GEORG-AUGUST-UNIVERSITÄT
GÖTTINGEN

University of Göttingen

Working Paper in Economics

Daily oil price shocks and their uncertainties

Shu Wang

December 2024

No. 436

ISSN: 2705-9141

Platz der Göttinger Sieben 5 | 37073 Göttingen

vwl-verwaltung@uni-goettingen.de | www.uni-goettingen.de/en/60864.html

Daily oil price shocks and their uncertainties

Shu Wang*

December 6, 2024

– Check for update –

Abstract

This paper presents a high-frequency structural VAR framework for identifying oil price shocks and examining their uncertainty transmission in the U.S. macroeconomy and financial markets. Leveraging the stylized features of financial data — specifically, volatility clustering effectively captured by a GARCH model — this approach achieves global identification of shocks while allowing for volatility spillovers across them. Findings reveal that increased variance in aggregate demand shocks increases the oil-equity price covariance, while precautionary demand shocks, triggering heightened investor risk aversion, significantly diminish this covariance. A real-time forecast error variance decomposition further highlights that oil supply uncertainty was the primary source of oil price forecast uncertainty from late March to early May 2020, yet it contributed minimally during the 2022 Russian invasion of Ukraine.

Keywords: Oil price, uncertainty, impulse response functions, structural VAR, forecast error variance decomposition, GARCH.
JEL Classification: Q43, Q47, C32, C58.

*Department of Econometrics, University of Göttingen, Humboldtallee 3, D-37073 Göttingen, Germany. Contact: shu.wang@uni-goettingen.de. The financial support of the Deutsche Forschungsgemeinschaft (HE 2188/17-1) is gratefully acknowledged. I thank Ralf Brüggemann, Christian Hafner, Helmut Herwartz, Lutz Kilian and Helmut Lütkepohl for insightful discussions and helpful comments on an earlier draft of this paper. Any errors and omissions should be regarded as those of the author.

1 Introduction

Structural vector autoregressions (SVARs) have become essential tools for identifying the structural shocks underlying oil price dynamics and examining their broader implications for macroeconomic activity, financial markets, and economic and environmental policies (see Kilian and Zhou 2023 for a literature review).[1] Classical models, incorporating oil production, real crude oil prices, economic activity indicators, and occasionally inventory levels, often relying on low-frequency monthly data, have long served as standard frameworks for empirical analysis and teaching in SVAR modeling (see Kilian and Lütkepohl 2017 for a textbook treatment). Recent work has explored the potential of high-frequency SVARs to capture the oil market's dynamics with greater temporal granularity. For example, Valenti et al. (2023) utilize weekly U.S. crude production and inventory data from the Energy Information Administration (EIA) to investigate short-run fluctuations in oil prices. Extending such analyses to even higher frequencies, such as daily data, is challenging due to the absence of data on key indicators like production, inventories, and aggregate economic activity. Nevertheless, daily observations of crude oil spot and futures prices – given oil's status as one of the most actively traded commodities – offer a promising high-frequency data source. Gazzani et al. (2024) demonstrate that asset prices can provide valuable insights for real-time structural analysis, enabling a refined examination of how different oil price shocks affect the macroeconomy and monetary policy decisions.

This paper contributes to the nascent yet rapidly growing literature on high-frequency SVAR models by presenting an econometric framework that provides granular insights into oil market shocks and their time-varying volatilities as they propagate through the U.S. economy and financial markets. Using a dataset of daily crude oil spot and futures prices alongside stock market indices covering more than 36 years, the model not only sheds light on recent post-pandemic developments but also allows for revisiting key historical episodes, such as the Gulf War in the

[1] Notable contributions to the SVAR literature on oil markets include Kilian (2008, 2009), Kilian and Murphy (2012, 2014), Baumeister and Hamilton (2019), Braun (2023). SVAR models have been employed to analyze the transmission and implications of oil price shocks on U.S. macroeconomic activity (e.g., Herrera and Rangaraju 2020), inflation and inflation expectations (e.g., Kilian and Zhou 2022, Baumeister 2023, Aastveit et al. 2023), equity markets (e.g., Kilian and Park 2009), monetary policy (e.g., Kilian and Lewis 2011), and climate policies (e.g., Herwartz et al. 2024).

early 1990s. A core contribution lies in demonstrating the unique advantages of high-frequency oil market models. While their granularity enables a detailed examination of market dynamics, the distinctive statistical properties of the data – such as volatility clustering and unconditional leptokurtosis, which align with a generalized autoregressive conditional heteroskedasticity (GARCH) structure – offer a valuable source for identification. These characteristics, when combined with the ability to capture volatility spillovers across equity, oil spot and futures markets, can lead to global identification of the shocks. Going beyond identification, high-frequency models provide a unique platform for studying the rich second-order moment dynamics of oil price shocks. To this end, I introduce several structural tools, including covariance impulse response functions (CIRFs), which quantify volatility spillovers across shocks and assess their effects on the variances and covariances of asset prices. Furthermore, recognizing that conditional covariances reflect the accuracy of the optimal predictor (in the L_2 sense) for oil and stock prices, I develop a time-varying forecast error variance decomposition (FEVD), which delivers real-time insights into the evolving sources of forecast uncertainty.

To estimate the model, I propose a quasi maximum likelihood (QML) estimator that is straightforward to implement and does not rely on strong distributional assumptions. I study its asymptotic properties by demonstrating consistency and asymptotic normality under relatively mild assumptions, while simulation exercises indicate its favorable finite-sample performance. Additionally, I discuss several Lagrange-multiplier (LM) and likelihood-ratio (LR) type statistics to test various competing specifications. For the oil market model under consideration, I examine alternative specifications, including a BEKK model and SVARs with differing volatility transmission schemes, and conduct tests for the alternative causal schemes as well as patterns of volatility spillovers. The selected benchmark model incorporates, among other features, the influence of aggregate demand and oil supply uncertainty on the variance of precautionary demand shock, and offers a more nuanced understanding of transmission of oil price shocks and their volatilities.

The findings demonstrate that high-frequency financial data contains relevant information on the dynamics of oil demand and supply, allowing the identification of three distinct oil price shocks: an aggregate demand shock, a precautionary demand shock, and an oil supply shock. While these shocks share many features with their low-frequency counterparts – for example, expectations about future oil demand and

supply can exert significant influence on oil prices and macroeconomic aggregates even in the absence of changes in current production, and oil supply shocks have historically contributed minimally to oil price variations (e.g., Kilian and Murphy 2014, Anzuini et al. 2015) – the high-frequency approach also reveal new properties in terms of transmission of their volatilities. Specifically, aggregate demand uncertainty emerges as the primary driver of equity market volatility, whereas oil supply uncertainty predominantly affects spot market volatility, and precautionary demand uncertainty exerts the strongest influence on futures markets. A variance increase in aggregate demand shocks raises the covariance between stock and oil prices, while precautionary demand shocks, which reflect heightened economic uncertainty and anticipated supply risks unrelated to current production, lower this covariance, reflecting a strong risk aversion. Precautionary demand shocks cause spot and futures prices to rise in tandem, similar to Gazzani et al. (2024), and significantly tighten financial conditions. The excess bond premium (EBP) from Gilchrist and Zakrajšek (2012) rises markedly, the global factor in risky asset prices (Miranda-Agrippino and Rey 2020) declines, and measures such as the VIX and bond spreads increase. Heightened risk aversion also leads to stock market declines, a rise in gold futures, and falling Treasury yields. By contrast, an unexpected oil supply shortfall triggers an immediate decline and gradual recovery in global oil production, pushing spot prices more than 1.5% above baseline, while futures prices adjust more modestly and with delay, resulting in a sharp reduction of the futures-spot spread. Examining low-frequency variables further reveals that inventory drawdowns help smooth demand and raise convenience yields, which raises the value of holding physical oil over futures, in line with the results found in Valenti (2022).

