG} on oil inventories	none

4.3 Matrix form

Following the specification of a SVAR the model is rewritten in matrix form as it follows:

$$\begin{pmatrix} 1 & 0 & -\gamma_{qp} & 0 & 0 \\ 0 & 1 & -\alpha_{yp} & 0 & 0 \\ 1 & -\alpha_{py} & -\alpha_{qp} & -1 & 1 \\ -\psi_1 & 0 & -\psi_2 & 1 & -\psi_3 \\ 0 & 0 & -\beta_{qp} & 0 & 1 \end{pmatrix} \begin{pmatrix} q_t^{pip} \\ ip_t \\ p_t \\ s_t \\ q_t^{pip} \end{pmatrix} = \mathbf{B}x_{t-1} + \begin{pmatrix} u_t^{\text{pipeline supply}} \\ u_t^{\text{aggregate demand}} \\ u_t^{\text{flow demand}} \\ u_t^{\text{precautionary shock}} \\ u_t^{\text{LNG supply}} \end{pmatrix}$$

Where the structural shocks are assumed to be normally distributed $u_t \sim N(0_{5x1}, \mathbf{D})$ and the contemporaneous elasticities in $\mathbf{A}$ satisfy the following 5 sign restrictions:

$$\gamma_{qp} > 0 \quad \beta_{qp} < 0 \quad \alpha_{yp} < 0 \quad \alpha_{qy} > 0 \quad \alpha_{qp} < 0$$

The way to implement these restrictions is to incorporate them in the prior of the contemporaneous matrix, $p(\mathbf{A})$, assigning truncated distributions in each case. The purpose of this is twofold, first it makes easier to incorporate prior knowledge of the elements in $\mathbf{A}$ that is not directly provided by the likelihood of the data. On top of that it helps to ensure that the posterior of the elasticities $p(\mathbf{A} \mid Y_t)$ gives most of the probability to values of $\mathbf{A}$ that make the impact matrix $\mathbf{A}^{-1}$ to satisfy the following signs.

$$\begin{pmatrix} \varepsilon_{q^{pip},t} \\ \varepsilon_{ip,t} \\ \varepsilon_{p,t} \\ \varepsilon_{s,t} \\ \varepsilon_{q^{LNG},t} \end{pmatrix} = \begin{pmatrix} + & + & + & ? & ? \\ + & + & - & ? & + \\ - & + & + & + & - \\ ? & ? & + & ? & ? \\ ? & + & + & ? & + \end{pmatrix} \begin{pmatrix} u_t^{\text{pipeline supply}} \\ u_t^{\text{aggregate demand}} \\ u_t^{\text{flow demand}} \\ u_t^{\text{precautionary shock}} \\ u_t^{\text{LNG supply}} \end{pmatrix}$$

4.4 Prior knowledge about these elasticities

The structural elasticities embedded in matrix A, which govern the contemporaneous interactions between quantities, prices, and real economic activity, can be disciplined using a well-established empirical literature on energy markets, particularly natural gas and oil. In set-identified SVAR models that rely on inequality or sign restrictions, these elasticities play a central role in shaping inference. Unlike in fully identified models, the posterior distribution does not converge to the likelihood even asymptotically, implying that prior information does not wash out with increasing sample size (Baumeister and Hamilton, 2015). As a consequence, the choice of prior distributions is not innocuous and must be grounded in economically and empirically plausible values.

The short-run price elasticity of pipeline natural gas supply, denoted by γ_{qp}, measures the responsiveness of pipeline gas imports to price changes. We impose the theoretical restriction $\gamma_{qp} > 0$ and center the prior at 0.1, reflecting the well-established view that natural gas supply is highly inelastic in the short run. Empirical studies consistently find small elasticities, typically close to zero or below 0.1, due to technological constraints, adjustment costs, and contractual rigidities (Al-Sahlawi, 1989; Dahl and Duggan, 1996; Arora, 2014; Rubaszek et al., 2021; Mason and Roberts, 2018). Centering the prior at a small positive value therefore captures an upward-sloping but steep supply curve, while the heavy-tailed Student-t prior allows for uncertainty without admitting implausibly large elasticities.

A similar prior is imposed on the short-run price elasticity of LNG supply, β_{qp}. Although LNG-specific elasticity estimates are scarce, short-run LNG supply is constrained by fixed liquefaction capacity, shipping availability, and regasification infrastructure, implying only limited responsiveness to price signals. We therefore impose $\beta_{qp} > 0$ and center the prior at 0.1, allowing the data to assess whether LNG supply is more flexible than pipeline supply while anchoring inference in a plausible range.

On the demand side, the short-run price elasticity of natural gas consumption, α_{qp}, is expected to be negative. We center the prior at -0.3 and truncate the distribution to enforce $\alpha_{qp} < 0$. This calibration reflects the broad consensus that gas demand is highly inelastic in the short run. Early evidence summarized by Al-Sahlawi (1989) documents short-run elasticities clustered around values close to -0.3, while more recent studies typically estimate elasticities between -0.04 and -0.16 (Joutz, 2009; Bernstein and Madlener, 2011; Arora, 2014), with demand remaining far from elastic even when allowing for heterogeneity **?**. Centering the prior at -0.3, with a scale implying a one-sided 90% interval approximately spanning $(-0.86, 0)$, therefore provides a conservative and empirically grounded characterization of short-run demand responsiveness, consistent with priors commonly adopted in the oil-market SVAR literature.

