The effects of gas and oil prices on European industrial production: a Growth-at-Risk approach

Lorenzo Rossi[1,2]

[1] *Department of Economics, Management, and Quantitative Methods, University of Milan, Milan, Italy*

[2] *Department of Economics and Management, University of Pavia, Pavia, Italy*

September 18, 2026

Abstract

This paper investigates how oil and gas price shocks affect European industrial production using a Growth-at-Risk (GaR) framework. By focusing on the conditional distribution of industrial production rather than its mean, the analysis captures asymmetric and nonlinear effects of energy price fluctuations on economic activity. Using monthly data from 2004 to 2023, the results show that gas and oil price increases primarily amplify downside risks in aggregate industrial production, while net gas price spikes (NGPI) also constrain the upper-tail growth potential of the Energy sector. These effects are most pronounced in Germany and Italy, where industrial activity is highly dependent on natural gas, but are largely absent in France, reflecting its more diversified and nuclear-based energy system. The post-2020 period—marked by the COVID-19 pandemic and the Russian invasion of Ukraine—intensifies these vulnerabilities, making energy shocks more persistent and systemic.

1 Introduction

Industrial production is one of the pillars of the European Union (EU) economy. accounting for over 15% of its gross domestic product (GDP) according to World Bank data[1]. IEA (2022b) reports that the energy needs of this sector are equally significant: more than a quarter of total energy consumption in the EU stems from industrial activity, with over 70% of this demand met by natural gas (31%), oil (10%) and electricity (34%) and despite being the largest energy source, approximately a third of European electricity is itself generated from gas[2]. Therefore, the industrial system is tied to fossil fuel

[1] World Bank
[2] European Council

markets both directly and indirectly. A disturbance originating in these markets reaches industrial users twice: once through the fuel they burn, and once again through the electricity they buy. In recent years, this dependence and exposure to fossil fuels has been tested repeatedly. The COVID-19 pandemic produced sudden demand shifts and a temporary collapse of energy consumption from firms' side due to lockdowns; the Russian invasion of Ukraine then delivered persistent supply-side disruptions of a different kind, particularly in natural gas availability. As noted by (IEA (2022a)), Russia supplied close to 40% of EU gas imports in 2021, but by the end of 2022 the supply fell to roughly one fifth of the pre-war level. European Countries reacted by diversifying their sources and activating emergency storage measures, yet these changes came at the cost of higher price levels and price volatility. According to (ECB (2023)), between mid-2021 and late 2022, wholesale gas prices in Europe rose more rose dramatically, and electricity prices reached record highs in different member states (e.g. Germany), becoming a source of uncertainty for European producers. Between mid-2021 and late 2022 wholesale gas prices in Europe . Industrial output moved with these developments. Eurostat figures show industrial production in the EU falling by approximately 2.5% in 2022, with the contraction concentrated precisely where energy enters the production function most directly: output in electricity and gas supply declined by more than 7%, while gas-dependent manufacturing sectors such as chemicals, paper and metallurgy lost more than 5% on average[3].

Exposure to energy market conditions is far from uniform across the Union, however, and national energy systems differ enough to make the aggregate picture a poor guide to individual experiences. According to IEA estimates, German industry absorbs about a quarter of total final energy consumption, with electricity covering roughly 28% of industrial energy use and natural gas a further 32%. Italy displays a similar but more gas-intensive profile: industry accounts for around 31% of national final energy demand, of which electricity supplies approximately 35% and natural gas 44%. France is the natural counterexample. Thanks to its nuclear-dominated generation mix, electricity accounts for more than 40% of industrial energy consumption while natural gas represents less than 20% (IEA, 2022b). The same movement in international gas or oil prices will thus reach these three economies along different channels and at different speeds: where industrial processes rely on gas, price increases translate almost mechanically into higher marginal costs and compressed operating margins, whereas a nuclear-based power system leaves domestic industry, at least in the short run, less directly linked to fossil fuel quotations.

The French case requires one qualification, since the 2022 episode is often read as evidence against the insulation argument. Over that year French nuclear availability fell to its lowest level in decades, and France became a net importer of electricity for the first time in more than four decades. The cause,

[3] Eurostat

however, was domestic and technical rather than macroeconomic: stress corrosion cracking detected in the safety injection circuits of several reactors from late 2021 onwards forced an extensive programme of inspections and repairs across the fleet. This disruption is unrelated to the pandemic, to the war in Ukraine and to the international fossil fuel market more generally, and it operated through physical generation capacity rather than through the price of gas or oil. It is therefore a source of variation that our conditioning variables (TTF gas and Brent quotations) do not span by construction, and it does not undermine the interpretation of France as an economy whose industrial output is comparatively insulated from international fossil fuel prices. If anything, it makes the French evidence conservative: French industry faced a severe domestic energy supply disturbance during the same window, and still displays no detectable tail sensitivity to gas and oil price movements.

Industrial production and energy prices move together over the cycle (the Industrial Production Index contracts sharply in 2008–2009), in 2020 and again in 2022, while gas prices reach a level after the invasion of Ukraine that has no precedent in the sample, but the co-movement is not stable. It appears considerably tighter around the crisis episodes than in the calm periods that separate them, and an average correlation is not well equipped to describe a relationship of that kind. Figure 4 therefore reports exceedance correlation profiles between monthly changes in the Industrial Production Index and monthly changes in gas and oil prices, a diagnostic that measures co-movement conditional on both series exceeding a common threshold and thus isolates what happens in the tails rather than on average.

The profiles are markedly higher at the extremes than in the centre of the distribution: energy and industrial cycles co-move most closely when industrial production growth is at its lowest or its highest, and the pattern is asymmetric, with large price increases associated with deeper downside outcomes more than with stronger upside ones. Two implications follow. First, energy price dynamics have asymmetric effects on industrial production output. Second, since energy dependence varies so widely across countries and sectors, the aggregate correlation conflates situations that are economically quite different.

These considerations lead naturally to a quantile-based approach. This paper adopts a Growth-at-Risk (GaR) framework to study how the conditional distribution of industrial production growth in the European Union varies with natural gas and oil price conditions. The choice of these two energy carriers reflects their role in European industry: natural gas is the dominant fossil input for both manufacturing processes and power generation, and therefore carries the regional component of energy price risk, while oil remains the global benchmark for energy and transport costs. Alongside ordinary monthly price changes we use the net price increase transformations introduced by Hamilton (2003), which retain only the amount by which the current price exceeds its twelve-month trailing maximum and are zero

Figure 1: Industrial Production Index (2004-2023).

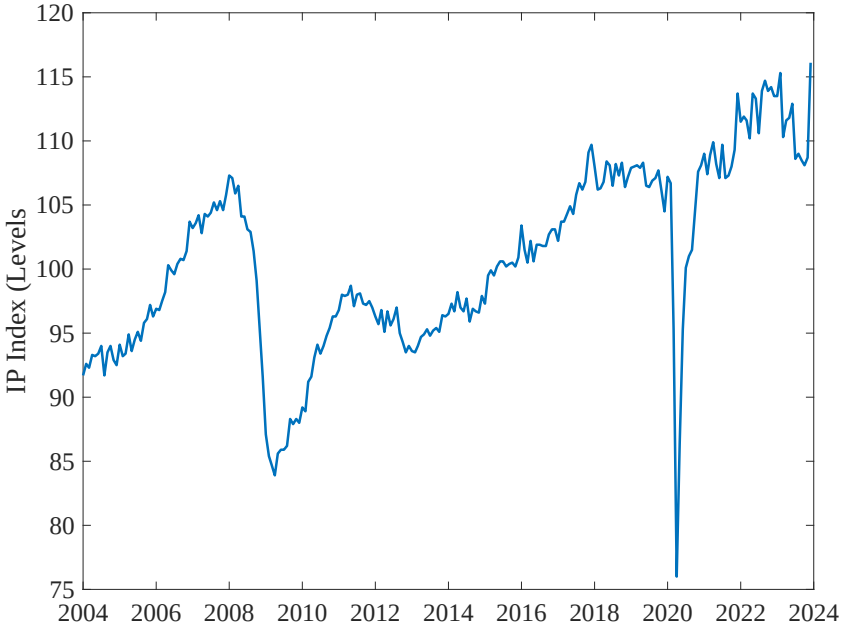
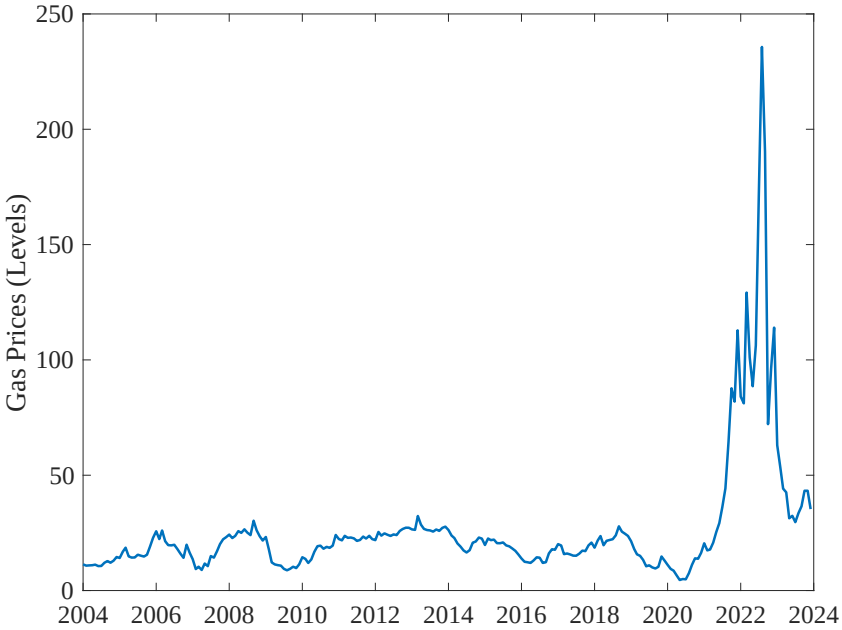
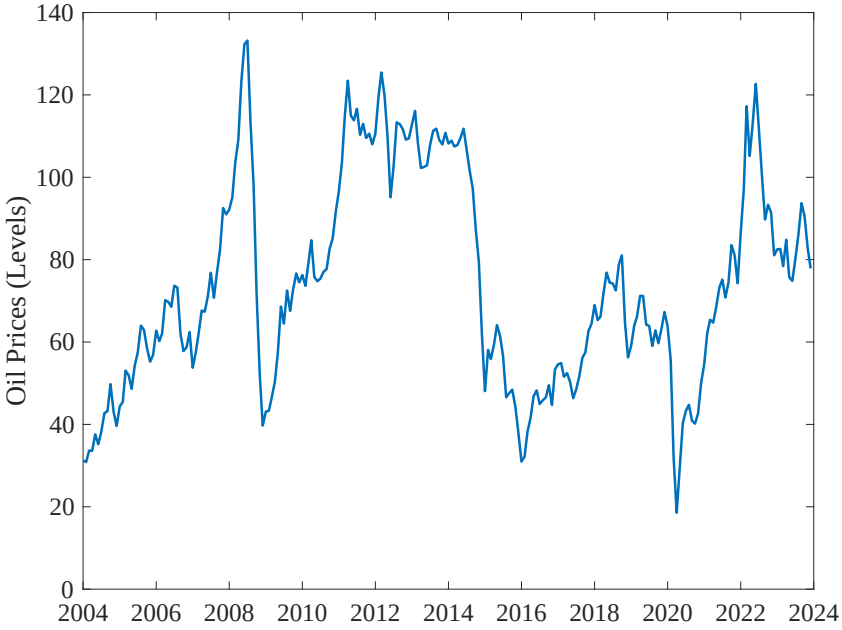


Figure 2: Natural Gas Prices, 2004-2023.



otherwise. These indicators, the Net Gas Price Increase (NGPI) and Net Oil Price Increase (NOPI), isolate the episodes of sustained upward price pressure that the energy–macroeconomy literature has

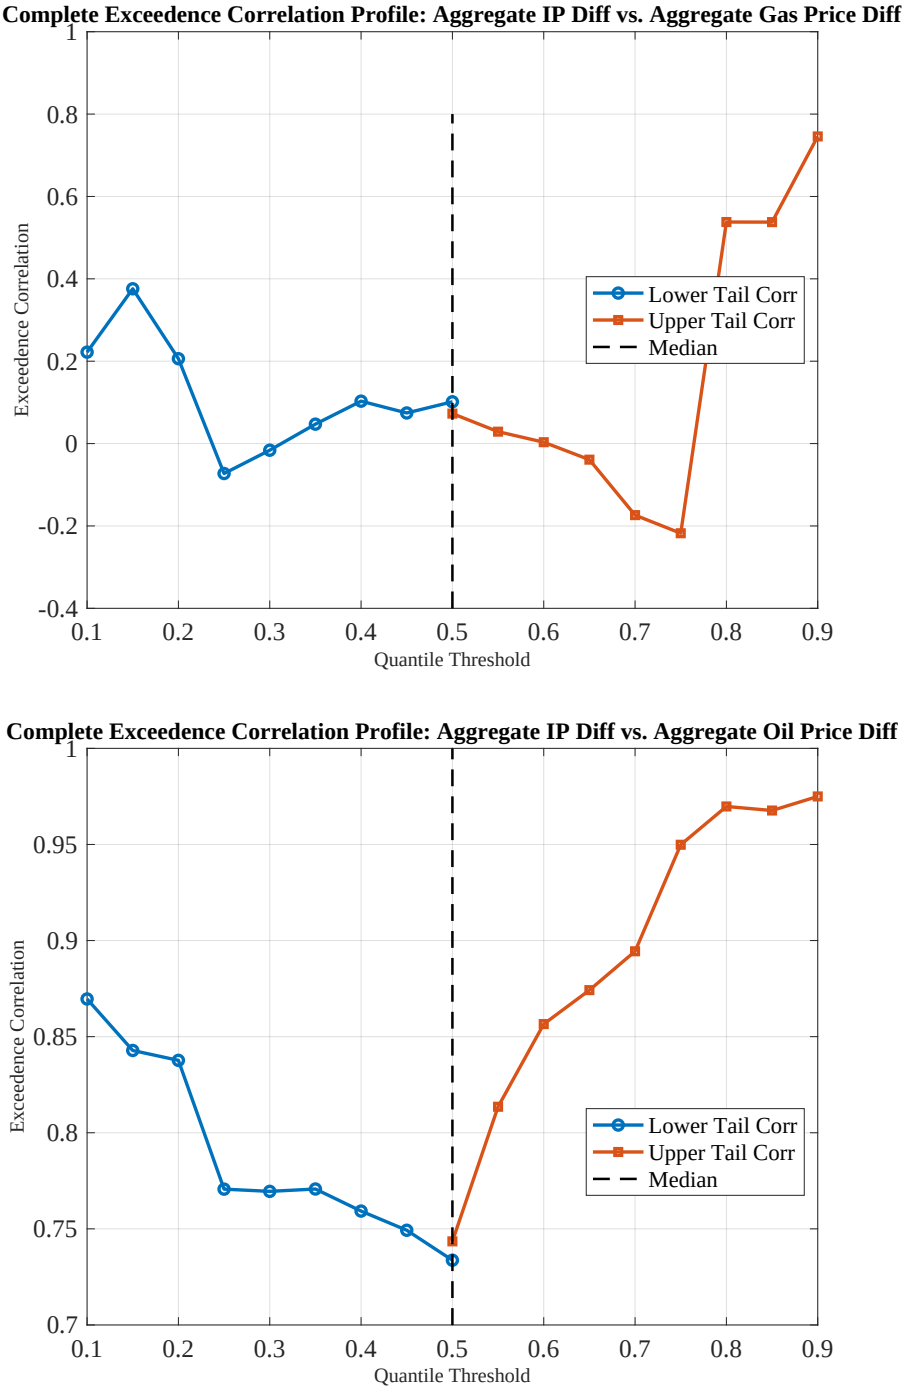
Figure 3: Brent Prices, 2004-2023.



found to matter most.

Following Adrian et al. (2019), we treat industrial production as an indicator whose conditional distribution changes shape with a set of conditioning variables, with energy conditions playing the role that financial conditions play in their analysis. Two things should be stated plainly about how the resulting estimates are to be read. First, the exercise is predictive rather than structural: energy prices are jointly determined with European industrial activity, and no attempt is made here to identify exogenous energy supply or demand disturbances in the sense of the structural oil literature. The coefficients below describe how the conditional quantiles of future industrial production growth vary with observed energy price conditions, not the causal response of output to an identified disturbance, and the terminology of the paper is chosen accordingly: we refer throughout to energy price movements, increases and conditions rather than to energy price shocks, a term we reserve for structurally identified objects. The net price increase transformation mitigates part of the concern, since it removes the ordinary two-way variation and retains only large directional breakouts (of which the 2021-2022 episode, driven by an unambiguously non-economic event, is by far the largest in the sample), but it does not amount to identification. Second, the horizon h indexes the lead between current conditions and subsequent industrial production growth within the estimation sample, not a genuine real-time forecast; the conditional distributions reported are fitted projections, consistent with the in-sample baseline of Adrian et al. (2019). Estimating conditional quantiles at several horizons allows the two edges of the

Figure 4: Exceedance correlation profiles of Industrial Production Index vs Natural Gas prices and Industrial Production Index vs Oil prices.



distribution to be examined separately: the accumulation of downside risk in the lower quantiles and the compression of upside potential in the upper ones. Neither is visible in a conditional mean.

Empirically, the analysis proceeds on two levels. At the aggregate level we assess how gas and oil price conditions map into the distribution of industrial production growth for the European Union; we

then disaggregate by Main Industrial Grouping, with particular attention to the Energy sector, and finally turn to three major euro area economies (Germany, Italy and France) whose energy structures differ in the ways documented above. A comparison of the pre-2020 subsample with the full sample completes the picture by asking whether the recent crises changed the strength and persistence of these relationships.

The results support a distributional reading of the energy: industry link and locate it more precisely than the motivating discussion would suggest. At the aggregate European level the relationship runs through the net gas price increase rather than through ordinary gas price variation. NGPI enters the 5th, 10th and 25th quantiles with significantly negative coefficients from $h = 6$ onwards, stabilising around -0.2 at the one-year horizon, while the upper quantiles are indistinguishable from zero at every horizon. Ordinary monthly gas price changes, by contrast, carry no significant information for any quantile or horizon, which suggests that routine price variability is absorbed by the production system and that vulnerability is associated with exceptional, sustained price breakouts rather than with volatility as such. The net oil price increase is largely uninformative, with only a short-lived negative coefficient at the 5th quantile around $h = 8$–10; ordinary Brent price changes reach the left tail only at the longest horizon. Differently for the Energy sector, the exposure lies in the opposite tail: both gas price changes and NGPI compress the *upper* quantiles, with coefficients at the 90th and 95th turning significantly negative as the horizon lengthens and reaching roughly -0.10 and -0.15 respectively by $h = 12$, while oil-based variables play no discernible role. The remaining Main Industrial Groupings show weak or insignificant effects throughout, so that the energy price–growth nexus is concentrated in the Energy sector, with any transmission to the rest of industry operating as a second-round effect. The fitted conditional densities make the same point in distributional terms. The expected shortfall of aggregate industrial production growth is persistently negative and highly volatile, falling to about -12% in 2008–2009 and to between -27% and -30% in 2020, whereas its upside counterpart remains positive and remarkably stable across the sample: the action is almost entirely in the left tail. Conditioning on energy prices also improves the fit of the predictive distribution during the months when it matters most, and the improvement is concentrated at short horizons, dissipating by $h = 12$. Cross-country results confirm that these regularities are not uniform. Germany reproduces the aggregate pattern most closely, with NGPI compressing the Energy sector's upper tail by about -0.2 at $h = 12$. Italy displays a milder version of the same configuration, with gas variables compressing the upper tail of both total and Energy sector production by around -0.1 at the one-year horizon. France shows no significant relationship for any energy variable, quantile or horizon — a null result consistent with its electricity-based industrial energy mix. Finally, the comparison between the pre-2020 subsample and the full sample indicates that these relationships predate the recent crises but

have become larger and more persistent since.

The remainder of the paper is organised as follows. Section 2 reviews the literature on Growth-at-Risk models and their extensions to macroeconomic and energy applications. Section 3 describes the dataset and the econometric methodology. Section 4 presents the empirical results at the aggregate, sectoral, subsample and country level, and Section 5 concludes by discussing what these findings imply for industrial resilience and for the European energy transition.

2 Literature Review

Growth-at-Risk (GaR) models have become essential tools for evaluating macroeconomic tail risks and understanding the nonlinear transmission of shocks across financial and real sectors. Unlike traditional econometric models that primarily estimate the conditional mean of macroeconomic indicators—thereby focusing on average outcomes—GaR emphasizes the full conditional distribution of growth, with special attention to its lower quantiles. These lower quantiles correspond to rare but highly disruptive economic outcomes such as recessions or sharp contractions in output. This shift from mean-based to distribution-based inference represents a fundamental methodological innovation, aligning macroeconomic risk modeling with concepts long established in finance, such as Value-at-Risk (VaR) and Expected Shortfall.

The foundational work by Adrian et al. (2019) introduced the GaR framework to quantify the entire conditional distribution of GDP growth as a function of prevailing financial conditions. By using quantile regressions, they demonstrated that deteriorating financial conditions systematically compress the lower tail of the growth distribution, increasing the likelihood of extreme negative realizations, while the upper tail remains comparatively stable. This asymmetry suggests that financial stress primarily operates through risk-amplifying mechanisms rather than by uniformly shifting the entire growth distribution. Their findings marked a significant advance in macro-financial analysis, bridging the gap between predictive econometrics and systemic risk monitoring.

Building on this foundation, Adrian, Grinberg, et al. (2022) extended the framework to study the intertemporal dimension of risks through the so-called "term structure of Growth-at-Risk." They documented that loose financial conditions tend to mitigate short-term downside risks but increase medium-term vulnerabilities, particularly during phases of rapid credit expansion. This result introduced a dynamic trade-off in financial stability policy: short-term stimulus may foster future fragility by encouraging leverage and risk-taking. The term-structure perspective thus provided a powerful interpretative tool for central banks and macroprudential authorities, who increasingly rely on GaR to assess the resilience of economic growth over different horizons.

Subsequent research has explored the nonlinear and multimodal nature of macro-financial dynamics. Adrian et al. (2021) showed that the joint distribution of economic and financial conditions can deviate substantially from Gaussianity, particularly during episodes of financial tightening. In such periods, the conditional distribution of growth becomes bimodal, reflecting multiple equilibrium paths—one associated with normal expansions and another with crisis regimes. This notion of "multimodality" emphasizes that financial constraints and nonlinear propagation mechanisms can give rise to abrupt shifts in macroeconomic states, challenging the assumptions of linear or symmetric adjustment processes.

Applications to the euro area have reinforced the relevance of the GaR methodology in multi-country and monetary union settings. Figueres & Jarociński (2020) applied the GaR framework using the Composite Indicator of Systemic Stress (CISS) and showed that surges in financial stress substantially elevate downside risks to growth, even when mean forecasts appear stable. Their findings provided early empirical evidence that financial vulnerabilities can silently accumulate beneath seemingly tranquil macroeconomic conditions. More recently, Lang et al. (2025) employed IMF data to analyze medium-term euro area risks, finding that financial shocks exert persistent, nonlinear effects that extend well beyond the initial impact period. These results collectively underscore that the effects of financial conditions on output are not only asymmetric but also state-dependent, with stronger transmission during fragile or leveraged periods.

A growing body of work has refined the empirical implementation of GaR models to improve real-time risk monitoring and cross-country comparability. Korobilis & Schröder (2025a) introduced the Quantile Factor-Augmented VAR (QFAVAR) framework, which integrates high-dimensional macro-financial information into quantile regressions, enabling a more coherent assessment of systemic risk spillovers across economies. Their model captures the interdependence of tail risks among European countries, a crucial step toward developing an early-warning system for synchronized downturns. Similarly, Korobilis & Schröder (2025b) proposed a probabilistic quantile factor model that allows for uncertainty quantification and posterior inference, thereby enhancing the interpretability of cross-country risk factors within a Bayesian context.

Further methodological innovation comes from quantile-based VAR approaches such as those of Chavleishvili & Manganelli (2024), who developed a Quantile Vector Autoregression (QVAR) to trace structural relationships across quantiles rather than just in the mean. This approach captures how shocks in financial or real variables propagate differently depending on the part of the conditional distribution being considered—for example, how a tightening shock may have a stronger effect during low-growth regimes than during expansions. These extensions demonstrate the flexibility of quantile-based methodologies in addressing heterogeneity and asymmetry—features that are central to understanding

the macroeconomic impact of shocks, including those arising from energy markets.

GaR-type methods have been applied to various macroeconomic risks. For instance, López-Salido & Loria (2024) introduced the Inflation-at-Risk (IaR) framework, while Brownlees & Souza (2021) performed global backtesting of GaR measures. Adams et al. (2021) and Mitchell et al. (2024) refined the construction of predictive densities from quantile regressions, improving forecast accuracy and allowing multimodal distributions. Clark et al. (2024) proposed a multicountry nonparametric quantile factor model, while Corradi et al. (2024) developed robust out-of-sample quantile coverage tests. Recent advances by Adrian et al. (2025) introduced robust econometric estimators that adapt to time-varying tail heaviness, enhancing the stability of GaR estimates during systemic episodes.

The quantile regression literature has evolved rapidly since the seminal contribution of Koenker & Bassett (1978), which introduced a robust framework for estimating conditional quantiles of a dependent variable given a set of covariates. Unlike classical least-squares estimation that targets conditional means, quantile regression (QR) enables researchers to study the entire conditional distribution of a variable and to quantify heterogeneous effects across its different regions. This approach has proven particularly valuable in macroeconomic and financial applications, where responses to shocks are often nonlinear and asymmetric. The methodology provides an essential foundation for frameworks such as Growth-at-Risk, Inflation-at-Risk, and related tail-risk models.

Over the last two decades, a rich body of research has extended the basic QR approach to capture additional sources of complexity in economic data. One major direction has been the modeling of cross-sectional and temporal dependence through latent factors and dynamic parameters. Chen et al. (2021) introduced Quantile Factor Models (QFM), which generalize the concept of factor structures beyond mean relationships by allowing latent factors to influence multiple quantiles of the distribution simultaneously. Their Quantile Factor Analysis (QFA) framework consistently estimates both quantile-dependent factors and loadings, accounting for heavy-tailed errors and heteroskedasticity that are typical in macro-financial data. This line of research has bridged the gap between quantile methods and high-dimensional factor modeling, forming the theoretical backbone of many multi-country or multi-sector GaR extensions.

Building on these ideas, Bernardi et al. (2024) proposed a Bayesian Dynamic Quantile Model Averaging (DQMA) approach, which explicitly models the instability of quantile relationships over time. By combining sequential Markov Chain Monte Carlo (MCMC) estimation with model averaging, their framework adapts dynamically to changing economic regimes, such as periods of high volatility or financial stress. Similarly, Candila et al. (2023) developed a mixed-frequency quantile regression framework for Value-at-Risk (VaR) and Expected Shortfall (ES) forecasting, integrating information from variables observed at different frequencies (daily, monthly, quarterly). This approach

is particularly relevant for energy and commodity markets, where high-frequency price information interacts with slower-moving macroeconomic fundamentals. Both contributions emphasize that quantile relationships are neither constant nor isolated across time and frequency domains, and that effective modeling must account for the layered nature of economic information.

Another important advancement involves the introduction of hidden state and regime-switching dynamics into quantile-based models. Merlo et al. (2022) developed Quantile Hidden Semi-Markov Models (QHSM), extending the traditional hidden Markov framework to accommodate conditional quantile dynamics within each latent regime. This allows for a probabilistic description of how economies transition between states characterized by distinct tail behaviors—such as expansion versus crisis phases—offering new insights into regime-dependent tail risks. In related work, Merlo et al. (2021) proposed a joint quantile regression model for VaR and ES, providing a coherent probabilistic structure that links pointwise quantile predictions to expected shortfall measures used in portfolio optimization. The earlier study by Laporta et al. (2018) contributes to this strand by comparing and selecting Value-at-Risk models for energy commodities, illustrating how quantile-based approaches outperform classical volatility models in capturing the asymmetric and fat-tailed nature of commodity price returns.

A parallel strand of research has emphasized semiparametric and nonparametric quantile models. Catania & Luati (2019) proposed a semiparametric framework that combines flexibility in functional form with the robustness of quantile estimation, allowing the data to inform the shape of the conditional distribution without imposing restrictive assumptions. Likewise, Abadir et al. (2023) developed a GARCH-based approach to density and functional forecasting that complements quantile models by providing full predictive distributions. Together, these approaches reflect a broader methodological shift toward flexible, data-driven representations of uncertainty that can accommodate nonlinearities, asymmetries, and regime dependence.

Bayesian methods have also played an increasingly central role in advancing quantile-based inference. Yu & Moyeed (2001) established the connection between quantile regression and the Asymmetric Laplace Distribution (ALD), showing that minimizing the quantile loss function is equivalent to maximizing the likelihood of an ALD-based model. This insight allowed quantile regression to be incorporated into a full Bayesian framework, enabling posterior inference, prior elicitation, and predictive simulation. Building on this foundation, Iacopini et al. (2022) proposed a Bayesian multivariate quantile regression with time-varying volatility, relaxing the common homoskedasticity assumption by embedding stochastic volatility and GARCH-type structures. Their model proves especially useful for capturing the evolving tail behavior of energy commodities, which experience volatility clustering and sudden regime shifts. Later, Iacopini et al. (2023b) extended this line of work to a mixed-frequency

quantile VAR (MF-QVAR), allowing for joint modeling of variables observed at heterogeneous frequencies—such as daily commodity prices and monthly macroeconomic indicators—thereby aligning quantile analysis more closely with real-time forecasting practices. More recently, Iacopini et al. (2023a) introduced the concept of quantile sensitivity, which measures how changes in one variable's quantiles affect the entire distribution of another, offering a new perspective on the transmission of monetary and inflationary shocks across the distribution of outcomes.

The literature has also made substantial progress in linking quantile regressions to the broader problem of density reconstruction. While quantile regressions yield point estimates for specific quantiles, policy and risk applications often require full predictive densities. Gaglianone & Lima (2012) proposed an econometric method for constructing conditional densities from sets of quantile estimates without assuming a parametric distribution. This nonparametric idea was later refined by Mitchell et al. (2024), who showed that density forecasts derived from quantile regressions can capture multimodal and asymmetric features in macro-financial data—properties often missed by Gaussian approximations. Complementary work by Lenza et al. (2025) employed quantile regression forests, a machine-learning approach that extends quantile estimation to large datasets, demonstrating notable gains in density forecasting for inflation risks. In parallel, Azzalini & Capitanio (2009) provided the theoretical underpinnings for modeling asymmetry through the multivariate skew-t distribution, a key building block in recent Growth-at-Risk implementations where tail heaviness and skewness vary over time and across economies.

While most GaR research focuses on financial and credit conditions, quantile-based methods have also gained increasing prominence in the analysis of energy markets and their macroeconomic repercussions. The early literature on energy and commodity price risks largely developed in parallel to the financial Value-at-Risk (VaR) tradition. Pioneering contributions such as Giot & Laurent (2003) and Laporta et al. (2018) employed quantile and VaR models to measure downside risk in oil, gas, and electricity futures markets. These studies showed that standard volatility models, including GARCH-type specifications, systematically underestimate tail risks during episodes of high price volatility, while quantile-based frameworks more accurately capture asymmetries and nonlinear dependence structures. Their results were instrumental in demonstrating that commodity prices, like financial returns, exhibit fat-tailed and regime-dependent dynamics that require distributional, rather than mean-based, modeling.

Building on these insights, recent contributions have extended quantile methods to analyze the complex interactions between energy prices, financial markets, and real economic activity. Ren et al. (2022) applied penalized quantile regressions (LASSO and SCAD) to forecast carbon return distributions, finding that incorporating nonlinear and asymmetric predictors improves the identification of

extreme price events. Similarly, Chuliá et al. (2024) employed quantile connectedness techniques to assess systemic vulnerabilities in European electricity markets, revealing that tail dependencies intensify during crises, leading to stronger cross-country and cross-sector contagion. Uribe et al. (2022) further explored the co-movement between electricity and natural gas prices across European markets, showing that these linkages become markedly nonlinear in times of distress — precisely the kind of state dependence that quantile models are designed to capture.

A particularly relevant methodological advance for macroeconomic applications comes from Apergis (2023), who developed a quantile-MIDAS framework to construct energy price risk measures that explicitly link the tails of oil and gas price distributions to GDP growth forecasts. Their findings highlight that periods of extreme volatility in energy prices are associated with elevated downside risks to output — a direct bridge between energy price uncertainty and macroeconomic vulnerability. Complementary work by Foroni et al. (2023) demonstrated that low-frequency macroeconomic fundamentals, such as industrial activity and inflation expectations, significantly influence high-frequency movements in electricity prices. This mixed-frequency evidence strengthens the case for quantile approaches that integrate multi-scale information, since shocks propagate differently depending on the time horizon and the distributional state of the economy.

These contributions resonate with the broader macroeconomic literature on the asymmetric effects of energy price shocks. A long line of empirical research beginning with Hamilton (2003) established that oil price increases have far more severe macroeconomic consequences than equivalent decreases — a pattern often referred to as the "asymmetric oil price response." Subsequent studies such as Hamilton (2011) and Bjørnland et al. (2018) reinforced this finding, showing that the magnitude and persistence of output contractions depend crucially on the state of the economy and the nature of the shock. Structural analyses by Blanchard & Gali (2007) linked this evolving sensitivity to changes in monetary policy credibility and the energy intensity of production. At a more disaggregated level, Herrera et al. (2011), Lee & Ni (2002), and Herrera & Karaki (2015) revealed that oil shocks affect sectors very differently, with energy-intensive industries—such as chemicals, basic metals, and transportation—displaying disproportionate output declines. These sectoral asymmetries provide a compelling rationale for moving beyond aggregate mean effects toward distributional frameworks that can capture heterogeneity across quantiles, sectors, and countries.

The convergence of these strands—quantile econometrics, energy price dynamics, and macroeconomic asymmetry—has opened the way for applying GaR-type models to energy-related systemic risks. However, such integration remains limited. Existing quantile-based energy models typically focus on return distributions or financial risk metrics, without explicitly linking them to real economic outcomes. A few emerging studies attempt to bridge this gap. For instance, Iacopini et al. (2022) and

Iacopini et al. (2023b) show that Bayesian multivariate quantile frameworks can capture the evolving joint tail behavior of energy prices and macro variables, while Candila et al. (2023) highlight the importance of mixed-frequency data structures in forecasting risk measures relevant to both financial and industrial sectors. Yet, comprehensive analyses that connect energy price asymmetries to the conditional distribution of industrial production across European economies remain scarce.

Beyond traditional energy and financial domains, GaR-type frameworks have been extended to environmental and climate risk assessments. Adrian, Boyarchenko, Giannone, et al. (2022) applied the GaR methodology to long-run paleoclimate data, uncovering how extreme climatic events shape macro-financial vulnerabilities through nonlinear feedback loops. Likewise, Vasseur & Aznarte (2021) used probabilistic quantile regression to forecast exceedance probabilities of air pollution thresholds, offering a new dimension of risk management for environmental policy. These studies underscore the versatility and cross-disciplinary relevance of quantile-based modeling: by focusing on the tails rather than the center of distributions, it becomes possible to understand how rare but consequential events—whether financial crises, energy shocks, or environmental disruptions—shape the broader economic landscape.

This paper contributes to the literature in three main ways. First, it extends the Growth-at-Risk framework beyond its traditional financial and macroeconomic domains to explicitly incorporate energy price dynamics, thereby linking fossil fuel shocks to real-sector tail risks. Second, by distinguishing between aggregate and sectoral industrial production — with a particular focus on the energy sector — it highlights how gas and oil price asymmetries transmit differently across layers of industrial activity, revealing both downside vulnerabilities and constraints on upside potential. Third, the paper integrates cross-country heterogeneity by comparing Germany, Italy, and France, three economies that differ markedly in their energy mixes and industrial dependence on natural gas. In doing so, it provides one of the first systematic assessments of how fossil fuel price shocks shape the conditional distribution of industrial output across Europe, offering novel empirical evidence and methodological insight into the asymmetric, nonlinear, and country-specific transmission of energy shocks to real economic performance.

3 Data and Model

3.1 Data

The analysis uses monthly data running from January 2004 to December 2023. Industrial production is measured by the Industrial Production Index (IPI) published by Eurostat,[4] which has the advantage of being harmonised across member states and available simultaneously at the aggregate level and

[4]Eurostat, volume index of production, calendar and seasonally adjusted, NACE Rev. 2 sections B–D, 2021 = 100.

by industrial grouping. The aggregate series covers the European Union as a whole (EU27); the country-level analysis is restricted to Germany, Italy and France.