The high-frequency structural model also provides an exceptionally granular perspective on sources of uncertainty perceived by markets as they compile information to forecast oil and asset prices. Notably, the composition of forecast uncertainty shows sizable time variation, often diverging markedly from results based on conventional FEVD using shocks' unconditional variances. For example, in early 2020, oil supply shocks contributed minimally to the one-day-ahead FEV of oil prices (Figure 1.1). Instead, demand shocks predominantly drove forecast uncertainty, with aggregate demand shocks temporarily accounting for approximately 44% of forecast variance on key dates – such as Centers for Disease Control and Prevention's (CDC) Covid-19 warnings and subsequent Covid policy announcements – over three times

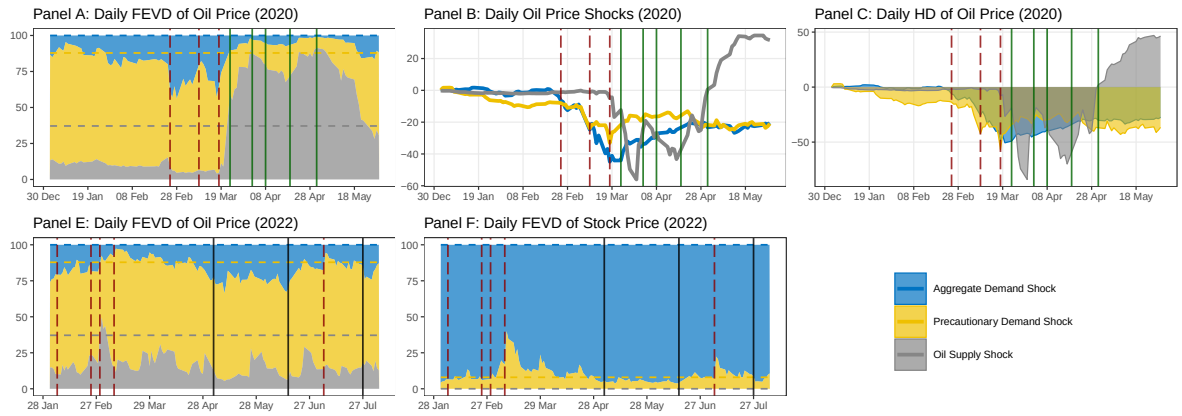

Figure 1.1: *Daily FEVD and HD for oil and stock price dynamics. First row: Panel A shows the FEVD of crude oil spot prices, while Panel B presents conditional heteroskedastic oil price shocks (cumulated) during early 2020. Panel C displays the HD for oil prices. Key events are marked by dashed and solid lines: red dashed lines indicate CDC's initial Covid-19 warning (Feb. 25), the Black Monday crash (Mar. 9), and major Covid-policy announcements (Mar. 18). Green solid lines denote significant oil market events, including intensification of the Russia-Saudi price war (Mar. 23), and anticipation of and agreements on the historic production cut (Apr. 2, 8, 19) effective on May 1. Second row: Panels E and F display FEVDs of oil and stock prices during the first half of 2022. Red dashed lines mark anticipation and onset of the Russian invasion of Ukraine (Feb. 5 and 24), U.S. legislative and executive actions banning Russian oil (Mar. 1 and 8), and Chinese lockdown news fueling recession concerns (Jul. 5). Black solid lines indicate dates of FOMC's announcements of rate hikes. Dashed horizontal lines in FEVD plots indicate the unconditional one-day-ahead FEV contributions of each shock.*

their unconditional level. However, from late March onward, amid the Russia-Saudi price war and rounds of negotiations culminating in a historic production cut effective May 1, the oil supply shock emerged as the dominant source of forecast uncertainty, at times contributing as much as 91%. In contrast, during the 2022 Russian invasion of Ukraine, oil supply shocks only modestly contributed to forecast uncertainty for oil prices, except in response to U.S. legislative actions banning Russian oil imports shortly after the invasion.

Related methodological literature Identification strategies in oil market SVARs have traditionally drawn on theoretical insights on the short-run supply and demand curves, particularly assumptions about their slopes (i.e., elasticities), often using zero or sign restrictions (e.g., Kilian 2009, Kilian and Murphy 2014, Baumeister and Hamilton 2019, Braun and Brüggemann 2023). Building on the literature on proxy SVAR and high-frequency identification, originally developed for monetary policy analysis, Känzig (2021) constructs a proxy for oil supply news shock by employing changes in oil futures prices around OPEC announcements (see also Kilian 2024 for discussion and references therein). Besides theory- and instrument-

based approaches, statistical identification methods leveraging heteroskedasticity (e.g., Lütkepohl and Netšunajev 2014, Bertsche and Braun 2022) and independence under non-Gaussianity (e.g., Braun 2023, Hafner et al. 2024) have also been used to identify monthly oil price shocks. Valenti et al. (2023) pursue set-identification in their weekly model using exclusion and sign restrictions following Baumeister and Hamilton (2019). In the daily oil market model, Gazzani et al. (2024) adopt a mixed identification strategy, using sign and magnitude restrictions based on observed correlation patterns between oil and stock prices combined with two narrative restrictions. They also construct a proxy capturing periods of marked oil price declines alongside spikes in market volatility, using indices such as the VIX and gold prices, to identify a forward-looking demand shock aimed at capturing shifts in expectations and uncertainty about future global demand.

Although exploiting conditional heteroskedasticity for identification appears especially well-suited for daily oil market models, this approach has not yet been previously explored. It should be noted, however, that using conditional heteroskedasticity for identification purposes has precedence dating back to King et al. (1994). Similarly, Sentana and Fiorentini (2001) and Lanne and Saikkonen (2007) developed identification approaches for factor models, though applications within SVAR frameworks are less common (exceptions are Normadin and Phaneuf 2004, Bouakez and Normandin 2010). Unlike these studies, this paper's approach relaxes the assumption of a diagonal GARCH structure, allowing volatility spillovers across shocks. Incorporating such spillovers is particularly pertinent here, as I find substantial evidence that uncertainty surrounding both aggregate demand and oil supply significantly influence the volatility of the precautionary demand shock. Surprisingly, relaxing the diagonal assumption leads to global identification, where permutations of columns in the structural impact multiplier or the rows/columns in the parameter matrices governing the volatility dynamics are no longer observationally equivalent.

This paper also speaks to the literature on structural multivariate GARCH (SM-GARCH) models, from which the concept of volatility IRFs originates (Hafner and Herwartz 2006, 2023a,b, Fengler and Polivka 2024).[2] Unlike SM-GARCH mod-

[2]Hafner and Herwartz (2006) pioneered the development of generalized IRFs for multivariate volatility and provided a foundational analysis of their properties. Hafner and Herwartz (2023a) extended this framework to accommodate asymmetric SM-GARCH models, while Hafner and Herwartz (2023b) proposed simulation techniques for computing correlation IRFs. Recently, Fengler and

els, which typically rely on higher-moment independence under non-Gaussianity (Hafner et al. 2022) or proxy variables (Fengler and Polivka 2024) for identification, the present model leverages conditional heteroskedasticity alone for identification. Moreover, while SM-GARCH models focus on modeling the covariance dynamics of observable variables, the GARCH structure here directly governs the law of motion of the shocks, endowing the model parameters with structural interpretations.

Outline and notations The paper proceeds as follows. Section 2 introduces the econometric framework for high-frequency oil market models, and Section 3 presents empirical results. Section 4 concludes. Appendix A contains proofs of main results and the online appendix (OA) provides additional proofs and lemmas, simulation details and further empirical results including model diagnosis and robustness check. Throughout, $\|x\|$ denotes the Euclidean norm for vectors, and $\|A\|$ denotes the operator norm for bounded linear operators, with $\|A\|^2 = \rho(A'A)$, where $\rho(.)$ is the spectral radius. For a square matrix A, $|A|$ represents its determinant. I use $\mathcal{B}(x, \delta)$ to indicate an open ball centered at x with radius $\delta > 0$. Conditional expectations are defined almost surely (a.s.), $\odot$ denotes the Hadamard product, $\mathbf{1}_N$ is a $N \times 1$ vector of ones.