The elasticity of gas demand with respect to aggregate economic activity, α_{qy}, captures the income elasticity of natural gas demand. We impose $\alpha_{qy} > 0$ and center the prior at 0.6, reflecting a strong but less than proportional response of gas consumption to changes in output. This calibration is motivated by survey evidence summarized by Al-Sahlawi (1989), which reports income elasticities ranging from near zero to above unity, as well as by more recent cross-country estimates for OECD economies that cluster around 0.7 Burke and Yang (2016). Centering the prior slightly below this benchmark provides a conservative yet empirically grounded characterization of demand responsiveness to economic activity.

Finally, the parameter α_{yp} captures the contemporaneous response of aggregate economic activity to changes in natural gas prices. We impose a negative sign restriction and center the prior at -0.05, reflecting the well-established view that higher energy prices act as a drag on output by increasing production costs and reducing real disposable income, while

Table 2: Priors and posteriors distributions for parameters in $\mathbf{A}$

	γ_{qp}	β_{qp}	α_{yp}	α_{qy}	α_{qp}	ψ_1	ψ_2	ψ_3
Priors affecting contemporaneous coefficients $\mathbf{A}$								
Location	0.1	0.1	-0.05	0.7	-0.1	0	0	0
Scale	0.2	0.2	0.1	0.2	0.2	0.5	0.5	0.5
Degrees of freedom	3	3	3	3	3	3	3	3
Sign restriction	t^+	t^+	t^-	t^+	t^-	t	t	t
Posterior affecting contemporaneous coefficients $\mathbf{A}$								
5%	0.033	0.339	-0.027	0.351	-2.566	-0.148	-0.112	0.010
50%	0.081	0.776	-0.007	0.881	-1.210	-0.036	-0.070	0.053
95%	0.131	1.474	-0.001	2.011	-0.654	0.075	-0.029	0.094
Mean	0.081	0.827	-0.010	0.993	-1.332	-0.036	-0.070	0.053

having only a modest immediate effect Kilian (2008). Similar calibrations are standard in Bayesian SVAR studies of energy markets, and posterior estimates are typically close to zero Baumeister and Hamilton (2019). Evidence from the European gas crisis likewise points to noticeable but non-catastrophic short-run output effects **?**, supporting a small negative prior mean.

In summary, all prior distributions are centered on values firmly supported by the empirical literature. Supply elasticities are assumed to be small and positive, demand price elasticities moderate and negative, income elasticities positive and sizable, and the contemporaneous output response to gas prices modestly negative. By adopting literature-based priors within the Bayesian framework of Baumeister and Hamilton (2015), the model achieves economically meaningful identification while avoiding improper posteriors. Although the data remain free to update these beliefs, the priors provide a necessary anchor that ensures sensible inference in a weakly identified SVAR setting.

After visual inspection of the matching of identified shocks with recent historical events, I would like to proceed with some further analysis such as a historical decomposition of the estimated variables or a counterfactual analysis of "No energy crisis", similar to Güntner, Reif & Wolters (2024) for the whole EU instead of only for Germany.

5 Results

In this section, we present the estimation results of the baseline model and discuss the dynamics of the European natural gas market. We begin by examining the posterior distributions of the model parameters, focusing in particular on the short-run price elasticities of demand and supply. Then, we analyze the posterior impulse response functions, which capture the dynamic responses of the endogenous variables to structural shocks. Next, we quantify both the short- and long-term effects of these shocks on the variables included in the VAR system and assess their historical decompositions, highlighting their contributions to the variation in natural gas prices. Finally, we conduct a counterfactual analysis based on a hypothetical scenario in which European LNG imports were not available during the post-energy-crisis recovery period.

5.1 Posterior knowledge of the elasticities

We begin by comparing the prior and posterior distributions of the parameters in the contemporaneous relations matrix **A**, with particular attention to the posterior medians of the short-run demand and supply elasticities. Figure 1 displays the prior distributions (red line) alongside the posterior ones (blue histograms), while the corresponding descriptive statistics are reported in the lower panel of Table 2. The main findings can be summarized as follows.

The posterior median of the short-run price elasticity of pipeline gas supply, γ_{qp}, is estimated at only 0.08, indicating a highly inelastic response of supply to price changes. This estimate is broadly consistent with the elasticities reported by Rubaszek et al. (2021) and Farag (2024) for the U.S. natural gas market, who find elasticities of approximately 0.01 and 0.02, respectively, and conclude that American gas supply is strongly price inelastic. It is also roughly half the elasticity estimated by Mason and Roberts (2018) for natural gas wells in Wyoming, who report a value of 0.026. Moreover, our estimate is comparable to the European supply elasticities obtained by Adolfsen et al. (2024) (around 0.859) and **?** (approximately $0.y$), further confirming the limited short-run responsiveness of natural gas supply to price movements.

Regarding demand elasticities, the posterior median of the short-run price elasticity, α_{qp}, is estimated at -1.21. In absolute terms, this value exceeds the prior and is more than twice as large as the estimates reported by Rubaszek et al. (2021) and Farag (2024) for the U.S. natural gas market, who find elasticities of -0.42 and -0.177, respectively. This result suggests that, in the short run, European natural gas demand is considerably more responsive to price changes what the American demand is. This was documented by **?**, suggesting that produce As for the income elasticity, α_{qy}, its posterior distribution remains relatively close to the prior, with a median of 0.881, broadly consistent with the estimate of 0.719 of Güntner et al. (2024) or 0.725 reported by Baumeister and Hamilton (2019) for crude oil. This similarity indicates that shifts in aggregate economic activity translate into comparable changes in demand across both energy commodities and across different markets.

For the remaining parameters, the estimates indicate that, in the short run, the effect of natural gas prices on economic activity, α_{yp}, is almost negligible, with a posterior median of -0.009. This value is well below both our prior belief and the estimates reported by **?**, yet it is close to the -0.002 obtained by Baumeister and Hamilton (2019) for the crude oil market. Regarding the parameters that capture the effect of changes in supply and prices on natural gas inventories, ψ_1 and ψ_3, the contemporaneous posterior medians are negative, amounting to -0.170 and -0.177, respectively. This suggests that natural gas inventories decrease in response to positive shocks in supply or prices.