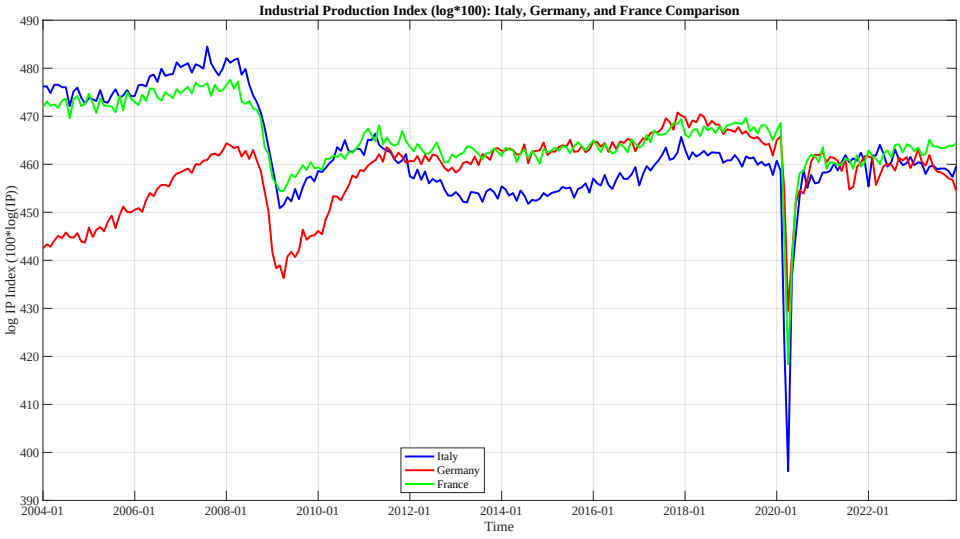
The choice of these three economies is dictated by their energy profiles rather than by size alone. Germany and Italy are among the most gas-intensive industrial systems in Europe, relying heavily on imported gas (particularly from Russia before 2022) as a primary energy source in manufacturing: natural gas covers approximately 32% of industrial energy consumption in Germany and 44% in Italy. France sits at the other end of the spectrum. Its industry runs predominantly on electricity, more than 70% of which is generated from nuclear plants, and its direct consumption of natural gas is below 20% of industrial energy use. The three countries therefore span the range of exposure within the euro area, which makes them informative about whether gas and oil price movements are associated with shifts in the whole distribution of industrial production growth in proportion to how much energy enters production, or whether the relationship is common to European industry regardless of its structure.

Figure 5 plots the logarithmic IPI, scaled by 100, for the European Union and for the three countries. The series share the two large cyclical episodes of the sample, which are the contraction of 2008–2009 and the collapse and rebound of 2020, but they separate after 2021. The German index (green) becomes visibly more volatile and declines through the 2022 energy crisis; the Italian recovery (red) stalls and remains uneven; French output (blue) fluctuates comparatively little. These divergences are consistent with the structural differences described above, although they need not be attributed to energy dependence alone: national fiscal and energy policy responses over 2022–2023 differed substantially in scale and design, and the figure cannot separate the two explanations. What it does establish is that a single aggregate series conceals a good deal of cross-country variation, which is what motivates the country-level estimates below.

Alongside the aggregate index, the analysis makes use of Eurostat's Main Industrial Groupings, and in particular of the Energy grouping, which collects the production, transformation and distribution of energy[5]. This sector occupies a special position in the argument. It is the part of industry whose costs are most directly tied to wholesale input prices, and it supplies an intermediate input to every other branch of production, so that any constraint on its output has the potential to propagate downstream. On the price side, the conditioning variables are the TTF natural gas price and the Brent crude oil price. Electricity prices are excluded from the baseline specification because they represent a potential transmission channel of natural gas prices to industrial production and are highly correlated with gas prices. Including both simultaneously would therefore complicate the interpretation of the gas price coefficients. All price and quantity series are transformed into natural logarithms and differenced, so that the dependent variable is the monthly growth rate of industrial production and the ordinary price

[5] Eurostat, Main Industrial Groupings (MIG) classification.

Figure 5: Industrial Production Index for Italy, Germany, and France, 2004–2023.



regressors are monthly rates of change, both expressed in percentage points

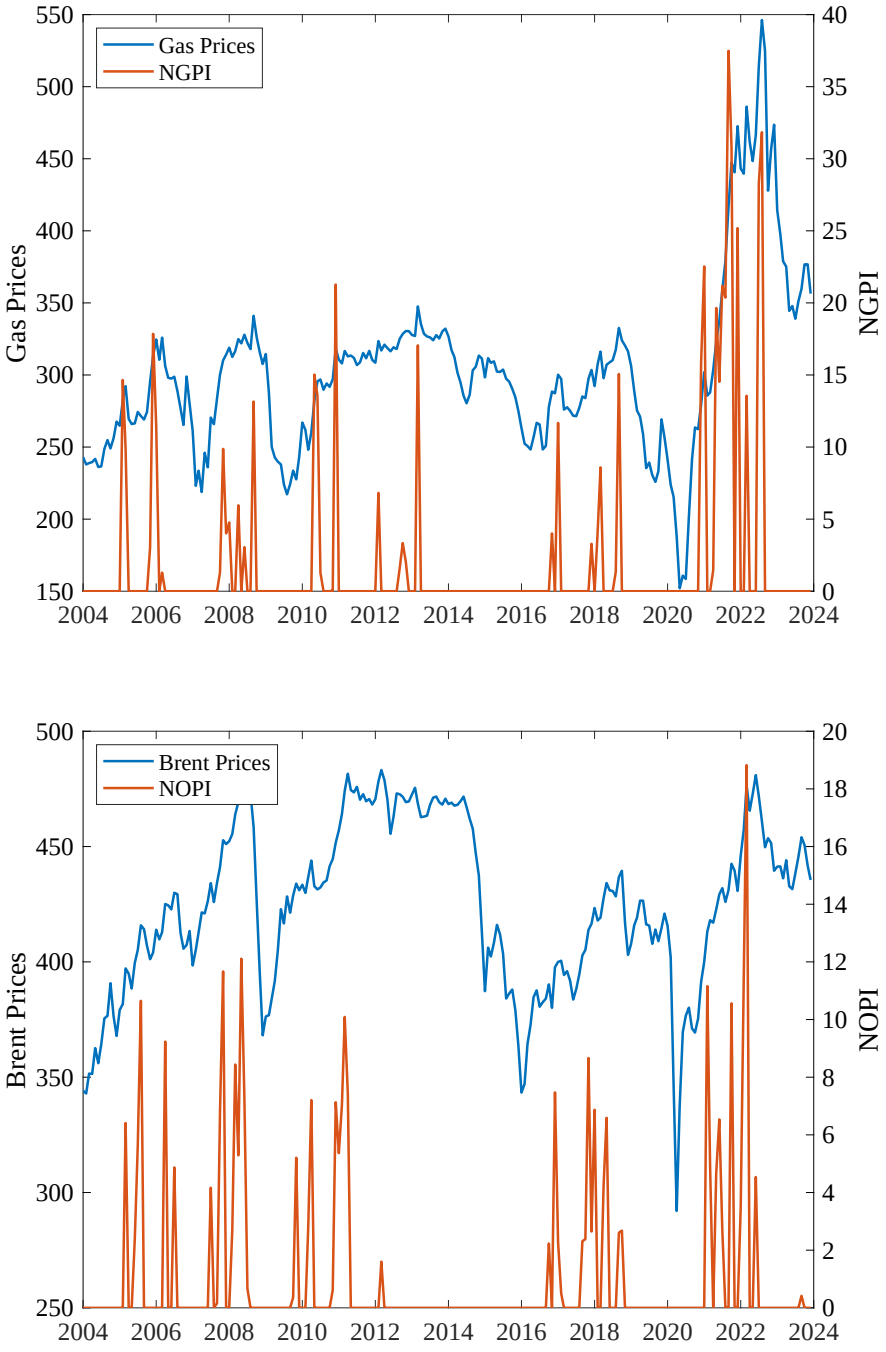
Ordinary price changes, however, are unlikely to be the object of interest. The evidence surveyed in Section 2 suggests that what disturbs production is not volatility as such but the arrival of price levels outside recent experience, and that increases and decreases are not mirror images of one another. To capture this we construct the net price increase measures introduced by Hamilton (2003). For oil, the Net Oil Price Increase at month t is

$$\text{NOPI}_t = \max(0,\ p_t - \max\{p_{t-1}, p_{t-2}, \ldots, p_{t-12}\})\,, \tag{1}$$

where p_t denotes the logarithm of the Brent price in month t, scaled by 100 so that the resulting index is expressed in percentage points. An analogous transformation applied to the TTF series yields the Net Gas Price Increase (NGPI). The measure is zero whenever the current price fails to exceed its twelve-month trailing maximum, and equal to the excess otherwise. It therefore discards mean-reverting fluctuations and retracements and retains only the amount by which the price has moved into territory not visited in the previous year. These indicators are designed to capture the effect of sudden and and persistent increases in fuel prices that the empirical literature has shown to be associated with disproportionately negative consequences for macroeconomic activity (Hamilton, 2003, 2011).

Figure 6 compares the behaviour of NGPI and NOPI with the underlying logarithmic price series. The construction changes the character of the regressor considerably. Where the raw series fluctuate continuously in response to seasonal demand, inventory cycles and geopolitical expectations, the net increase measures are zero during periods of stability or retracement. The result is a form of nonlinear

Figure 6: NGPI against log of Gas Prices and NOPI against log of Brent Prices. Orange spikes represent the difference between the current value and the maximum value over the previous 12 months, indicating episodes of net positive price shocks.



filtering that retains only those episodes of sustained price acceleration most likely to translate into real disruption. For gas, the nonzero episodes cluster in three periods: the commodity boom of 2008, the post-crisis tensions of 2011, and above all 2021–2022, when the magnitude of the spike dwarfs

anything else in the sample, reflecting both the size and the persistence of the surge as European supply arrangements were abruptly restructured. The oil index is sparser still, with meaningful realisations essentially confined to 2008 and mid-2022. One would therefore expect gas breakouts to be associated with larger and more persistent shifts in the distribution of European industrial production growth than oil breakouts of comparable size, and the estimates in Section **??** bear this out.

3.2 Model

The econometric framework follows the Growth-at-Risk methodology of Adrian et al. (2019), which extends mean-based models by focusing on the conditional distribution of future outcomes, and particularly on its tails. Rather than modelling the expected value of future industrial production growth, we model its conditional quantiles, and then use those quantiles to recover a full conditional distribution from which tail risk measures can be computed. The interest lies in whether energy conditions change the shape of that distribution — widening its left side, compressing its right — and not only its location.

Let ΔIPI_{t+h} denote the growth rate of industrial production h months ahead and let X_t be the vector of exogenous conditioning variables, that is, variables dated t and treated as predetermined with respect to ΔIPI_{t+h}. For each quantile level $\tau \in (0, 1)$ the conditional quantile function is specified as linear in the conditioning set,

$$Q_{\Delta IPI_{t+h}}(\tau \mid X_t) = \alpha_\tau + X_t'\beta_\tau, \tag{2}$$

where α_τ and β_τ are quantile-specific parameters.

The conditioning vector always contains one lag of industrial production growth, which carries the persistence of the series, together with a single energy variable: the monthly change in the gas price, the monthly change in the Brent price, NGPI, or NOPI.The four energy variables are estimated in separate specifications rather than jointly. This is a deliberate restriction. NGPI is a nonlinear function of the same gas price whose first difference is the alternative regressor, and the two are mechanically related; entering them together would make the resulting coefficients difficult to interpret and would spend degrees of freedom that the extreme quantiles cannot spare. Reporting the specifications side by side has the further advantage of making visible which representation of the energy price actually carries information, which turns out to be a substantive result rather than a technicality. A baseline specification conditioning on lagged industrial production alone serves as the point of comparison throughout.

The estimators $(\hat{\alpha}_\tau, \hat{\beta}_\tau)$ are obtained by solving the minimisation problem introduced by Koenker

& Bassett (1978),

$$(\hat{\alpha}_\tau, \hat{\beta}_\tau) = \arg\min_{\alpha,\beta} \sum_{t=1}^{T} \rho_\tau(\Delta IPI_{t+h} - \alpha - X_t'\beta)\,, \tag{3}$$

where $\rho_\tau(u)$ is the asymmetric quantile loss function, defined as

$$\rho_\tau(u) = \begin{cases} \tau u & \text{if } u \geq 0, \\ (\tau - 1)u & \text{if } u < 0. \end{cases} \tag{4}$$

Weighting positive and negative residuals differently is what makes the fitted line a quantile rather than a mean: at the optimum a proportion τ of the observations lies below the estimated hyperplane,

$$\mathbb{P}(\Delta IPI_{t+h} \leq \alpha_\tau + X_t'\beta_\tau \mid X_t) = \tau. \tag{5}$$

The problem is convex and piecewise linear, and is solved by linear programming. Estimation is conducted separately for each quantile τ, so that the coefficients are free to differ across the distribution; it is precisely this freedom that allows an energy variable to matter in the left tail and not at the median, or in the right tail of one sector and the left tail of another.

Under standard regularity conditions the estimator is consistent and asymptotically normal,

$$\sqrt{T}\left(\hat{\beta}_\tau - \beta_\tau\right) \xrightarrow{d} \mathcal{N}(0, V_\tau)\,, \qquad V_\tau = \tau(1-\tau)\, D^{-1}\Omega D^{-1}, \tag{6}$$

with $D = \mathbb{E}\big[f_{\Delta IPI|X}(Q_{\Delta IPI}(\tau \mid X))\, XX'\big]$ and $\Omega = \mathbb{E}[XX']$. This expression is of limited practical use here for two reasons. The sandwich form requires an estimate of the conditional density of the dependent variable at the quantile of interest, which is exactly the object that is hardest to estimate in the tails; and the asymptotic result assumes independent observations, whereas the direct multi-horizon design used below induces overlap. Projecting ΔIPI_{t+h} on information dated t for $h > 1$ means that consecutive residuals share $h-1$ periods, so that they follow a moving average process of order $h-1$ by construction, and conventional standard errors would be understated.

Inference is therefore based on a moving-block bootstrap. The procedure is applied separately to every (τ, h) pair and to every specification: for each of them, 1,000 bootstrap samples are drawn by resampling overlapping blocks of consecutive observations of $(\Delta IPI_{t+h}, X_t)$, the quantile regression is re-estimated on each sample, and the reported bands are the 68% and 95% quantiles of the resulting distribution of $\hat{\beta}_{\tau,h}$. The block length is set to $\ell = 12$, which covers the overlap generated at the longest forecast horizon.

The intercept α_τ locates the conditional quantile, while β_τ measures how changes in the regressors

displace it. A large negative coefficient at the lower quantiles ($\tau = 0.05$ or 0.10) indicates that increases in gas or oil prices are associated with a deterioration of the left tail of the industrial production growth distribution, consistent with the asymmetric patterns documented by Hamilton (2003) and Hamilton (2011). Note that the regressors are price *changes* and net price *increases*, not price levels, so the coefficients describe the association between the recent movement of the price and the subsequent distribution of output growth.

Although quantile regression does not assume a specific error distribution, it can be interpreted as a maximum likelihood estimator under an asymmetric Laplace distribution (ALD) for the residuals,

$$f(u \mid \tau, b) = \frac{\tau(1-\tau)}{b} \exp\Big[-\rho_\tau\Big(\frac{u}{b}\Big)\Big], \tag{7}$$

where $b > 0$ is a scale parameter. The ALD interpretation is not required for estimation, but it clarifies that the passage from a collection of quantile regressions to a single density is a change of representation rather than a new set of assumptions.

Following Adrian et al. (2019), once the conditional quantiles have been estimated at several quantile levels, the full conditional distribution of industrial production growth is reconstructed by fitting a skewed-t distribution through them. The skewed-t of Azzalini & Capitanio (2009) provides a flexible parametric form capable of capturing both fat tails and asymmetry, the two features characteristic of macroeconomic and energy-related series, and it does so with four interpretable parameters: location μ_t, scale σ_t, skewness λ_t and tail thickness ν_t. For each period t these four parameters are *estimated*, not imposed: they are chosen to minimise the distance between the fitted and the estimated quantiles, so that the parameter vector $\theta_t = (\mu_t, \sigma_t, \lambda_t, \nu_t)$ solves

$$\min_{\theta_t} \sum_{i=1}^{N_\tau} \Big[Q^{QR}_{\Delta IPI_{t+h}}(\tau_i \mid X_t) - Q^{ST}_{\Delta IPI_{t+h}}(\tau_i \mid \theta_t)\Big]^2, \tag{8}$$

where $Q^{ST}(\tau_i \mid \theta_t)$ denotes the theoretical quantile function of the skewed-t distribution,

$$Q^{ST}(\tau_i \mid \theta_t) = \mu_t + \sigma_t F^{-1}_{ST}(\tau_i; 0, 1, \nu_t, \lambda_t), \tag{9}$$

where $F^{-1}_{ST}(\cdot)$ is the quantile function of the standardised skewed-t distribution. Since N_τ exceeds the four parameters, the system is overidentified and (8) is a genuine nonlinear least-squares problem rather than an exact match; details of the minimisation procedure are provided in Appendix **??**.

The fitted distribution serves two purposes. It provides a smooth and continuous approximation of the conditional density implied by the discrete quantile estimates, which also resolves any crossing between adjacent estimated quantiles, and it makes tail-based risk measures computable. Growth-at-

Risk is defined as the fifth percentile of the fitted conditional distribution and Expected Shortfall as the mean outcome conditional on falling below it,

$$\mathrm{GaR}_{t+h}^{(0.05)} = F_{ST}^{-1}\Big(0.05;\, \hat{\mu}_t, \hat{\sigma}_t, \hat{\nu}_t, \hat{\lambda}_t\Big)\,, \qquad \mathrm{ES}_{t+h}^{(0.05)} = \mathbb{E}\Big[y_{t+h} \mid y_{t+h} < \mathrm{GaR}_{t+h}^{(0.05)}\Big]\,, \tag{10}$$

where y_{t+h} corresponds to ΔIPI_{t+h}. For symmetry of exposition we also report the upper-tail counterpart, the expected outcome conditional on exceeding the ninety-fifth percentile, which we label Expected Longrise. Comparing the two over time is the most direct way of seeing whether energy conditions are associated with movements in downside risk, in upside potential, or in both. The four parameters carry a straightforward reading in this context: μ_t is the location parameter of the distribution, σ_t its dispersion, λ_t its asymmetry, with negative values indicating a longer left tail, and ν_t the weight of the extremes.

Quantiles are estimated at $\tau = 0.05, 0.10, 0.25, 0.50, 0.75, 0.90, 0.95$ and horizons at $h = 1, 3, 6, 12$ months. The grid is deliberately concentrated in the tails, where the questions of interest lie, while retaining enough central mass to pin down the location of the fitted density: the lower quantiles describe periods of industrial contraction, the upper ones the loss of upside growth potential. The horizons span the range over which energy costs plausibly work their way into production decisions, from the immediate response to the annual contracting cycle. Each (τ, h) pair is estimated by a separate direct regression,

$$(\hat{\alpha}_{\tau,h}, \hat{\beta}_{\tau,h}) = \arg\min_{\alpha,\beta} \sum_{t=1}^{T-h} \rho_\tau(\Delta IPI_{t+h} - \alpha - X_t'\beta)\,, \tag{11}$$

with the sum truncated at $T - h$ to account for the horizon. The direct approach avoids the accumulation of specification error that afflicts iterated multi-step forecasts and lets the coefficients vary freely with the horizon, which is what makes it possible to trace the profile of $\beta_{\tau,h}$ against h and to observe, for instance, that the association between a net gas price increase and the lower tail builds gradually rather than arriving on impact.

A clarification about the meaning of h is in order, since the word "horizon" suggests forecasting. The regression in (11) relates conditioning information observed in month t to industrial production growth realised in month $t+h$, and its coefficients are estimated using the entire 2004–2023 sample. At no point is the model estimated on data up to some date and then evaluated on data after it. Consequently, the conditional distributions and the risk measures derived from them are *in-sample fitted values*: they answer the question "given the energy conditions prevailing in month t, how was the distribution of industrial production growth h months later configured, on average over this sample?", and not the question "what would a forecaster observing month t have predicted?". The distinction matters because

the second question requires estimating the model recursively with only the information available at each date, which we do not do here. This follows the baseline specification of Adrian et al. (2019), who characterise the historical conditional distribution over the full sample and treat out-of-sample performance as a separate exercise; a rolling-window evaluation of predictive accuracy is left for future work.

Two features of this framework deserve emphasis. The use of both country-level and sector-specific data allows heterogeneous patterns to be identified and structural vulnerabilities located within particular industrial segments rather than assumed to be uniform. And the distinction between linear price movements and net upward increases lets the specification accommodate the asymmetric responses that the empirical literature on nonlinear energy–macroeconomy linkages has documented, instead of imposing that a price rise and a price fall of equal size carry the same information.

This modeling framework is advantageous for several reasons. First, the use of both country-level and sector-specific data allows for the identification of heterogeneous transmission mechanisms and the detection of structural vulnerabilities across industrial segments. Second, by distinguishing between linear price movements and net upward shocks (NGPI and NOPI), the model incorporates asymmetric responses that are consistent with empirical evidence on nonlinear energy-macroeconomy linkages. Third, the explicit focus on the distribution tails rather than the conditional mean permits a more policy-relevant analysis of extreme events, enabling the quantification of worst-case scenarios with direct implications for risk management and macroprudential oversight.

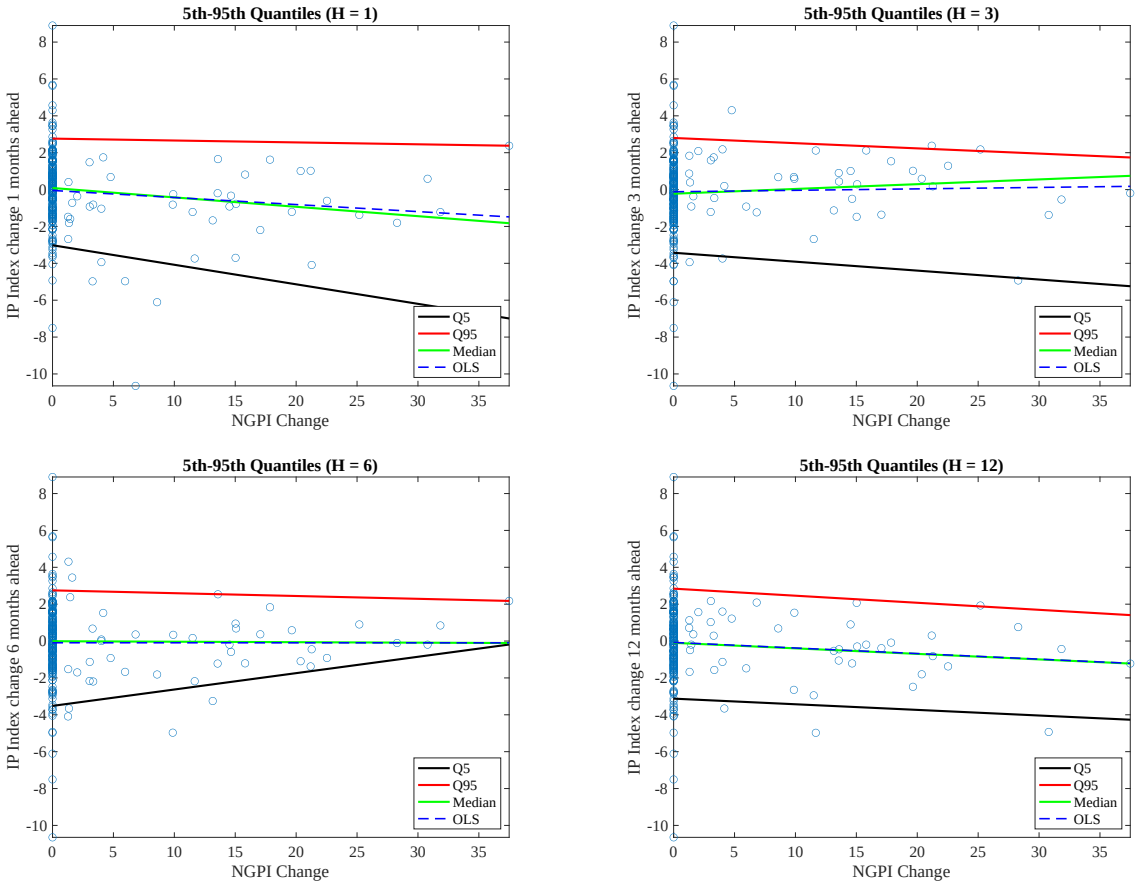
4 Results

This section presents the results of the quantile regression analysis, examining how gas and oil price dynamics affect the conditional distribution of industrial production growth across multiple dimensions of the European economy. The analysis proceeds in four complementary stages.

First, we evaluate the effects of energy price shocks at the aggregate European Union level, assessing how both regular price fluctuations and large net increases shape the entire distribution of industrial production outcomes. Second, we extend the analysis to a sectoral level, distinguishing the heterogeneous responses of the Energy sector. Third, we perform a subsample comparison, contrasting the pre-2020 period with the full sample, to identify how recent crises (namely, the COVID-19 pandemic and the Russian invasion of Ukraine) have altered the magnitude and persistence of these effects. Finally, we conduct a country-level analysis for Germany, Italy, and France, comparing how structural differences in energy dependence and industrial composition influence the transmission of price dynamics.

Together, these layers of analysis provide a comprehensive view of how energy market volatility translates into macroeconomic tail risks, highlighting both common vulnerabilities and structural asymmetries across the European industrial landscape.

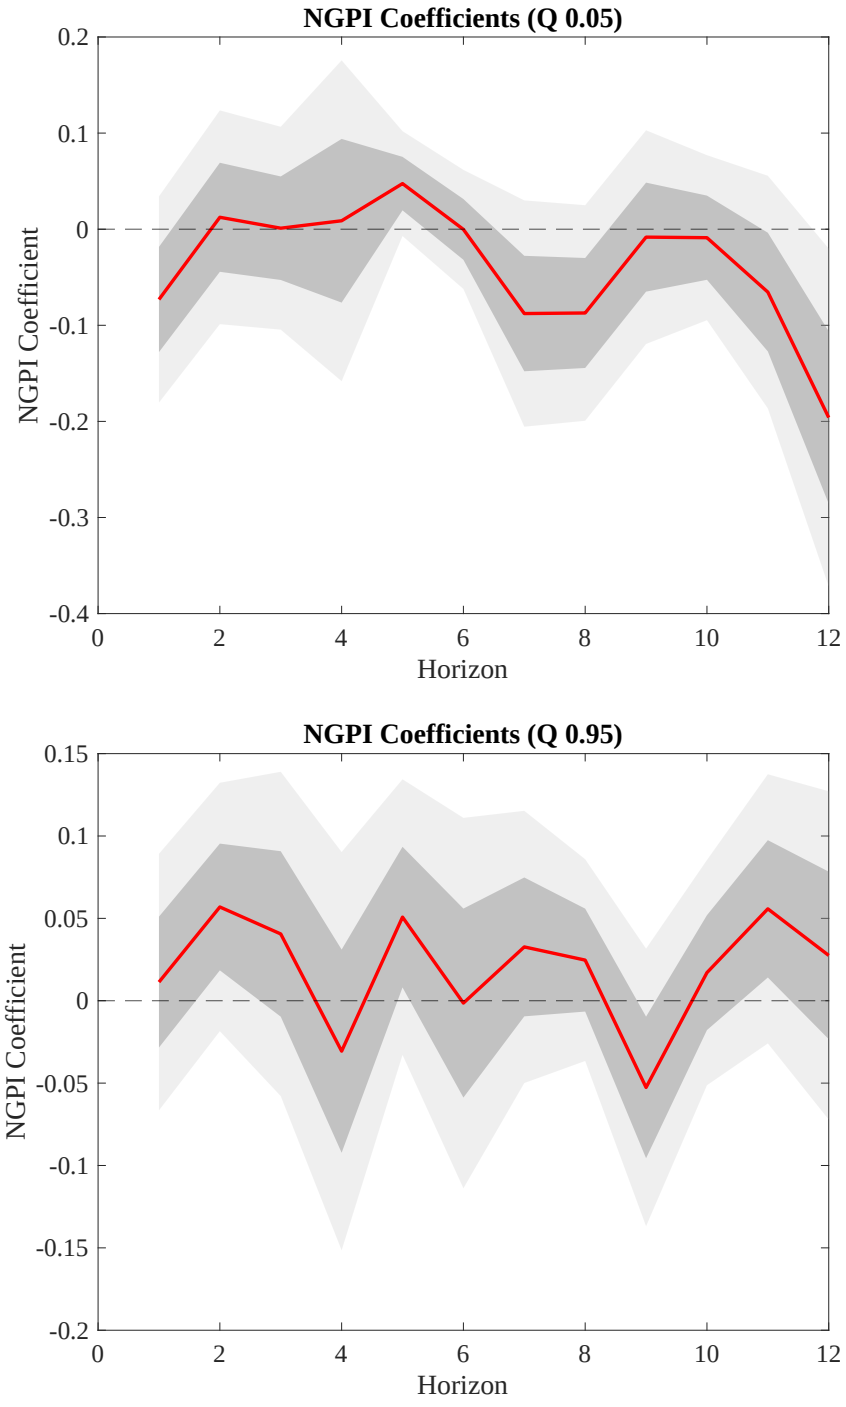
Figure 7: Scatter plot of IPI Change "h steps ahead" relative to current NGPI ($h = 1, 3, 6, 12$) for the IPI Change of the energy sector. The lines represent the slopes of the quantile regression for different quantiles.



Figures 8 and 24 display the estimated coefficients of the Net Gas Price Increase (NGPI) variable across quantiles and forecast horizons, while Figure 7 provides a visual representation of the quantile regression fit between NGPI and IPI Change at the tails ($\tau = 0.05$ and $\tau = 0.95$). The results reveal a clear asymmetric pattern in the conditional distribution of IPI Change: NGPI exerts a persistent negative effect concentrated in the lower quantiles, while the upper quantiles remain statistically insignificant across all horizons.

The effects of NGPI gradually intensify over time, suggesting that gas price shocks have delayed but substantial impacts on downside risk. In particular, the coefficients for the lower quantiles ($\tau = 0.05, 0.10, 0.25$) become significantly negative from $h = 6$ months onward, reaching their maximum

Figure 8: Coefficients of NGPI on the total aggregate IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.



magnitude around the one-year horizon. This finding implies that sharp and persistent gas price spikes progressively increase the likelihood of low-output outcomes, as captured by the left tail of the industrial production distribution. The estimated coefficients stabilize around -0.2 at the twelve-

month horizon, confirming the buildup of downside risk. These dynamics are consistent with the energy shock asymmetry described by Hamilton (2003) and Hamilton (2011), who show that large and sustained energy price increases have delayed but pronounced real effects, primarily through higher input costs and production constraints. In this context, NGPI captures these nonlinear adjustments by filtering out symmetric price fluctuations and isolating the adverse effects of abrupt gas price surges.

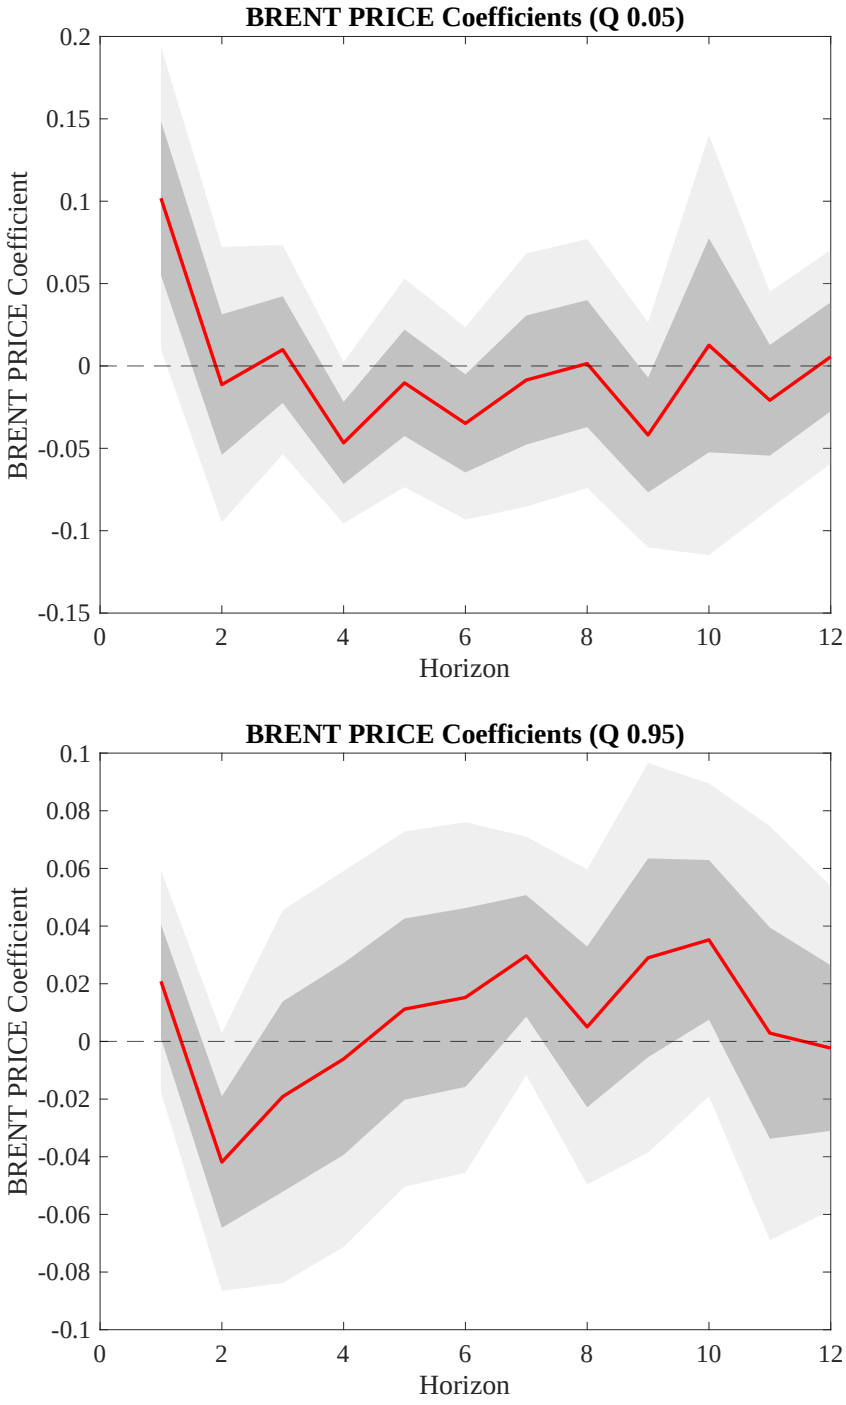
In contrast, when using the standard monthly percentage change in gas prices as a regressor, the results show no statistically significant impact on aggregate industrial production. As illustrated in Figures 21 and 23, the coefficients for ΔGP_t fluctuate around zero across all quantiles and horizons. Similarly, the scatter plots in Figure 16 show no discernible pattern between gas price changes and IPI Change. This lack of significance suggests that regular gas price variability is largely absorbed by the production system, whereas exceptional price spikes—captured by NGPI—are the true drivers of vulnerability.

Turning to oil prices, Figures 27 and 28 display the estimated coefficients of Brent price changes across quantiles and forecast horizons, while Figure 19 shows the corresponding quantile regression fits for the distribution tails ($\tau = 0.05$ and $\tau = 0.95$). The results indicate that oil price fluctuations have a delayed but significant effect on the lower tail of industrial production growth. Specifically, the coefficients for the 5th quantile become negative by $h = 12$, indicating that higher oil prices amplify the likelihood and magnitude of extreme negative outcomes in industrial production. At shorter horizons, a mild positive response can be observed in the lower quantiles, possibly reflecting firms' anticipatory behavior—such as accelerating production to hedge against future energy cost increases or fulfilling existing contracts under fixed production schedules. However, as horizons lengthen, these temporary effects are outweighed by the cumulative impact of rising energy costs, resulting in a clear buildup of downside risk over time.

In contrast, the Net Oil Price Increase (NOPI) variable exhibits weaker and less persistent effects. As shown in Figures 30 and 31, the NOPI coefficients fluctuate around zero for most quantiles and horizons, with only a short-lived negative effect at the 5th quantile—approximately -0.2 at $h = 8, 10$—which fades and reverts toward zero by $h = 12$. These results suggest that isolated oil price spikes during the sample period did not translate into sustained industrial contractions.