2 Econometrics framework

2.1 The dynamic model

The dynamic model for a vector of endogenous variables $y_t \in \mathbb{R}^N$ is a finite order VAR(P) process. Let $\{\mathcal{F}_t\}_{t\in\mathbb{N}}$ denote a sequence of increasing σ-fields generated by $\{y_s : s \leq t\}$ and let L represent the lag-operator. Apart from deterministic and exogenous terms, the structural form is given by

$$A(L)y_t = \left(I_N - \sum_{i=1}^{P} A_i L^i\right) y_t = B\xi_t, \text{ with } \mathbb{E}[\xi_t|\mathcal{F}_{t-1}] = 0, \ \mathbb{E}[\xi_t\xi_t'|\mathcal{F}_{t-1}] = \Sigma_t, \tag{2.1}$$

where the structural shocks in vector ξ_t, which form a martingale difference sequence (MDS), are mapped to the observable system via the non-singular matrix B and have a diagonal conditional covariance matrix Σ_t with diagonal elements $\sigma_t = (\sigma_{1t}, \ldots, \sigma_{Nt})'$. Define the reduced-form residuals as $u_t = A(L)y_t = B\xi_t$. Under the normalization $\mathbb{E}[\Sigma_t] = I_N$, the vector u_t has zero mean and a non-diagonal

Polivka (2024) investigate asymptotic properties of volatility IRFs.

unconditional covariance matrix[3]

$$\Omega := \mathbb{E}[u_t u_t'] = B\mathbb{E}[\mathbb{E}[\xi_t \xi_t' | \mathcal{F}_{t-1}]]B' = B\mathbb{E}[\Sigma_t]B' = BB'. \tag{2.2}$$

Suppose y_t is causal and second-order stationary (with precise assumptions specified later) with invertible reverse characteristic polynomial, i.e., $|A(z)| \neq 0$ for $|z| \leq 1$, it has a Wold moving average (MA) representation:

$$y_t = \Theta(L)\xi_t = \sum_{i=0}^{\infty} \Theta_i \xi_{t-i}, \text{ where } \Theta(L) = A(1)^{-1}B. \tag{2.3}$$

The structural MA representation is particularly interesting, since it traces the dynamic effects of structural shocks on the conditional expectation of the observable variables through the impulse response functions (IRFs). To distinguish this concept from the covariance impulse response functions (CIRFs) introduced later, I refer to these as mean IRFs (MIRFs), defined as follows:

Definition 2.1. For a sequence of random vectors $\{X_{t+h}\}_{h\geq 0}$, integrable and measurable with respect to $\mathcal{F}_{t-1}$, the MIRFs are defined as

$$\mathcal{M}^X_{t+h}(\xi_t | \mathcal{F}_{t-1}) := \mathbb{E}[X_{t+h} | \mathcal{F}_{t-1}, \xi_t] - \mathbb{E}[X_{t+h} | \mathcal{F}_{t-1}], \quad \xi_t \in \mathbb{R}^N, \ h \in \mathbb{N}.$$

MIRFs describe the path of X_{t+h} expected to occur given its history, after a shock at time point t, relative to the potential path, where X_t is already at its 'steady state' absent any current or future shock. For y_t with an MA representation (2.3), the MIRF at horizon $h \in \mathbb{N}$ is given directly by the coefficient Θ_h:

$$\mathcal{M}^y_{t+h}(\xi^* | \mathcal{F}_{t-1}) = \Theta_h \xi^*, \ \forall t. \tag{2.4}$$

Here, $\xi^* \in \mathbb{R}^N$ is a particular point in the support of the distribution of ξ_t that represents the 'dose' of the shock. A common choice of the dose is $\xi^* := e_j$ with e_j being the j-th column of the identity matrix I_N, so the j-th column of Θ_h represents the effect of a 'unit' shock to the j-th structural shock, with other shocks being muted.

Notably, the MIRFs depend on both the reduced-form parameters A_j and the structural parameters in B, with $\Theta_0 = B$. Thus, the impact multipliers in B describe the

[3]The imposition of unit variance for the latent shocks is made to normalize the size of the shocks and is applied without any loss of generality. Alternative normalizations have also been proposed. For example, one approach involves restricting the diagonal elements of matrix B to be unity, denoted as $B^\dagger$, while allowing $\mathbb{E}[\Sigma_t] = \Sigma$ to have flexible diagonal elements. In this case, $B = B^\dagger \Sigma^{1/2}$ (see Gouriéroux et al. 2017 and references therein for a discussion).

contemporaneous response of y_t to an exogenous change in ξ_t. However, it is well known that matrix B cannot be identified from the unconditional second-moment condition (2.2). Specifically, any $B^* = BQ$, where Q is an orthogonal matrix, satisfies the condition $\Omega = B^*B^{*\prime}$ equally wel.

Identifications for low-frequency oil market models often build upon economic theory and institutional knowledge on price elasticity and various motives for holding oil, which guide zero and sign restrictions, or construction of proxies. Meanwhile, statistical identification approaches that exploit dependence of higher-order moments or time-varying volatilities have yielded valuable insights. Despite these advances, statistical identification of high-frequency models remains underexplored. Modeling daily oil shocks as GARCH processes is particularly appealing, as high-frequency financial asset prices exhibit well-documented empirical features (correlation of squared-returns, volatility clustering, unconditional leptokurtosis) that align well with a GARCH model. I outline the specification of GARCH shocks next.

2.2 Conditional heteroskedasticity

A strong GARCH(p, q) specification for the marginal variance processes of the structural shocks is defined by

$$\xi_t = \Sigma_t^{1/2}\eta_t,\ \eta_t \overset{iid}{\sim} (0, I_N) \tag{2.5}$$

$$\sigma_{k,t} = \gamma_{k0} + \sum_{i=1}^{q}\sum_{n=1}^{N} G_{kn,i}\xi_{n,t-i}^2 + \sum_{j=1}^{p}\sum_{n=1}^{N}\Gamma_{kn,j}\sigma_{n,t-j},\ k = 1, \dots, N, \tag{2.6}$$

where $\gamma_{k0} = 1 - \sum_{i=1}^{q}\sum_{n=1}^{N} G_{kn,i} - \sum_{j=1}^{p}\sum_{n=1}^{N}\Gamma_{kn,j}$, and $G_{kn,i}, \Gamma_{kn,j} \geq 0$ for $k, n \in \{1, \dots, N\}$, $i, j \in \{1, \dots, \max(p, q)\}$. These parameter constraints ensure that the variances remain positive a.s. under suitable initial conditions, and satisfy the unconditional normalization $\mathbb{E}[\sigma_t] = \mathbf{1}_N$ imposed earlier. Combined with the conditional mean equation (2.1), the GARCH innovations imply that the conditional covariance of the observable variables is $H_t := \mathrm{Cov}[y_t|\mathcal{F}_{t-1}] = B\Sigma_t B'$. Focusing on a GARCH(1,1) specification, the conditional variance process can be compactly written as

$$\sigma_t = \gamma_0 + G_1(\xi_{t-1} \odot \xi_{t-1}) + \Gamma_1\sigma_{t-1}, \tag{2.7}$$

where $\gamma_0 = (I_N - G_1 - \Gamma_1)\mathbf{1}_N$. By allowing non-zero off-diagonal elements in the matrices G_1 and Γ_1, specification (2.7) captures not only volatility clustering within

individual shocks but also spillovers across distinct shocks. If there is sufficient heteroskedasticity in the latent shocks, provided that the $T \times N$ matrix $\boldsymbol{S}_T = (\sigma_1, \ldots, \sigma_T)'$ of sample variance paths has full rank, the matrix B is identified up to column permutations and sign changes:

Lemma 2.2. Let $\mathcal{P}$ be a permutation matrix and $\mathcal{D}$ a diagonal matrix with diagonal elements taken from $\{1, -1\}$. If rank $\boldsymbol{S}_T = N$, then the structural impact multiplier matrix B in (2.1) is identified up to the set $\{B^* : B^* = B\mathcal{P}\mathcal{D}\}$ and coefficient matrices G_i, Γ_j, in (2.6) are identified up to the set $\{G_i^* : G_i^* = \mathcal{P}'G_i\mathcal{P}\}$ and $\{\Gamma_j^* : \Gamma_j^* = \mathcal{P}'\Gamma_j\mathcal{P}\}$ with $i \in \{1, \ldots, q\}$, $j \in \{1, \ldots, p\}$, respectively.