5.2 Impulse Response Functions

We decompose the dynamics of the European natural gas market into five structural shocks, following the structural oil market SVAR literature Kilian (2009); Kilian and Murphy (2012,

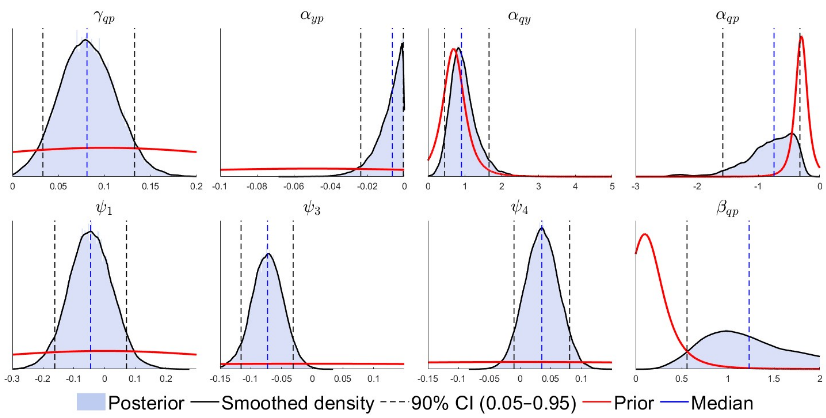

Figure 1: Prior and posterior of contemporaneous elasticities for the baseline model. Note:The baseline prior is represented using solid red lines, whereas the posterior is depicted using blue histograms. These distributions concern the contemporaneous coefficients in matrix $\mathbf{A}$ in the baseline model.

2014); Baumeister and Hamilton (2019). The model distinguishes two supply shocks, pipeline gas and LNG, and three demand shocks, aggregate demand, flow demand, and precautionary or storage demand. We trace their cumulative impulse responses for the set of endogenous variables.

First, consider the IRFs to a pipeline supply shock. An adverse pipeline shock lowers pipeline imports on impact and raises the real natural gas price, while industrial production declines in the short run. Inventories fall for a few months, consistent with an initial drawdown that partially cushions the physical shortfall, and LNG imports increase as supply is reallocated toward alternative inflows. In our estimates, the price effect is transitory and dissipates within about five months after the shock.

Next, consider the IRFs to a positive LNG supply shock that raises total gas availability, leading to a sustained decline in the real natural gas price and a persistent improvement in market conditions. Inventories increase for an extended period as the additional supply supports storage accumulation, while the adjustment in other inflows is gradual. Conversely, an adverse LNG shock reduces LNG availability, generates a prolonged tightening of the market, and raises prices in a persistent manner, reflecting the central role of LNG as the medium-run balancing margin in Europe.

These dynamics align with the gas literature documenting temporary price effects of pipeline disruptions. For example, Adolfsen et al. (2024) also find that pipeline supply shocks raise natural gas prices only temporarily, although their price response persists for up to nine months. A natural explanation is that their supply shock aggregates pipeline and LNG supply, so the estimated response reflects a mixture of short-lived pipeline disruptions and more persistent LNG-related adjustment. By distinguishing the two supply channels, our framework isolates the shorter-lived nature of pipeline shocks more clearly. The European experience in 2022, marked by the sharp reduction in Russian pipeline deliveries, provides a salient illustration of this mechanism, with pronounced price increases accompanied by comparatively limited immediate output losses as inventories and substitute imports partially mitigated the supply shortfall.

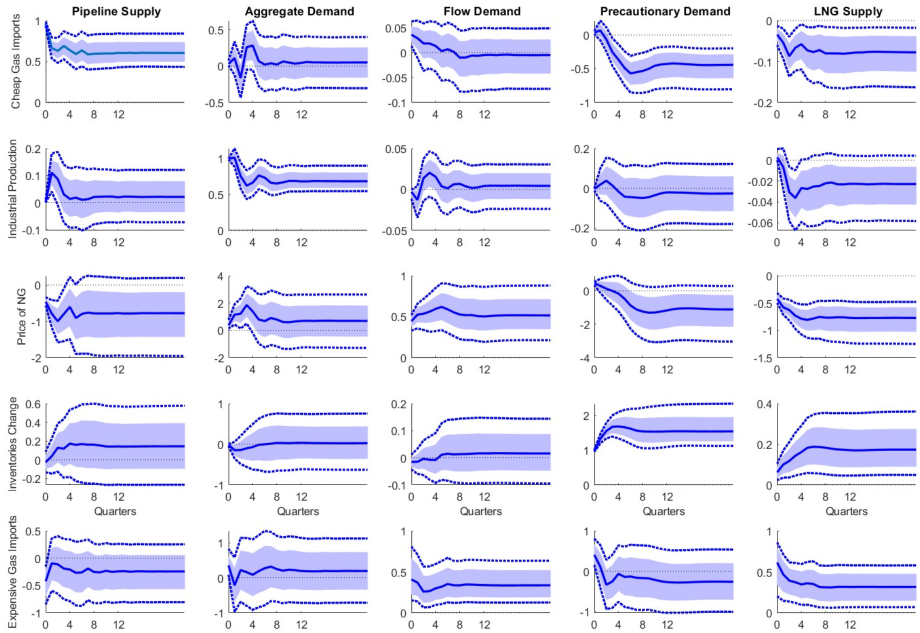

Figure 2: Cumulative Impulse Responses of the endogenous variables to structural shocks. Solid line represent the median estimate, blue shaded area and blue dashed lines represent the 68% and 90% confidence bands respectively

Compared with pipeline shocks, LNG shocks tend to be more persistent because LNG is the marginal balancing source in Europe and therefore shapes the medium-run market-clearing price. A pipeline disruption is typically a corridor-specific shortfall whose price impact is partly offset through inventory withdrawals, demand adjustment, and substitution toward other inflows, so the price effect can fade relatively quickly. An LNG supply shock instead reflects a change in Europe's access to flexible global cargoes, which relaxes the constraint on replacement supply and shifts the marginal supply schedule for longer. As a result, LNG shocks affect prices more durably than pipeline shocks in the impulse responses.