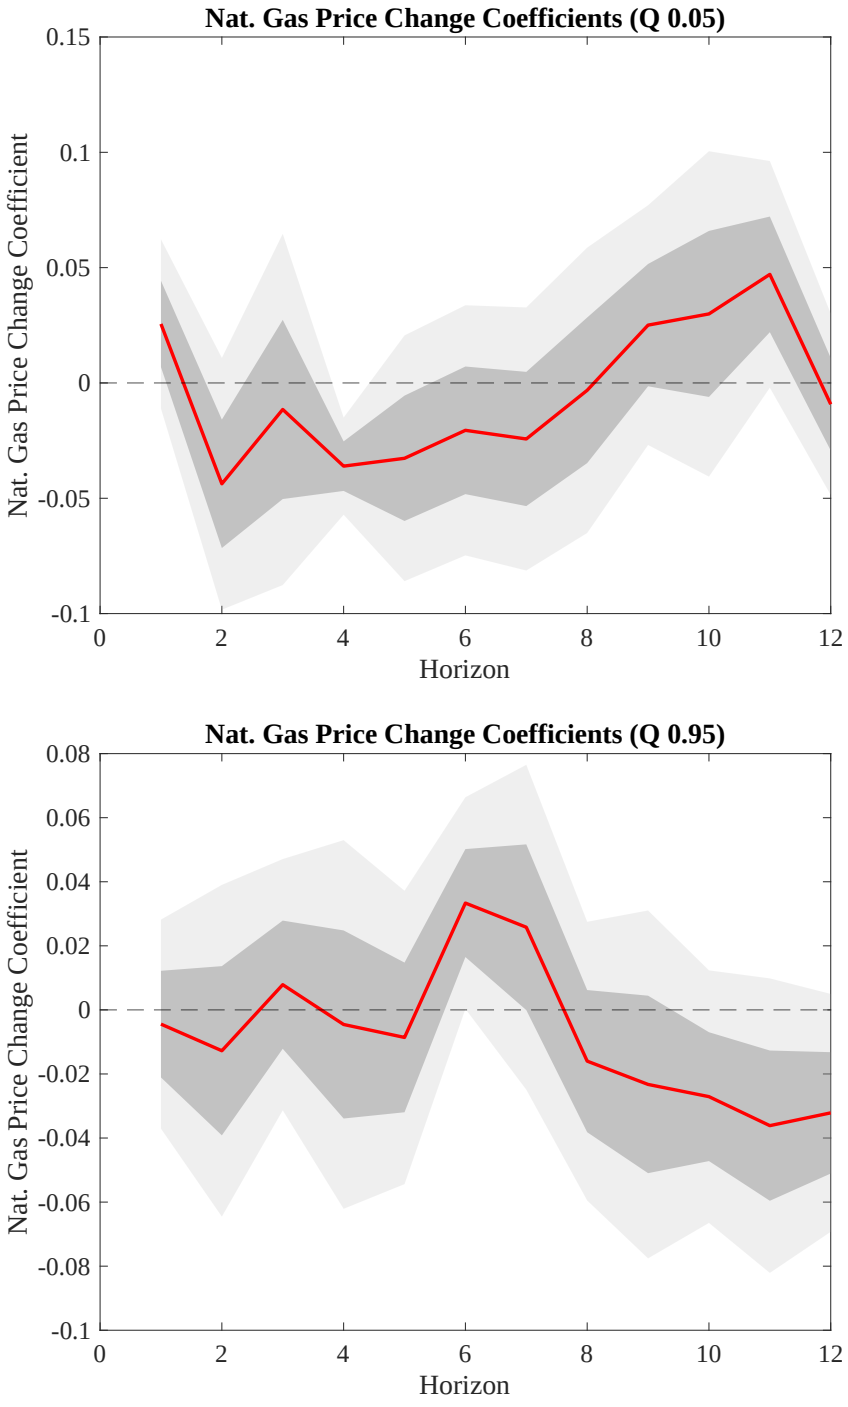
Figure 11 displays the time-varying Expected Shortfall and Expected Longrise of aggregate industrial production growth at the one-month horizon, derived from the fitted skewed t distributions conditional on NGPI (left panel) and Brent Price Change (right panel). The Expected Shortfall captures the average growth outcome conditional on realizations falling below the 5th percentile, while the Expected Longrise captures the symmetric upside counterpart above the 95th percentile. Both panels exhibit a striking asymmetry: the Expected Shortfall is persistently negative and highly volatile,

Figure 9: Coefficients of Brent Prices on the Energy sector's IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.



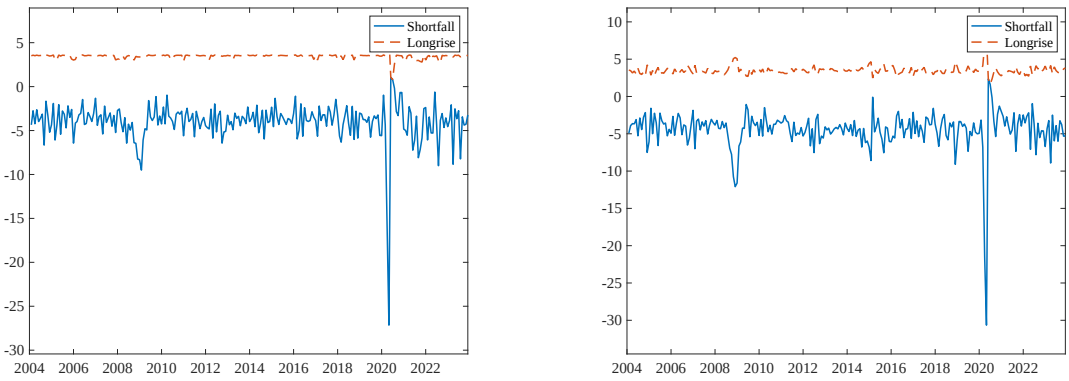
whereas the Expected Longrise remains positive and remarkably stable throughout the sample, showing minimal sensitivity to crisis episodes. Two major deteriorations in downside risk are visible in both specifications (the 2008–2009 Global Financial Crisis, where the Shortfall drops to approximately

Figure 10: Coefficients of Gas Price Change on the Energy sector's IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.



-12%, and the COVID-19 shock of 2020, where it plunges to roughly -27% to -30%). The similarity between the two panels is notable: despite capturing distinct energy price channels, both NGPI and Brent Price Change yield nearly identical Shortfall dynamics at $h = 1$. This convergence suggests

Figure 11: Expected Shortfall and Expected Longrise of aggregate industrial production growth at $h = 1$, conditional on NGPI (left) and Brent Price Change (right).

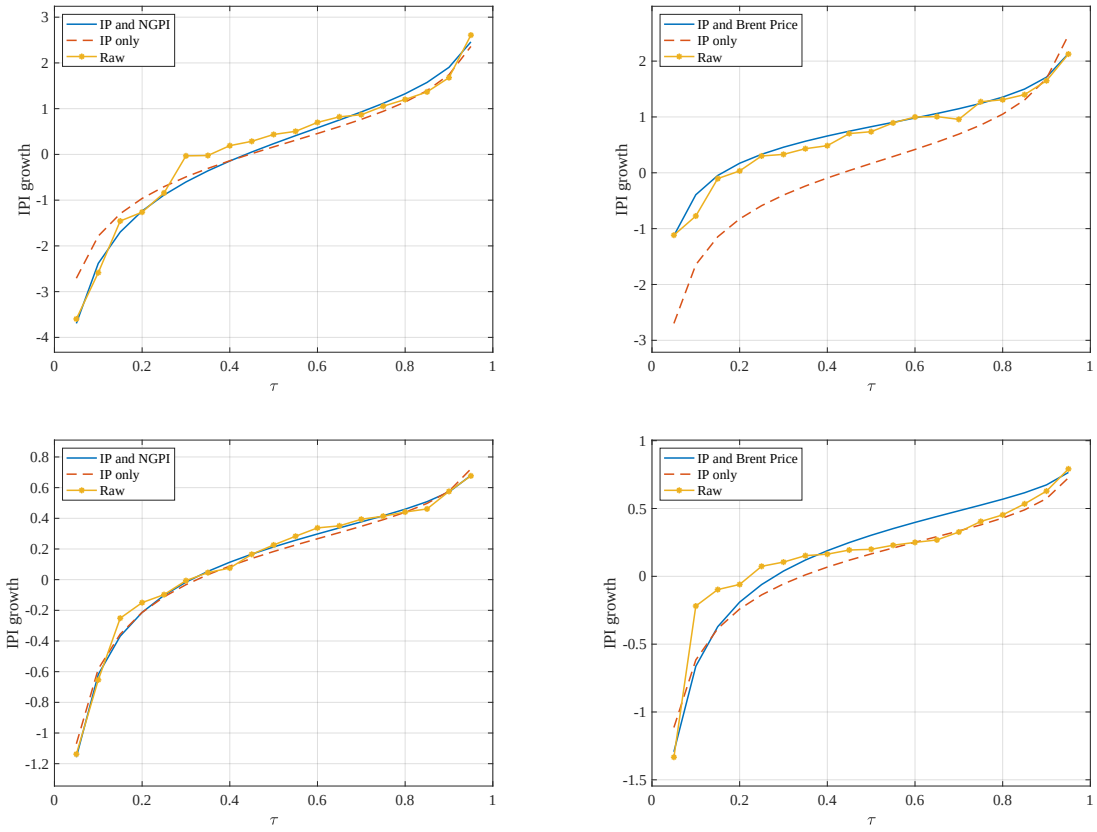


that at very short horizons, the conditional distribution of industrial production is dominated by common macroeconomic disturbances, such as pandemic lockdowns or systemic financial stress, rather than by the specific energy price channel. The distinguishing role of gas versus oil shocks becomes more apparent at longer horizons, where the persistence and magnitude of the Shortfall diverge across specifications (see Appendix 5 for the corresponding figures at $h = 3, 6$, and 12 months).

To evaluate whether the inclusion of energy price information improves the characterization of the conditional distribution during periods of acute stress, Figure 12 compares the fitted inverse CDF of aggregate industrial production growth for March 2022, the month right after the Russian invasion of Ukraine, under three specifications: a model conditioning on lagged industrial production and the energy price variable (IP and NGPI, left column; IP and Brent Price, right column), a baseline model conditioning on lagged industrial production alone (IP only), and the empirical quantiles computed from the raw data (Raw). The top row displays the results at $h = 1$, the bottom row at $h = 12$. Results for intermediate horizons ($h = 3$ and $h = 6$) are reported in Appendix **??**.

At the one-month horizon, both energy-augmented models track the empirical quantile function far more closely than the baseline. The improvement, however, operates through different regions of the distribution depending on the energy variable. For the NGPI specification (top left), the IP-only model diverges primarily in the middle and upper quantiles ($\tau \approx 0.3$–0.8), systematically underestimating growth outcomes. For the Brent specification (top right), the distortion is concentrated in the left tail: the IP-only model predicts growth as low as -3% at the 5th percentile, while both the Brent-augmented model and the raw quantiles indicate values closer to -1%. In other words, the baseline model without oil price information substantially overestimates downside risk during the initial months of the energy crisis. These complementary patterns suggest that gas and oil prices carry distinct informational content for different parts of the conditional distribution (NGPI primarily refines the characterization

Figure 12: Fitted inverse CDF of aggregate industrial production growth for March 2022, comparing energy-augmented models (blue), baseline IP-only model (red dashed), and empirical quantiles (orange). Top row: $h = 1$; bottom row: $h = 12$. Left column: NGPI specification; right column: Brent Price Change specification.



of the center and right tail, while Brent prices sharpen the left tail where downside risk is measured).

At the twelve-month horizon (bottom row), the three curves largely converge in both specifications: the energy-augmented and IP-only models produce nearly indistinguishable quantile functions, both closely aligned with the empirical quantiles. This convergence indicates that the informational content of the contemporaneous energy price shock is gradually absorbed into the conditional distribution as the horizon lengthens, and that the marginal contribution of energy price conditioning is strongest at short to medium horizons—precisely where timely risk assessment is most valuable. Together, these results provide direct evidence that energy price conditioning materially improves the fitted conditional density during episodes of severe market stress, reinforcing the empirical relevance of the GaR framework adopted in this study

Sectoral analysis provides further insight into the sources of heterogeneity across industrial activity. Figures 10 and 23 present the estimated coefficients of Gas Price Change across quantiles and horizons for the Energy sector, while Figure 17 displays the corresponding quantile regression fit at

the distribution extremes. The results show a clear and persistent contraction in the sector's upper tail: coefficients for the 95th quantile ($\tau = 0.95$) become significantly negative by $h = 12$, reaching values around -0.1. This pattern indicates that increases in gas prices dampen the Energy sector's upside potential (its capacity for high-growth outcomes over the medium term).

A similar pattern is observed for the NGPI variable, as shown in Figures 25 and 26. Here, the right tail of the Energy sector's IPI Change distribution contracts notably, with coefficients for the 90th and 95th quantiles turning significantly negative as horizons extend. The estimated effects reach approximately -0.15 by $h = 12$, confirming that large upward movements in gas prices suppress long-run expansion potential within the sector. This finding reveals a distinct form of vulnerability: while gas shocks increase downside risks for aggregate industrial production, they simultaneously constrain upside opportunities within the Energy sector itself.

The reduction in the Energy sector's upper-tail growth potential carries broader implications. Because the Energy sector provides critical inputs to all other branches of industrial activity, its limited capacity to expand may indirectly restrict production growth elsewhere in the economy. Such spillover effects are likely to materialize with a lag, as energy supply constraints lead to higher costs or reduced availability for downstream industries.

In contrast, oil price dynamics appear to play a negligible role in shaping the Energy sector's growth distribution. The results for Brent Price Change and NOPI (reported in the Appendix) show coefficients that are statistically insignificant across all quantiles and horizons. This asymmetry between gas- and oil-driven effects underscores the European industrial system's structural dependence on natural gas as a direct production input rather than a secondary energy source.

For the other industrial groupings—such as Capital Goods, Consumer Goods (both durable and non-durable), and Intermediate Goods—the estimated effects of both Gas Price Change and NGPI are weak or statistically insignificant. In these sectors, price shocks may be absorbed through anticipatory production adjustments, cost pass-through mechanisms, or longer production cycles that mitigate short-term exposure to energy volatility. As a result, the energy price–growth nexus appears to be concentrated within the Energy sector, with only potential second-round effects diffusing to the rest of the industrial system.

4.1 Subsample Comparison

To evaluate how recent crises have altered the relationship between fossil fuels prices and industrial production, the quantile regressions were re-estimated on a restricted sample covering January 2004 – December 2019, thus excluding the COVID-19 pandemic and the Russian invasion of Ukraine. Figures 38 and 39 compare the coefficients from the pre-COVID sample with those from the full sample

for the 5th quantile, focusing on downside risks to industrial production.

The comparison reveals that the negative effects of both NGPI and Oil Price Change were already present in the earlier period, indicating that industrial production had been vulnerable to energy price shocks even under normal conditions. However, the crises of the 2020s have amplified and prolonged these effects. For both variables, the negative coefficients become more pronounced and persist over longer horizons in the full sample. In the pre-COVID sample, coefficients tend to turn negative earlier in the forecast horizon but fade relatively quickly, whereas in the full sample they remain negative and significant throughout, suggesting more durable downside risk. This persistence is particularly evident for NGPI, where the negative response of industrial production deepens and remains statistically significant up to twelve months ahead.

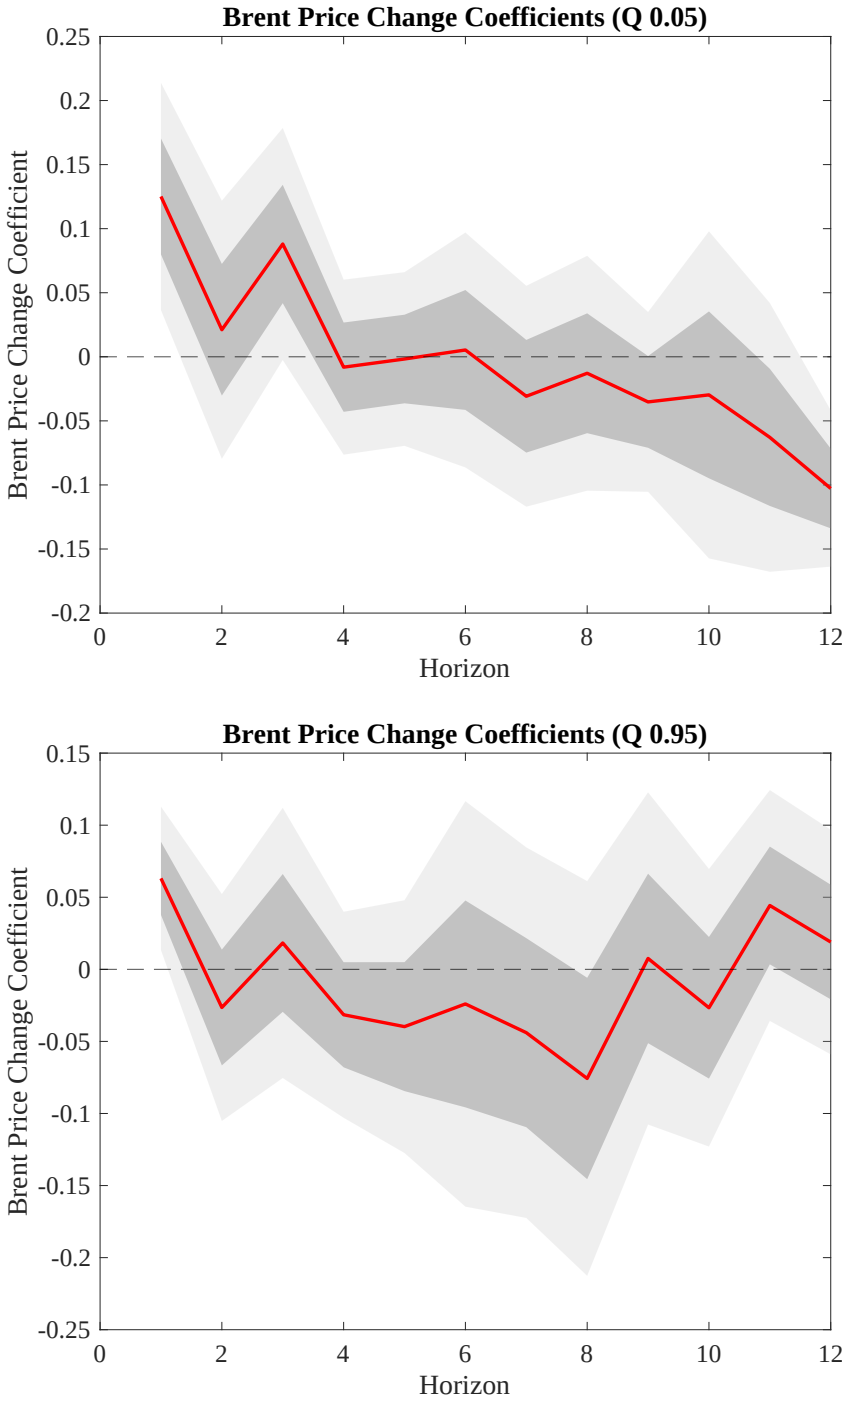
These results suggest that the pandemic and the subsequent energy-supply disruptions have modified the persistence of industrial vulnerability to energy price shocks. Prior to 2020, the effects of gas and oil price increases were more transient, reflecting partial resilience and faster mean reversion. In contrast, in the full sample the shocks exhibit more cumulative and lasting effects, consistent with heightened fragility in the energy–industry nexus. The crises therefore did not reverse the direction of the relationship between fossil fuels prices and industrial output but intensified its magnitude and persistence, making downside risks more entrenched in the post-2020 environment.

A similar pattern emerges for the Energy sector itself. Figure 40 compares the estimated coefficients for the 95th quantile of the sector's industrial production between the two samples. The results show that both Gas Price Change and NGPI reduce the sector's upper-tail growth potential in both periods, but the magnitude of this effect increases substantially in the post-2020 sample. In the pre-COVID subsample, the negative coefficients appear only marginally and dissipate by $h = 8$, whereas in the full sample they become larger—reaching roughly -0.15 by $h = 12$—and remain significant over the entire horizon. This indicates that recent crises have not altered the sign of the relationship but have amplified the degree to which gas price shocks constrain the Energy sector's capacity for high growth. As the Energy sector provides essential inputs to other industries, this heightened sensitivity may have indirect consequences for overall industrial performance in the European economy.

4.2 Country Analysis

This section examines country-specific results for Germany, Italy, and France, comparing how energy price shocks transmit to industrial production at both the aggregate and sectoral levels. While each economy displays distinctive structural characteristics, the general patterns observed at the EU level—strong downside risks from oil prices and constrained upside potential from gas price shocks—are largely confirmed within national data.

Figure 13: Coefficients of Brent Prices on the total aggregate IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.



For Germany, Figures 13 and 45 illustrate the estimated coefficients and quantile regression fits for the effects of Brent Price Changes on the 5th quantile of aggregate industrial production growth. The results closely mirror the aggregate EU findings: increases in oil prices exert a negative and

persistent effect on the lower tail of the German IPI distribution, indicating a heightened probability of adverse growth outcomes. The coefficients become significantly negative at medium to long horizons, suggesting that oil price shocks propagate gradually through cost channels, amplifying downside risks to industrial activity.

Turning to the Energy sector, Figures 14 and 44 display the estimated coefficients and quantile regression fits for the effects of the Net Gas Price Increase (NGPI) on the 95th quantile of sectoral IPI Change. The results indicate that NGPI significantly reduces the upper-tail growth potential of the German Energy sector. Coefficients for the right tail become negative and statistically significant over the forecast horizon, reaching magnitudes of approximately -0.2 by $h = 12$. This suggests that large gas price spikes not only depress aggregate industrial growth through input-cost pressures but also limit the Energy sector's capacity to expand, reinforcing the asymmetric impact observed at the EU level.

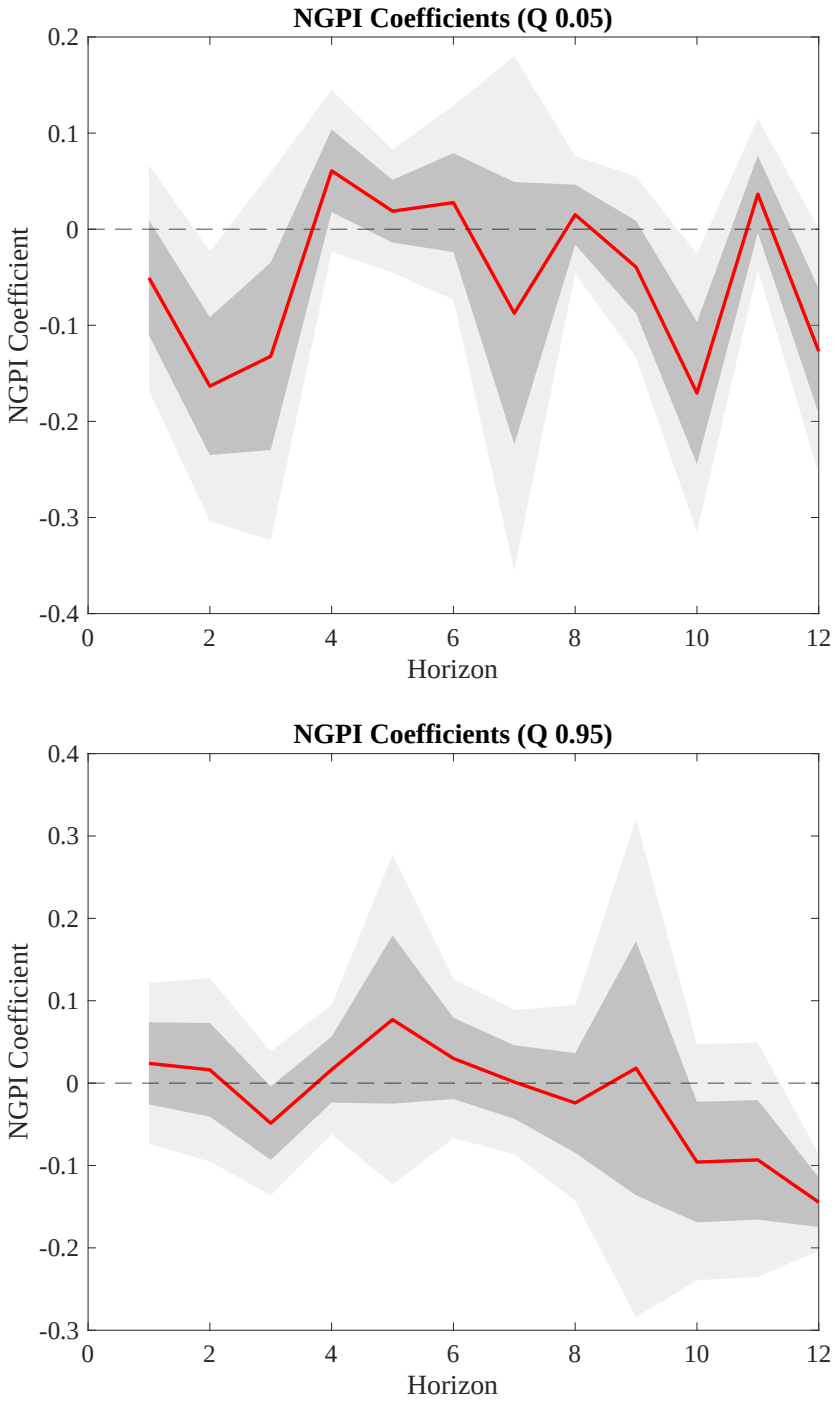
Italy exhibits a related but distinct pattern. Figures 57 show the estimated coefficients for the effects of Brent Price Changes on the 5th quantile of aggregate IPI Change. Unlike in Germany, oil price fluctuations do not exert a statistically significant effect on Italian industrial production. The coefficients remain close to zero across all horizons, suggesting that oil price shocks are not a key driver of downside risk for Italy's industrial sector. This relative insensitivity likely reflects the country's diversified energy mix and comparatively lower dependence on oil-intensive manufacturing processes.

In contrast, gas price dynamics play a more pronounced role in shaping Italy's industrial outcomes. As shown in Figures 54 and 15, the standard Gas Price Change variable (ΔGP_t) exerts a mild but statistically significant negative effect on the 95th quantile of both total and Energy sector IPI Change by $h = 12$. The corresponding scatter plots in Figure 50 confirm this pattern, showing a gradual compression of the upper tail of the conditional distribution. Rising gas prices therefore appear to dampen Italy's upside growth potential rather than exacerbate downside risk. Although the effect is moderate (around -0.1 at the one-year horizon), it points to a structural constraint on high-growth scenarios when energy costs rise persistently.

Figures 56 and 52 report the results for the NGPI variable. Consistent with the EU-level findings, NGPI negatively affects the energy sector's right tail, confirming that unexpected gas price surges constrain high-growth outcomes. The coefficients become significantly negative by $h = 12$, but the decline in the right tail is more gradual and less steep than in Germany, suggesting some resilience or slower transmission of gas shocks in Italy's industrial structure.

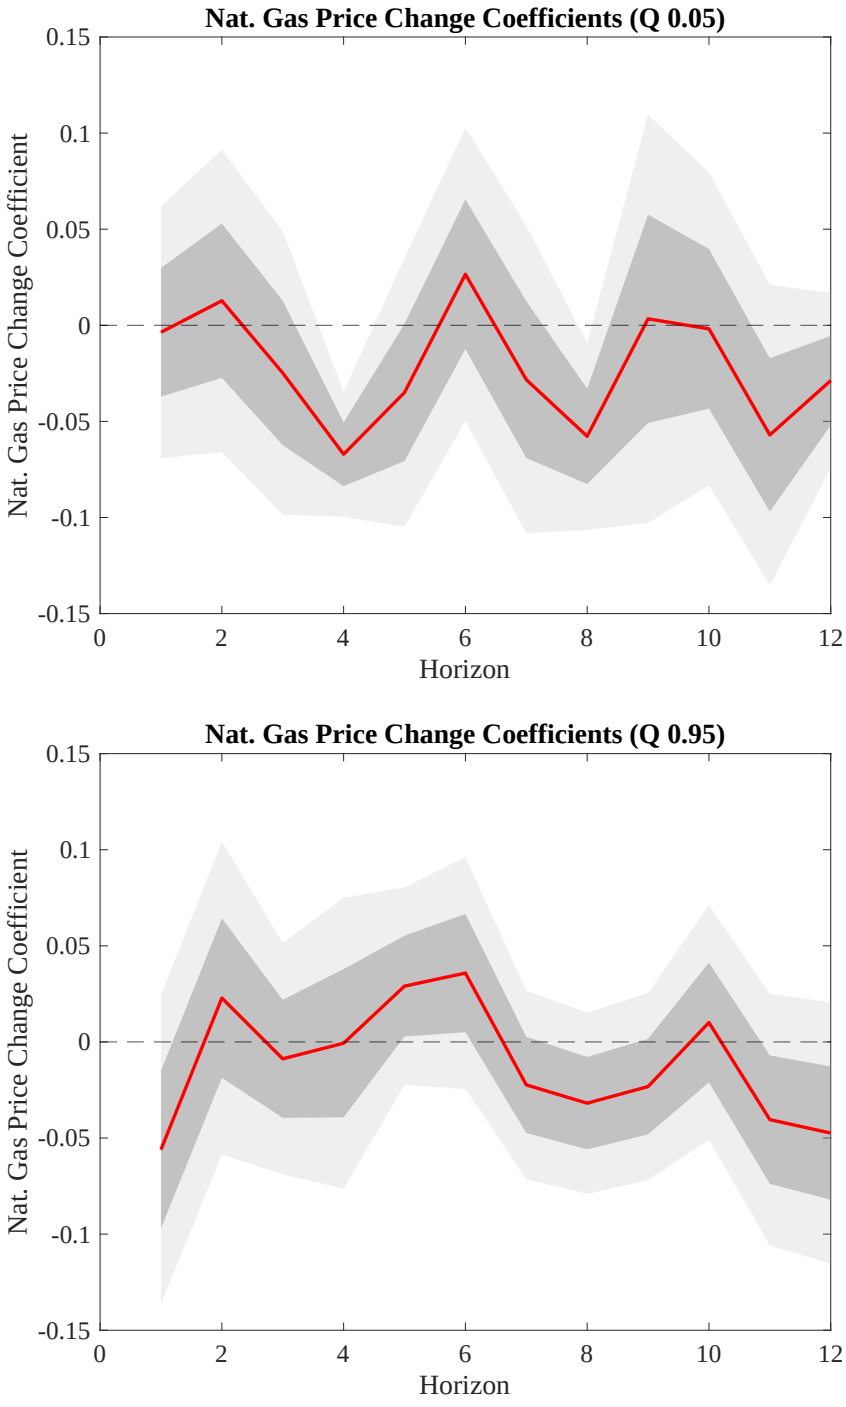
For France, the results do not indicate any statistically significant effects of energy price dynamics on either aggregate or sectoral industrial production. Neither Brent Price Changes nor gas-related

Figure 14: Coefficients of NGPI on the Energy sector's IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.



variables (Gas Price Change and NGPI) show robust coefficients across quantiles or horizons. This suggests that French industrial production has been comparatively less sensitive to fluctuations in international energy markets, likely reflecting its lower dependence on natural gas and greater reliance

Figure 15: Coefficients of Gas Price Change on the Energy sector's IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.



on nuclear and renewable energy sources. As a result, NGPI does not exhibit significant effects on the energy sector's right tail, with coefficients remaining close to zero across the forecast horizons.

Given the absence of significant relationships, detailed figures and estimations for France are re-

ported in the Appendix for completeness.

5 Conclusions

This paper investigated the effects of gas and oil price dynamics on European industrial production through a Growth-at-Risk (GaR) framework. By focusing on the conditional distribution of industrial output, rather than its mean, the analysis captured the asymmetric and nonlinear nature of energy price shocks across aggregate, sectoral, and national dimensions.

The results indicate that energy price changes—particularly those related to natural gas—have predominantly negative effects on industrial production in the European Union. The impacts are asymmetric, concentrating on the left tail of the aggregate industrial production distribution and the right tail of the Energy sector distribution. In other words, energy price increases simultaneously elevate the risk of sharp industrial downturns and reduce the potential for high-growth outcomes within the energy-producing sector itself. This pattern is consistent with Hamilton's seminal findings Hamilton, 2003, 2011, which emphasize that large and sustained energy price spikes depress output growth primarily through nonlinear cost and supply channels.

At the aggregate level, the Net Gas Price Increase (NGPI) emerges as a particularly powerful driver of downside risk, with its effects intensifying and persisting over the forecast horizon. Ordinary gas price fluctuations (ΔGP_t) show limited significance, suggesting that it is the occurrence of exceptional and sustained gas price surges—rather than routine volatility—that underpins industrial vulnerability. Similarly, oil price changes exert negative effects on the lower tail of industrial production, although their magnitude and persistence are generally weaker than those of gas shocks.

Sectoral disaggregation reveals that the Energy sector plays a pivotal role in transmitting these dynamics. Both Gas Price Change and NGPI reduce the upper-tail growth potential of the Energy sector, confirming that energy price surges limit its capacity for long-term expansion. In contrast, consumer and capital goods sectors are largely unaffected, likely due to differences in energy intensity, the ability to pass through costs, and longer production cycles that smooth the impact of short-term price fluctuations.

The cross-country analysis further highlights structural asymmetries within the euro area. In Germany and Italy, NGPI significantly reduces the high-growth potential of the Energy sector, with coefficients on the right tail becoming notably negative by $h = 12$. These effects are stronger and more persistent in Germany, reflecting its heavy reliance on industrial gas use, while Italy shows a milder but still discernible contraction of high-growth outcomes. France, by contrast, exhibits negligible sensitivity to gas price shocks, a result likely attributable to its more diversified and nuclear-intensive

energy system, which buffers industrial output from international energy market volatility.

Finally, the comparison between pre- and post-2020 samples suggests that the COVID-19 pandemic and the Russian invasion of Ukraine have amplified these vulnerabilities. While the asymmetric effects of energy price shocks were already present before 2020, they have since become stronger and more persistent, indicating an enduring shift in the energy–industry nexus toward greater fragility and prolonged downside risk.

References

Abadir, K. M., Luati, A., & Paruolo, P. (2023). GARCH density and functional forecasts. *Journal of Econometrics*, 235(2), 470–483. https://doi.org/10.1016/j.jeconom.2022.04.010

Adams, P. A., Adrian, T., Boyarchenko, N., & Giannone, D. (2021). Forecasting macroeconomic risks. *International Journal of Forecasting*, 37(3), 1173–1191. https://doi.org/10.1016/j.ijforecast.2021.01.003

Adrian, T., Boyarchenko, N., & Giannone, D. (2019). Vulnerable Growth. *American Economic Review*, 109(4), 1263–1289. https://doi.org/10.1257/aer.20161923

Adrian, T., Boyarchenko, N., & Giannone, D. (2021). MULTIMODALITY IN MACROFINANCIAL DYNAMICS. *International Economic Review*, 62(2), 861–886. https://doi.org/10.1111/iere.12501

Adrian, T., Boyarchenko, N., Giannone, D., Prasad, A., Seneviratne, D., & Xiao, Y. (2022). 800,000 Years of Climate Risk. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4213744

Adrian, T., Grinberg, F., Liang, N., Malik, S., & Yu, J. (2022). The Term Structure of Growth-at-Risk. *American Economic Journal: Macroeconomics*, 14(3), 283–323. https://doi.org/10.1257/mac.20180428

Adrian, T., Sasaki, Y., & Wang, Y. (2025). Robust Econometrics for Growth-at-Risk [Version Number: 1]. https://doi.org/10.48550/ARXIV.2508.00263

Apergis, N. (2023). Forecasting energy prices: Quantile-based risk models. *Journal of Forecasting*, 42(1), 17–33. https://doi.org/10.1002/for.2898

Azzalini, A., & Capitanio, A. (2009). Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t distribution [Version Number: 1]. https://doi.org/10.48550/ARXIV.0911.2342

Bernardi, M., Casarin, R., Maillet, B. B., & Petrella, L. (2024). Bayesian dynamic quantile model averaging. *Annals of Operations Research*. https://doi.org/10.1007/s10479-024-06378-7

Bjørnland, H. C., Larsen, V. H., & Maih, J. (2018). Oil and Macroeconomic (In)Stability. *American Economic Journal: Macroeconomics*, 10(4), 128–151. https://doi.org/10.1257/mac.20150171

Blanchard, O., & Gali, J. (2007, September). *The Macroeconomic Effects of Oil Shocks: Why are the 2000s So Different from the 1970s?* (Tech. rep. No. w13368). National Bureau of Economic Research. Cambridge, MA. https://doi.org/10.3386/w13368

Brownlees, C., & Souza, A. B. (2021). Backtesting global Growth-at-Risk. *Journal of Monetary Economics*, 118, 312–330. https://doi.org/10.1016/j.jmoneco.2020.11.003

Candila, V., Gallo, G. M., & Petrella, L. (2023). Mixed-frequency quantile regressions to forecast value-at-risk and expected shortfall. *Annals of Operations Research*. https://doi.org/10.1007/s10479-023-05370-x

Catania, L., & Luati, A. (2019). Semiparametric Modeling of Multiple Quantiles. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3494995

Chavleishvili, S., & Manganelli, S. (2024). Forecasting and stress testing with quantile vector autoregression. *Journal of Applied Econometrics*, 39(1), 66–85. https://doi.org/10.1002/jae.3009

Chen, L., Dolado, J. J., & Gonzalo, J. (2021). Quantile Factor Models. *Econometrica*, 89(2), 875–910. https://doi.org/10.3982/ECTA15746

Chuliá, H., Klein, T., Muñoz Mendoza, J. A., & Uribe, J. M. (2024). Vulnerability of European electricity markets: A quantile connectedness approach. *Energy Policy*, 184, 113862. https://doi.org/10.1016/j.enpol.2023.113862

Clark, T. E., Huber, F., Koop, G., Marcellino, M., & Pfarrhofer, M. (2024). Investigating Growth-at-Risk Using a Multicountry Nonparametric Quantile Factor Model. *Journal of Business & Economic Statistics*, 42(4), 1302–1317. https://doi.org/10.1080/07350015.2024.2310020

Corradi, V., Fosten, J., & Gutknecht, D. (2024). Reprint of: Out-of-sample tests for conditional quantile coverage: An application to Growth-at-Risk. *Journal of Econometrics*, 244(2), 105746. https://doi.org/10.1016/j.jeconom.2024.105746

ECB. (2023). *Economic Bulletin Issue 1, 2023* (tech. rep.).