Evidently, imposing a diagonal structure on G_i and Γ_j, as in Normadin and Phaneuf (2004), Lanne and Saikkonen (2007), and Bouakez and Normandin (2010), precludes volatility spillovers across shocks and leads only to local identification of both B and the GARCH parameters themselves. By contrast, allowing for non-zero off-diagonal elements can result in global identification, with the only permissible permutation matrix $\mathcal{P}$ in Lemma 2.2 being the identity matrix. A another key distinction of the Lemma from previous results (see, e.g., Sentana and Fiorentini 2001) lies in the structural identification of and interpretation afforded to the parameters governing the law of motion of the variance processes. While many existing specifications of time-varying variances, including non-parametric approaches (see Lewis 2021), offer distinct advantages in flexibility and generality in terms of identification, variance parameters are often treated as nuisance elements. Yet, scrutinizing these parameters can yield insights into the rich second-order dynamics of the shocks and facilitate closed-form structural tools for analyzing volatility transmission patterns and decomposing forecast uncertainty. I next introduce two structural devices for studying volatility transmission in the model.

2.3 Structural analysis of uncertainty transmission

Covariance impulse response functions To quantify the effect of an innovation in the shock variance on the subsequent covariances of a random vector of interest, such as the structural shocks ξ_t or the observable variables y_t, I define the covariance impulse response functions (CIRFs) as follows:

Definition 2.3. For a sequence of random vectors $\{X_{t+h}\}_{h\geq 1}$, square-integrable and

measurable with respect to $\mathcal{F}_{t-1}$, the CIRFs are defined as

$$\mathcal{V}^X_{t+h}(\eta_t|\mathcal{F}_{t-1}) := \text{Cov}[X_{t+h}|\mathcal{F}_{t-1}, \eta_t] - \text{Cov}[X_{t+h}|\mathcal{F}_{t-1}], \quad \eta_t \in \mathbb{R}^N, \ h \in \mathbb{N}^+.$$

Analogous to the MIRFs, CIRFs measure the difference, conditional on the past, between a shocked path of covariance and its potential path, where the covariance remains at its 'steady state' in the absence of any current or future volatility shocks.[4]

Theorem 2.4. For y_t with an MA representation (2.3) and structural shock variances following the GARCH process (2.7), the CIRF of ξ_t at horizon $h \in \mathbb{N}^+$ is given by

$$\mathcal{V}^{\xi}_{t+h}(\eta^*|\mathcal{F}_{t-1}) = \left((G_1 + \Gamma_1)^{h-1} G_1 \Sigma_t [(\eta^* \odot \eta^*) - \mathbf{1}_N]\mathbf{1}'_N\right) \odot I_N, \tag{2.8}$$

where $\eta^* \in \mathbb{R}^N$ is a point in the support of distribution of η_t. By setting $\mathcal{V}^{\xi}_t(\eta^*|\mathcal{F}_{t-1}) := 0$, the CIRF of y_t can be expressed as

$$\mathcal{V}^{y}_{t+h}(\eta^*|\mathcal{F}_{t-1}) = \sum_{i=0}^{h-1} \Theta_i \mathcal{V}^{\xi}_{t+h-i}(\eta^*|\mathcal{F}_{t-1})\Theta'_i. \tag{2.9}$$

Noting the elementary relation $\Sigma_t = (\sigma_t \mathbf{1}'_N) \odot I_N$, the CIRFs of the shocks themselves take a familiar form as in Hafner and Herwartz (2006). Specifically, at horizon $h = 1$,

$$\mathcal{V}^{\xi}_{t+1}(\eta^*|\mathcal{F}_{t-1}) = \left(\mathbb{E}\left[\frac{\partial \sigma_{t+1}}{\partial(\eta_t \odot \eta_t)}\middle|\mathcal{F}_{t-1}, \eta_t = \eta^*\right]\mathbf{1}'_N\right) \odot I_N, \tag{2.10}$$

while CIRFs at higher horizons can be computed recursively. For the specification in (2.7), the conditional expectation in (2.10) becomes $G_1\Sigma_t[(\eta^* \odot \eta^*) - \mathbf{1}_N]$ and for $h > 1$, with diag(.) being a vector of diagonal elements of a square matrix,

$$\text{diag}\left(\mathcal{V}^{\xi}_{t+h}(\eta^*|\mathcal{F}_{t-1})\right) = (G_1 + \Gamma_1)\,\text{diag}\left(\mathcal{V}^{\xi}_{t+h-1}(\eta^*|\mathcal{F}_{t-1})\right).$$

Detailed derivations are provided in the proof. The second part of the Theorem extends their results to a vector MA process. Notably, unlike MIRFs (2.4), which are time-invariant, the CIRFs in (2.9) depend on t through the initial conditional variance of the underlying shocks Σ_t. The deviation between the shocked and potential paths arises if the squared variance innovation $\eta_t \odot \eta_t$ diverges from its unconditional

[4] In our framework of conditional heteroskedasticity, the shocked covariance paths for ξ_t at horizon $h = 0$ all begin at zero for any η_t, as $\xi_t|\mathcal{F}_{t-1}, \eta_t = \Sigma_t^{1/2}\eta_t$ a.s. This complicates the interpretation of the difference $-\Sigma_t$. Therefore, $h = 0$ is excluded from the definition, and where appropriate, we set $\mathcal{V}^{\xi}_t(\eta_t|\mathcal{F}_{t-1}) = 0$ for all η_t.

expectation $\mathbf{1}_N$. A common approach is to set the shock dose at a quantile, say 99%, from the empirical distribution of η_t (see, e.g., Hafner and Herwartz 2023a). Evidently, the persistence of the covariance response is governed by the spectral norm $\rho(G_1+\Gamma_1)$, and under stationarity assumptions (to be detailed later), the CIRFs decay exponentially. Importantly, if both G_1 and Γ_1 were restricted to be diagonal, a shock in the variance of the j-th component, with others fixed at unity, would only induce responses in the j-th diagonal element of $\mathcal{V}_{t+1}^{\xi}$. Variances of other shocks will remain unaffected across all horizons.

Forecast error variance decomposition The least MSE h-step forecast ($h \geq 1$) given information available at time t is the conditional expectation $y_{t,h} = \mathbb{E}[y_{t+h}|\mathcal{F}_t]$. Let $u_{t,h} = y_{t+h} - y_{t,h} = \sum_{i=0}^{h-1} \Theta_i \xi_{t+h-i}$ denote the h-step forecast error, the forecast error covariance is given by

$$\mathrm{Cov}[u_{t,h}|\mathcal{F}_t] = \sum_{i=0}^{h-1} \Theta_i \mathbb{E}[\xi_{t+h-i}\xi_{t+h-i}'|\mathcal{F}_t]\Theta_i' = \sum_{i=0}^{h-1} \Theta_i \mathbb{E}[\Sigma_{t+h-i}|\mathcal{F}_t]\Theta_i',$$

where the last equality holds since, for any $h \geq 1$, $\mathcal{F}_t$ is a sub-$\boldsymbol{\sigma}$-field of $\mathcal{F}_{t+h-1}$, thus

$$\mathbb{E}[\xi_{t+h}\xi_{t+h}'|\mathcal{F}_t] = \mathbb{E}[\mathbb{E}[\xi_{t+h}\xi_{t+h}'|\mathcal{F}_t, \mathcal{F}_{t+h-1}]|\mathcal{F}_t] = \mathbb{E}[\mathbb{E}[\xi_{t+h}\xi_{t+h}'|\mathcal{F}_{t+h-1}]|\mathcal{F}_t] = \mathbb{E}[\Sigma_{t+h}|\mathcal{F}_t].$$