Turning to aggregate and flow demand shocks (second and third columns), both generate upward pressure on natural gas prices, but they differ in the persistence of the price response. Aggregate demand shocks raise prices only temporarily, consistent with a gradual adjustment of quantities as the economy reverts toward baseline. Flow demand shocks, in contrast, trigger a sharper contemporaneous tightening of market conditions, reflected in an immediate price increase and a pronounced drawdown of inventories, and they may be mildly contractionary for industrial production as higher energy costs and short-run scarcity constrain activity. Over time, however, the price effect of flow demand shocks fades as consumption normalizes and inventories are replenished.

Finally, precautionary demand shocks capture shifts in desired inventory holdings driven by expectations of future supply scarcity. A positive precautionary shock induces a simultaneous increase in inventories and natural gas prices, consistent with stockpiling behavior. In response, LNG imports remain broadly unchanged, whereas pipeline imports decline and settle at a permanently lower level after about three months. This pattern is suggestive of forward-looking inventory management, whereby storage holders anticipate persistent disruptions in pipeline supply, as in the 2022 episodes, and adjust their sourcing and inventory decisions accordingly.

Overall, the cumulative impulse responses indicate that European natural gas prices reflect the interaction of demand conditions, supply rigidities, and storage behavior. In the short run, price fluctuations are driven by both supply- and demand-side shocks, whereas medium-run dynamics are mostly associated with flow demand and LNG supply adjustments. These patterns are consistent with evidence from the oil and gas literatures and highlight the particular sensitivity of gas prices to infrastructure constraints and concerns about supply security.

5.3 Variance Decomposition

A Forecast Error Variance Decomposition (FEVD) quantifies the relative importance of each structural shock in explaining the uncertainty of gas price fluctuations over time, thereby providing a summary measure of the dominant drivers of price variability. Table 3 reports the median forecast error variance decomposition (FEVD) of the natural gas price at different horizons.

Table 3: Median FEVD of the Price

	Cheap. Sup.	**Aggr. Dem.**	**Flow Dem.**	**Prec. Dem.**	**Exp. Sup.**
1 month	22.7	15.0	21.1	21.1	19.1
2 months	18.4	42.9	11.0	13.9	10.9
3 months	19.3	43.1	9.4	15.1	9.5
6 months	16.9	55.8	5.2	14.6	5.5
1 year	15.7	53.7	4.4	19.5	4.7
2 years	15.6	53.7	4.3	19.8	4.7

Notes: This table reports descriptive statistics for the main variables in the sample. Data are monthly from 2010–2024.

At very short horizons, price fluctuations are driven by a relatively balanced mix of all the shocks in the model. At the one-month horizon, every shock explain equally the price volatility, while LNG supply explains a slightly smaller share. This pattern reflects the fact that in the immediate aftermath of shocks, gas prices react to a wide range of disturbances affecting both supply and demand conditions. As the horizon increases, aggregate demand shocks emerge as the dominant source of price variance. From three months onward, aggregate demand explains more than 40 percent of the forecast error variance, rising to nearly 55 percent at horizons of six months and beyond. This indicates that medium- to long-run movements in natural gas prices are strongly linked to broader macroeconomic conditions and demand forces, rather than to short-lived supply shocks.

Turning to the relative roles of supply and demand shocks, the variance decomposition reveals clear differences in their persistence and horizon dependence. On the demand side, flow demand shocks, which capture short-run adjustments in industrial and residential consumption, play an important role at very short horizons, accounting for more than 20 percent of price variance at the one-month horizon. Their contribution, however, declines rapidly and falls below 5 percent beyond one year, indicating that immediate consumption adjustments have largely transitory effects on price volatility. In contrast, precautionary demand shocks display a more persistent pattern. While their short-run contribution is moderate, their importance increases with the forecast horizon, explaining close to 20 percent of price variance at one- and two-year horizons. Unlike flow demand shocks, which dissipate once short-run consumption adjusts, precautionary motives affect prices through expectations and inventory decisions that operate over longer horizons. This suggests that uncertainty-driven behavior and expectations about future supply conditions generate lasting effects on prices, in line with the evidence from the energy crisis.

On the supply side, both pipeline and LNG supply shocks contribute more modestly to overall price variance when averaged over the full sample. Pipeline supply shocks explain a non-negligible share of short-run price variability but stabilize at around 15 percent at longer horizons, pointing to a limited yet persistent influence. Meanwhile, LNG, supply shocks account for an even smaller fraction of price variance, typically below 10 percent beyond the short run. This indicates that while supply disruptions, particularly in LNG markets, can trigger sharp price movements during specific episodes, their contribution to overall price uncertainty is relatively limited in the long run.

Taken together, these results highlight a clear horizon-dependent structure in which short-run price fluctuations reflect a combination of supply disturbances and immediate demand adjustments, whereas medium- and long-run volatility is increasingly shaped by aggregate demand conditions and persistent precautionary shocks rather than by supply shocks forces.