Figueres, J. M., & Jarociński, M. (2020). Vulnerable growth in the euro area: Measuring the financial conditions. *Economics Letters*, 191, 109126. https://doi.org/10.1016/j.econlet.2020.109126

Foroni, C., Ravazzolo, F., & Rossini, L. (2023). Are low frequency macroeconomic variables important for high frequency electricity prices? *Economic Modelling*, 120, 106160. https://doi.org/10.1016/j.econmod.2022.106160

Gaglianone, W. P., & Lima, L. R. (2012). Constructing Density Forecasts from Quantile Regressions. *Journal of Money, Credit and Banking*, 44(8), 1589–1607. https://doi.org/10.1111/j.1538-4616.2012.00545.x

Giot, P., & Laurent, S. (2003). Market risk in commodity markets: A VaR approach. *Energy Economics*, 25(5), 435–457. https://doi.org/10.1016/S0140-9883(03)00052-5

Hamilton, J. D. (2003). What is an oil shock? *Journal of Econometrics*, 113(2), 363–398. https://doi.org/10.1016/S0304-4076(02)00207-5

Hamilton, J. D. (2011). NONLINEARITIES AND THE MACROECONOMIC EFFECTS OF OIL PRICES. *Macroeconomic Dynamics*, 15(S3), 364–378. https://doi.org/10.1017/S1365100511000307

Herrera, A. M., & Karaki, M. B. (2015). The effects of oil price shocks on job reallocation. *Journal of Economic Dynamics and Control*, 61, 95–113. https://doi.org/10.1016/j.jedc.2015.08.006

Herrera, A. M., Lagalo, L. G., & Wada, T. (2011). OIL PRICE SHOCKS AND INDUSTRIAL PRODUCTION: IS THE RELATIONSHIP LINEAR? *Macroeconomic Dynamics*, 15(S3), 472–497. https://doi.org/10.1017/S1365100511000290

Iacopini, M., Poon, A., Rossini, L., & Zhu, D. (2023a). Money Growth and Inflation: A Quantile Sensitivity Approach [Version Number: 3]. https://doi.org/10.48550/ARXIV.2308.05486

Iacopini, M., Poon, A., Rossini, L., & Zhu, D. (2023b). Bayesian mixed-frequency quantile vector autoregression: Eliciting tail risks of monthly US GDP. *Journal of Economic Dynamics and Control*, 157, 104757. https://doi.org/10.1016/j.jedc.2023.104757

Iacopini, M., Ravazzolo, F., & Rossini, L. (2022). Bayesian Multivariate Quantile Regression with alternative Time-varying Volatility Specifications [Version Number: 2]. https://doi.org/10.48550/ARXIV.2211.16121

IEA. (2022a). *Gas Market Report Q1 2022* (tech. rep.).

IEA. (2022b). *World Energy Outlook 2023* (tech. rep.). IEA.

Koenker, R., & Bassett, G. (1978). Regression Quantiles. *Econometrica*, 46(1), 33. https://doi.org/10.2307/1913643

Korobilis, D., & Schröder, M. (2025a). Monitoring multi-country macroeconomic risk: A quantile factor-augmented vector autoregressive (QFAVAR) approach. *Journal of Econometrics*, 249, 105730. https://doi.org/10.1016/j.jeconom.2024.105730

Korobilis, D., & Schröder, M. (2025b). Probabilistic Quantile Factor Analysis. *Journal of Business & Economic Statistics*, 43(3), 530–543. https://doi.org/10.1080/07350015.2024.2396956

Lang, J. H., Rusnák, M., & Greiwe, M. (2025). Medium-Term Growth-at-Risk in the Euro Area. *IMF Economic Review*. https://doi.org/10.1057/s41308-024-00272-9

Laporta, A. G., Merlo, L., & Petrella, L. (2018). Selection of Value at Risk Models for Energy Commodities. *Energy Economics*, 74, 628–643. https://doi.org/10.1016/j.eneco.2018.07.009

Lee, K., & Ni, S. (2002). On the dynamic effects of oil price shocks: A study using industry level data. *Journal of Monetary Economics*, 49(4), 823–852. https://doi.org/10.1016/S0304-3932(02)00114-9

Lenza, M., Moutachaker, I., & Paredes, J. (2025). Density forecasts of inflation: A quantile regression forest approach. *European Economic Review*, 178, 105079. https://doi.org/10.1016/j.euroecorev.2025.105079

López-Salido, D., & Loria, F. (2024). Inflation at risk. *Journal of Monetary Economics*, 145, 103570. https://doi.org/10.1016/j.jmoneco.2024.103570

Merlo, L., Maruotti, A., Petrella, L., & Punzo, A. (2022). Quantile hidden semi-Markov models for multivariate time series. *Statistics and Computing*, 32(4), 61. https://doi.org/10.1007/s11222-022-10130-1

Merlo, L., Petrella, L., & Raponi, V. (2021). Forecasting VaR and ES using a joint quantile regression and its implications in portfolio allocation. *Journal of Banking & Finance*, 133, 106248. https://doi.org/10.1016/j.jbankfin.2021.106248

Mitchell, J., Poon, A., & Zhu, D. (2024). Constructing density forecasts from quantile regressions: Multimodality in macrofinancial dynamics. *Journal of Applied Econometrics*, 39(5), 790–812. https://doi.org/10.1002/jae.3049

Ren, X., Duan, K., Tao, L., Shi, Y., & Yan, C. (2022). Carbon prices forecasting in quantiles. *Energy Economics*, 108, 105862. https://doi.org/10.1016/j.eneco.2022.105862

Uribe, J. M., Mosquera-López, S., & Arenas, O. J. (2022). Assessing the relationship between electricity and natural gas prices in European markets in times of distress. *Energy Policy*, 166, 113018. https://doi.org/10.1016/j.enpol.2022.113018

Vasseur, S. P., & Aznarte, J. L. (2021). Comparing quantile regression methods for probabilistic forecasting of NO2 pollution levels. *Scientific Reports*, 11(1), 11592. https://doi.org/10.1038/s41598-021-90063-3

Yu, K., & Moyeed, R. A. (2001). Bayesian quantile regression. *Statistics & Probability Letters*, 54(4), 437–447. https://doi.org/10.1016/S0167-7152(01)00124-9

Appendix

Numerical Solution of the Skewed-t Fitting Problem

The minimization problem in Equation **??** constitutes a nonlinear least squares problem, since the skewed-t quantile function $F_{ST}^{-1}(\tau_i;\theta_t)$ is a nonlinear function of the parameter vector θ_t with no closed-form expression. The problem is solved numerically using a Levenberg-Marquardt algorithm, as implemented in MATLAB's `lsqnonlin` routine.

Define the residual vector $r(\theta_t) \in \mathbb{R}^{N_\tau}$ with components:

$$r_i(\theta_t) = Q^{\mathrm{QR}}_{\Delta IPI_{t+h}}(\tau_i|X_t) - F_{ST}^{-1}(\tau_i;\theta_t), \quad i = 1,\ldots,N_\tau. \tag{12}$$

The algorithm proceeds iteratively. At each step k, the residual vector is linearized around the current estimate $\theta_t^{(k)}$ using its Jacobian $J(\theta_t^{(k)}) = \partial r/\partial\theta_t$, an $N_\tau \times 4$ matrix approximated via finite differences. The parameter update is then given by:

$$\theta_t^{(k+1)} = \theta_t^{(k)} - \left(J^T J + \lambda^{(k)} I\right)^{-1} J^T r(\theta_t^{(k)}), \tag{13}$$

where $\lambda^{(k)} > 0$ is an adaptive damping parameter. When $\lambda^{(k)}$ is large, the step approximates a gradient descent direction, ensuring stability far from the optimum; as $\lambda^{(k)}$ decreases, the update converges to a Gauss-Newton step, which exploits the local curvature of the objective and accelerates convergence near the solution. The damping parameter is adjusted at each iteration based on the realized reduction in the objective function $\|r(\theta_t)\|^2$.

Bound constraints are imposed to ensure that the estimated parameters remain in the admissible region of the parameter space—specifically, $\sigma_t > 0$ and $\nu_t > 2$, the latter guaranteeing the existence of finite variance. When bound constraints are active, the algorithm switches to a trust-region-reflective method, which restricts each update step to a region consistent with the imposed bounds. Iteration continues until the change in the objective function or the parameter vector falls below a pre-specified numerical tolerance.

Figures

EU Aggregate

Figure 16: Scatter plot of IPI Change "h steps ahead" relative to current Nat. Gas Prices Change ($h = 1, 3, 6, 12$) for the total aggregate IPI Change. The lines represent the slopes of the quantile regression for different quantiles.

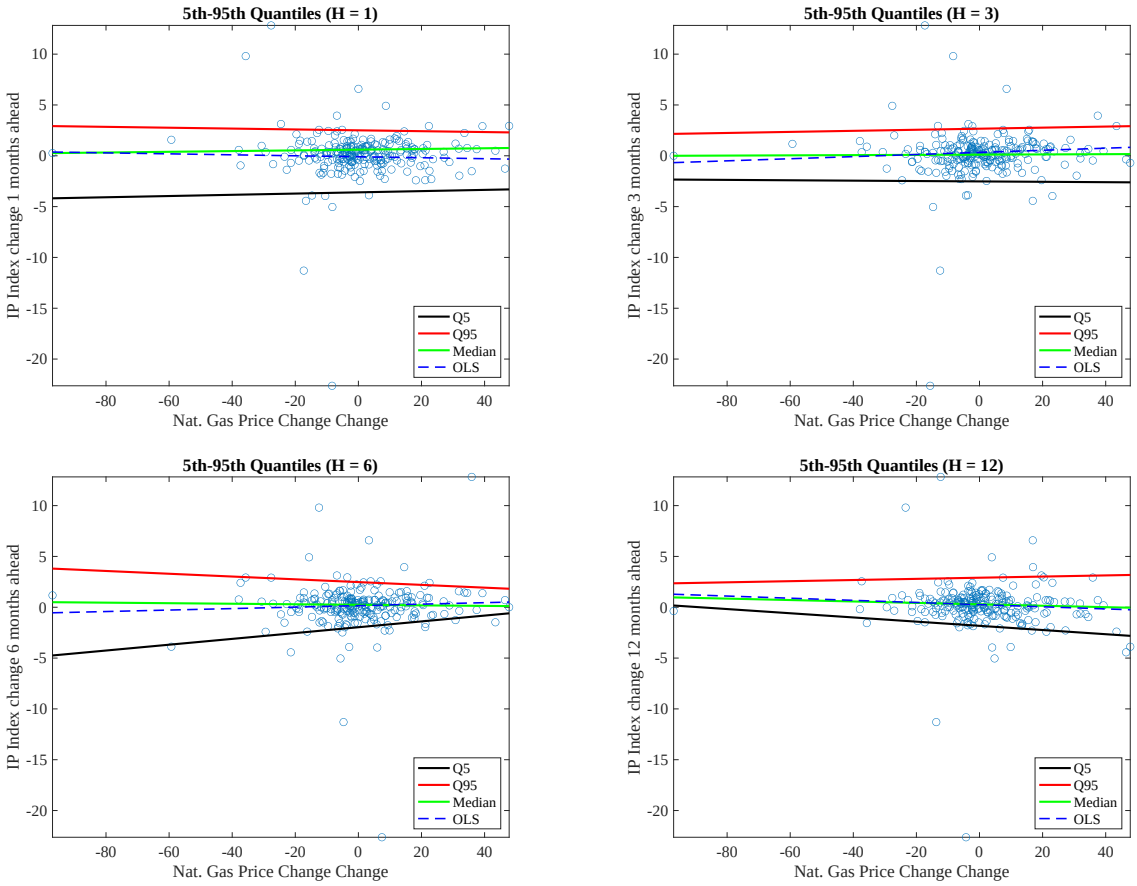


Figure 17: Scatter plot of IPI Change "h steps ahead" relative to current Nat. Gas Prices Change ($h = 1, 3, 6, 12$) for the IPI Change of the energy sector. The lines represent the slopes of the quantile regression for different quantiles.

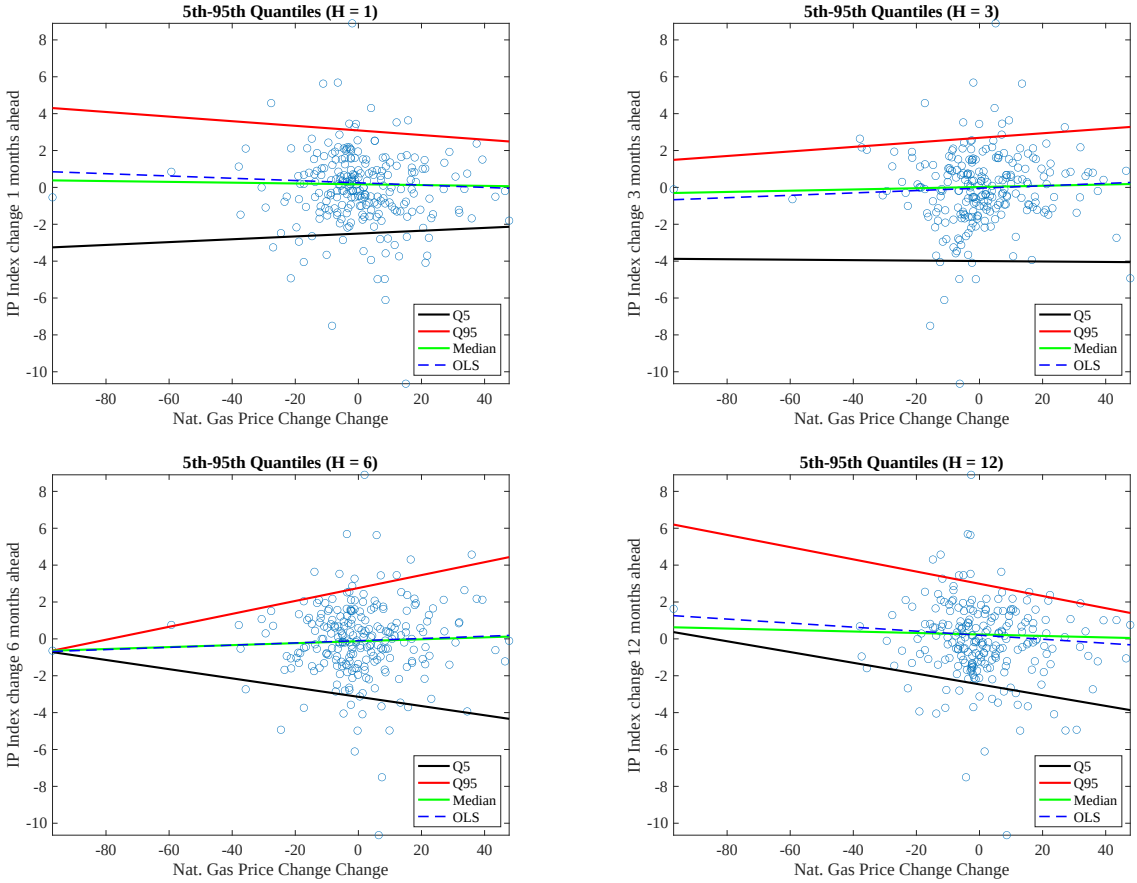


Figure 18: Scatter plot of IPI Change "h steps ahead" relative to current NGPI ($h = 1, 3, 6, 12$) for the total aggregate IPI Change. The lines represent the slopes of the quantile regression for different quantiles.

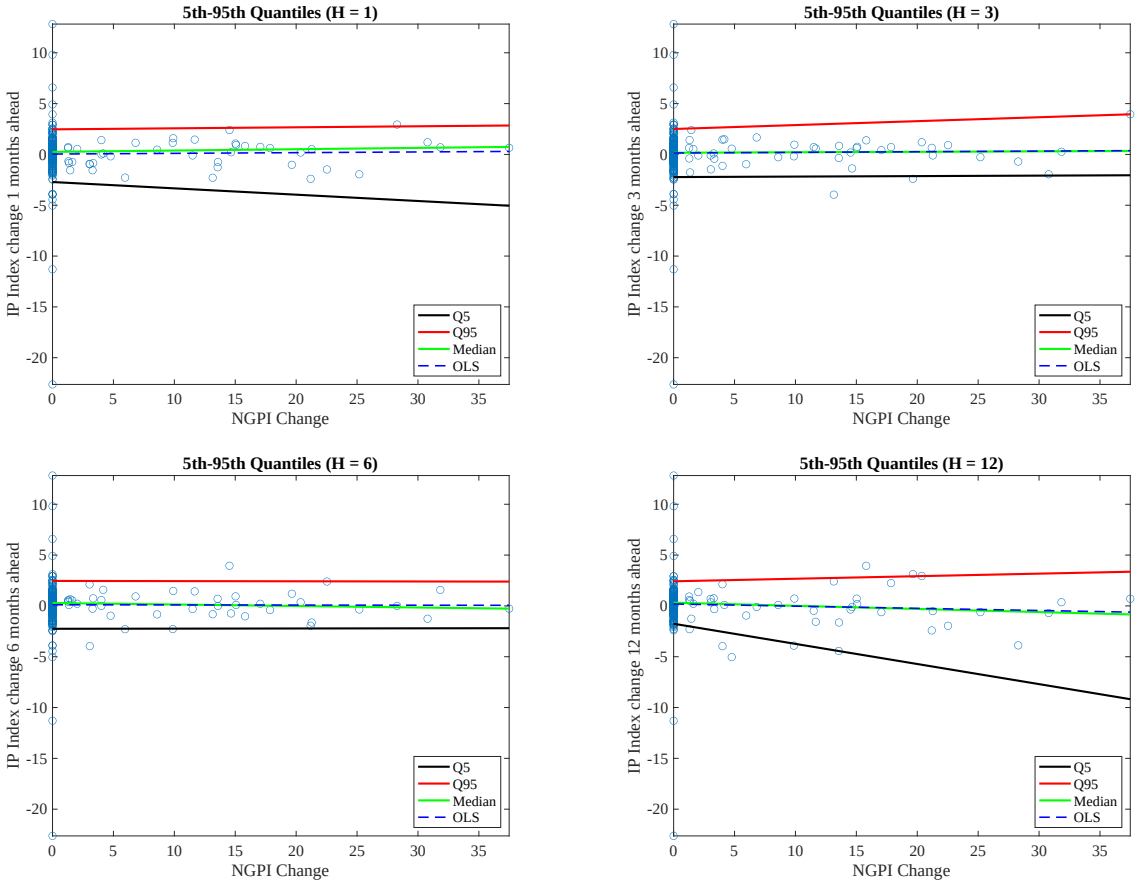


Figure 19: Scatter plot of IPI Change "h steps ahead" relative to current Oil Prices Change ($h = 1, 3, 6, 12$) for the total aggregate IPI Change. The lines represent the slopes of the quantile regression for different quantiles.

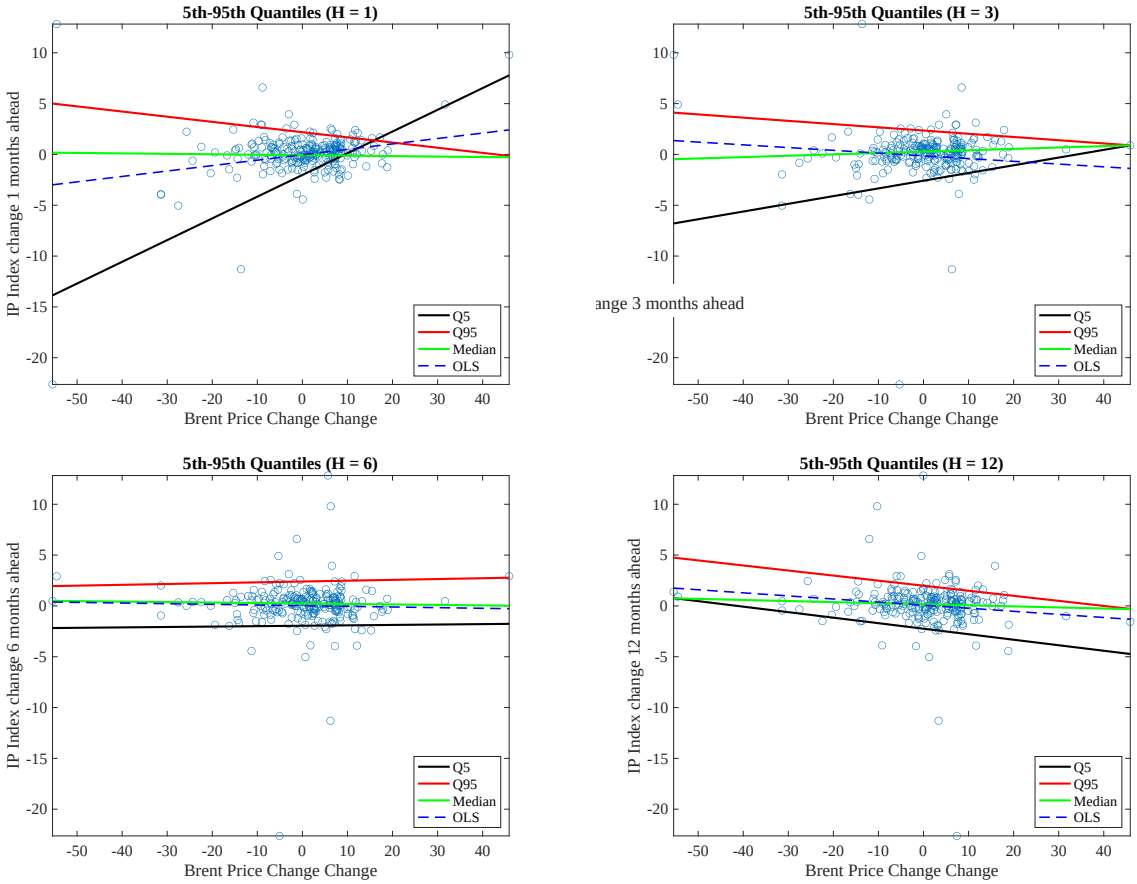


Figure 20: Scatter plot of IPI Change "h steps ahead" relative to current NOPI ($h = 1, 3, 6, 12$) for the total aggregate IPI Change. The lines represent the slopes of the quantile regression for different quantiles.

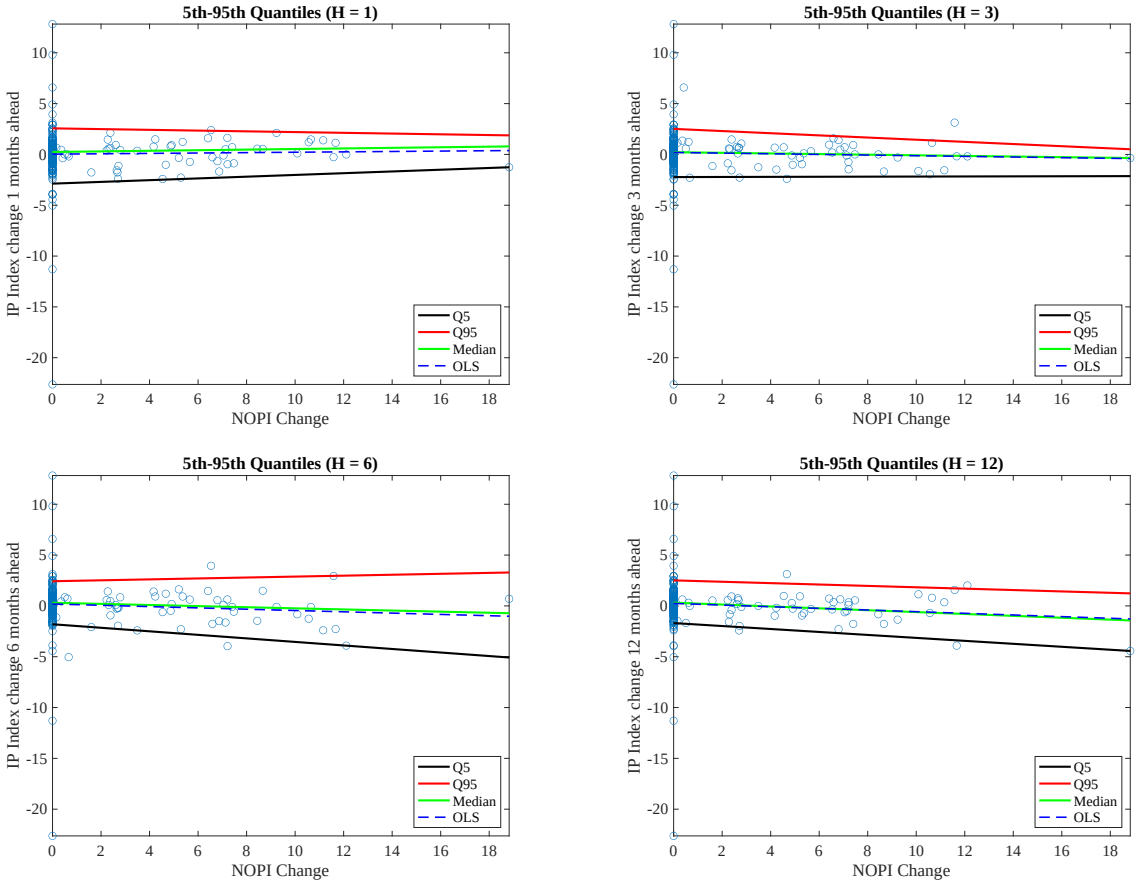


Figure 21: Coefficients of Gas Price Change on the total aggregate IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

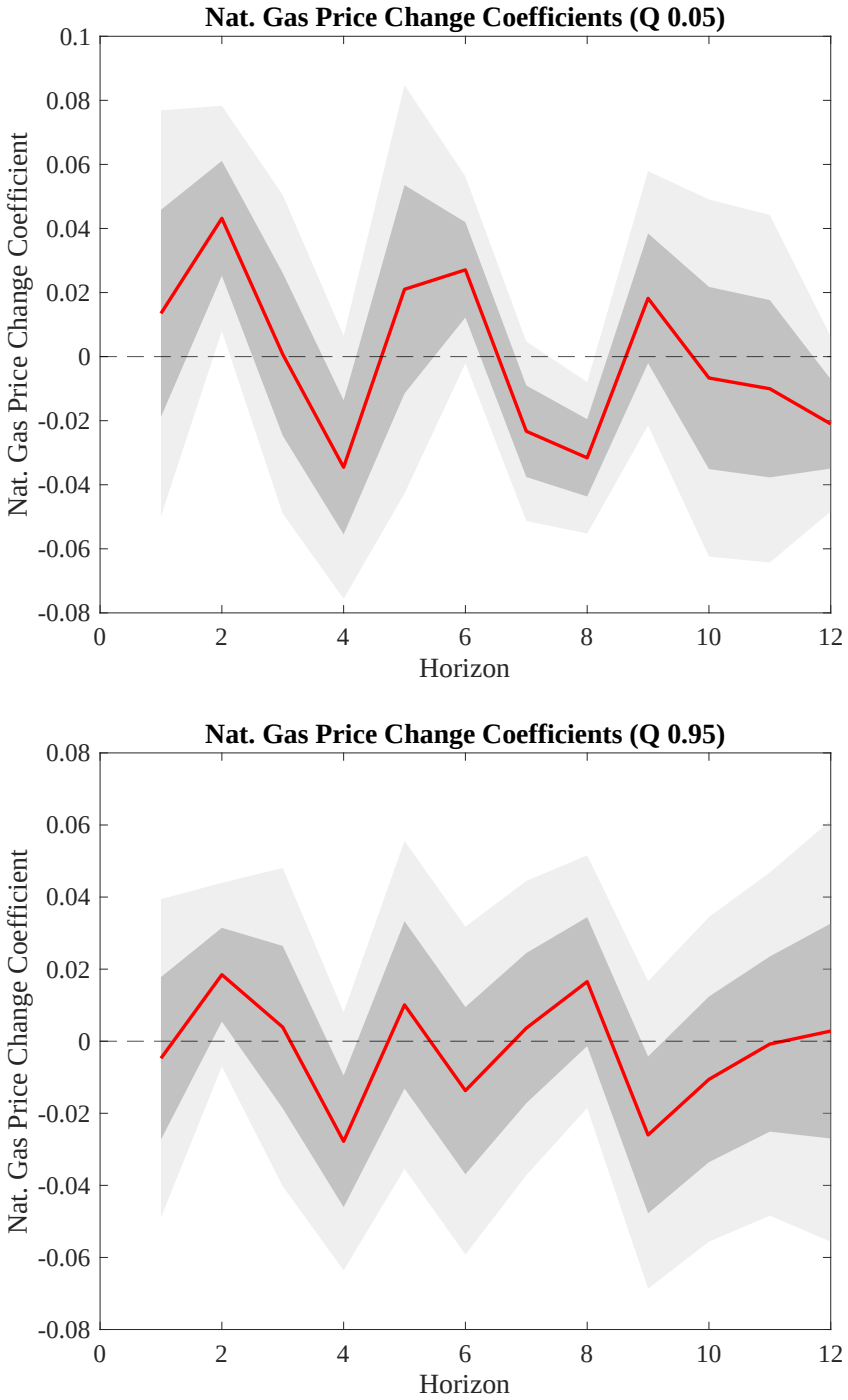


Figure 22: Coefficients of Gas Price Change on the total aggregate IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

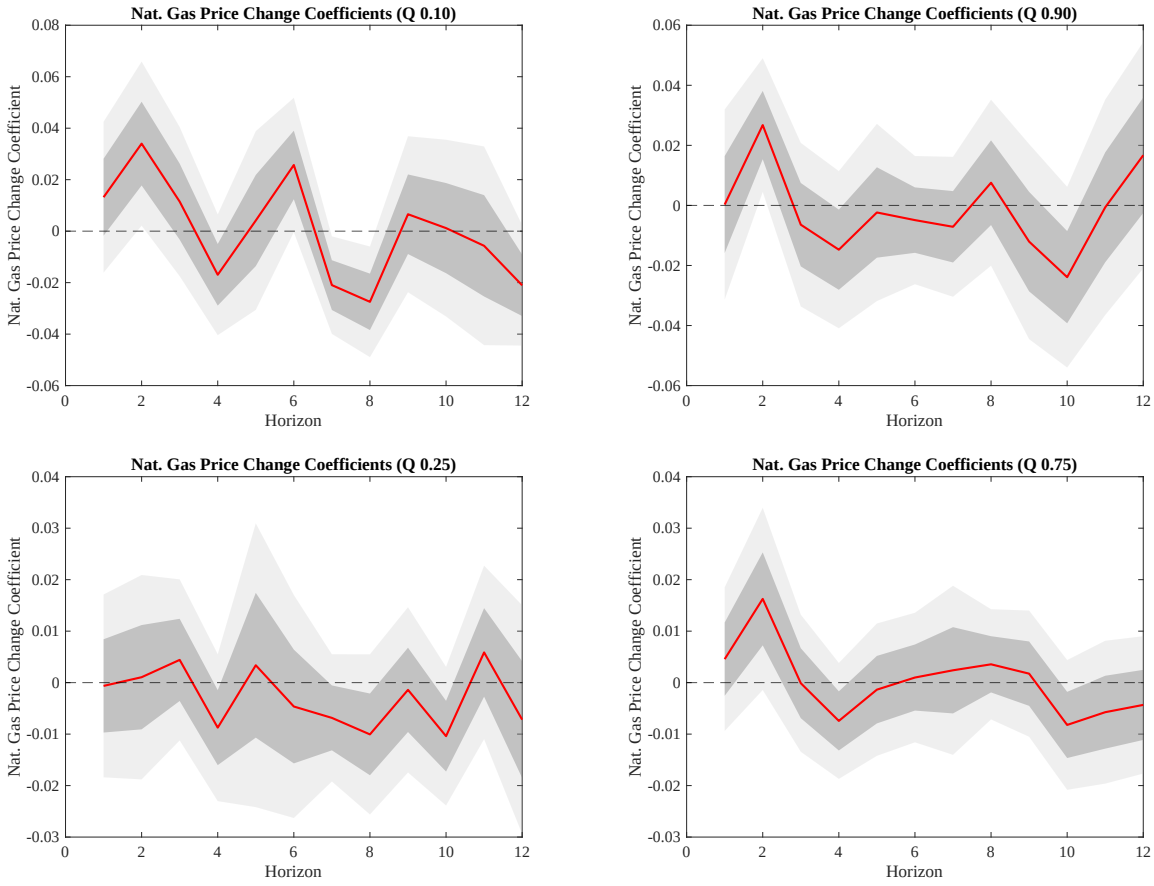


Figure 23: Coefficients of Gas Price Change on the Energy sector's IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

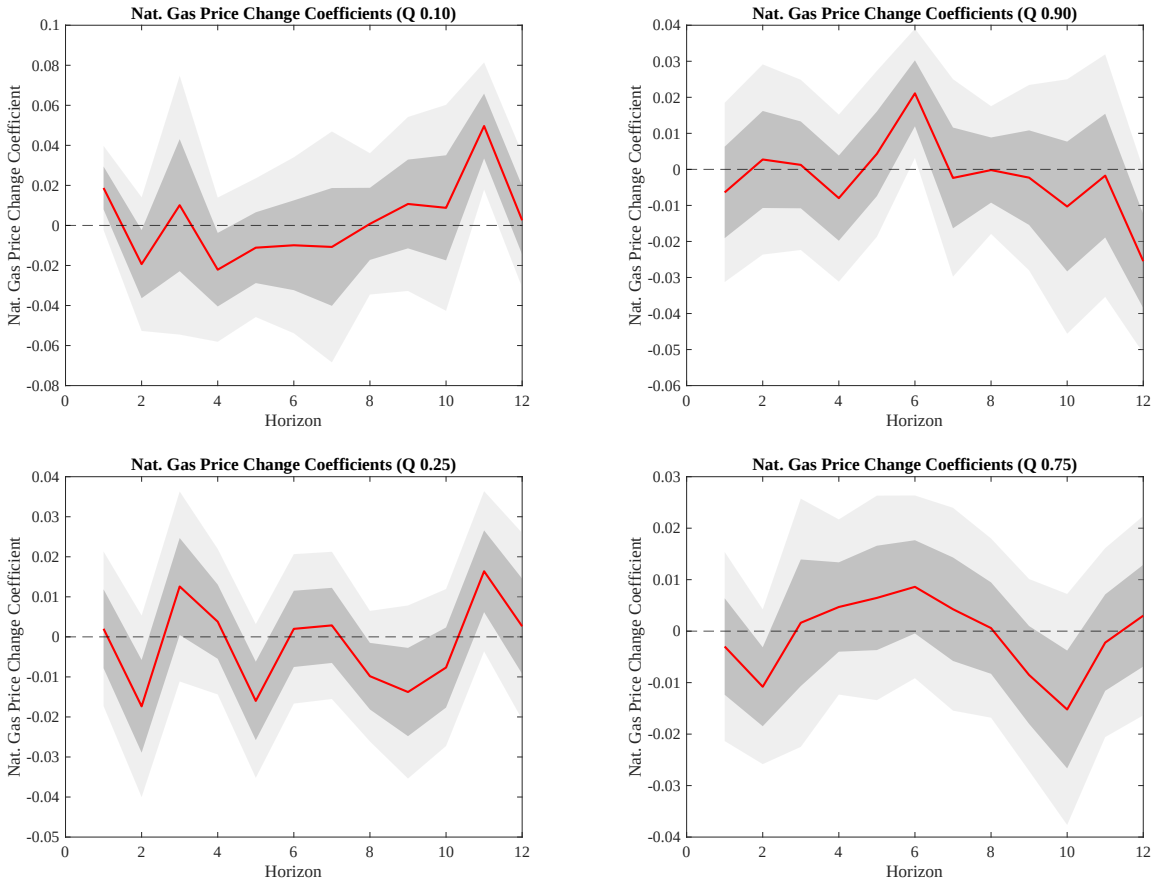


Figure 24: Coefficients of NGPI on the total aggregate IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