As shown in the proof of Theorem 2.5, the h-step prediction of structural shock variances follows the recursive relation

$$\mathbb{E}[\sigma_{t+h}|\mathcal{F}_t] = \gamma_0 + (G_1+\Gamma_1)\mathbb{E}[\sigma_{t+h-1}|\mathcal{F}_t] = \mathbf{1}_N + (G_1+\Gamma_1)^{h-1}(\sigma_{t+1} - \mathbf{1}_N), \tag{2.11}$$

where $\sigma_{t+1} = \mathbb{E}[\sigma_{t+1}|\mathcal{F}_t] = \gamma_0 + G_1(\xi_t \odot \xi_t) + \Gamma_1\sigma_t$. With $\rho(G_1+\Gamma_1) < 1$, the long-term forecast error variance converges to the unconditional variance, $\lim_{h\to\infty} \mathbb{E}[\sigma_{t+h}|\mathcal{F}_t] = \mathbf{1}_N$. We can now derive the time-varying FEVD as follows:

Theorem 2.5. For y_t with an MA representation (2.3) and structural shock variances following the GARCH process (2.7), the contribution of the j-th shock to the h-step forecast error variance of the k-th variable, with $h \in \mathbb{N}^+$, is given by

$$\frac{\sum_{i=0}^{h-1} \mathbb{E}[\sigma_{j,t+h-i}|\mathcal{F}_t] e_k' \Theta_i e_j}{\sum_{j=1}^{N}\sum_{i=0}^{h-1} \mathbb{E}[\sigma_{j,t+h-i}|\mathcal{F}_t] e_k' \Theta_i e_j}, \tag{2.12}$$

where $\mathbb{E}[\sigma_{j,t+h-i}|\mathcal{F}_t] = e_j'\mathbb{E}[\sigma_{t+h-i}|\mathcal{F}_t]$ and for $i = 0, \ldots, h-1$,

$$\mathbb{E}[\sigma_{t+h-i}|\mathcal{F}_t] = \mathbf{1}_N + (G_1+\Gamma_1)^{h-i-1}(\sigma_{t+1} - \mathbf{1}_N),$$

with $\sigma_{t+1} = \gamma_0 + G_1(\xi_t \odot \xi_t) + \Gamma_1\sigma_t$.

Since all elements in G_1 and Γ_1 are non-negative, the predictor variance positively depends on the difference between the conditional variance of the shocks at the forecast origin, σ_{t+1}, and their unconditional variances, $\mathbf{1}_N$. When the conditional variance of the j-th shock at the forecast origin substantially exceeds one, this shock becomes the primary source of forecast uncertainty; conversely, a shock with a negative difference improves predictive accuracy. Thus, this real-time decomposition identifies the key drivers of prediction error variance at any given point, offering insights into what the market perceives as sources of uncertainty when assimilating information to forecast oil and asset prices. Next, I describe estimation of the model.

2.4 Quasi maximum likelihood estimation

Let $\vartheta = (\boldsymbol{\alpha}', \boldsymbol{\beta}', \phi')' \in \Theta$ denote the parameter vector, where $\boldsymbol{\alpha} = \text{vec}(\boldsymbol{A})$ with $\boldsymbol{A} = (A_1, \ldots, A_P)$, $\boldsymbol{\beta} = \text{vec}(B)$, $\phi = \text{vec}\left(G_1, \ldots, G_q, \Gamma_1, \ldots, \Gamma_p\right)$. Define $\boldsymbol{y} = \text{vec}(y_1, \ldots, y_T)$, $Z_t = \text{vec}(y_t, \ldots, y_{t-P+1})$, $Z = (Z_0, \ldots, Z_{T-1})$. One way to jointly estimate the reduced-form AR coefficients $\boldsymbol{\alpha}$, structural parameters $\boldsymbol{\beta}$, and GARCH parameters ϕ is to maximize the quasi log-likelihood function based on a Gaussian conditional density

$$f(y_t|\vartheta, \mathcal{F}_{t-1}) = (2\pi)^{-N/2}|H_t(\boldsymbol{\beta}, \phi)|^{-1/2} \exp\left(-\frac{1}{2}u_t'(\boldsymbol{\alpha})H_t^{-1}(\boldsymbol{\beta}, \phi)u_t(\boldsymbol{\alpha}))\right),$$

where $u_t(\boldsymbol{\alpha}) = y_t - (Z_{t-1}' \otimes I_N)\boldsymbol{\alpha}$ and $H_t(\boldsymbol{\beta}, \phi) = B(\boldsymbol{\beta})\Sigma_t(\phi)B(\boldsymbol{\beta})'$. The one-step QML estimator, obtained as $\hat{\vartheta}_T^{\dagger} := \text{argmax}_{\vartheta\in\Theta} \frac{1}{T}\sum_{t=1}^{T} l_t^{\dagger}(\vartheta)$, with $l_t^{\dagger}(\vartheta) = \log f(y_t|\vartheta, \mathcal{F}_{t-1})$, is consistent and asymptotically normal. If η_t is Gaussian, its asymptotic covariance will attain the parametric lower bound. However, an important drawback is that the scale of the structural shocks is unidentified. See Theorem B.1 in OA B for further details.[5] Because the scales of the shocks play a crucial role in analyzing volatility transmission, I propose a two-step estimation procedure where the scale is identified up to a sign flip. The ambiguity of the sign can be resolved based on whether the analyst is interested in, for instance, an unexpected increase or decrease in oil supply.

Using the parameterization $B = \Omega^{1/2}(\boldsymbol{\omega})Q(\varrho)$ with $\Omega^{1/2}$ obtained by spectral decomposition, $\boldsymbol{\omega} = \text{vech}(\Omega)$ and $Q(\varrho)$ as the product of a sequence of Givens rotation matrices, it is ensured that $B(\boldsymbol{\omega}, \varrho)$ is non-singular and satisfies the unconditional

[5]Specifically, denote $W = B^{-1}$ and $w_j = e_j'W$. There exists a non-zero constant $c \in \mathbb{R}$ and $\boldsymbol{\beta}^*$, ϕ^* such that $w_j(\boldsymbol{\beta}^*) = cw_j(\boldsymbol{\beta})$ and $\sigma_{jt}(\phi^*) = c^2\sigma_{jt}(\phi)$ for some $j \in \{1, \ldots, N\}$, and $\vartheta^* = (\boldsymbol{\alpha}', \boldsymbol{\beta}^{*\prime}, \phi^{*\prime})'$ has the same likelihood as ϑ.

moment condition in (2.2), as $|Q(\varrho)| = \|Q(\varrho)\| = 1$, $\forall \varrho$ (see OA C for further details).[6] Define $\vartheta = (\vartheta_r', \vartheta_s')'$, where $\vartheta_r = (\boldsymbol{\alpha}', \boldsymbol{\omega}')'$ and $\vartheta_s = (\varrho', \phi')'$ contains the reduced-form, structural and GARCH parameters, respectively. The first-step estimator of the reduced-form parameters is obtained by solving the optimization problem $\hat{\vartheta}_{r,T} = \text{argmax}_{\vartheta_r \in \Theta_r} \frac{1}{T} \sum_{t=1}^{T} g_t(\vartheta_r)$, where

$$g_t(\vartheta_r) = -\left(\log|\Omega| + u_t'(\boldsymbol{\alpha})\Omega^{-1}u_t(\boldsymbol{\alpha})\right).$$

In the second step, ϑ_s is estimated by QML (conditional on fixed initial values) using the log-likelihood function given by $\mathcal{L}_T(\vartheta_s|\vartheta_r) = -\frac{1}{T}\sum_{t=1}^{T} l_t(\vartheta_s|\vartheta_r)$, where

$$l_t(\vartheta_s|\vartheta_r) = \log|\Sigma_t(\phi)| + \varepsilon_t'(\vartheta_r)Q(\varrho)\Sigma_t^{-1}(\phi)Q(\varrho)'\varepsilon_t(\vartheta_r), \tag{2.13}$$

with $\varepsilon_t(\vartheta_r) = \Omega^{-1/2}(\omega)(y_t - (Z_{t-1}' \otimes I_N)\boldsymbol{\alpha})$. The first-step estimator is the standard ML estimator, with the familiar closed-form:

$$\hat{\boldsymbol{\alpha}}_T = ((ZZ')^{-1}Z \otimes I_N)\boldsymbol{y}, \quad \hat{\Omega}_T = \frac{1}{T}\sum_{t=1}^{T}(y_t - (Z_{t-1}' \otimes I_N)\hat{\boldsymbol{\alpha}}_T)(y_t - (Z_{t-1}' \otimes I_N)\hat{\boldsymbol{\alpha}}_T)', \tag{2.14}$$

which is readily implemented in common software packages, making the proposed procedure adaptable to a wide range of applications and thus particularly appealing to practitioners. Nonetheless, ignoring conditional heteroskedasticity does reduce the first-step estimator's efficiency. The two-step estimator is thus defined as $\hat{\vartheta}_T = (\hat{\vartheta}_{r,T}'\hat{\vartheta}_{s,T}')'$ with $\hat{\vartheta}_{r,T}$ given by (2.14) and $\hat{\vartheta}_{s,T} = \text{argmax}_{\vartheta_s \in \Theta_s} \mathcal{L}_T(\vartheta_s|\hat{\vartheta}_{r,T})$. The asymptotic properties of $\hat{\vartheta}_{r,T}$ are well established and have been extensively discussed in Lütkepohl (2005) and Brüggemann et al. (2016). Next, I discuss the asymptotic properties of the QML estimator for the structural and GARCH parameters.