5.4 Historical Decomposition

A historical decomposition expresses the evolution of an observed time series as the cumulative contribution of the structural shocks identified in the model. In a structural VAR framework, this tool is particularly useful for attributing large price movements to distinct economic mechanisms and for assessing how the relative importance of supply, demand, and precautionary forces varies over time. Rather than focusing on the response to isolated shocks, the historical decomposition provides a narrative interpretation of realized outcomes, allowing one to identify which shocks were most relevant during specific historical episodes.

The dynamics of European natural gas real prices are best understood as the outcome of three interacting forces: (i) very low short-run demand elasticity, (ii) binding supply and infrastructure constraints, and (iii) the pivotal role of storage and expectations when supply security deteriorates. Because gas is hard to substitute in heating and many industrial processes in the short run, even moderate physical imbalances translate into large price movements. In Europe, these imbalances are magnified by the fact that pipeline supply is tied to specific corridors and contracts, while LNG is globally traded but constrained by liquefaction capacity, shipping availability, and regasification infrastructure. Hence, price spikes emerge when the market must simultaneously (a) replace lost pipeline volumes, and (b) rebuild inventories ahead of winter, especially under heightened uncertainty.

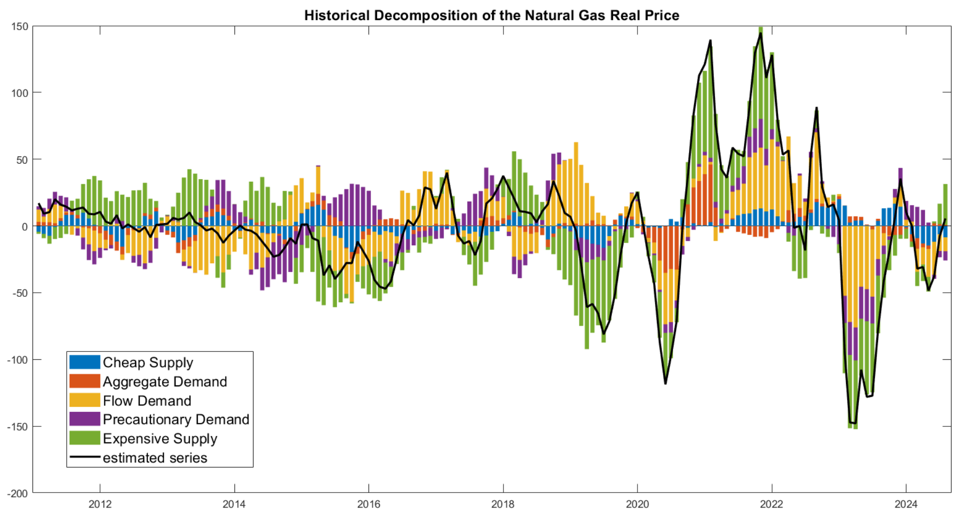

Figure 3: Impulse responses of the European gas price to supply shocks.

Figure 3 presents the historical decomposition of the European real natural gas price based on the five structural shocks identified in the model. The black line depicts the estimated log-level price, while the colored bars show the contemporaneous contribution of cheap supply, aggregate demand, flow demand, precautionary demand, and expensive supply shocks. Several episodes align closely with well-documented market developments. Following the 2014 Crimea invasion and the deterioration in EU Russia energy relations, the decomposition suggests an increase in storage-related and flow pressures, although price effects remained contained, consistent with the gradual nature of the adjustment and the buffering role of contracts and infrastructure. The COVID-19 shock in 2020 appears as a demand-driven episode: the sharp decline in prices coincides with large negative contributions from flow demand and precautionary demand, consistent with the collapse in economic activity and energy use documented for global gas markets during the pandemic International Energy Agency (2020). In the post-COVID recovery of 2021, prices start rising well before the full-scale 2022 crisis. This pattern is consistent with the combination of a rapid rebound in demand and supply-side frictions in global LNG markets and shipping [1], together with unusually low European storage entering the winter season Popkostova (2022). The most pronounced regime shift occurs during the 2022 energy crisis, when the sudden reduction in

[1]Specially a stronger power and gas demand in Asia pulled LNG away from Europe Reuters (2021). This has combined with a global supply crunch and limited injections, resulting in record-low stocks in the second half of 2021

Russian pipeline deliveries forced Europe to rely on scarce flexible LNG volumes under tight global capacity conditions International Energy Agency (2023). In this period, the decomposition is dominated by positive contributions from flow demand, precautionary demand, and expensive supply shocks, consistent with scarcity concerns and storage motives playing a central role in price formation. Aggregate demand contributes comparatively little during the spike, which is also consistent with contemporaneous assessments that emphasize supply insecurity and risk premia rather than macroeconomic overheating Furtuna et al. (2022). Finally, the normalization in 2023 coincides with the unwinding of these forces and with improved fundamentals, including strong storage buffers and the gradual expansion of LNG import capacity and market resilience EuropeanCommission (2024).

Overall, the historical decomposition clarifies how the same structural mechanisms operate with very different intensity across episodes. Periods of stability are characterized by offsetting shocks that keep prices anchored, whereas crisis episodes arise when multiple supply- and risk-related forces reinforce each other. The analysis shows that large price movements emerge not from a single dominant shock, but from the interaction between constrained supply, expectations, and limited adjustment margins. This perspective helps rationalize both the severity of the 2022 price spike and the speed of its subsequent reversal.