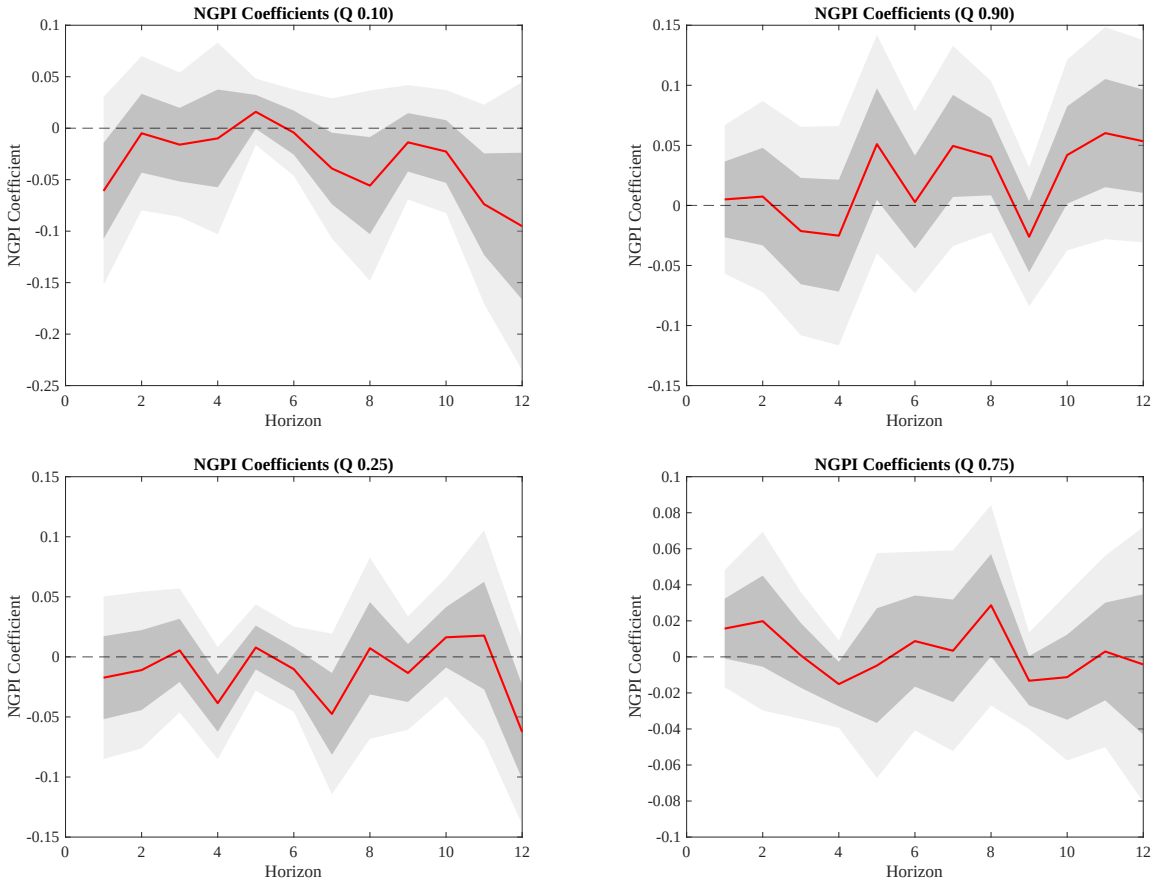


Figure 25: Coefficients of NGPI on the Energy sector's IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

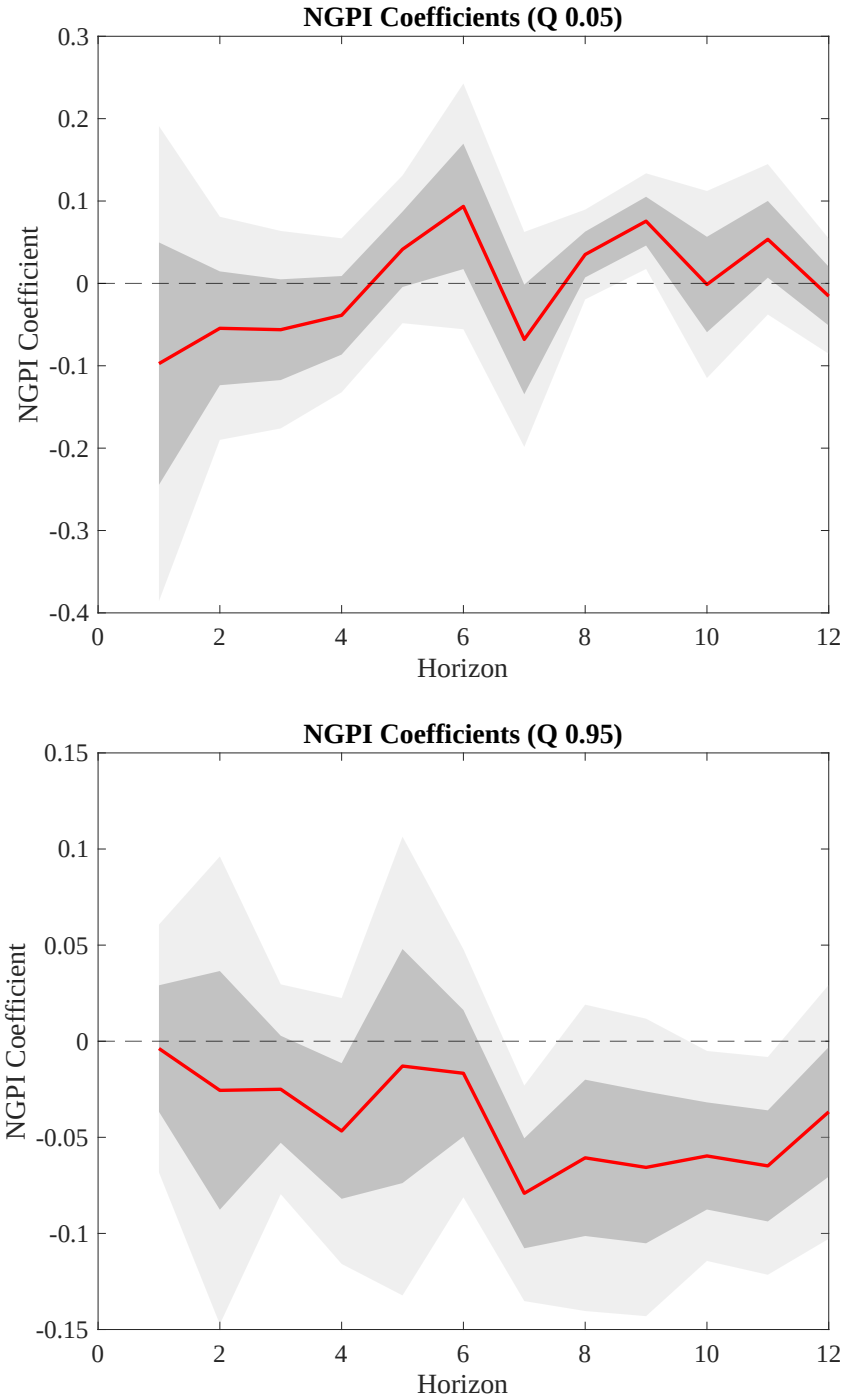


Figure 26: Coefficients of NGPI on the Energy sector's IPI Change across the horizons (Q10, Q25, Q75 and Q90). The bands represent the 95% and 68% confidence interval.

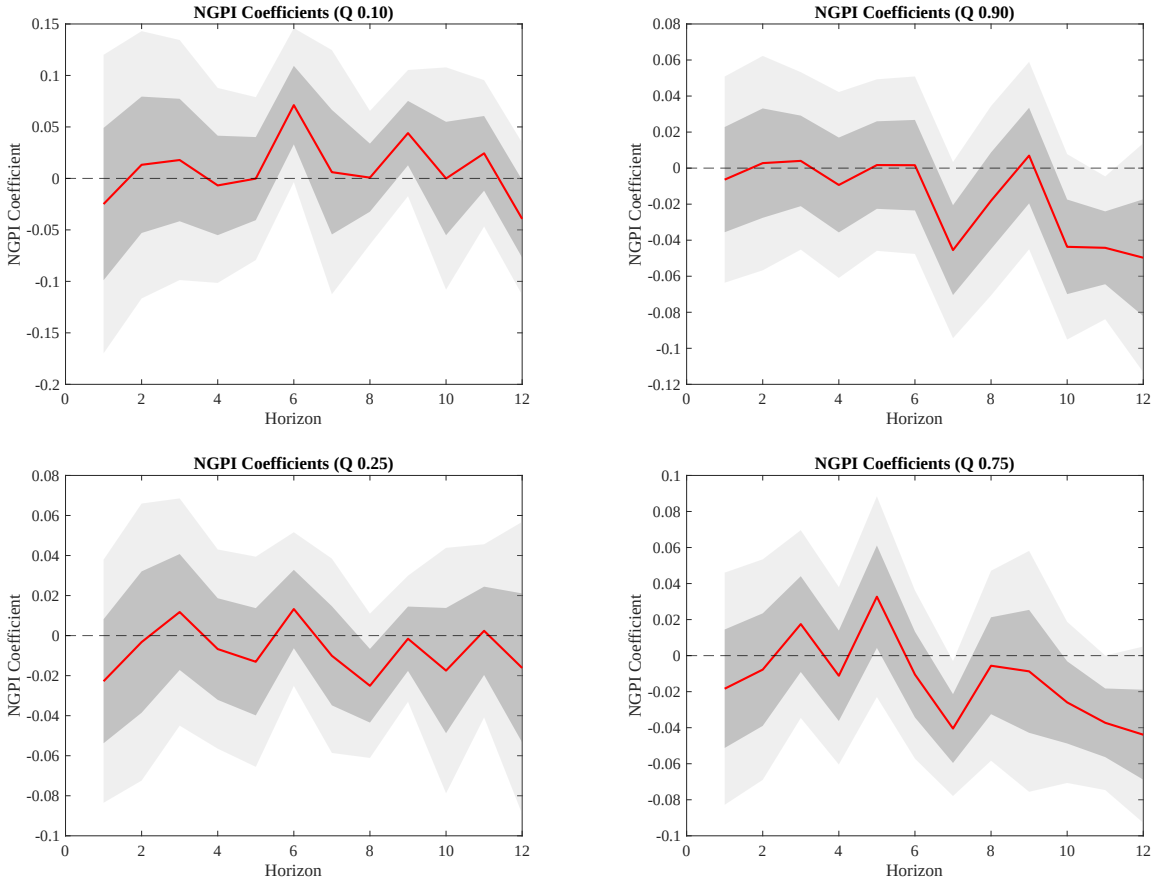


Figure 27: Coefficients of Brent Prices on the total aggregate IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

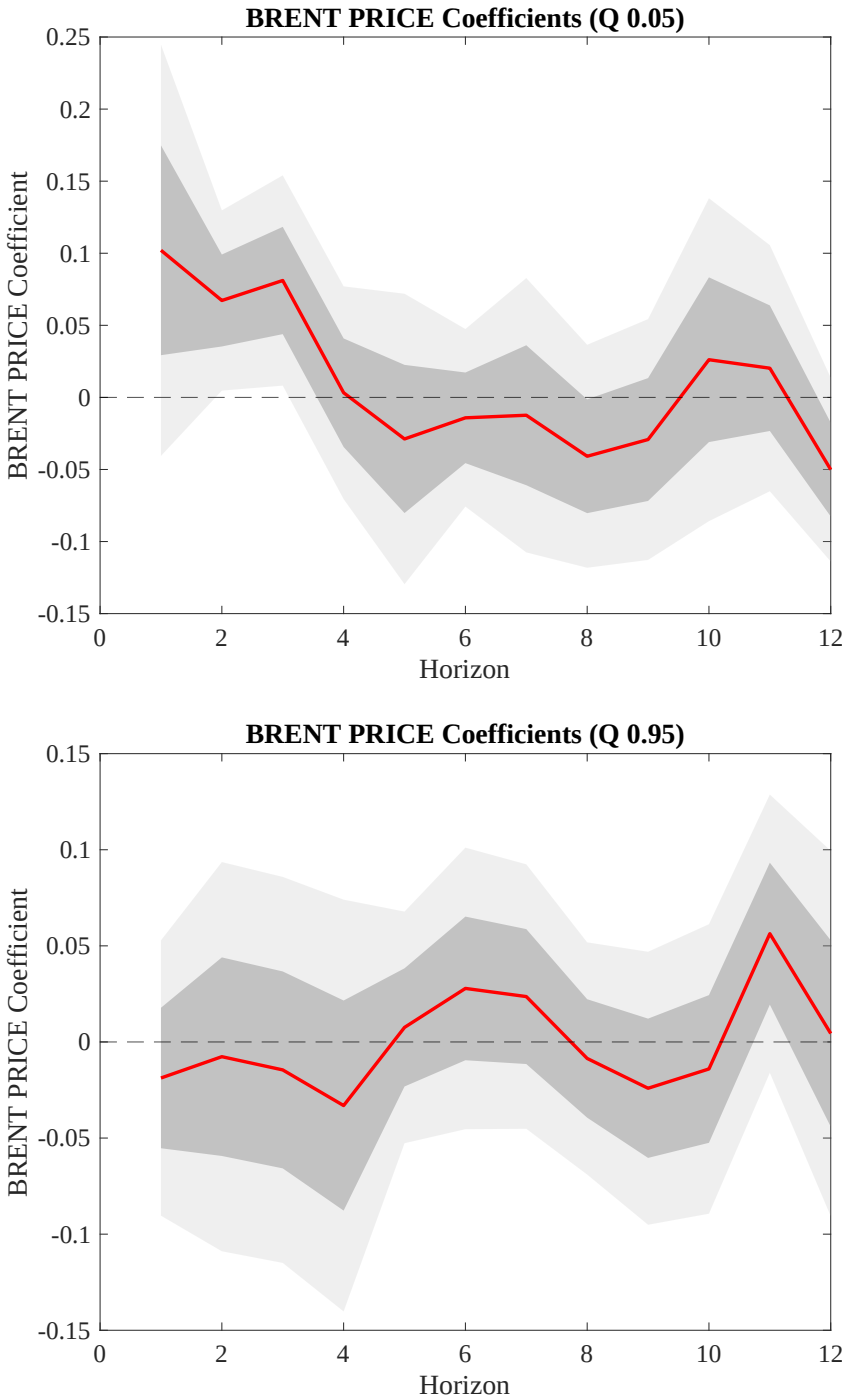


Figure 28: Coefficients of Brent Prices on the total aggregate IPI Change across the horizons (Q10, Q25, Q75 and Q90). The bands represent the 95% and 68% confidence interval.

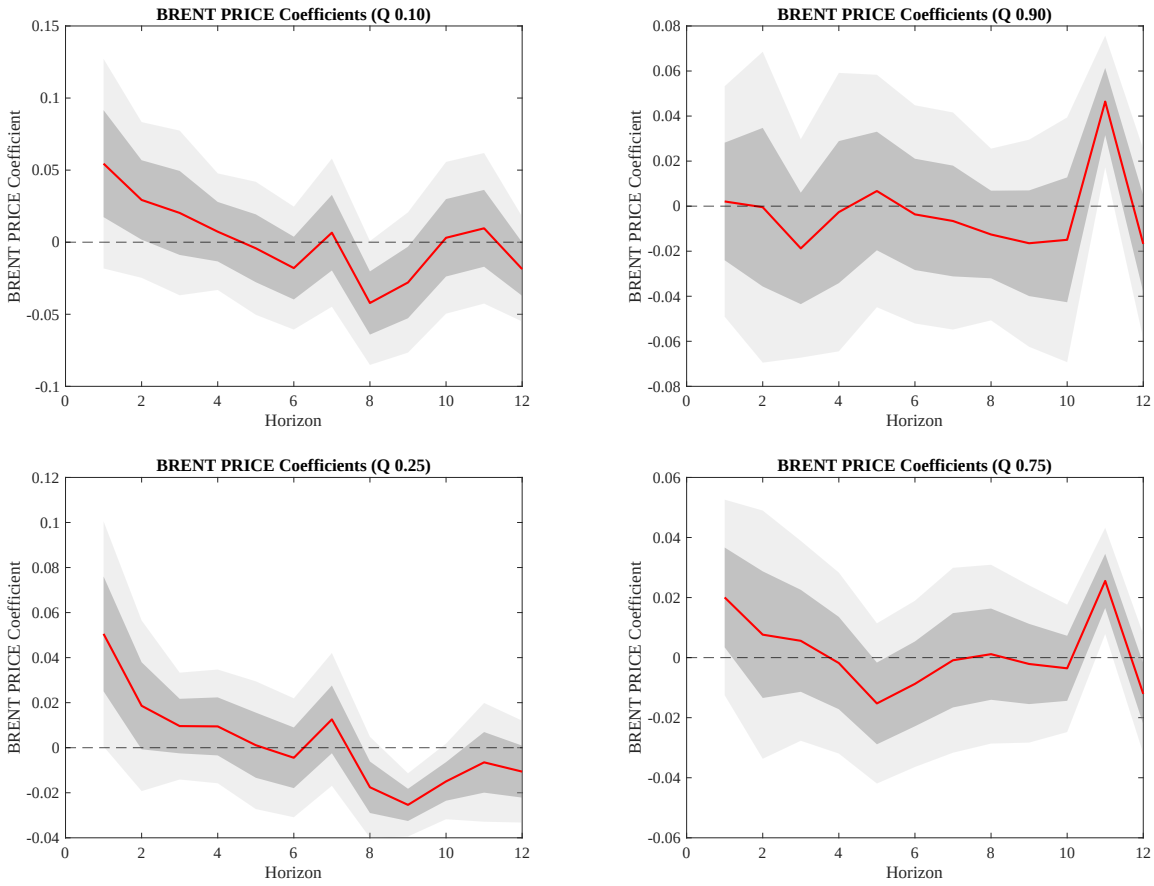


Figure 29: Coefficients of Brent Prices on the Energy sector's IPI Change across the horizons (Q10, Q25, Q75 and Q90). The bands represent the 95% and 68% confidence interval.

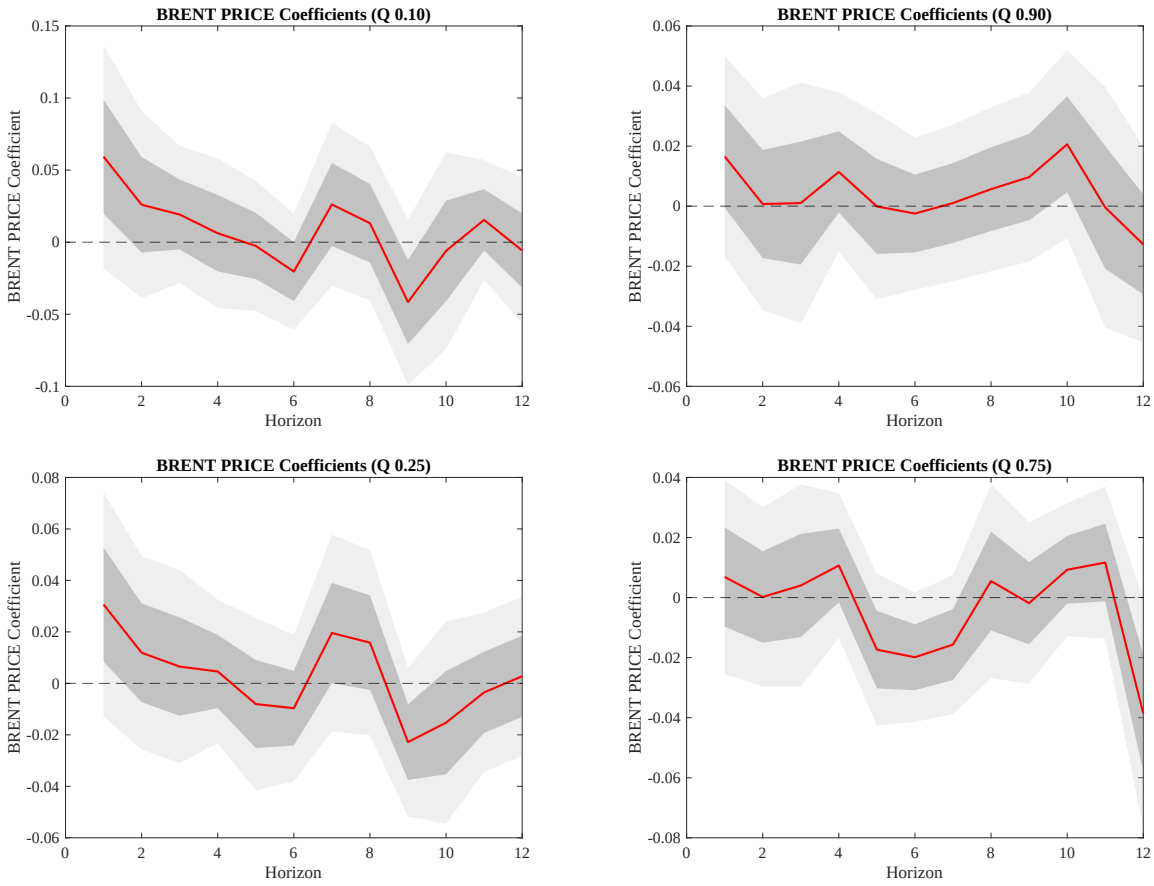
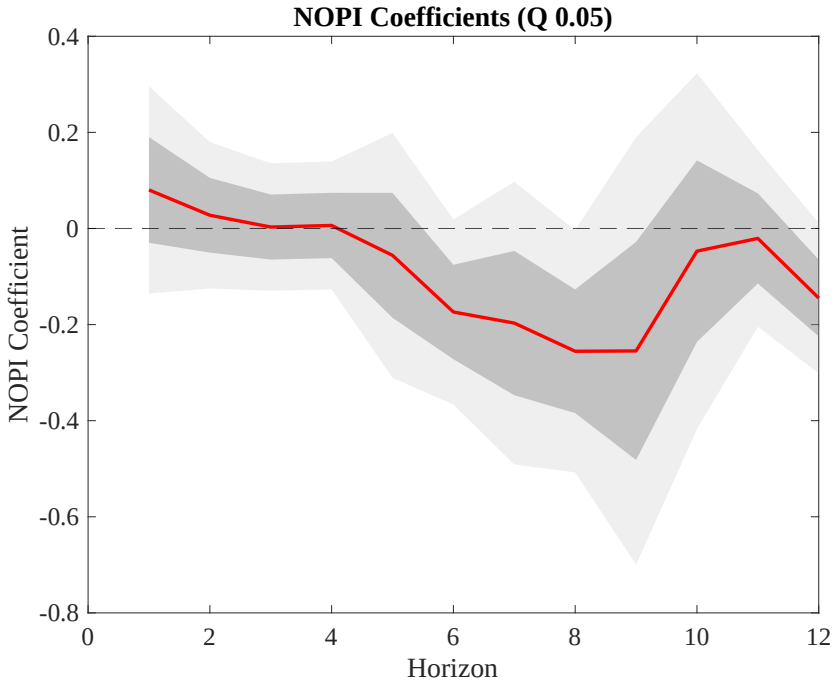


Figure 30: Coefficients of NOPI on the total aggregate IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.



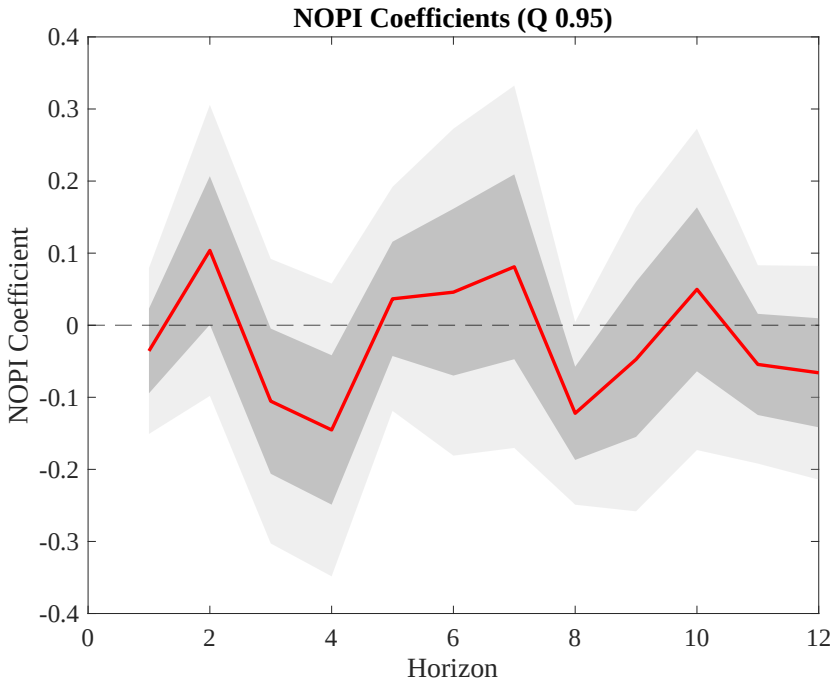


Figure 31: Coefficients of NOPI on the total aggregate IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

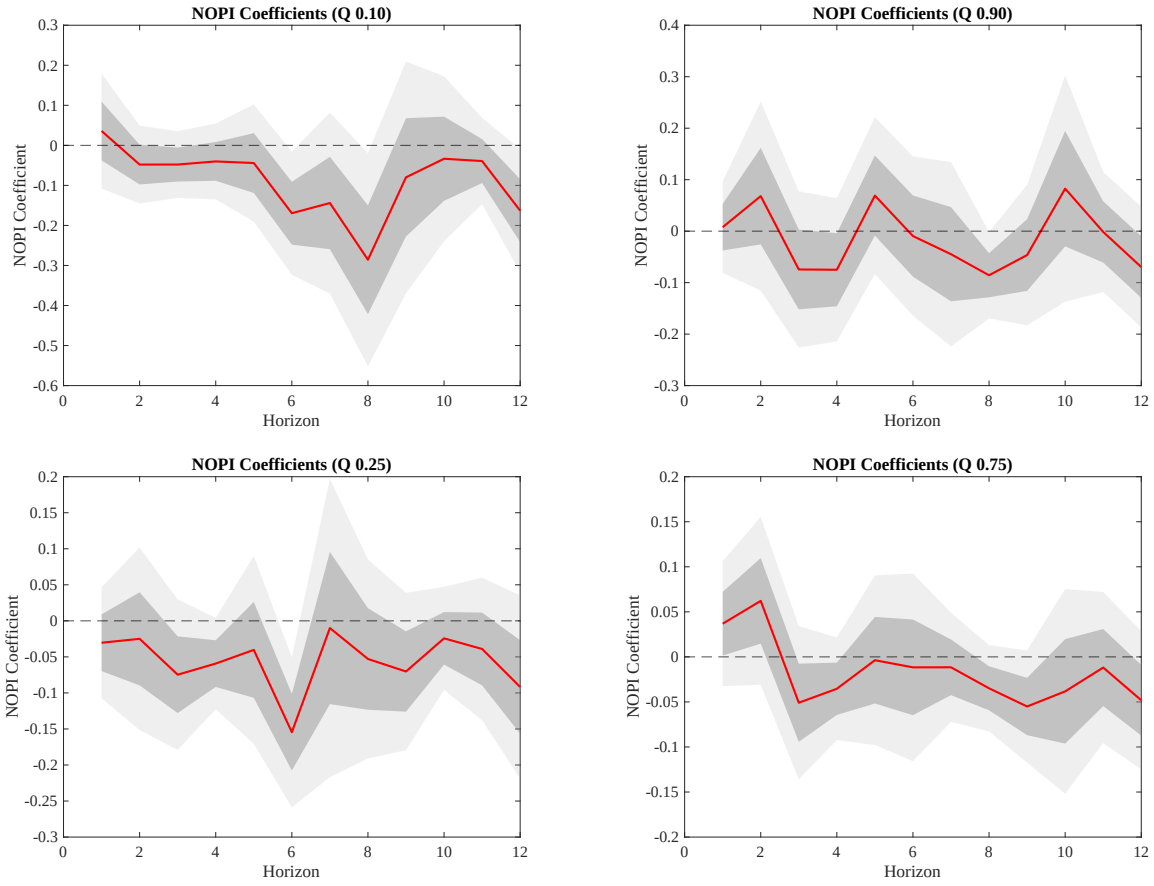


Figure 32: Coefficients of NOPI on the Energy sector's IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

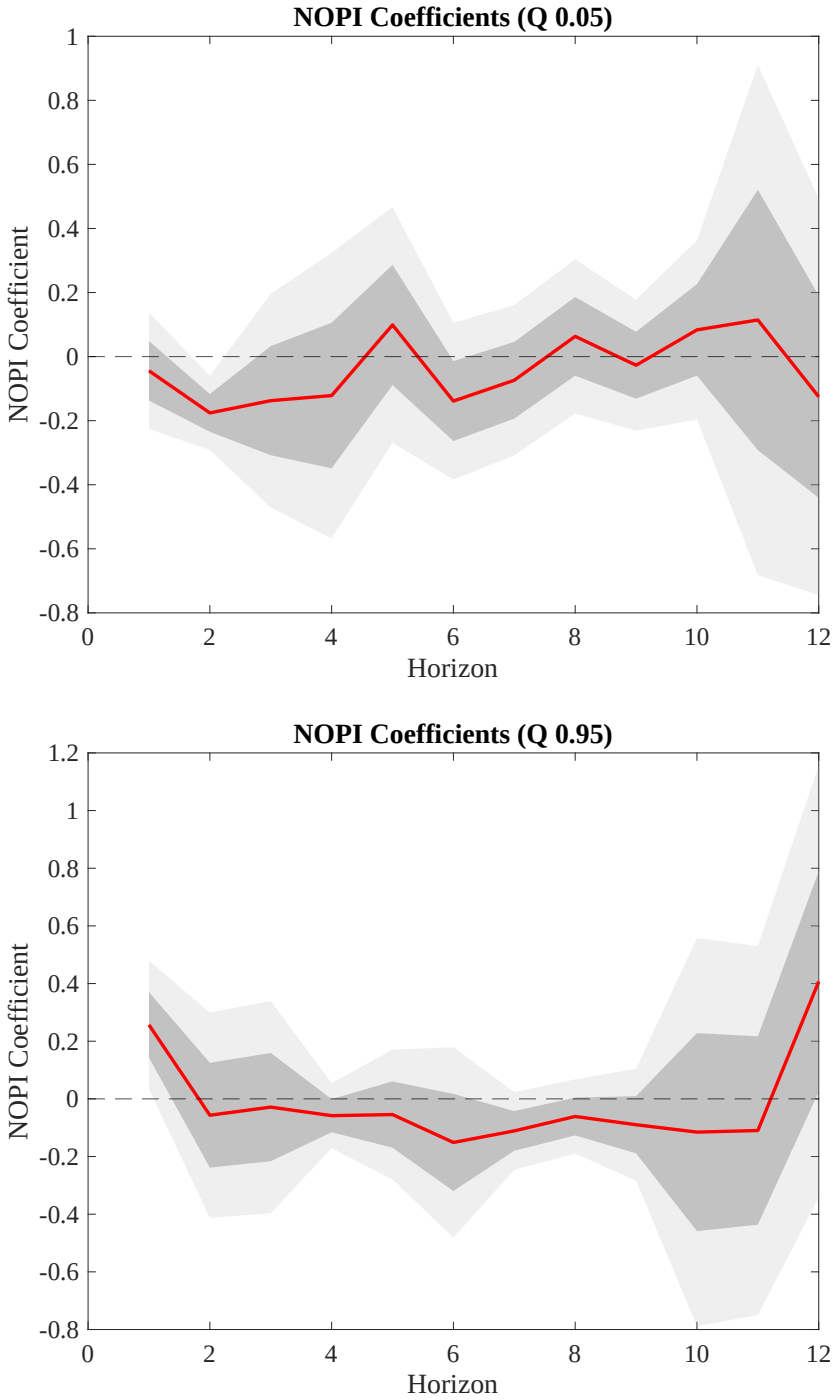


Figure 33: Coefficients of NOPI on the Energy sector's IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

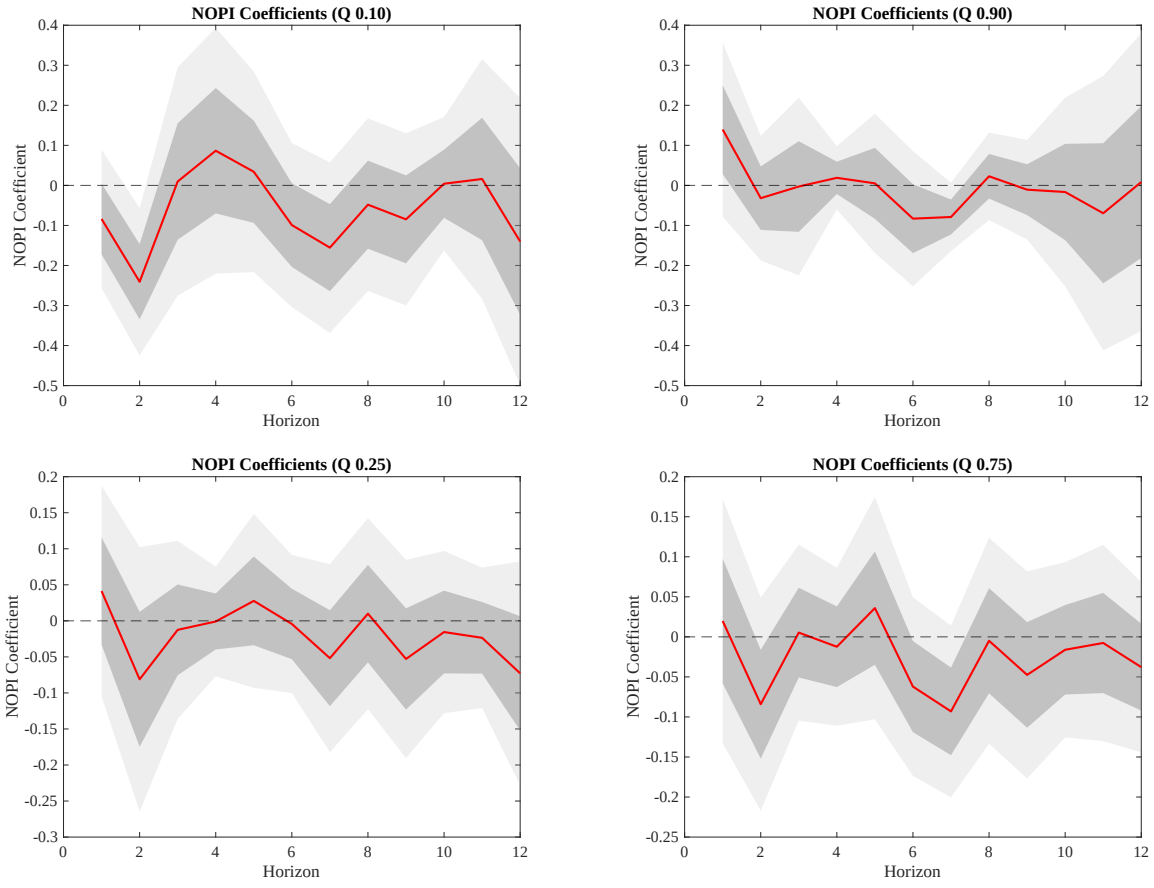


Figure 34: Expected Shortfall and Expected Longrise conditional on NGPI, $h = 3, 6, 12$.

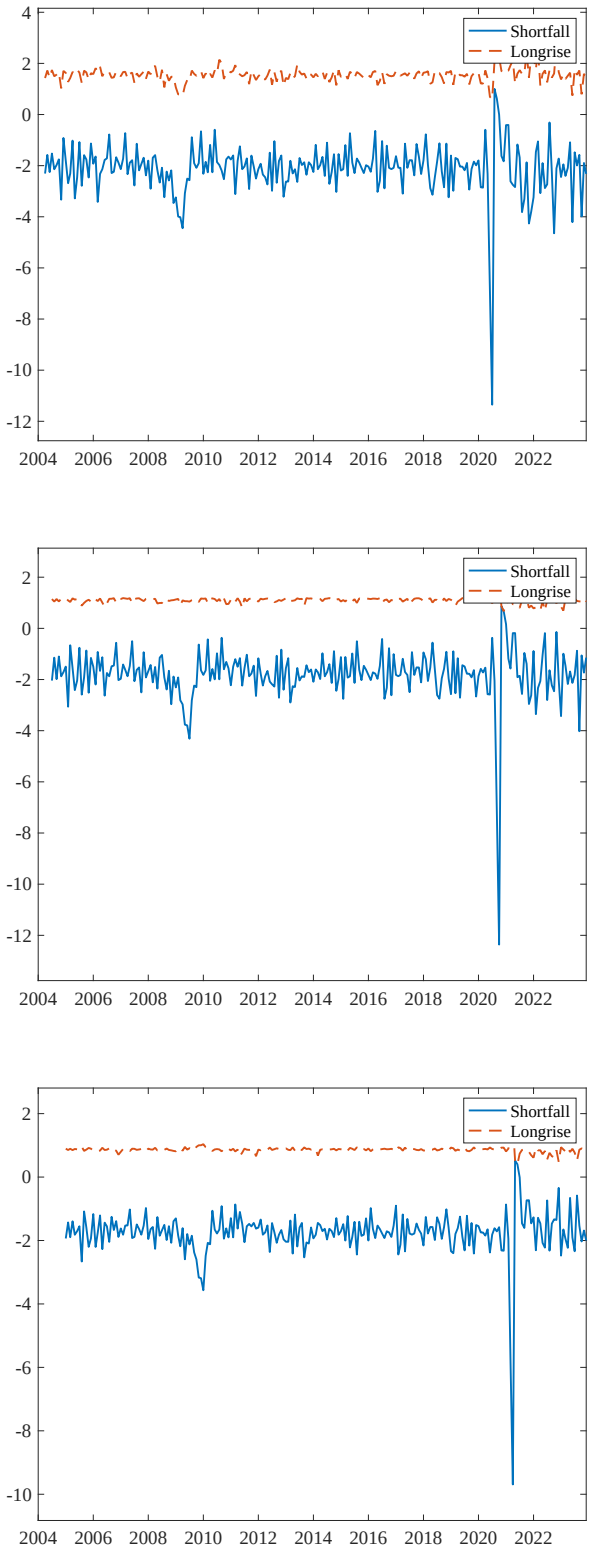


Figure 35: Expected Shortfall and Expected Longrise conditional on Brent Price Change, $h = 3, 6, 12$.