2.5 Asymptotic properties

For model y_t defined in (2.1), with structural shocks ξ_t following the GARCH process (2.5) and (2.7), and parameter vectors as defined above, I establish consistency of the estimator under the following assumptions:

Assumption (A1). The true parameter value $\vartheta_0 \in \Theta$ with $\Theta \subset \mathbb{R}^{d_\vartheta}$ being a compact subspace of $\mathbb{R}^{d_\vartheta}$, on which the identification condition holds.

[6] As also detailed in OA C, this orthogonal rotation parameterization is also advantageous arithmetically, since it is infinitely differentiable in ϱ, and its r-th partial derivative with respect to each element in ϱ is obtained by rotating the corresponding axes by 90 degrees clockwise r times.

Common compactification techniques can be applied and condition rank $\boldsymbol{S}_T = N$ leads to local identification by Lemma 2.2. While global identification on Θ may be achieved by imposing mechanical restrictions (e.g., Lanne et al. 2017) or adopting specific labelling strategies (e.g., Lewis 2021) to fix a permutation. As previously mentioned, however, relaxing the assumption of diagonal GARCH parameters can actually lead to global identification. Several examples illustrating this will be provided later in this section. Define $N \times N^2$ matrix Δ_σ such that $\sigma_t = \Delta_\sigma \text{vec}(\Sigma_t)$ and $\text{vec}(\Sigma_t) = \Delta_\sigma^+ \sigma_t$, with the explicit formula and properties of Δ_σ discussed in OA A.

Assumption (A2). The reverse characteristic polynomial has no roots on or within the complex unit circle, i.e., $|A(z, \boldsymbol{\alpha})| \neq 0$, for any $|z| \leq 1$ and for all $\boldsymbol{\alpha}$, and $\sup_{\phi \in \Phi} \mathbb{E}\left[\log \| \left((\eta_t \odot \eta_t)' \otimes G_1(\phi)\right) \Delta_\sigma' + \Gamma_1(\phi) \| \right] < 0$.

Since GARCH parameters are non-negative, a sufficient condition is $\sup_{\phi \in \Phi} \rho(G_1(\phi) + \Gamma_1(\phi)) < 1$. When G_1 and Γ_1 are restricted to be diagonal, as in previous literature, it suffices to assume $G_{jj,1} + \Gamma_{jj,1} < 1$ for all j.

Assumption (A3). $\sup_{\phi \in \Phi} \rho(\Gamma_1(\phi)) < 1$, $\mathbb{E}\|\eta_t\|^{2s} < \infty$ for some $s > 0$.

This assumption ensures that any fixed initialization of the variance process becomes asymptotically negligible. When G_1 and Γ_1 are diagonal, this assumption is redundant once stationarity is assumed, as all diagonal elements of Γ_1 are constrained to be less than unity. The following theorem establishes the consistency of the QML estimator.

Theorem 2.6. Under Assumptions (A1) – (A3), as $T \to \infty$,

$$\hat{\vartheta}_T \xrightarrow{p} \vartheta_0.$$

Define $\overline{\sigma_t} := \max_{j \in \{1,\ldots,N\}} \sigma_{jt}$ and $\underline{\sigma_t} := \min_{j \in \{1,\ldots,N\}} \sigma_{jt}$. Additional assumptions are introduced to establish asymptotic normality:

Assumption (A4). The true parameter $\vartheta_0 \in \text{interior}(\Theta)$.

Assumption (A5). $\mathbb{E}\|\xi_t\|^6 < \infty$ and $\sup_{\phi \in \Phi} \mathbb{E}[\overline{\sigma_t}^s(\phi)] < \infty$ for some $s > 1$.

While the existence of fourth-order moments is necessary for inference concerning the reduced-form covariance $\boldsymbol{\omega}$ (or any functions thereof), the slightly stronger Assumption (A5) is required to bound derivatives of the variance processes.

Theorem 2.7. Under Assumptions (A1) – (A5), as $T \to \infty$,

$$\sqrt{T}(\hat{\vartheta}_{s,T} - \vartheta_{s,0}) \xrightarrow{d} \mathcal{N}(\mathbf{0}, V).$$

The detailed expression of the asymptotic covariance matrix is presented in Appendix A, where it is shown that estimation uncertainty in the reduced-form slope parameter $\boldsymbol{\alpha}$ does not affect the asymptotic variance of the structural rotation and variance parameters in the second step.

2.6 Simulation performance

To investigate the finite-sample properties of the QML estimator, a series of simulation experiments is conducted. For data generation, I employ a three-dimensional VAR(1) model with parameters calibrated based on the theoretical framework of An and Schorfheide (2007), with structural impact multipliers (with different reduced-form correlations and structural rotations) randomly sampled. Regarding the GARCH parameters, I consider three alternative patterns of variance transmission:[7]

$$\text{Type a: } \begin{bmatrix} * & 0 & 0 \\ * & * & * \\ 0 & * & * \end{bmatrix}, \quad \text{Type b: } \begin{bmatrix} * & 0 & * \\ * & * & * \\ 0 & 0 & * \end{bmatrix}, \quad \text{Type c: } \begin{bmatrix} * & 0 & * \\ 0 & * & * \\ * & 0 & * \end{bmatrix}, \tag{2.15}$$

where $*$ denotes unrestricted elements in both matrices G_1 and Γ_1 and are randomly sampled. Importantly, volatility spillovers among different shocks are permitted due to the presence of non-zero off-diagonal elements, and it can be readily verified that no permutation $\mathcal{P}'G_1\mathcal{P}$ or $\mathcal{P}'\Gamma_1\mathcal{P}$ replicates the transmission pattern, except when $\mathcal{P} = I_3$. I explore multiple sample sizes T and consider three distinct (centered and standardized) distributions for the variance innovations η_t: Gaussian, Student's t, and χ^2, with various degrees of freedom. Further details on the simulation design and performance assessment criteria are provided in OA D.1.

Estimator Table 2.1 presents the means of the estimation errors (measured as mean squared errors, with detailed definitions in OA D.1) for all model parameters, structural shocks, their variances, and the variance innovations across 1,000 simulation

[7] While I focus on three specific variance transmission patterns that achieve global identification, other configurations, such as distinct structures on G_1 and Γ_1, may also lead to global identification. Importantly, this flexibility allows for economic theories tailored to specific variance transmission schemes to be effectively leveraged for shock identification within the proposed framework. Exploring this direction, however, extends beyond the scope of this paper and is left for future research.