5.5 Counterfactual Analysis

The structural VAR framework employed in this paper enables a rich set of counterfactual and scenario-based analyses, in line with Güntner et al. (2024) and Antolín-Díaz et al. (2021). By specifying alternative realizations of selected structural shocks, we can simulate how the system would have evolved under different hypothetical conditions and trace the resulting paths of the endogenous variables. Exploiting this feature, we construct a three different counterfactual scenarios to evaluate the effects of a hypothetical events. The first scenario consists on the absence of LNG imports starting in February 2022 which we saw had a substantial effect lowering the prices specially in 2023. The second scenario quantifies the contribution of reduced flow demand resulting from the European Union's energy-saving measures to the stabilization of natural gas prices toward the end of 2022. A final third counterfactual scenario is drawn, in which we impose no bottleneck in the trade of LNG and, therefore, we assume that every lost unit of pipeline gas after the beginning of the energy crisis is immediately substituted by the same unit provided by LNG suppliers.

5.5.1 First Counterfactual: No Availability of LNG

In response to the looming gas crisis in 2022, European policymakers moved swiftly to diversify gas supply by expanding imports of liquefied natural gas (LNG). Germany, which had relied entirely on pipeline gas until then, fast-tracked the construction of floating LNG terminals – the first of which became operational in December 2022. The EU also negotiated deals with alternative suppliers (for example, a joint EU–US task force secured additional LNG deliveries of at least 15 bcm for 2022) and maximized use of existing import facilities. As a result, Europe's LNG imports surged by over 60% in 2022, adding roughly 66 bcm to supply, primarily from the United States, Qatar and other partners.

This counterfactual assumes a zero contribution of LNG to natural gas prices from February 2022 to December 2023. The left panel of Figure 4 presents the counterfactual scenario (red dashed line), alongside the baseline median forecast (solid line) and the observed data for the log level of natural gas prices. For ease of comparison, both model-generated series are expressed in log levels by accumulating log growth rates over time, starting from the realized price level in January 2022.

The right panel of Figure 4 displays the statistical difference between the counterfactual and the baseline model-generated log prices. While the initial gap is modest, the most pronounced divergence emerges from the autumn of 2022 onward, coinciding with the period in which LNG inflows began exerting downward pressure on European natural gas prices. As expected, the difference remains zero for all dates prior to February 2022. Since both the baseline and counterfactual series are expressed in log levels, their statistical difference can be interpreted as the percentage change in natural gas prices under the counterfactual scenario of no LNG supply arrivals during the energy crisis. The results indicate that LNG supply played a substantial role in mitigating price pressures in Europe. In the absence of investments in LNG infrastructure, natural gas prices would have been more than 400% higher throughout 2024.

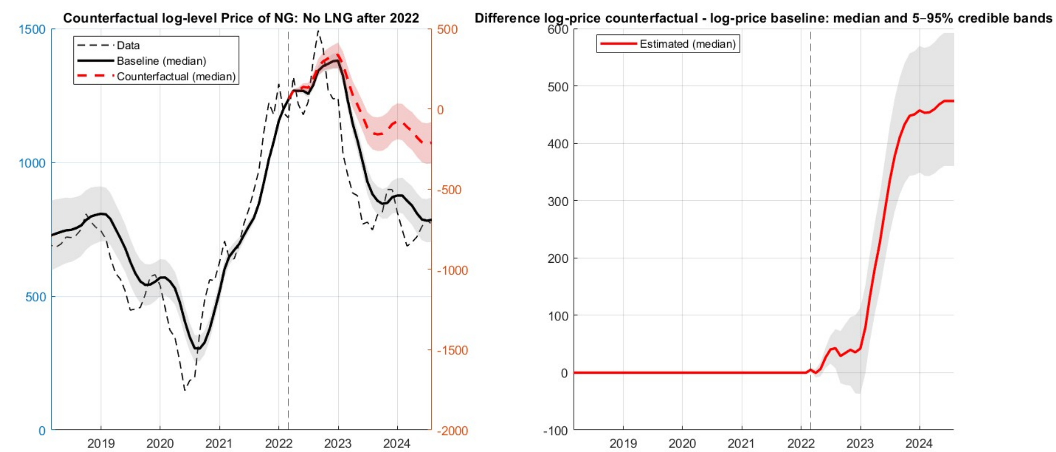

Figure 4: .

5.5.2 Second Counterfactual: No energy-saving measures

In parallel with supply-side responses, the European Union also took unprecedented steps to curb natural gas demand during the crisis. In July 2022, EU member states agreed to a coordinated plan ("Save Gas for a Safe Winter") to voluntarily reduce gas consumption by 15% between August 2022 and March 2023 IEEFA.org. National governments implemented measures such as public awareness campaigns, thermostat limits in buildings, fuel switching in industry, and other energy-saving regulations. High gas prices likewise incentivized consumers and firms to conserve energy. These efforts – combined with a milder winter – proved effective: EU gas consumption fell by roughly 18% in late 2022 compared to the five-year average, exceeding the target (about 56 bcm less gas used)IEEFA.org.

Now, we consider a counterfactual scenario in which these conservation policies had an even more pronounced impact on gas demand. Suppose that, from autumn 2022 onward, European gas consumers cut back usage significantly more (through stricter enforcement of rationing or exceptionally strong compliance with saving measures). This counterfactual creates an scenario in which the flow demand structural shocks were zero from October 2022 to December 2023. In the SVAR model, this can be captured as a one-off negative demand shock (or equivalently, a positive shock to gas inventories) in late 2022, reflecting an abrupt drop in consumption relative to baseline. We again use full-sample parameter estimates and the actual 2022/23 winter temperatures. By comparing this conditional forecast (with extra demand reduction) to the unconditional forecast and actual outcomes, we can analyze how much the EU's energy-saving measures (had they been even more aggressive) might have alleviated the gas supply strain during the crisis.