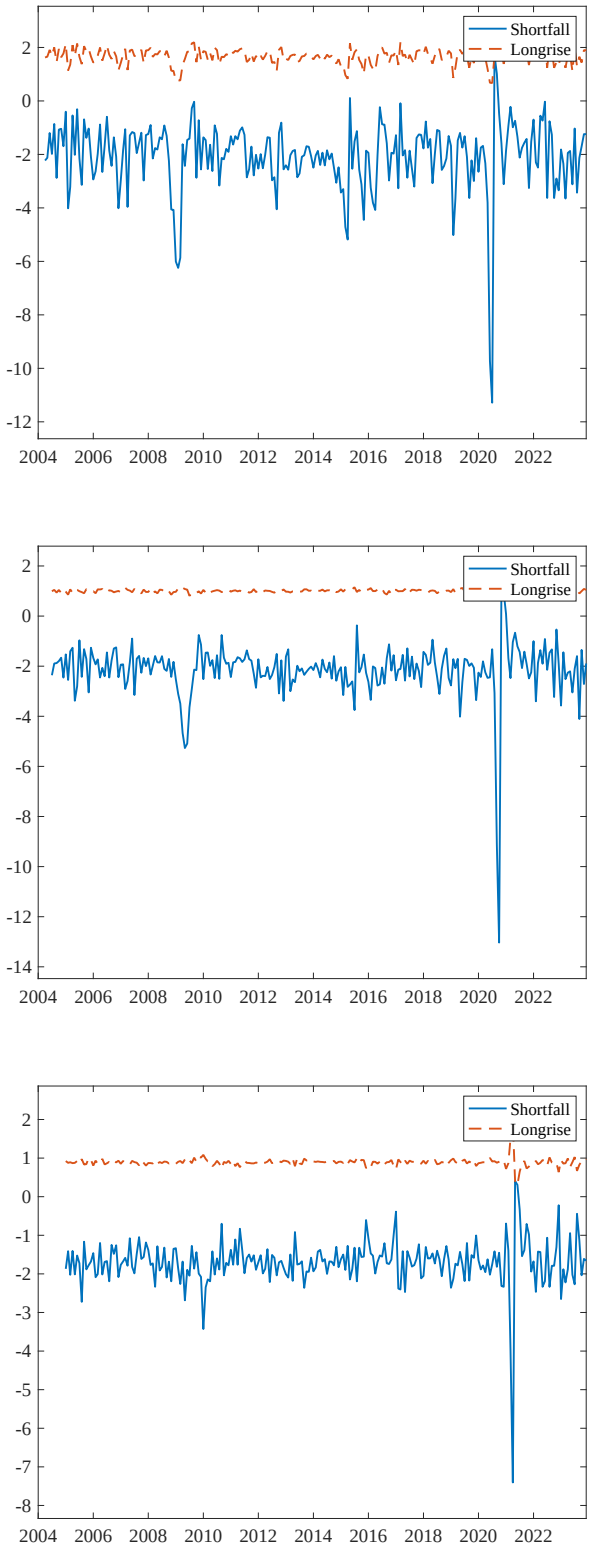


Figure 36: Fitted inverse CDF, March 2022: NGPI specification, $h = 3$ and $h = 6$.

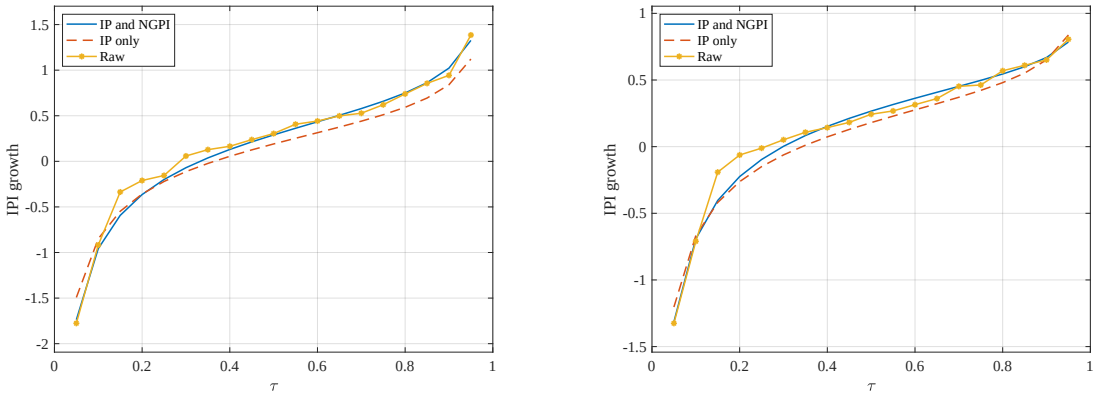
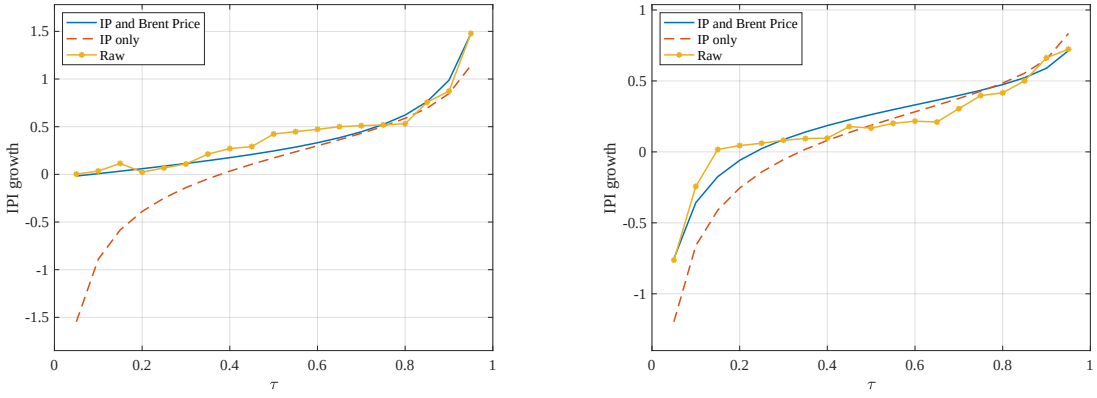


Figure 37: Fitted inverse CDF, March 2022: Brent Price Change specification, $h = 3$ and $h = 6$.



Subsample Comparison

Figure 38: Comparison of coefficients of NGPI on the total aggregate IPI Change across the horizons (Q5 and Q95). On the left, we observe the coefficient of the subsample (2004-2019), on the right, the full sample. The bands represent the 95% and 68% confidence interval.

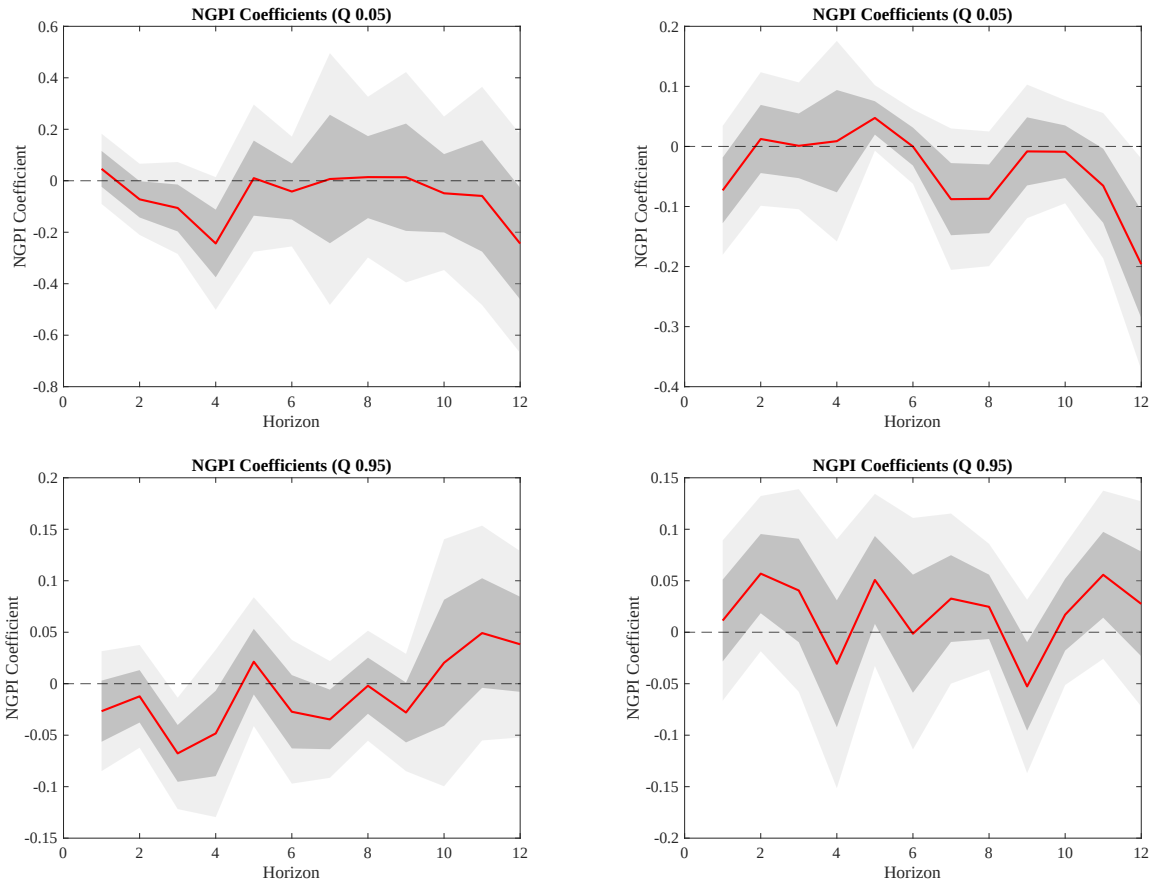


Figure 39: Comparison of coefficients of Brent Prices on the total aggregate IPI Change across the horizons (Q5 and Q95). On the left, we observe the coefficient of the subsample (2004-2019), on the right, the full sample. The bands represent the 95% and 68% confidence interval.

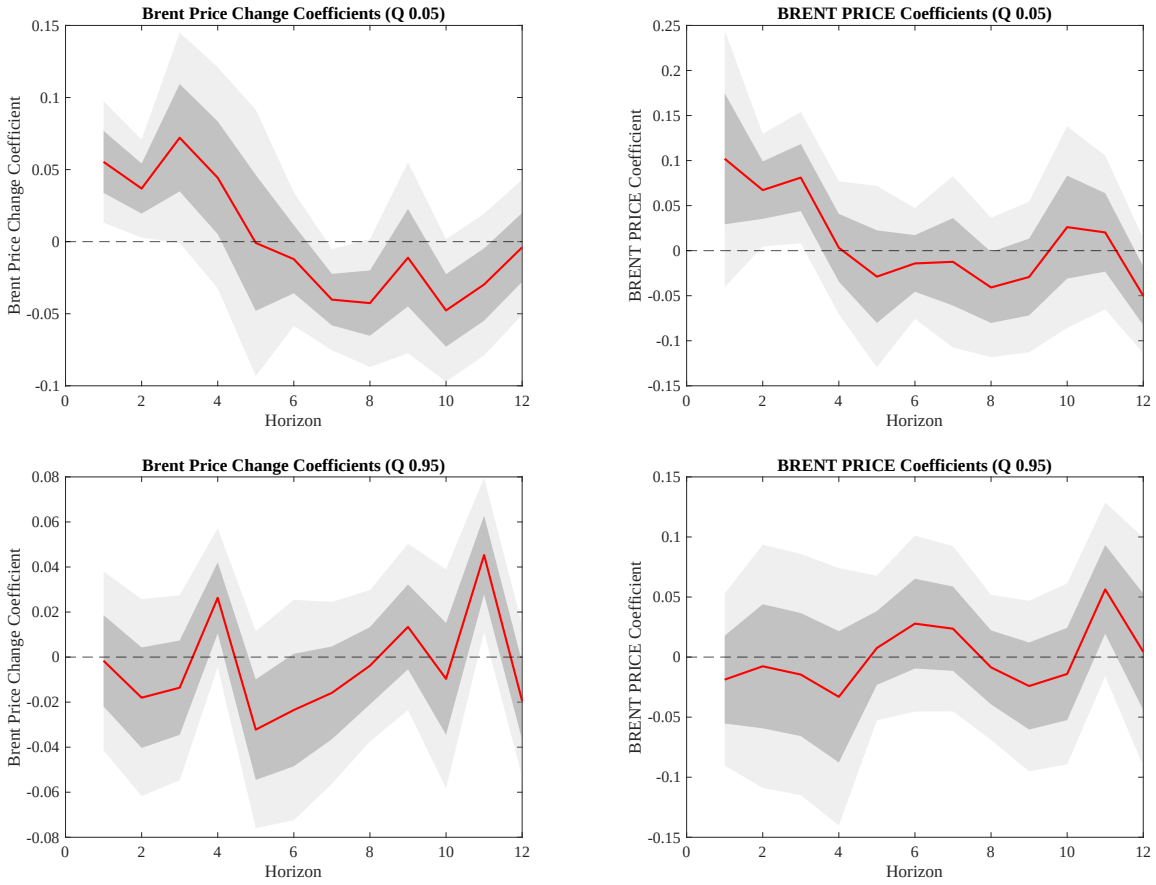
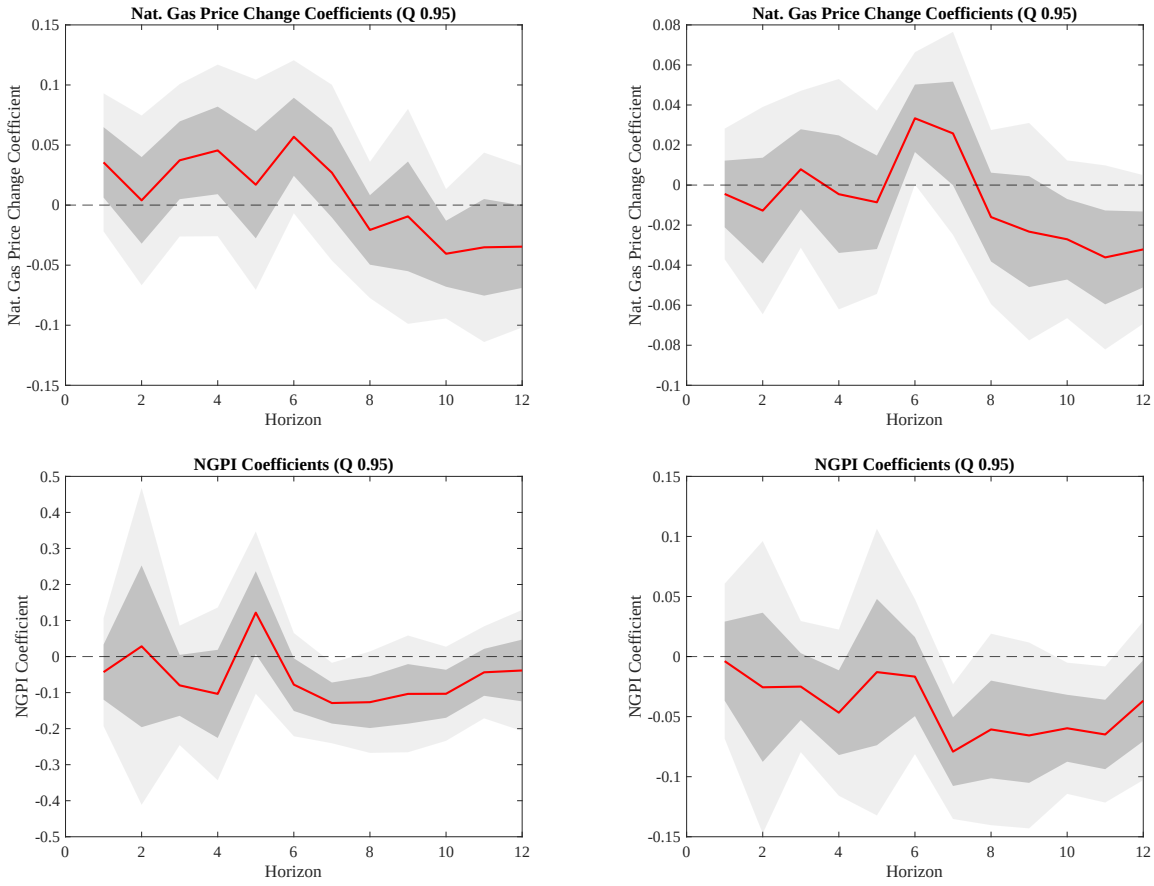


Figure 40: Comparison of coefficients of Gas Prices Change (above) and NGPI (below) on the Energy sector's IPI Change on the right tail (Q95). On the left, we observe the coefficient of the subsample (2004-2019), on the right, the full sample. The bands represent the 95% and 68% confidence interval.



Germany

Figure 41: Scatter plot of IPI Change "h steps ahead" relative to current Nat. Gas Prices Change ($h = 1, 3, 6, 12$) for the total aggregate IPI Change. The lines represent the slopes of the quantile regression for different quantiles.

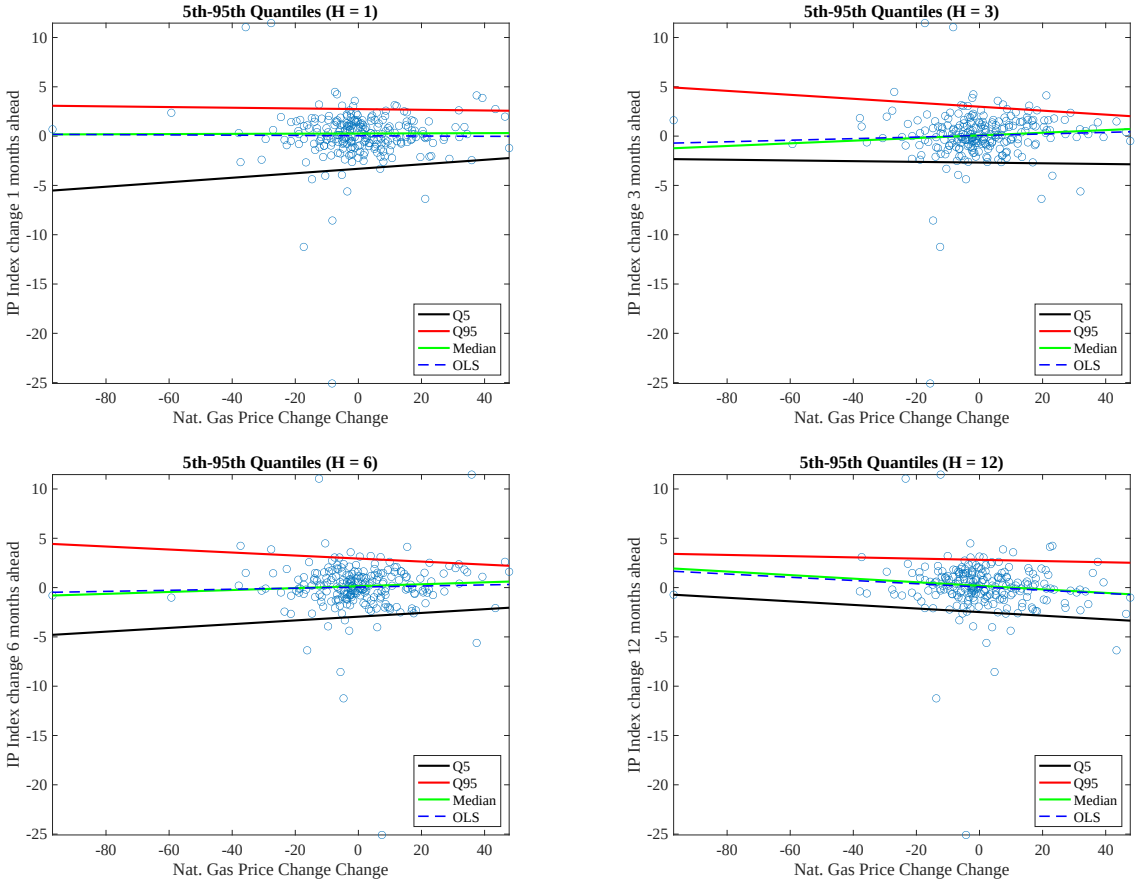


Figure 42: Scatter plot of IPI Change "h steps ahead" relative to current Nat. Gas Prices Change ($h = 1, 3, 6, 12$) for the IPI Change of the energy sector. The lines represent the slopes of the quantile regression for different quantiles.

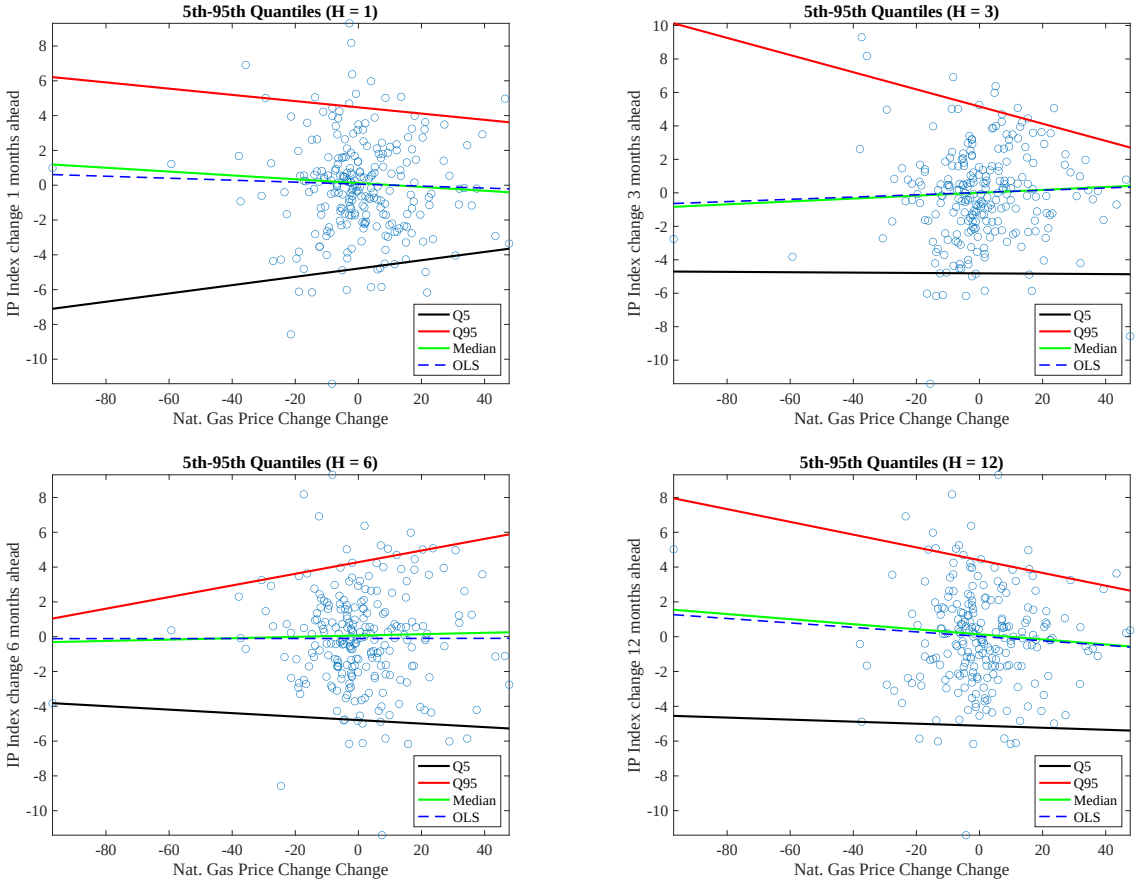


Figure 43: Scatter plot of IPI Change "h steps ahead" relative to current NGPI ($h = 1, 3, 6, 12$) for the total aggregate IPI Change. The lines represent the slopes of the quantile regression for different quantiles.

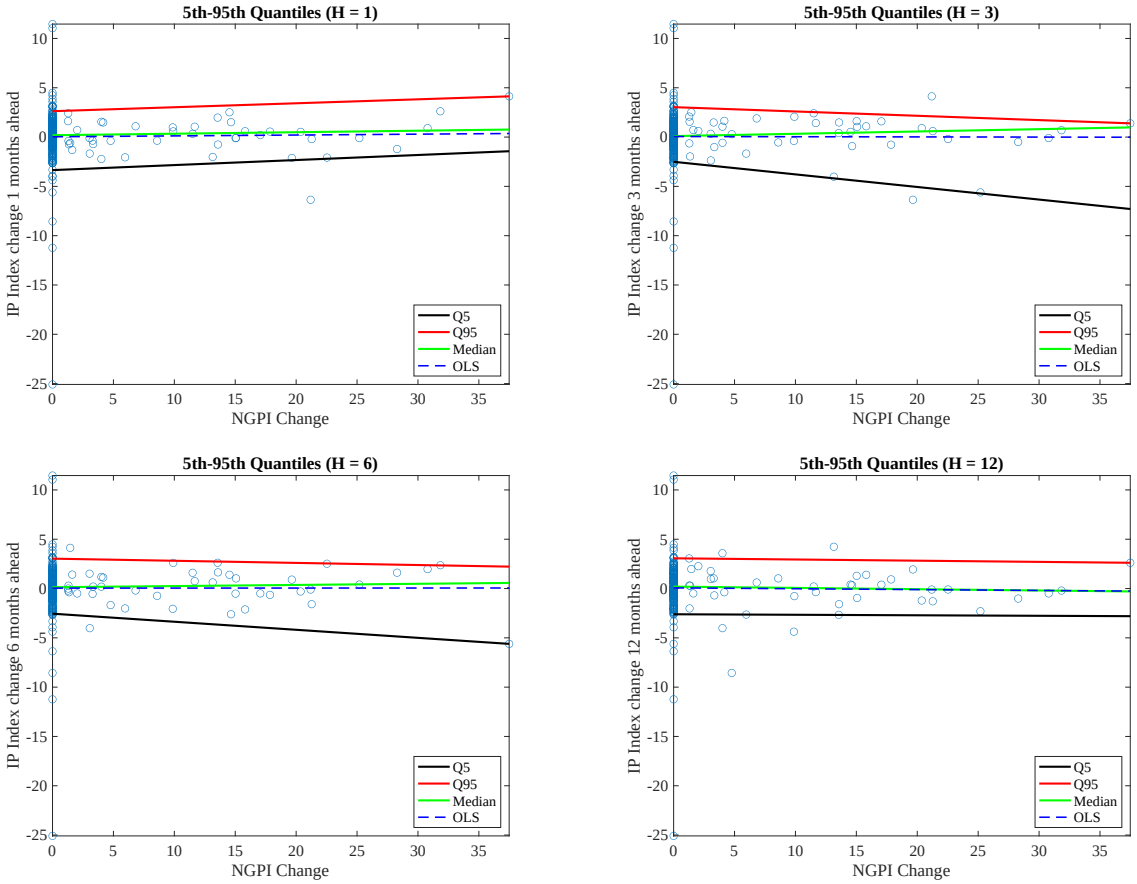


Figure 44: Scatter plot of IPI Change "h steps ahead" relative to current NGPI ($h = 1, 3, 6, 12$) for the IPI Change of the energy sector. The lines represent the slopes of the quantile regression for different quantiles.

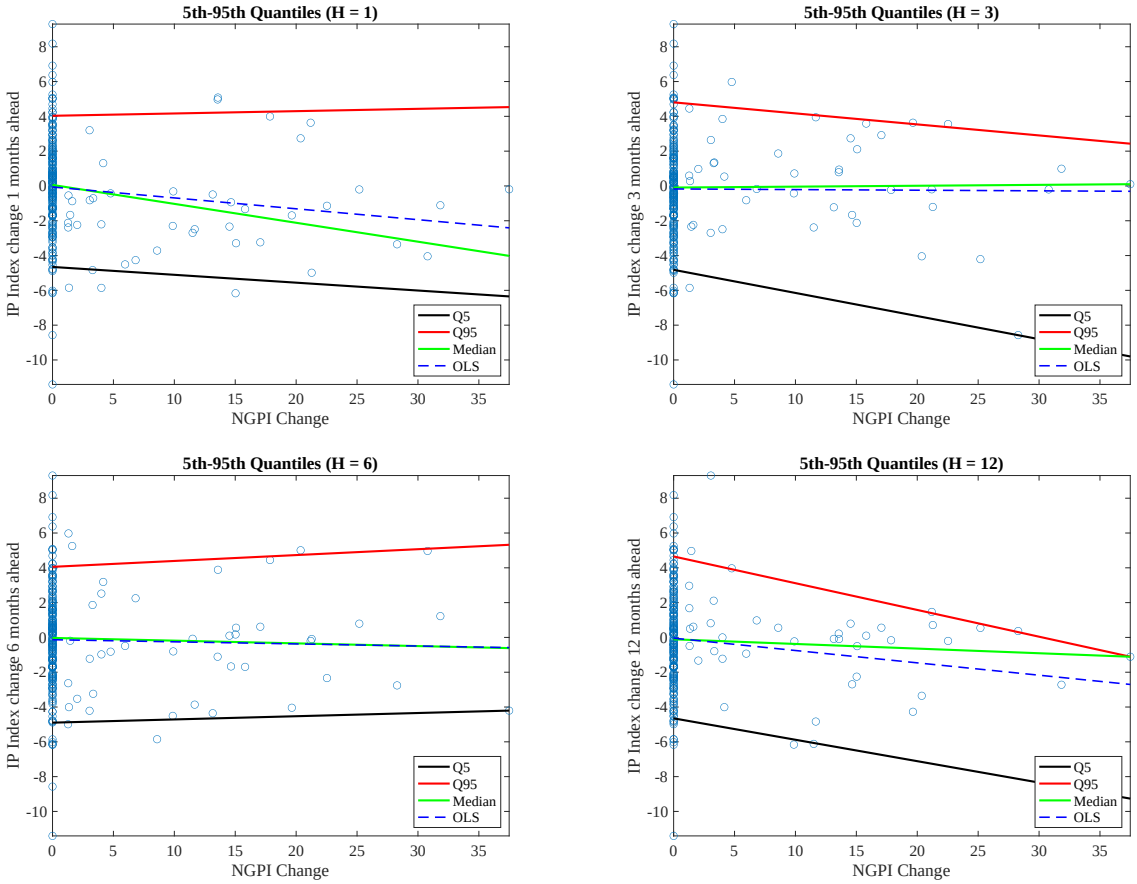


Figure 45: Scatter plot of IPI Change "h steps ahead" relative to current Oil Prices Change ($h = 1, 3, 6, 12$) for the total aggregate IPI Change. The lines represent the slopes of the quantile regression for different quantiles.