		$T = 2,000$							$T = 8,000$						
Dist	Type	$\boldsymbol{\alpha}$	$\boldsymbol{\beta}$	G_1	Γ_1	$\{\xi_t\}_{t=1}^T$	$\{\eta_t\}_{t=1}^T$	$\{\sigma_t\}_{t=1}^T$	$\boldsymbol{\alpha}$	$\boldsymbol{\beta}$	G_1	Γ_1	$\{\xi_t\}_{t=1}^T$	$\{\eta_t\}_{t=1}^T$	$\{\sigma_t\}_{t=1}^T$
Gaussian	a	0.004	0.216	0.011	0.126	0.117	0.112	0.064	0.001	0.042	0.002	0.051	0.026	0.025	0.015
		0.009	2.402	0.027	0.508	0.746	0.745	0.170	0.001	0.024	0.003	0.173	0.020	0.019	0.045
	b	0.003	0.276	0.011	0.134	0.130	0.125	0.048	0.001	0.120	0.003	0.064	0.050	0.049	0.014
		0.008	2.860	0.024	0.375	1.060	1.060	0.128	0.001	0.040	0.004	0.213	0.024	0.022	0.036
	c	0.004	0.424	0.014	0.117	0.179	0.173	0.120	0.001	0.151	0.004	0.052	0.060	0.058	0.030
		0.010	3.956	0.070	0.317	1.332	1.330	0.223	0.002	0.035	0.005	0.151	0.026	0.022	0.058
Stud. t	a	0.004	0.263	0.014	0.128	0.130	0.124	0.107	0.001	0.105	0.004	0.062	0.046	0.045	0.029
		0.009	2.866	0.043	0.425	1.136	1.124	0.339	0.002	0.037	0.006	0.192	0.026	0.022	0.086
	b	0.004	0.249	0.013	0.137	0.119	0.112	0.084	0.001	0.089	0.003	0.067	0.041	0.039	0.074
		0.009	1.232	0.033	0.426	0.433	0.424	0.202	0.002	0.033	0.006	0.211	0.023	0.021	0.052
	c	0.004	0.358	0.015	0.121	0.160	0.153	0.135	0.001	0.110	0.004	0.055	0.048	0.046	0.064
		0.011	3.622	0.039	0.321	1.272	1.268	0.412	0.002	0.037	0.007	0.170	0.027	0.022	0.112
Chi-square	a	0.004	0.296	0.014	0.126	0.144	0.138	0.150	0.001	0.043	0.003	0.058	0.027	0.026	0.027
		0.011	3.027	0.041	0.401	1.252	1.268	0.394	0.002	0.027	0.006	0.176	0.021	0.019	0.061
	b	0.004	0.330	0.014	0.133	0.148	0.142	0.078	0.001	0.112	0.004	0.064	0.050	0.049	0.021
		0.009	3.552	0.037	0.366	1.295	1.288	0.215	0.002	0.032	0.007	0.181	0.022	0.021	0.054
	c	0.005	0.374	0.017	0.126	0.169	0.162	0.244	0.001	0.108	0.005	0.058	0.046	0.044	0.036
		0.012	3.474	0.047	0.317	1.275	1.285	0.363	0.003	0.040	0.007	0.173	0.027	0.023	0.117

Table 2.1: *Estimation errors under alternative distributions for η_t and GARCH specifications. Odd-numbered (shaded) rows report the mean estimation errors (measured as element-wise mean squared errors), while even-numbered rows report the 95%-quantiles of the empirical distribution of estimation errors across 1,000 simulation replications. The table includes results for all model parameters, structural shocks ξ_t, their variances σ_t, and variance innovations η_t. See OA D.1 for further details.*

replications, along with the 95%-quantiles of their empirical distributions. I consider two sample sizes: $T = 2,000$ and $T = 8,000$, approximating roughly 8 and 32 years of daily observations (based on 250 trading days per year). In the smaller sample, the reduced-form slope parameters in $\boldsymbol{\alpha}$ are estimated with good accuracy, though the estimator for the structural impact multipliers in $\boldsymbol{\beta}$ shows sizable bias, with large biases occurring occasionally, as indicated by the high upper-quantile values. Among the GARCH parameters, G_1 is notably easier to estimate than Γ_1, with the estimation error for G_1 often an order of magnitude lower than for Γ_1. Although the variance process is generally well estimated, occasional large biases appear in the estimates for the shocks ξ_t and η_t, as indicated by the respective quantile statistics. As the sample size increases, the estimator's performance markedly improves across all parameters, as evident in the results presented in the right part of the table. As we expect from the consistency of the estimator, both the mean estimation errors and the 95%-quantiles converge to zero. In larger samples, biases in $\boldsymbol{\beta}$ are notably reduced, and estimation errors for Γ_1 also converge, further closing the gap with

		$T = 2,000$				$T = 8,000$				$T = 32,000$			
		Size	Power			Size	Power			Size	Power		
Dist	Test		Alt.1	Alt.2	Alt.3		Alt.1	Alt.2	Alt.3		Alt.1	Alt.2	Alt.3
Gaussian	$\mathcal{T}_T^{LM1}$	0.17/0.10/0.03	-	-	0.94	0.13/0.07/0.02	-	-	0.98	0.10/0.05/0.01	-	-	0.99
	$\mathcal{T}_T^{LM2}$	0.35/0.28/0.15	1.00	0.93	0.95	0.28/0.21/0.09	1.00	0.99	0.98	0.23/0.16/0.06	1.00	1.00	0.99
	$\mathcal{T}_T^{LM3}$	0.21/0.12/0.03	1.00	1.00	0.94	0.13/0.07/0.02	1.00	1.00	0.98	0.10/0.05/0.01	1.00	1.00	0.99
	$\mathcal{T}_T^{LR}$	0.24/0.14/0.04	-	0.98	-	0.10/0.07/0.02	-	1.00	-	0.09/0.05/0.01	-	1.00	-
Stud. t	$\mathcal{T}_T^{LM1}$	0.26/0.16/0.06	-	-	0.92	0.18/0.10/0.02	-	-	0.98	0.13/0.07/0.02	-	-	0.99
	$\mathcal{T}_T^{LM2}$	0.48/0.40/0.26	1.00	0.95	0.94	0.33/0.25/0.14	1.00	0.99	0.98	0.26/0.20/0.09	1.00	1.00	0.99
	$\mathcal{T}_T^{LM3}$	0.31/0.21/0.09	1.00	1.00	0.93	0.18/0.10/0.02	1.00	1.00	0.98	0.14/0.07/0.02	1.00	1.00	1.00
	$\mathcal{T}_T^{LR}$	0.23/0.16/0.07	-	0.98	-	0.18/0.11/0.04	-	1.00	-	0.13/0.08/0.03	-	1.00	-
Chi-square	$\mathcal{T}_T^{LM1}$	0.29/0.20/0.08	-	-	0.93	0.16/0.10/0.03	-	-	0.98	0.13/0.07/0.02	-	-	0.99
	$\mathcal{T}_T^{LM2}$	0.52/0.43/0.30	1.00	0.95	0.95	0.37/0.29/0.17	1.00	0.99	0.98	0.28/0.21/0.11	1.00	1.00	0.99
	$\mathcal{T}_T^{LM3}$	0.36/0.25/0.11	1.00	1.00	0.94	0.19/0.12/0.04	1.00	1.00	0.98	0.14/0.08/0.01	1.00	1.00	0.99
	$\mathcal{T}_T^{LR}$	0.19/0.13/0.07	-	0.98	-	0.18/0.11/0.04	-	1.00	-	0.14/0.09/0.03	-	1.00	-

Table 2.2: *Rejection frequencies of the LR and LM tests under different distributions for η_t and varying sample sizes, based on 1,000 simulation replications. Nominal sizes under the null hypothesis are set at 0.10/0.05/0.01. The alternatives used to evaluate test power include a VAR(2) process (Alt.1), a non-recursive causal scheme (Alt.2), and a misspecified GARCH structure (Alt.3), with a testing size of 0.01. Further details are provided in OA D.2.*

the more stably estimated G_1 parameter. Similar observations also hold for the estimated shocks. In scenarios where variance innovations are drawn from leptokurtic or skewed distributions, the estimator's performance shows only minimal distortion, which demonstrates robustness under non-Gaussian data. Overall, the simulation results indicate that the proposed QML estimator exhibits favorable finite-sample performance, with notable accuracy in estimating the structural parameters and variance transmission patterns, in samples of moderate size.