5.5.3 Third Counterfactual: No bottleneck in LNG

Now suppose as a counterfactual that the EU had managed to import enough LNG to offset a large share of the lost Russian pipeline gas immediately after the invasion. For instance, imagine that every lost bcm of natural gas imported by pipeline is instead added to the LNG imports. This creates a situation in which the scarcity of natural gas in 2022 did not really happen, but instead a different composition in the supply. In our structural VAR framework, we can represent this policy scenario as a one-off positive supply shock (a sudden increase in LNG gas availability) in mid-2022, boosting LNG net imports. We maintain the same parameter estimates than in the full sample. This allows us to compare the conditional forecast under the "LNG surge" scenario with the baseline unconditional forecast and the realized data, to gauge how effective an accelerated LNG substitution might have been.

Figure 5 displays the evolution of the log-level of natural gas prices, comparing the observed data with both the unconditional median forecast and the counterfactual scenario. In this counterfactual, it is assumed that LNG supply fully offset the pipeline supply disruptions experienced in 2022. This is implemented by reassigning the negative pipeline supply shocks as equivalent positive LNG supply shocks, effectively modeling a situation in which LNG entirely absorbed the shortfall from reduced pipeline inflows.

The left panel shows that, under this assumption, the counterfactual median price path (red dashed line) remains substantially below the baseline forecast (black solid line) and the realized data (black dashed line) throughout most of 2023 and 2024. This suggests that a perfectly elastic LNG substitution mechanism would have significantly mitigated the price surge observed during the energy crisis. The gap between the counterfactual and baseline series widens notably from mid-2022 onward, precisely when pipeline flows from Russia declined sharply. The right panel shows the log-level difference between the counterfactual and baseline scenarios, which becomes increasingly negative through late 2022 and 2023, eventually exceeding a gap of 200 log points. This implies that, in a scenario where LNG fully offset the pipeline supply losses, natural gas prices would have been significantly lower—by a margin exceeding 200 log points—throughout 2023 and into 2024. As expected, the difference remains exactly zero before February 2022, when both the baseline and counterfactual scenarios are identical by construction.

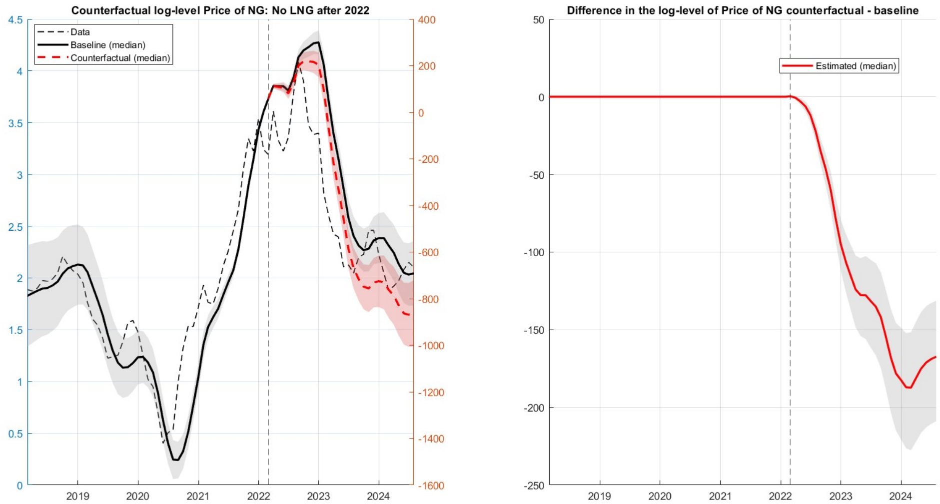

Figure 5: .

6 Conclusions

This paper revisits the fundamentals of the European natural gas market by explicitly distinguishing between pipeline gas and liquefied natural gas (LNG) within a Bayesian structural VAR framework identified through sign and elasticity restrictions. Motivated by the 2022 energy crisis, the analysis departs from the standard practice of modeling gas supply as a single aggregate and instead allows for substitution across heterogeneous supply channels. This distinction proves essential for understanding natural gas price dynamics, the transmission of supply disruptions, and the role of storage and expectations in shaping market outcomes.

The empirical results point to substantial heterogeneity in short-run supply elasticities. Pipeline gas supply is found to be extremely price inelastic, reflecting contractual rigidities and infrastructure constraints, whereas LNG supply is significantly more elastic and operates as the marginal adjustment mechanism in the European market. As a consequence, pipeline supply disruptions generate sharp but relatively transitory price increases, while LNG supply shocks exert more persistent effects on prices by shifting the medium-run market-clearing condition. On the demand side, short-run gas consumption is estimated to be highly price responsive, consistent with strong adjustment in industrial and residential usage during periods of stress, while income effects remain economically significant and stable. Storage and precautionary demand shocks emerge as an important amplification channel, particularly during episodes of heightened supply uncertainty.

Impulse response analysis and forecast error variance decompositions further highlight a clear horizon-dependent structure in gas price formation. Short-run price fluctuations are driven by a combination of supply shocks, flow demand adjustments, and inventory behavior, whereas medium- to long-run price dynamics are dominated by aggregate demand conditions and persistent precautionary motives rather than by supply disturbances alone. The historical decomposition underscores that major price spikes, such as those observed in 2022, arise from the interaction of multiple forces: binding supply constraints, low short-run substitution elasticities, and forward-looking storage behavior. No single shock is sufficient to explain these episodes in isolation.

The counterfactual exercises underscore the macroeconomic relevance of supply substitution. In the absence of LNG inflows during the crisis, European natural gas prices would have been several multiples higher well into the post-crisis period, highlighting the critical stabilizing role played by LNG infrastructure and global market integration. Conversely, scenarios with stronger demand reduction or frictionless LNG substitution indicate that both policy-driven conservation measures and supply-side flexibility can materially dampen price pressures, though neither fully eliminates the sensitivity of prices to supply security concerns.