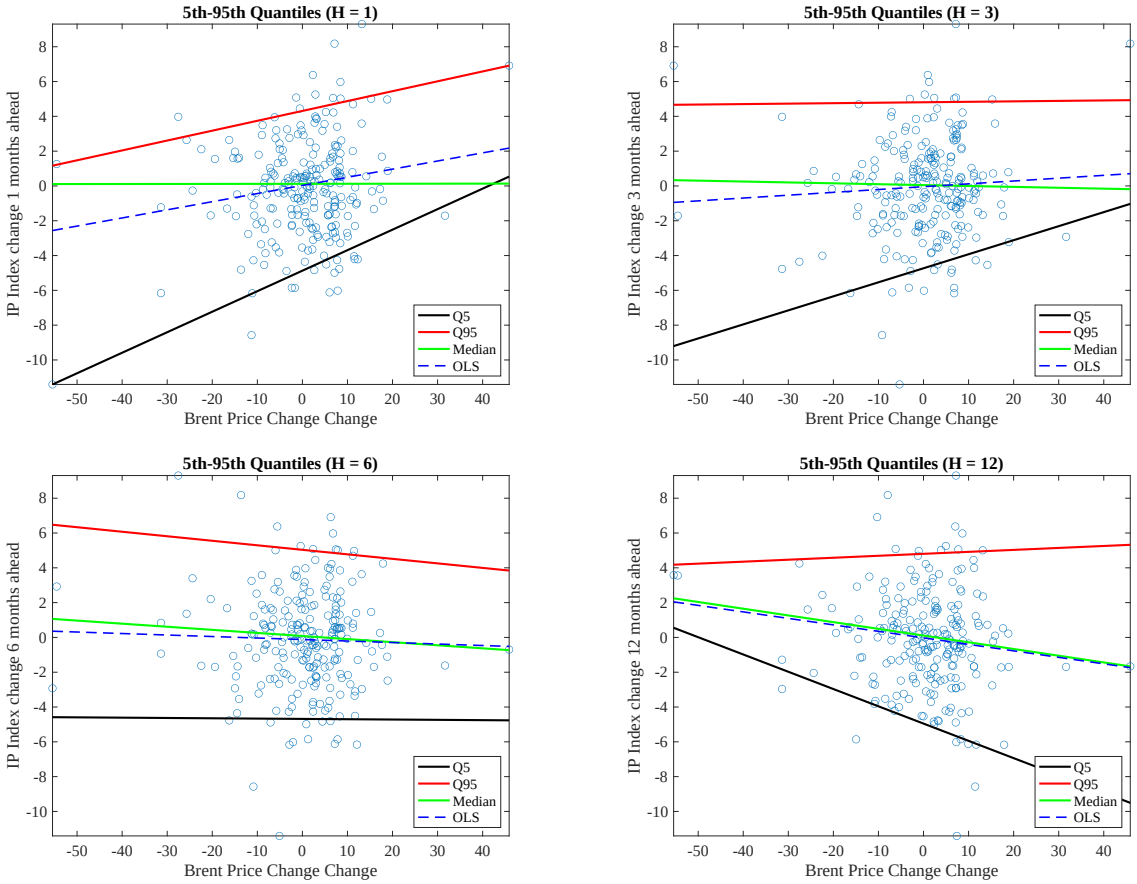


Figure 46: Coefficients of Gas Price Change on the total aggregate IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

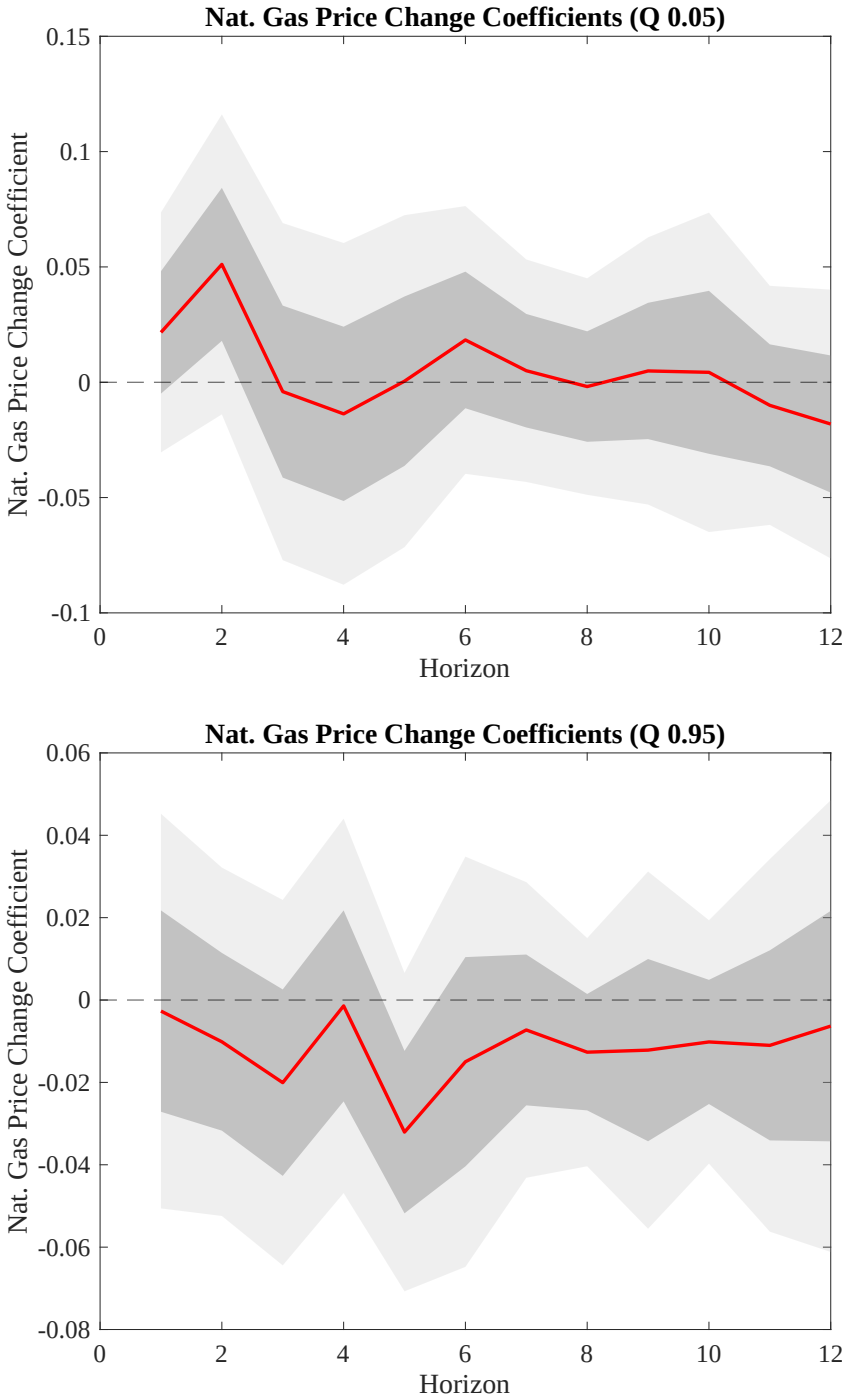


Figure 47: Coefficients of Gas Price Change on the Energy sector's IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

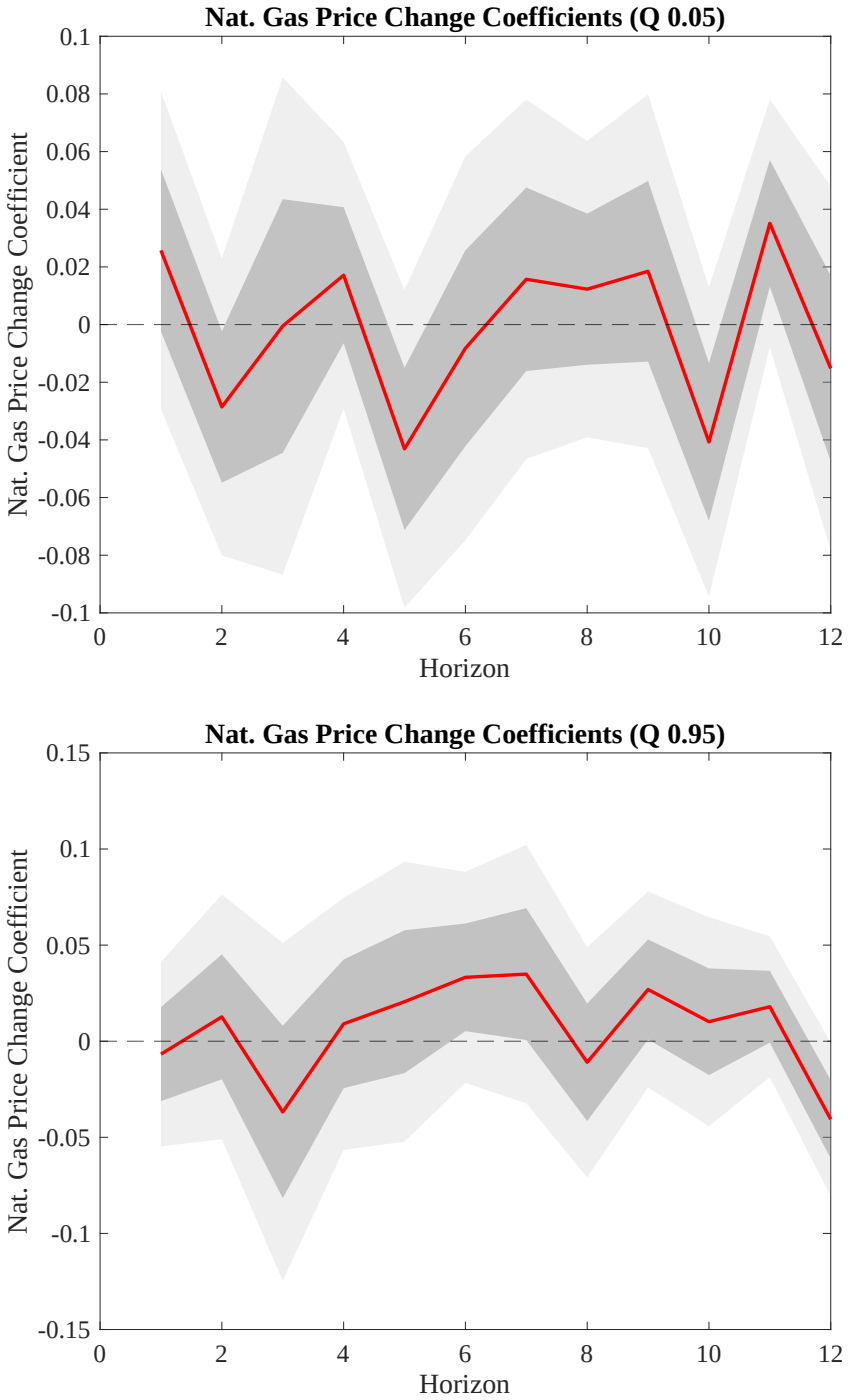
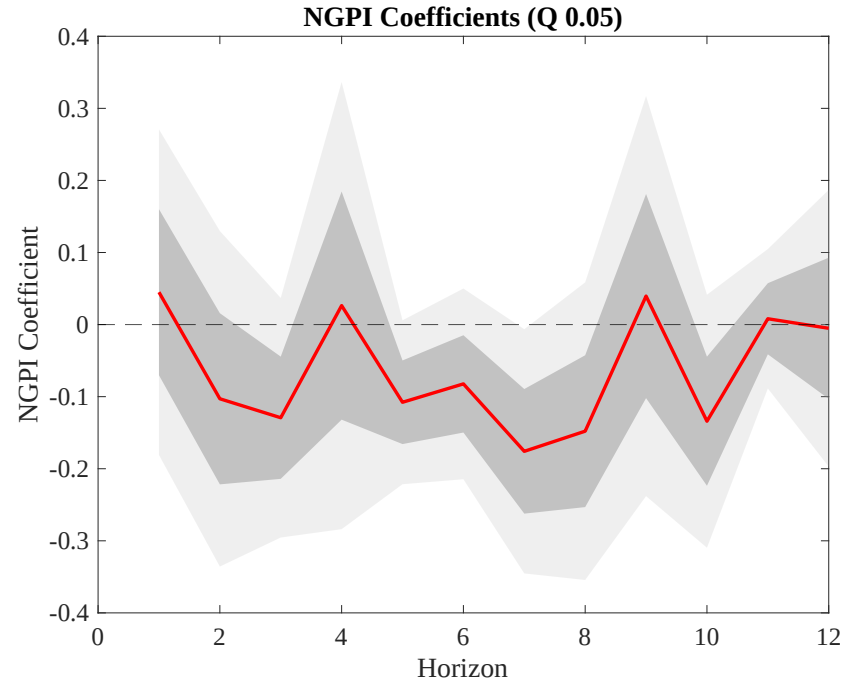
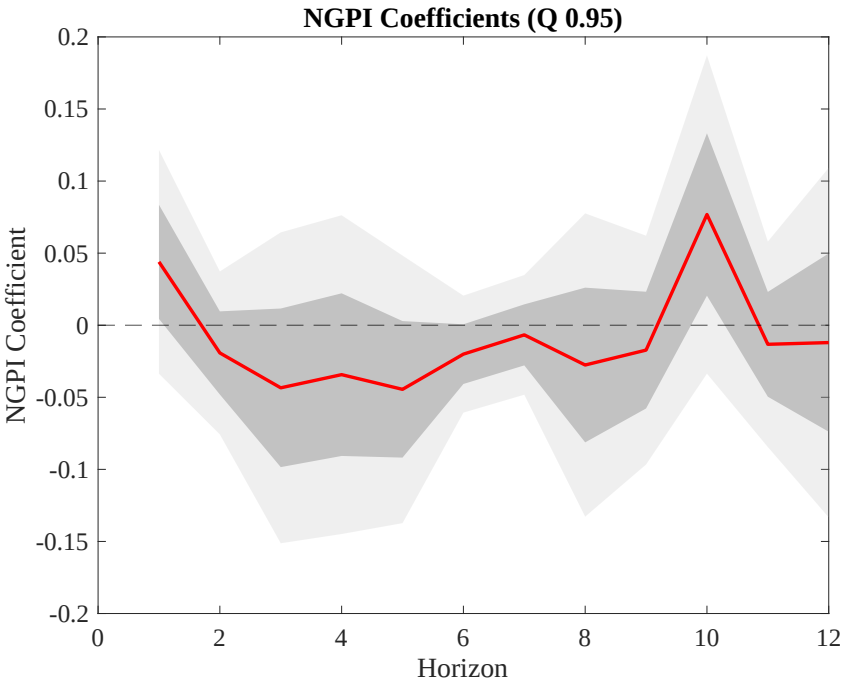


Figure 48: Coefficients of NGPI on the total aggregate IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.





Italy

Figure 49: Scatter plot of IPI Change "h steps ahead" relative to current Nat. Gas Prices Change ($h = 1, 3, 6, 12$) for the total aggregate IPI Change. The lines represent the slopes of the quantile regression for different quantiles.

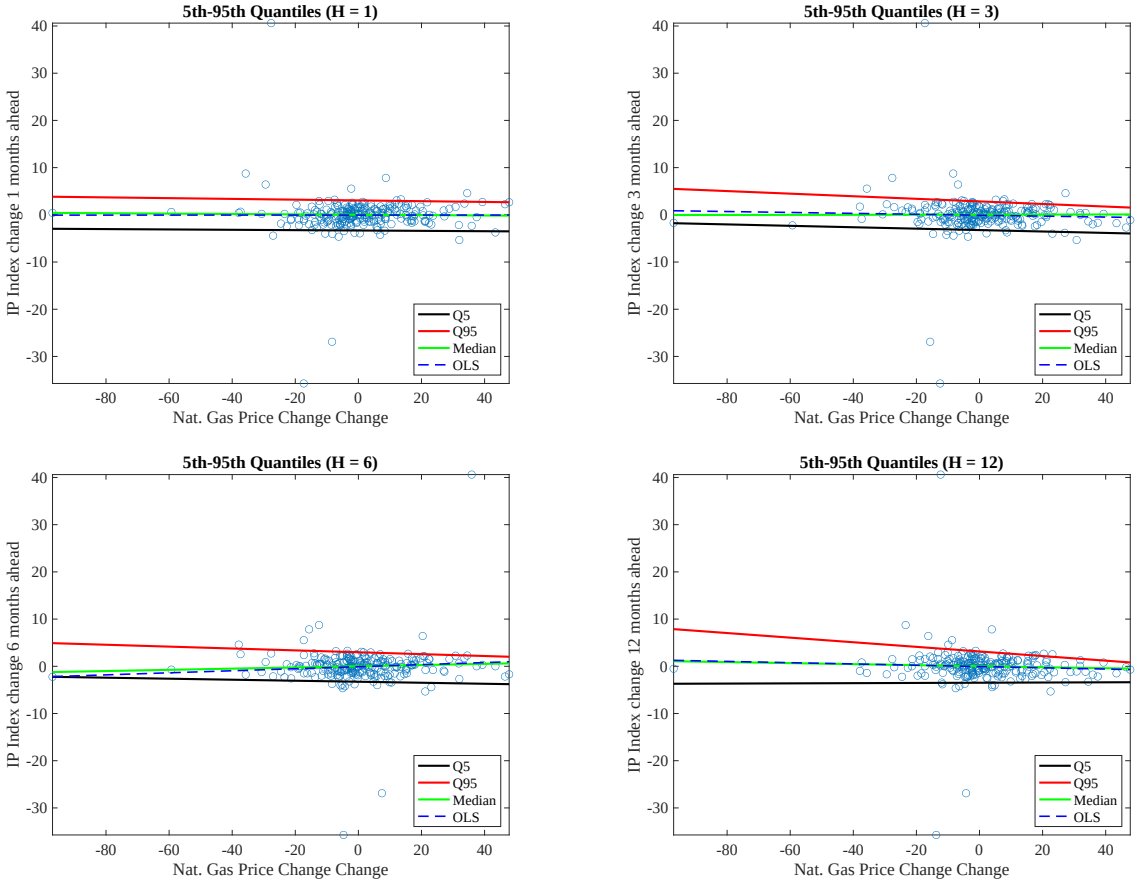


Figure 50: Scatter plot of IPI Change "h steps ahead" relative to current Nat. Gas Prices Change ($h = 1, 3, 6, 12$) for the IPI Change of the energy sector. The lines represent the slopes of the quantile regression for different quantiles.

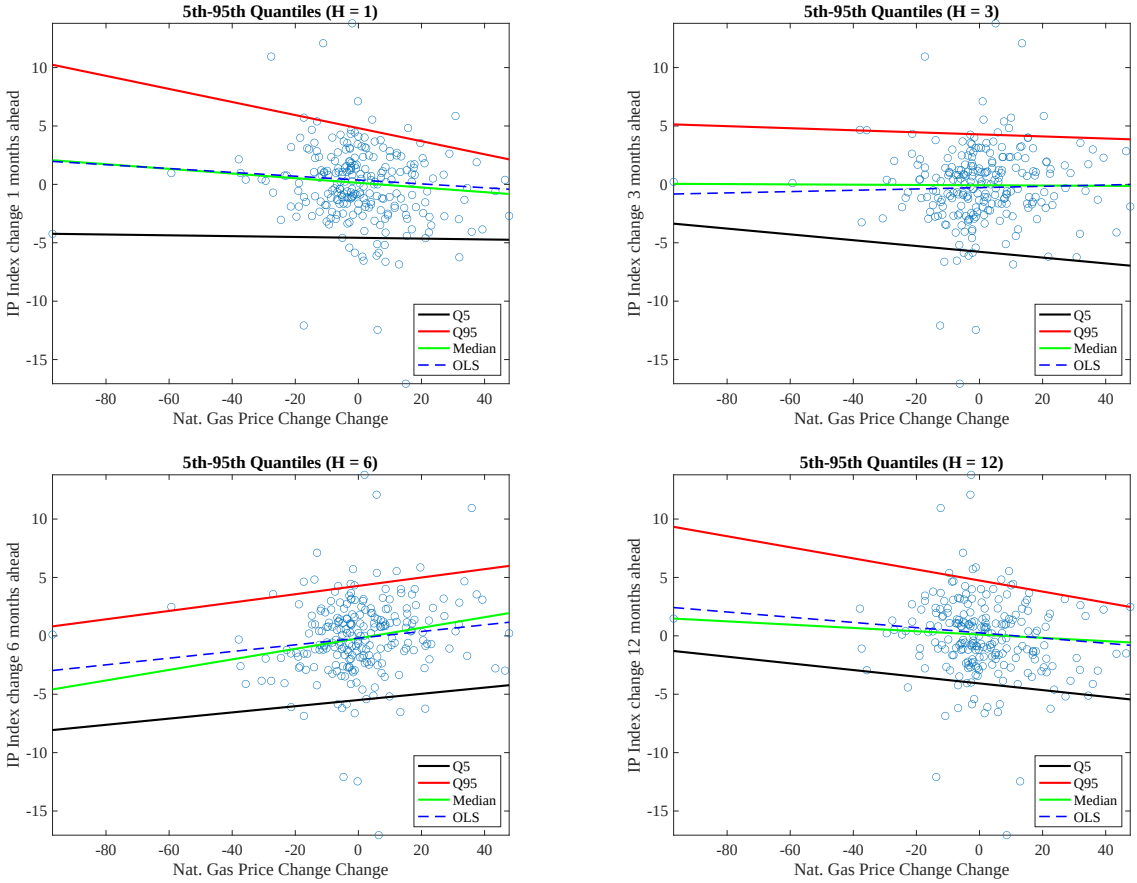


Figure 51: Scatter plot of IPI Change "h steps ahead" relative to current NGPI ($h = 1, 3, 6, 12$) for the total aggregate IPI Change. The lines represent the slopes of the quantile regression for different quantiles.

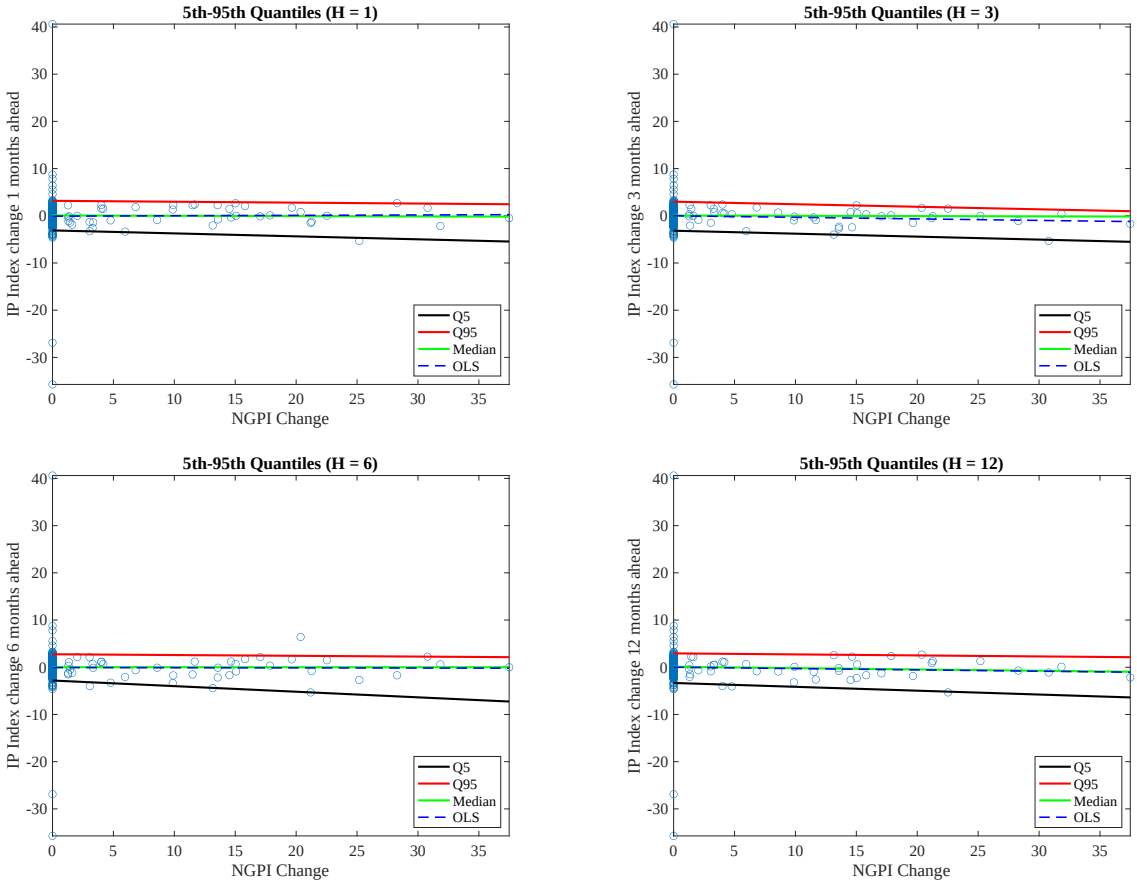


Figure 52: Scatter plot of IPI Change "h steps ahead" relative to current NGPI ($h = 1, 3, 6, 12$) for the IPI Change of the energy sector. The lines represent the slopes of the quantile regression for different quantiles.

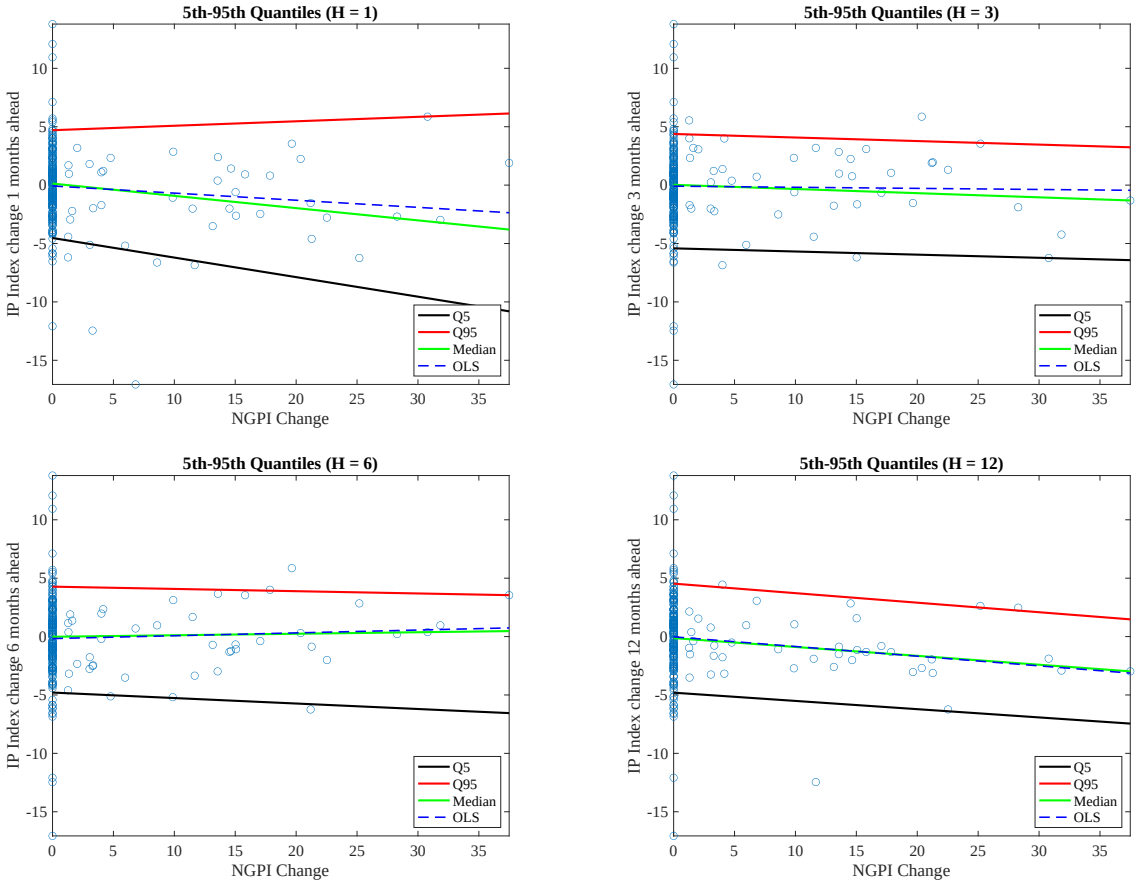


Figure 53: Scatter plot of IPI Change "h steps ahead" relative to current Oil Prices Change ($h = 1, 3, 6, 12$) for the total aggregate IPI Change. The lines represent the slopes of the quantile regression for different quantiles.

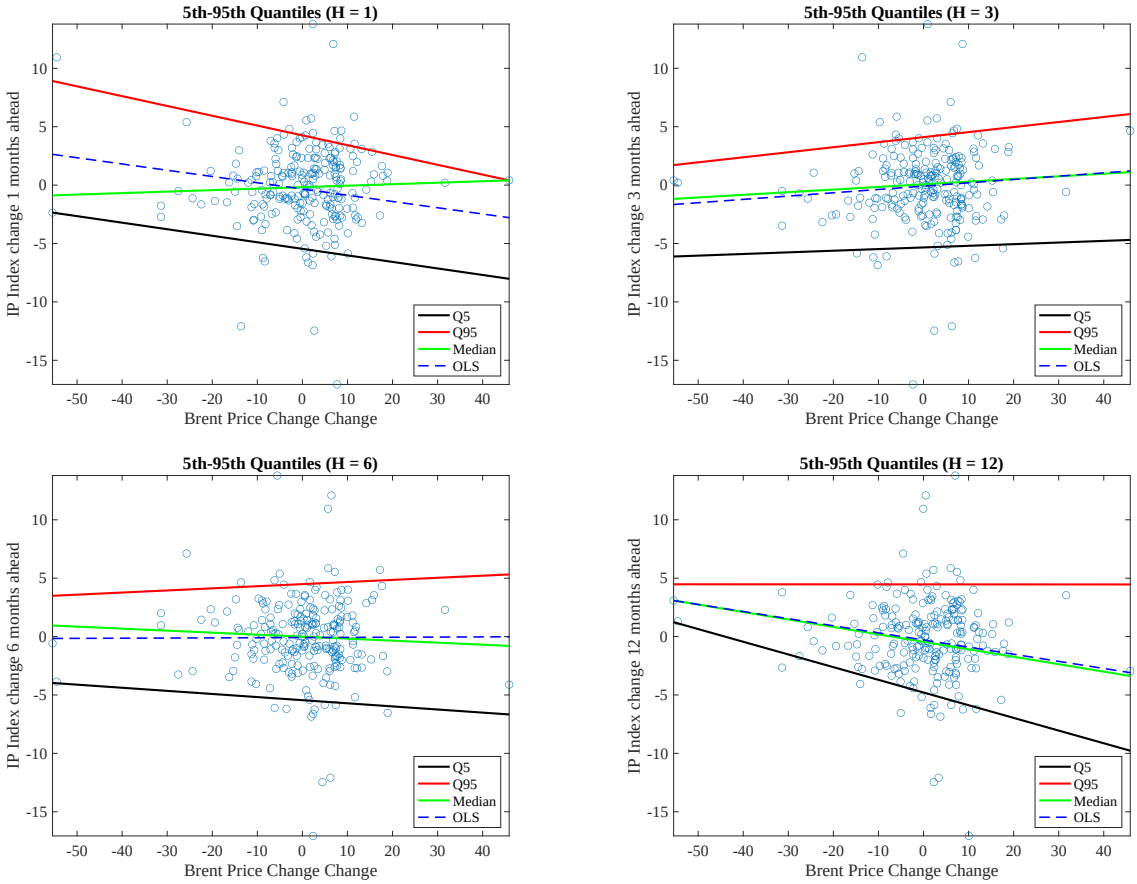


Figure 54: Coefficients of Gas Price Change on the total aggregate IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

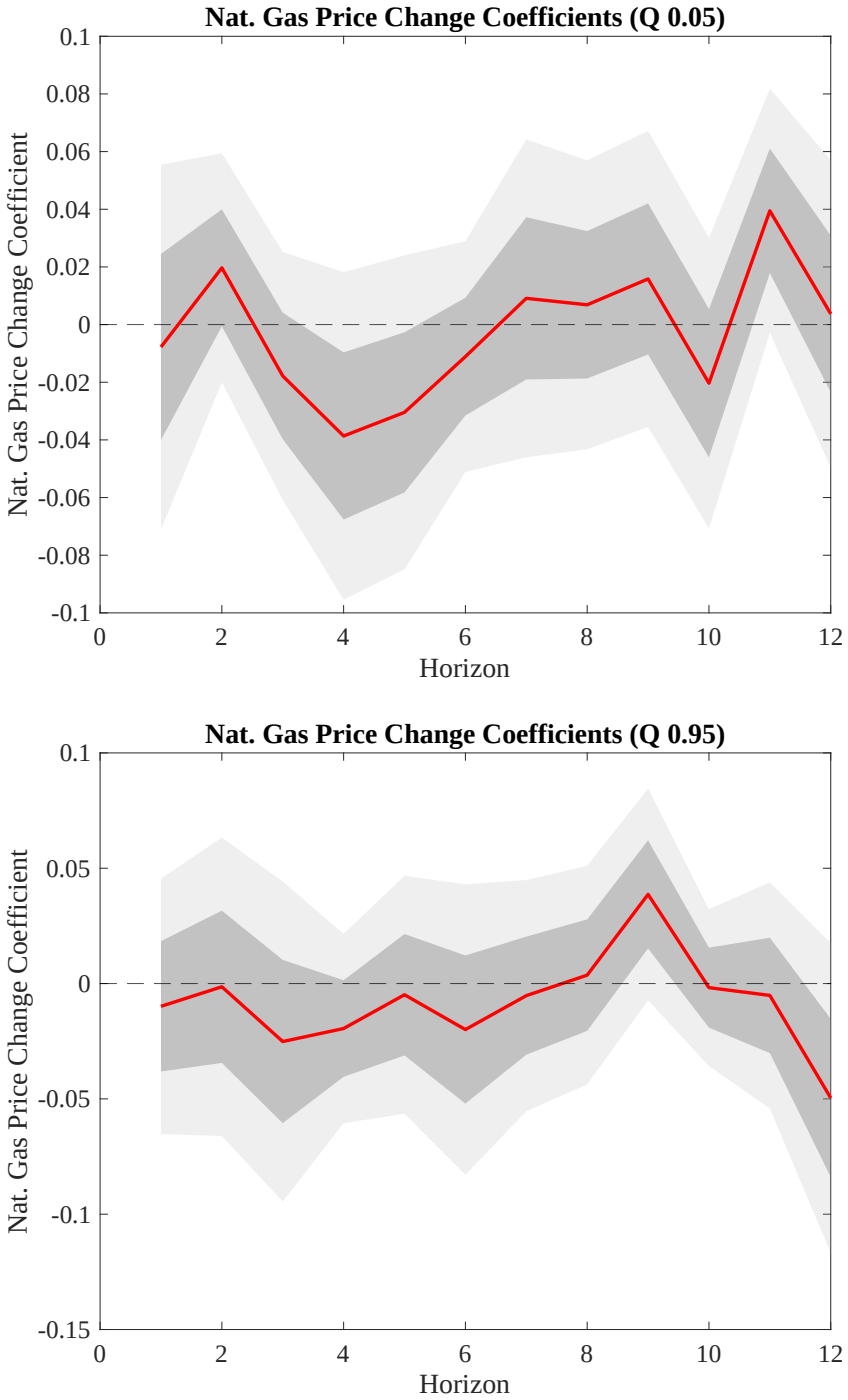
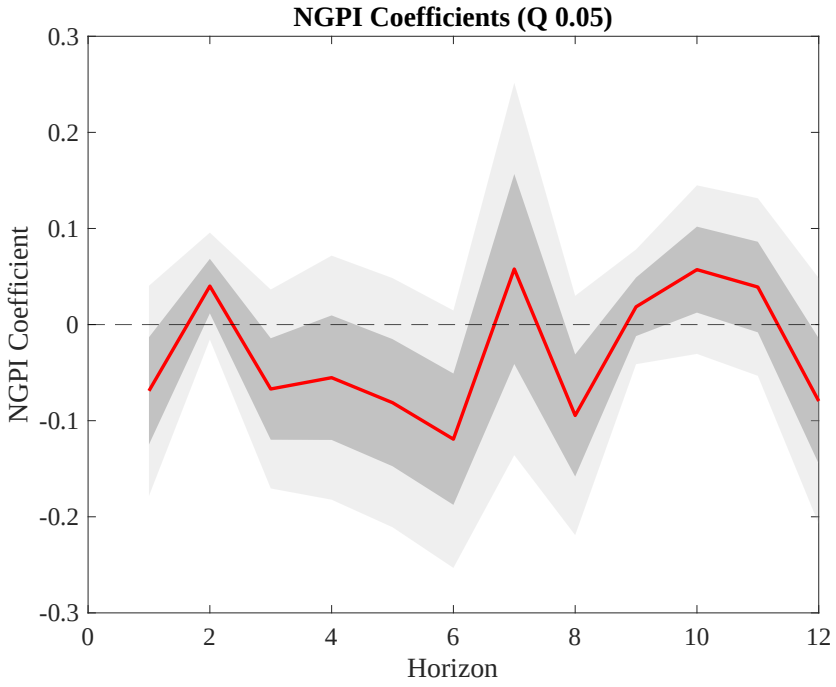


Figure 55: Coefficients of NGPI on the total aggregate IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.



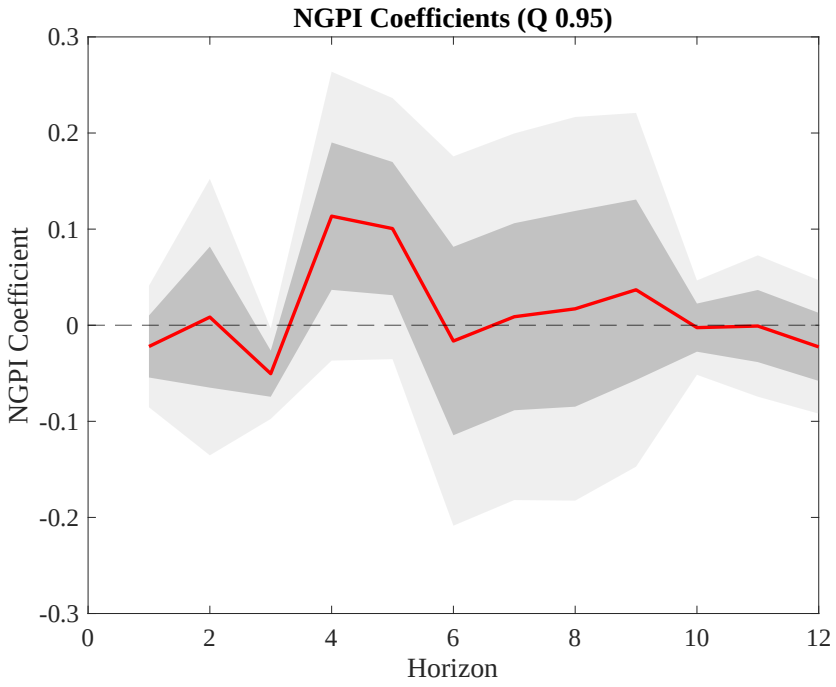
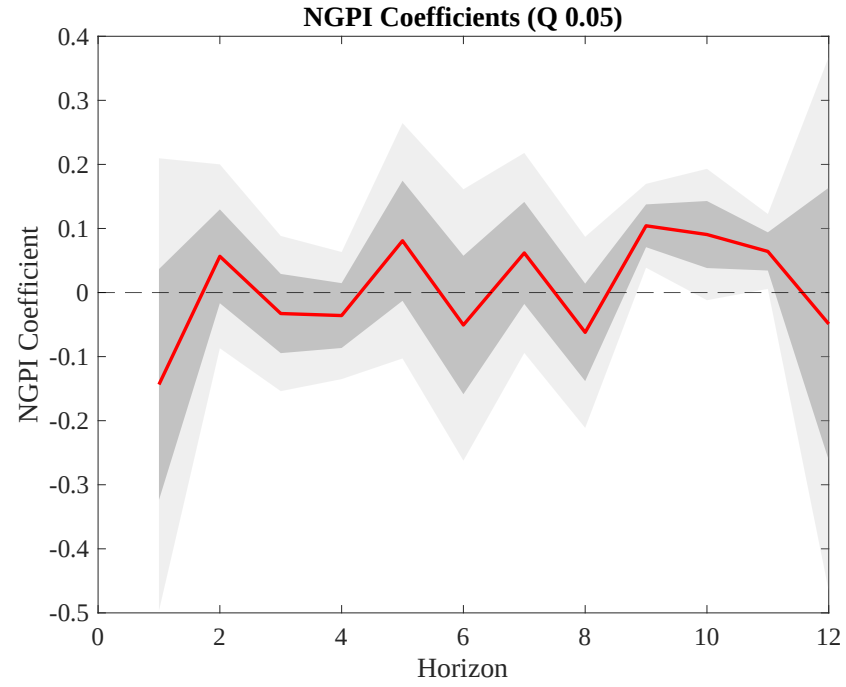


Figure 56: Coefficients of NGPI on the Energy sector's IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.