Taken together, these findings carry important policy implications. They suggest that energy security in Europe depends not only on the level of supply, but also on its composition and flexibility. Investments in LNG import capacity, storage infrastructure, and market integration enhance resilience by expanding the set of feasible adjustment margins following adverse shocks. At the same time, the results caution against over-reliance on aggregate measures of gas supply, as such aggregates can obscure critical substitution dynamics that shape prices and macro-financial outcomes.

Several avenues for future research emerge from this analysis. Extending the framework to explicitly model international spillovers through global LNG markets, incorporating forward prices and futures-based measures of expectations, or embedding the identified shocks into a broader macro-financial model could further enrich the understanding of energy-driven fluctuations. Moreover, exploiting higher-frequency data or disaggregating demand by sector may help to sharpen inference on adjustment mechanisms during periods of acute stress. Overall, the evidence presented here reinforces the view that accounting for supply heterogeneity is indispensable for both empirical modeling and policy analysis of European natural gas markets.

References

Adolfsen, J. F., Ferrari Minesso, M., Mork, J. E., and Van Robays, I. (2024). Gas price shocks and euro area inflation. *Journal of International Money and Finance*, 149:103183.

Al-Sahlawi, M. A. (1989). The demand for natural gas: A survey of price and income elasticities. *The Energy Journal*, 10(1):77–90.

Antolín-Díaz, J., Petrella, I., and Rubio-Ramírez, J. F. (2021). Structural scenario analysis with svars. *Journal of Monetary Economics*, 117(C):798–815.

Arora, V. (2014). Estimates of the price elasticities of natural gas supply and demand in the united states. MPRA Paper 54232, University Library of Munich, Germany.

Baumeister, C. and Hamilton, J. D. (2015). Sign restrictions in bayesian structural vars. *Journal of Business Economic Statistics*, 33(1):1–13.

Baumeister, C. and Hamilton, J. D. (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review*, 109(5):1873–1910.

Bernstein, M. A. and Madlener, R. (2011). Residential natural gas demand elasticities in oecd countries: An empirical survey and meta-analysis. *Energy Policy*, 39(5):2859–2876.

Burke, P. J. and Yang, H. (2016). The price and income elasticities of natural gas demand: International evidence. *Energy Economics*, 56:551–564.

Dahl, C. and Duggan, T. E. (1996). U.s. energy product supply elasticities: A survey and application to the u.s. oil market. *Resource and Energy Economics*, 18(3):243–263.

EuropeanCommission (2024). Quarterly reports confirm significant recovery on eu gas and electricity markets in 4th quarter 2023. European Commission, Directorate-General for Energy.

Farag, M. (2024). Revisiting the dynamics and elasticities of the u.s. natural gas market. EWI Working Papers 2024-8, Energiewirtschaftliches Institut an der Universitaet zu Koeln (EWI).

Furtuna, O., Grassi, A., Ianiro, A., Kallage, K., Koci, R., Lenoci, F., Sowiński, A., and Vacirca, F. (2022). Financial stability risks from energy derivatives markets. Financial stability review, special feature, European Central Bank.

Güntner, J., Reif, M., and Wolters, M. (2024). Sudden stop: Supply and demand shocks in the german natural gas market. *Journal of Applied Econometrics*, 39(7):1282–1300.

Hamilton, J. D. (1996). This is what happened to the oil price-macroeconomy relationship. *Journal of Monetary Economics*, 38(2):215–220.

Hamilton, J. D. (2023). Comment on "the economic effects of a stop of energy imports from russia". *Brookings Papers on Economic Activity*, Fall:456–481.

International Energy Agency (2020). Gas 2020. Technical report, International Energy Agency (IEA). Published 10 June 2020.

International Energy Agency (2023). Gas market lessons from the 2022-2023 energy crisis. Technical report, International Energy Agency (IEA). Report page.

Joutz, F. (2009). Estimating regional short-run and long-run price elasticities of residential natural gas demand in the u.s. Working paper, George Washington University. Available at SSRN.

Kilian, L. (2008). The economic effects of energy price shocks. *Journal of Economic Literature*, 46(4):871–909.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–69.

Kilian, L. and Murphy, D. P. (2012). Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market var models. *Journal of the European Economic Association*, 10(5):1166–1188.

Kilian, L. and Murphy, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29(3):454–478.

Mason, C. F. and Roberts, G. (2018). Price elasticity of supply and productivity: An analysis of natural gas wells in wyoming. *The Energy Journal*, 39(1_suppl):79–100.

Moll, B., Schularick, M., and Zachmann, G. (2023). The power of substitution: The great german gas debate in retrospect. *Brookings Papers on Economic Activity*, 2023(2):395–481.

Molnar, G. (2022). Economics of gas transportation by pipeline and lng. In *Handbook of Natural Gas Transmission and Processing*. Springer.

Popkostova, Y. (2022). Europe's energy crisis conundrum: Origins, impacts and way forward. EUISS Brief 2, European Union Institute for Security Studies (EUISS).

Reuters (2021). On the cusp of europe's winter season, gas storage hits 10-year low. Reuters.

Ritz, R. A. (2019). A strategic perspective on competition between pipeline gas and lng. Technical report, University of Cambridge.

Rubaszek, M., Szafranek, K., and Uddin, G. S. (2021). The dynamics and elasticities on the u.s. natural gas market. a bayesian structural var analysis. *Energy Economics*, 103:105526.

Sharples, J. (2023). A brave new world? lng contracts in the context of market turbulence and an uncertain future. OIES Paper NG 187, Oxford Institute for Energy Studies.

Szafranek, K. and Rubaszek, M. (2024). Have european natural gas prices decoupled from crude oil prices? evidence from tvp-var analysis. *Studies in Nonlinear Dynamics Econometrics*, 28(3):507–530.