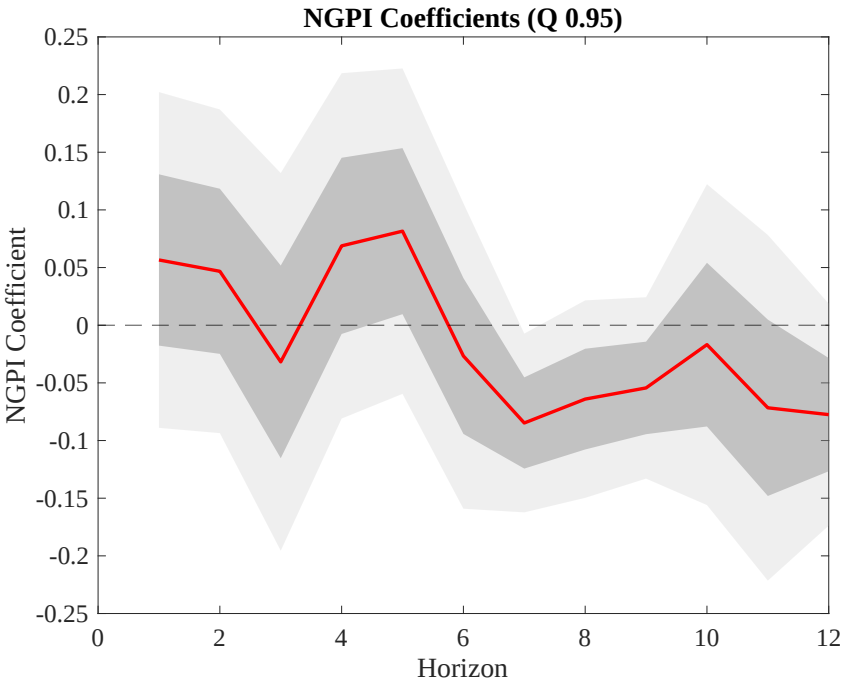
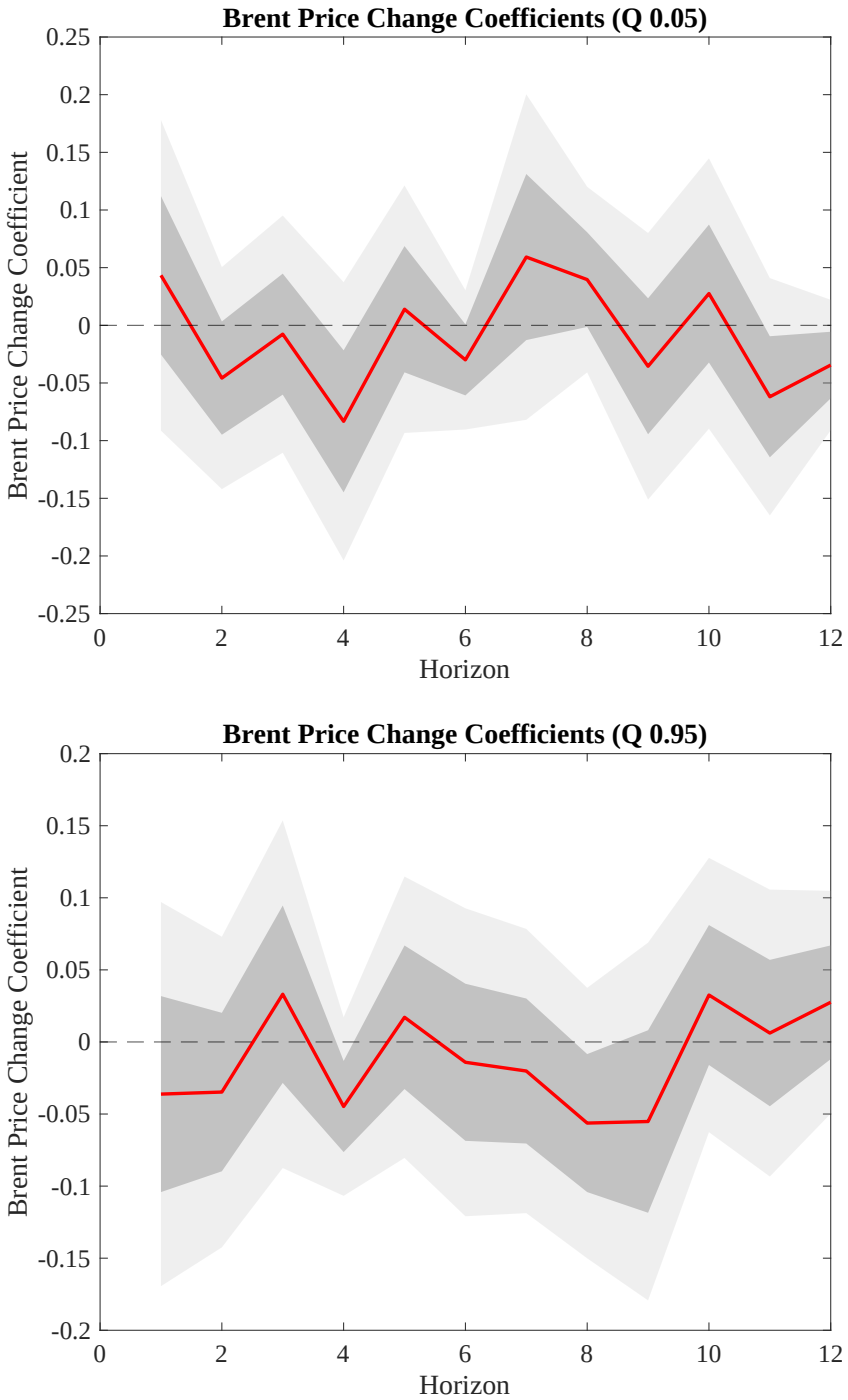


Figure 57: Coefficients of Brent Prices on the Energy sector's IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.



France

Figure 58: Scatter plot of IPI Change "h steps ahead" relative to current Nat. Gas Prices Change ($h = 1, 3, 6, 12$) for the total aggregate IPI Change. The lines represent the slopes of the quantile regression for different quantiles.

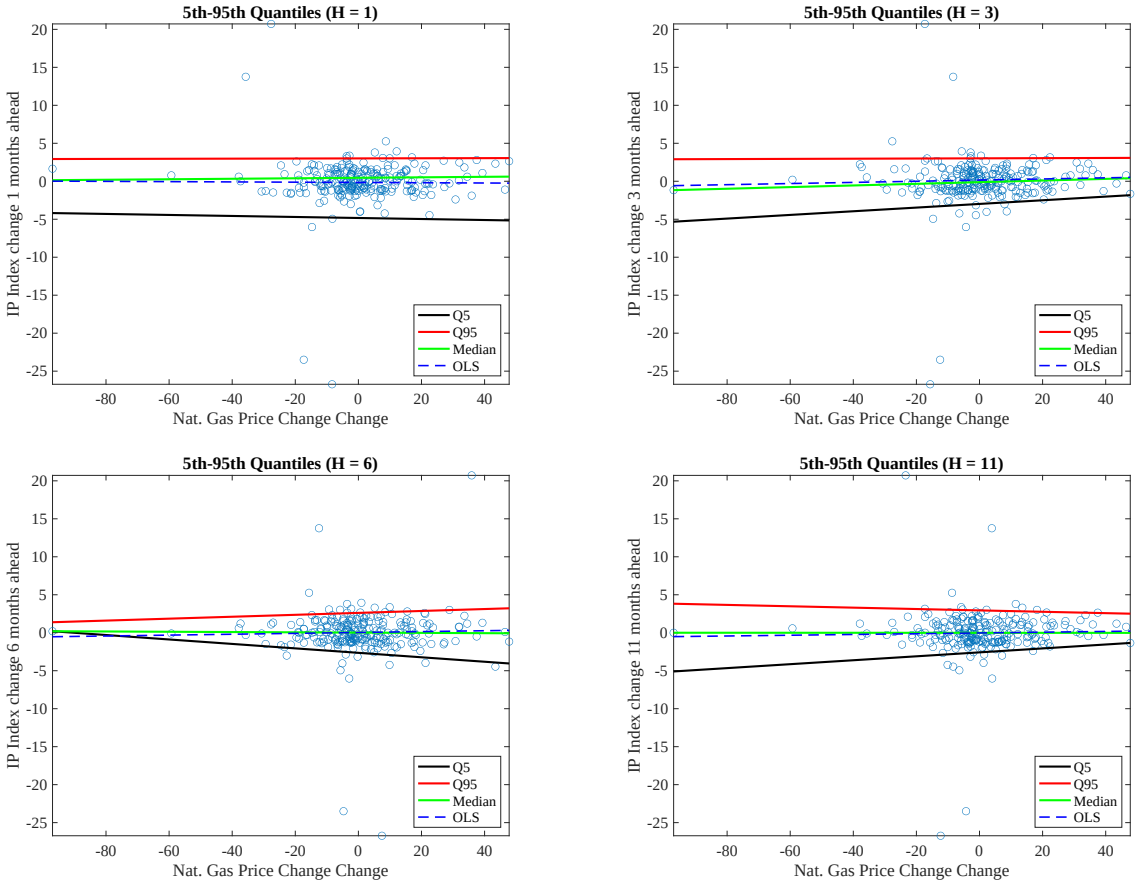


Figure 59: Scatter plot of IPI Change "h steps ahead" relative to current Nat. Gas Prices Change ($h = 1, 3, 6, 12$) for the IPI Change of the energy sector. The lines represent the slopes of the quantile regression for different quantiles.

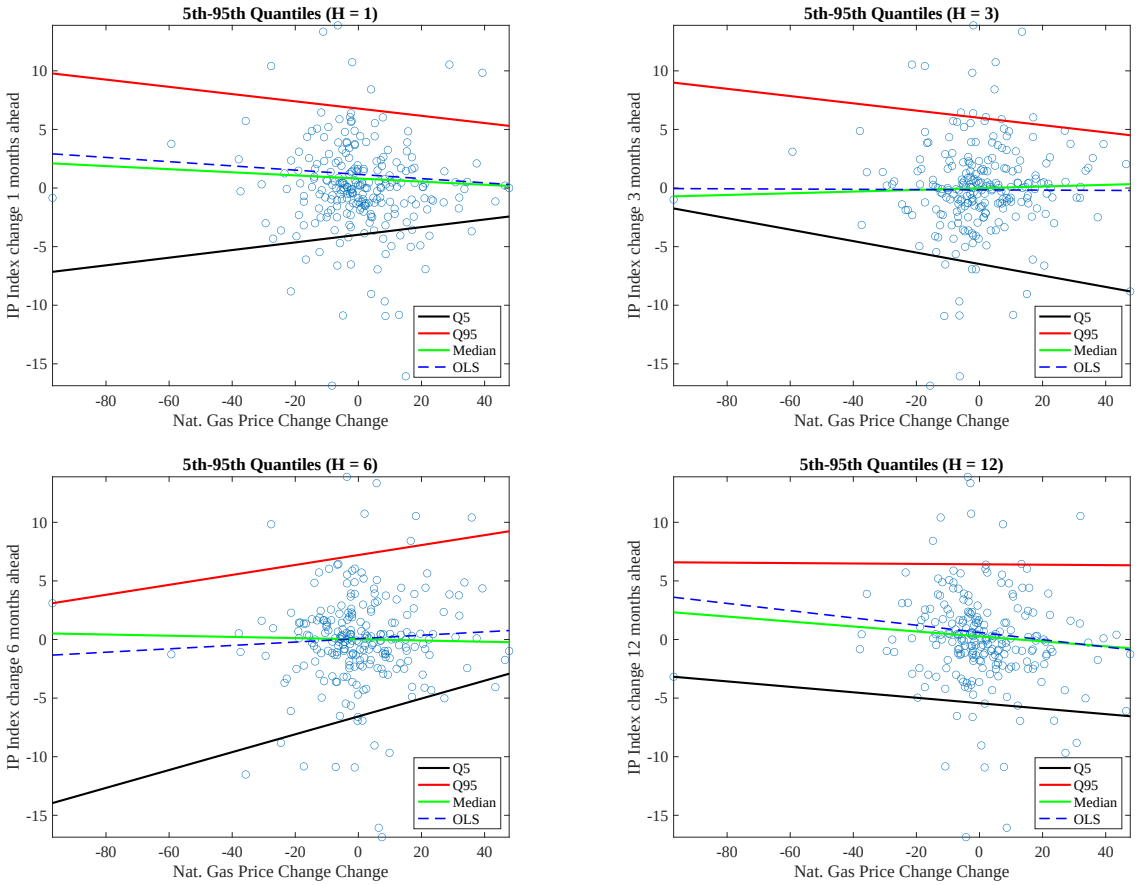


Figure 60: Scatter plot of IPI Change "h steps ahead" relative to current NGPI ($h = 1, 3, 6, 12$) for the total aggregate IPI Change. The lines represent the slopes of the quantile regression for different quantiles.

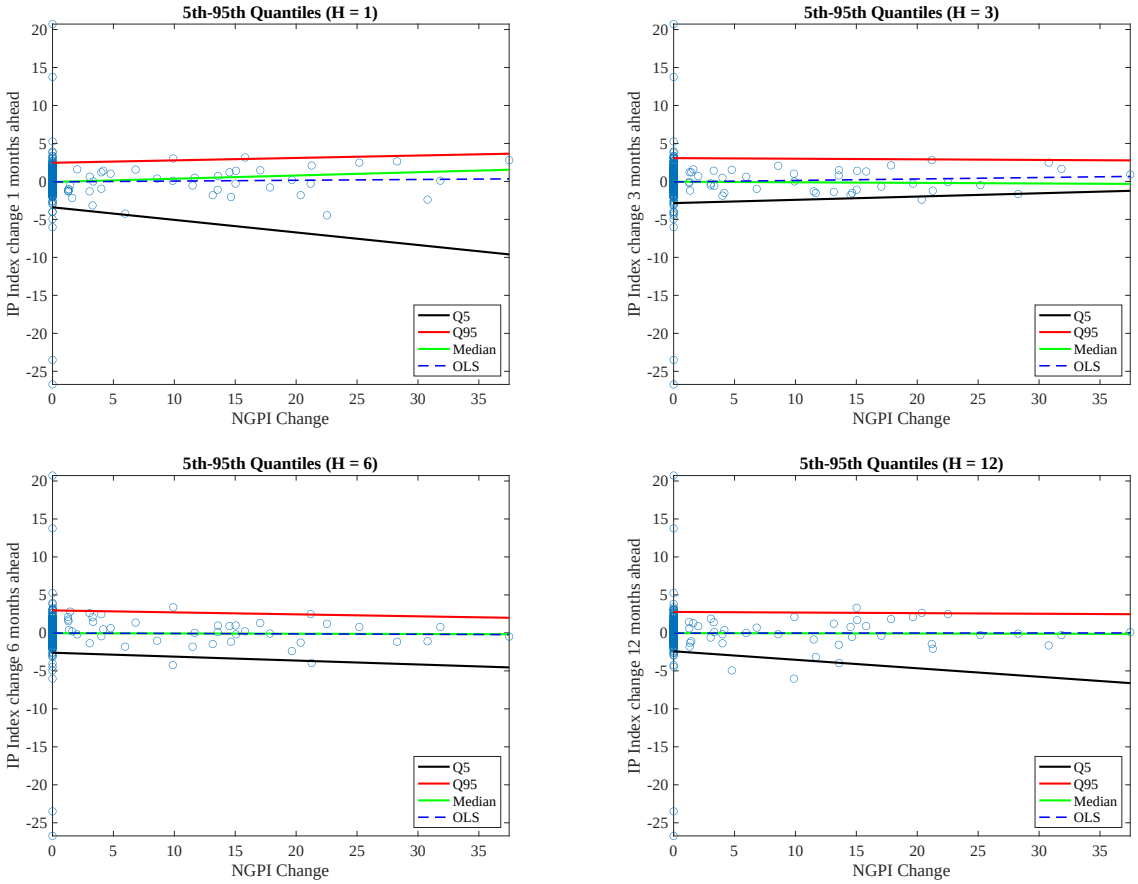


Figure 61: Scatter plot of IPI Change "h steps ahead" relative to current NGPI ($h = 1, 3, 6, 12$) for the IPI Change of the energy sector. The lines represent the slopes of the quantile regression for different quantiles.

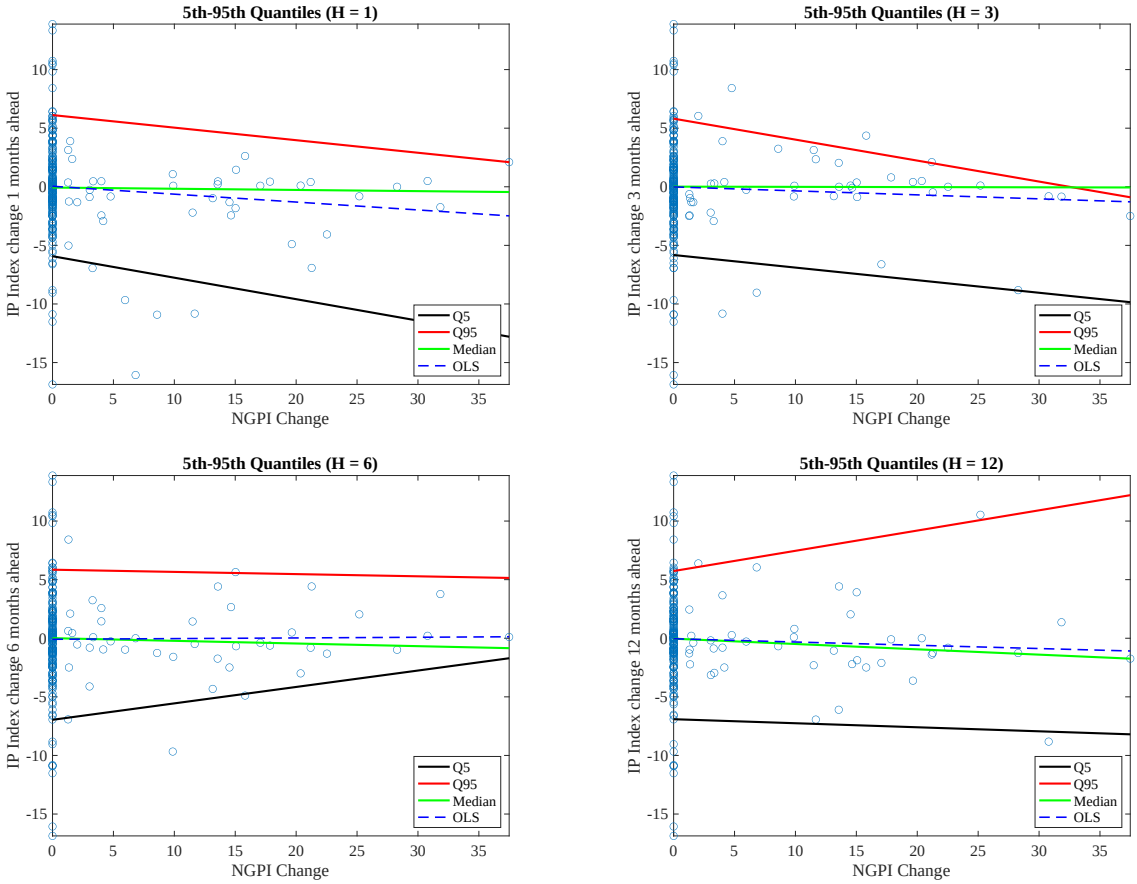


Figure 62: Scatter plot of IPI Change "h steps ahead" relative to current Oil Prices Change ($h = 1, 3, 6, 12$) for the total aggregate IPI Change. The lines represent the slopes of the quantile regression for different quantiles.

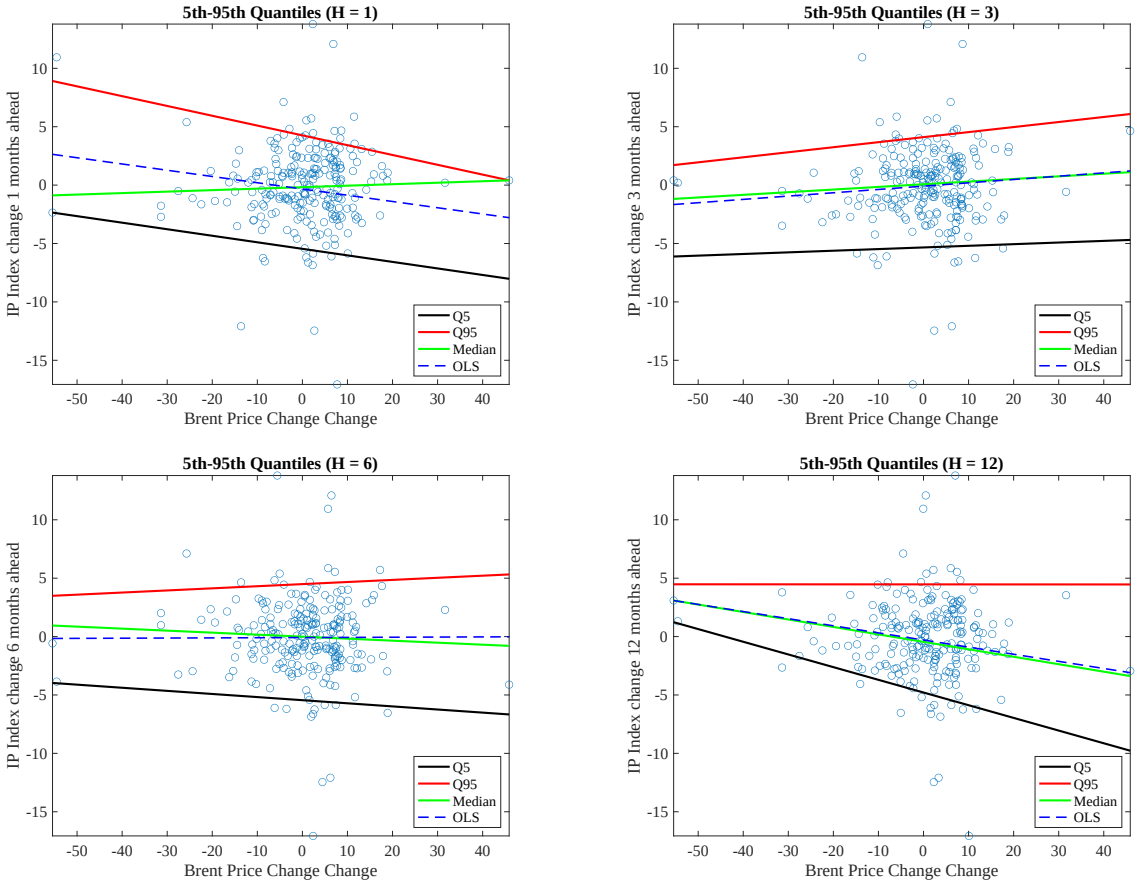


Figure 63: Coefficients of Gas Price Change on the total aggregate IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

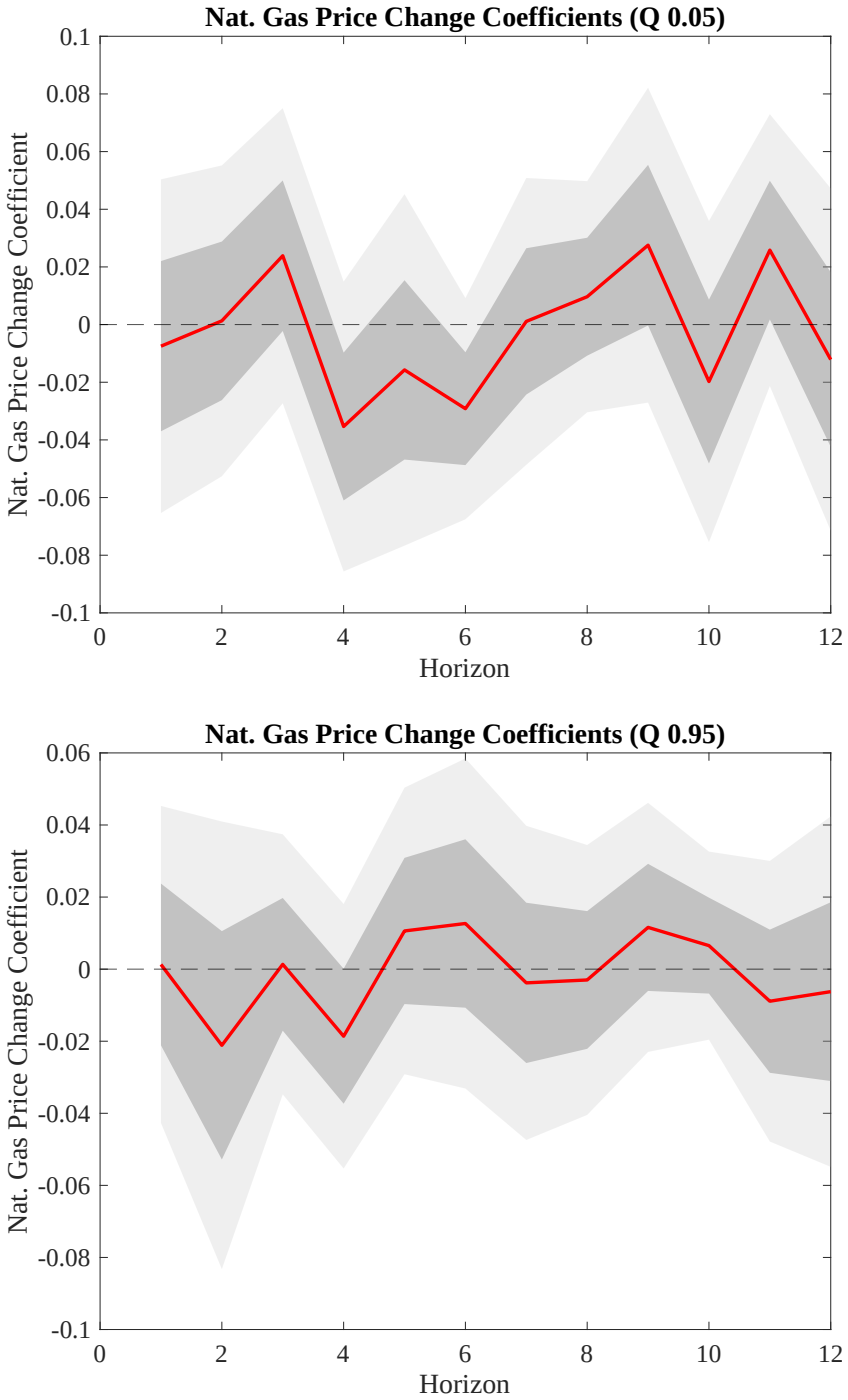


Figure 64: Coefficients of Gas Price Change on the Energy sector's IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

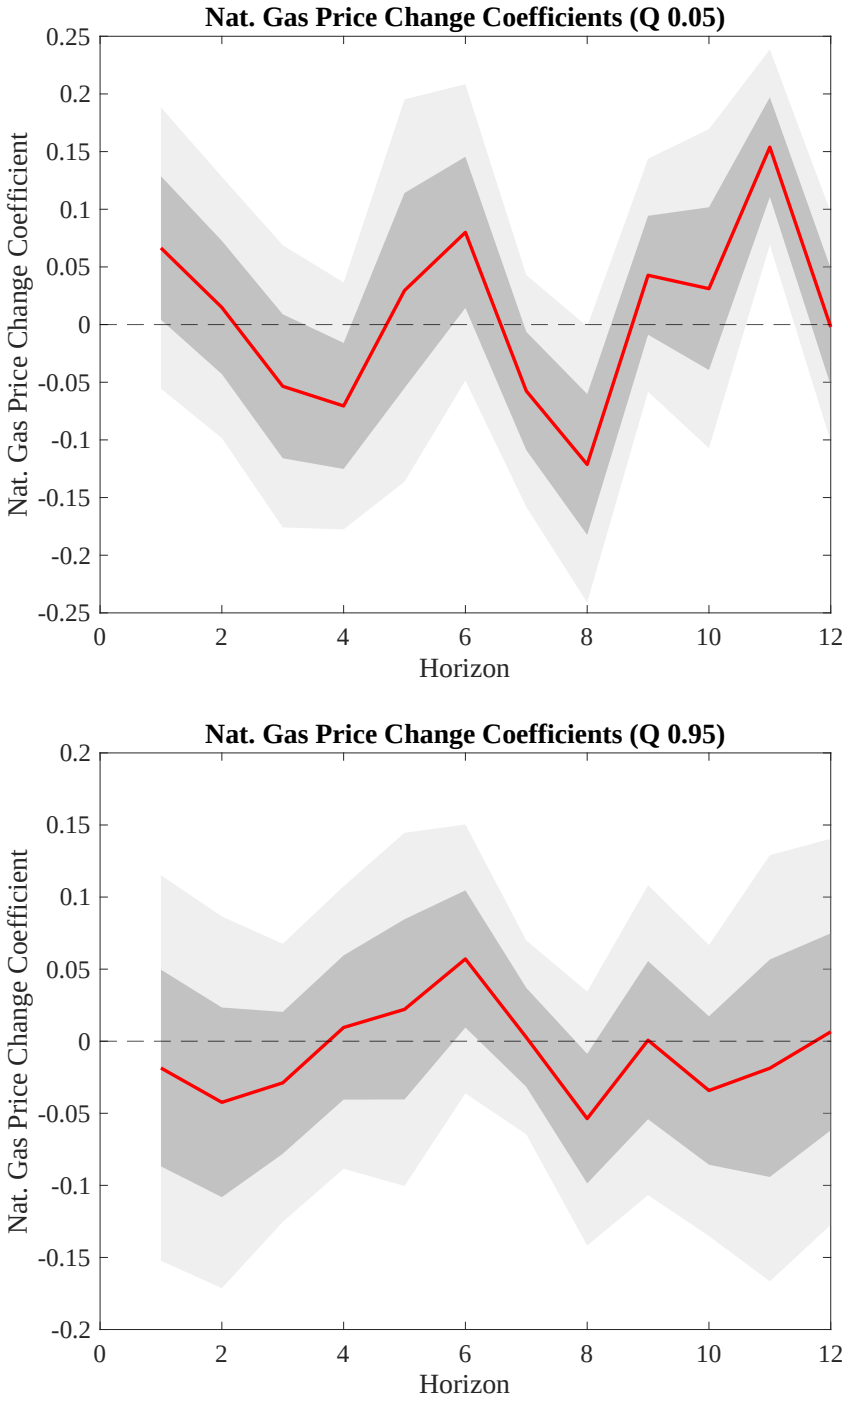
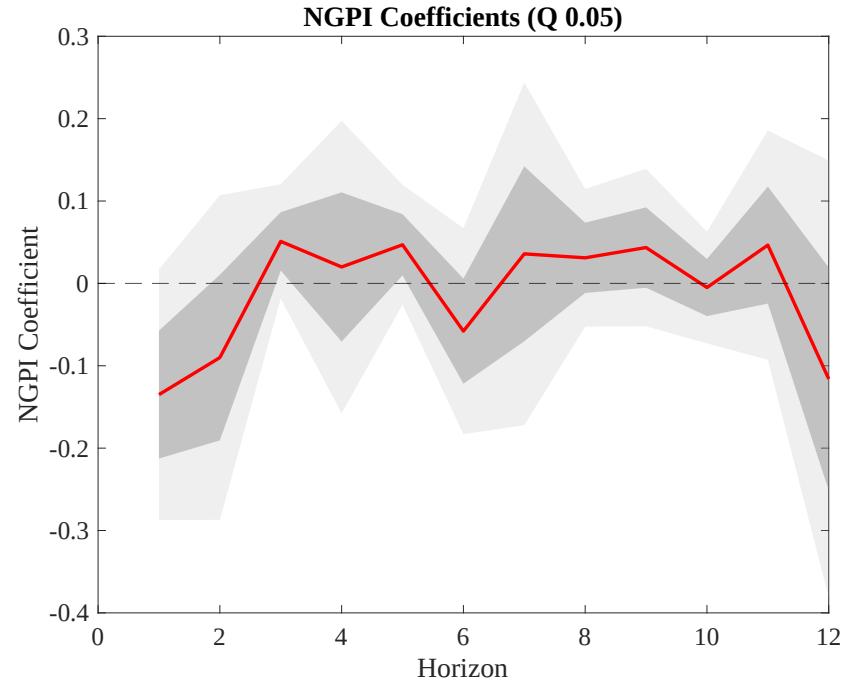


Figure 65: Coefficients of NGPI on the total aggregate IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.



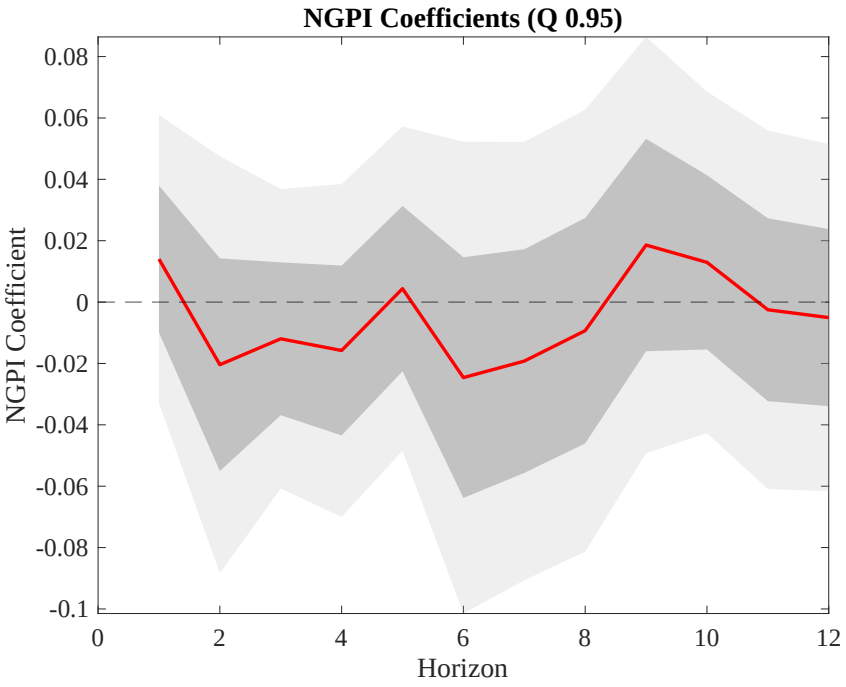


Figure 66: Coefficients of NGPI on the Energy sector's IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

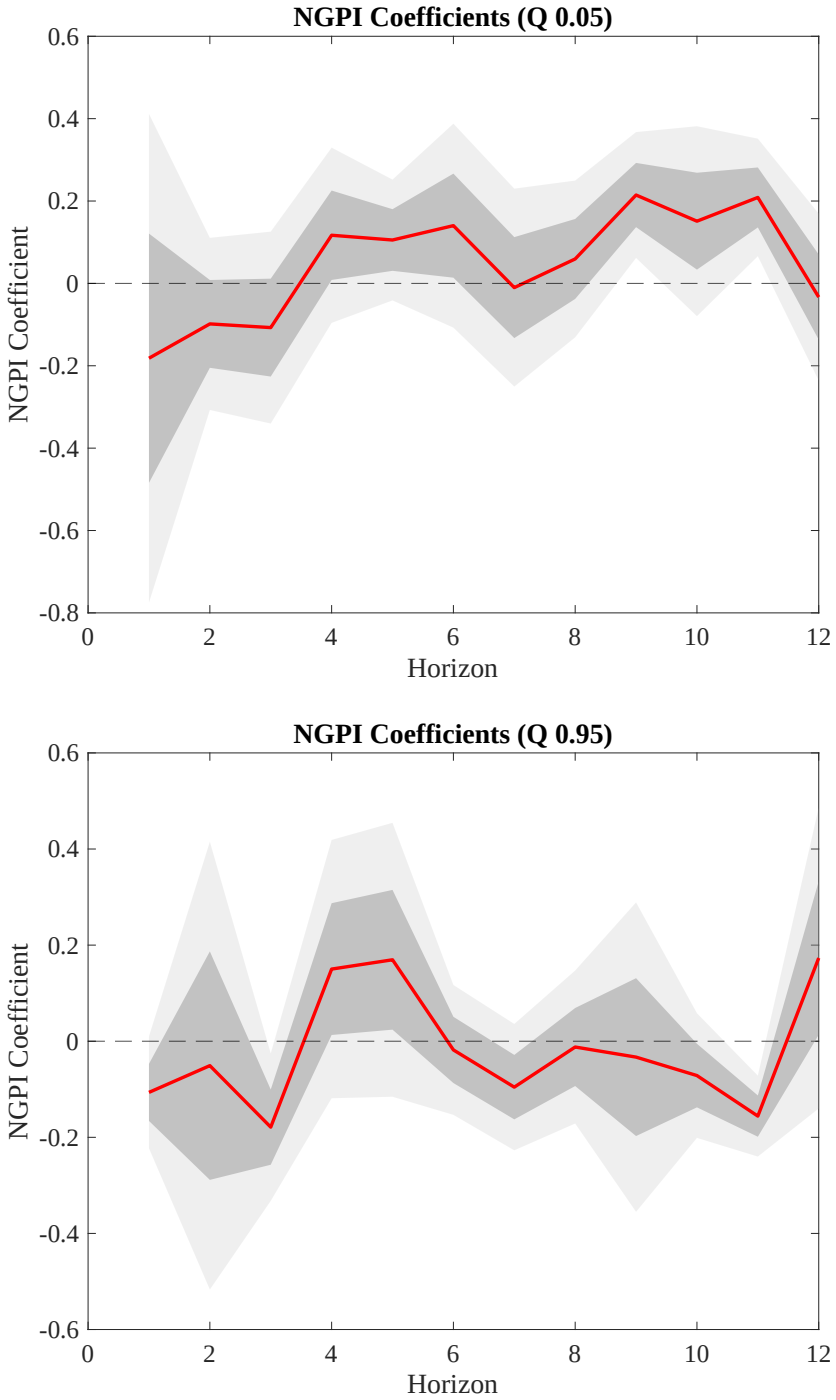


Figure 67: Coefficients of Brent Prices on the Energy sector's IPI Change across the horizons (Q5 and Q95). The bands represent the 95% and 68% confidence interval.