Specification tests As a corollary of Theorem 2.7, specific parameter constraints can be tested using the likelihood ratio (LR) test. Let $\vartheta_{s,0}$ denote the structural and GARCH parameter vector under the constraint. By Wilks' theorem, the LR test statistic has an asymptotic χ^2 distribution:

$$\mathcal{T}_T^{LR} = -2\left(l_t(\vartheta_{s,0}) - l_t(\hat{\vartheta}_{s,T})\right) \xrightarrow{d} \chi^2(\nu),$$

where ν is the number of parameters under constraint and $l_t(\vartheta_s) := l_t(\vartheta_s|\vartheta_{r,0})$, which may be replaced with the consistent estimator $l_t(\vartheta_s|\hat{\vartheta}_{r,T})$. This framework allows, for example, testing for specific causal schemes (e.g., upper/lower triangular or symmetric impact multiplier) through constraints on the structural rotations ϱ. However, the LR test cannot be directly applied to test the significance of GARCH parameters,

as the null hypothesis places $\vartheta_{s,0}$ on the boundary of the parameter space, which affects the asymptotic distribution of $\mathcal{T}_T^{LR}$ (see Ch. 8.3 Francq and Zakoian 2019 for details). In such cases, the Lagrange multiplier (LM) test can be applied, which retains a standard asymptotic χ^2 distribution even when the parameter lies on the boundary:

$$\mathcal{T}_T^{LM1} = T\mathcal{S}_T^{LM1}(\vartheta_{s,0})'(\hat{\mathcal{I}}_T^{LM1})^{-1}(\vartheta_{s,0})\mathcal{S}_T^{LM1}(\vartheta_{s,0}) \xrightarrow{d} \chi^2(\nu),$$

where $\mathcal{S}_T^{LM1}(\vartheta_{s,0}) = \frac{1}{T}\sum_{t=1}^{T}\frac{\partial}{\partial\vartheta_s}l_t(\vartheta_{s,0})$ and $\hat{\mathcal{I}}_T(\vartheta_{s,0})$ is a consistent estimator of the information matrix, e.g., $\hat{\mathcal{I}}_T^{LM1}(\vartheta_{s,0}) = \frac{1}{T}\sum_{t=1}^{T}\frac{\partial}{\partial\vartheta_s}l_t(\vartheta_{s,0})\frac{\partial}{\partial\vartheta_s'}l_t(\vartheta_{s,0})$. Additionally, one can consider joint LM tests involving both reduced-form and structural parameters, allowing for constraints on ϑ_r. Denoting the parameter vector under the null as $\vartheta_0 = (\vartheta_{r,0}', \vartheta_{s,0}')'$, we define two versions of the LM test statistics, denoted $\mathcal{T}_T^{\bullet}$ for $\bullet \in \{LM2, LM3\}$:

$$\mathcal{T}_T^{\bullet} = T\mathcal{S}_T^{\bullet}(\vartheta_0)'(\hat{\mathcal{I}}_T^{\bullet})^{-1}(\vartheta_0)\mathcal{S}_T^{\bullet}(\vartheta_0) \xrightarrow{d} \chi^2(\nu),$$

with $\mathcal{S}_T^{LM2}(\vartheta_0) = \left(\frac{1}{T}\sum_{t=1}^{T}\frac{\partial}{\partial\vartheta_r'}g_t(\vartheta_{r,0}), \frac{1}{T}\sum_{t=1}^{T}\frac{\partial}{\partial\vartheta_s'}l_t(\vartheta_{s,0})\right)'$ and $\mathcal{S}_T^{LM3}(\vartheta_0) = \frac{1}{T}\sum_{t=1}^{T}\frac{\partial}{\partial\vartheta}l_t^{\dagger}(\vartheta_0)$, $\hat{\mathcal{I}}_T^{\bullet}$ defined analogously. Analytical expressions for the score functions are provided in OA D.2.

To evaluate the properties of these tests, I simulate data from the DGP described earlier, randomly selecting a GARCH specification from the three variance transmission schemes with equal probability and adopting a recursive (i.e., lower-triangular) causal scheme. Table 2.2 documents the rejection frequencies of these tests under the null hypothesis, based on their asymptotic distribution and nominal sizes of 0.10, 0.05, and 0.01. While all tests exhibit elevated Type-I errors in smaller samples of 1,000 observations, the rejection frequencies align more closely with nominal sizes as T increases and are well-sized for $T = 32,000$. An exception is the LM test statistic $\mathcal{T}_T^{LM2}$, which tends to over-reject even in large samples. This result is expected, as the reduced-form likelihood does not account for the GARCH effect, leading to an underestimation of the covariance matrix. These trends hold across different distributions for η_t. To assess test power, I consider three alternative hypotheses. First, I test whether the data follows a VAR(2) process (Alt.1) using $\mathcal{T}_T^{LM2}$ and $\mathcal{T}_T^{LM3}$. Next, I test against a non-recursive causal scheme (Alt.2) using the same statistics, alongside the LR test, and finally, I test against alternative GARCH specifications using all three LM tests (Alt.3). As indicated by the rejection frequencies (based on a

nominal level of 0.01) in Table 2.2, all tests show favorable power across sample sizes and distributions considered.

3 A daily structural oil market model

3.1 Financial market data and model specifications

The model includes three endogenous variables: the (log) returns of the spot price of crude oil (ospot), U.S. stock market prices (sp), and the futures price of crude oil (ofuture). For the spot price, I use the current price of West Texas Intermediate (WTI) crude oil, while for the futures price, I use NYMEX WTI futures with a four-month delivery horizon. To construct U.S. stock market returns, I extract the first principal component from the returns of the S&P 500, NASDAQ, and Russell 2000 indices, which captures 91% of the variation in these composites.[8] The basic model structure aligns with that of Gazzani et al. (2024). However, prices for WTI are used due to their high liquidity ensuring arbitrage-free conditions, and their longer historical availability, allowing for an analysis that revisits important historical episodes, such as the Gulf War in the early 1990s. For similar reasons, I restrict the analysis to the U.S. stock market, which, as shown later, contains substantial information on expectations regarding future oil demand and supply that enables us to properly identify distinct oil price shocks. The data comprise $T = 9,056$ daily observations, spanning over 36 years, from January 1, 1988, to April 5, 2024. Addressing potential data issues related to the Covid-19 outbreak in 2020, I retain these observations instead of excluding or down-weighting them by variance, as suggested by Lenza and Primiceri (2022). The pandemic period coincides with notable oil market turbulence, including some historically unprecedented events, which are essential to our analysis. As shown later, the GARCH model effectively captures these dramatic changes in the volatility of the variables. To control for potential fixed effects of Covid-19 policies on the conditional means, I include the Covid-19 Stringency Index for the U.S., constructed by Hale et al. (2021), as an exogenous variable.[9] Monthly dummies are

[8] Data are obtained from the EIA website and FRED, with further details on sources and data transformations available in OA E.1. Comparable results are found when using shorter delivery horizons, beginning with the front-month contract provided by the EIA.

[9] This index is a composite measure based on nine government policy indicators. Since the original series ended in 2022, I extend it through polynomial extrapolation until the official end of the public health emergency as declared by the WHO in early May 2023.

	Alternative model specifications							
Vola.	VAR0-BEKK	SVAR-f	SVAR-d	SVAR-p-a	SVAR-p-b	SVAR-g-a	SVAR-g-b	SVAR-g-c
Scheme	$\begin{bmatrix} * & * & * \\ * & * & * \\ * & * & * \end{bmatrix}$	$\begin{bmatrix} * & * & * \\ * & * & * \\ * & * & * \end{bmatrix}$	$\begin{bmatrix} * & 0 & 0 \\ 0 & * & 0 \\ 0 & 0 & * \end{bmatrix}$	$\begin{bmatrix} * & 0 & 0 \\ 0 & * & * \\ 0 & * & * \end{bmatrix}$	$\begin{bmatrix} * & 0 & 0 \\ * & * & * \\ * & * & * \end{bmatrix}$	$\begin{bmatrix} * & 0 & 0 \\ * & * & * \\ 0 & * & * \end{bmatrix}$	$\begin{bmatrix} * & 0 & * \\ * & * & * \\ 0 & 0 & * \end{bmatrix}$	$\begin{bmatrix} * & 0 & * \\ 0 & * & * \\ * & 0 & * \end{bmatrix}$
AIC	101213.7	16381.66	13827.45	16296.56	14903.75	12586.00	12929.59	12948.32
BIC	101384.5	16531.00	13891.45	16389.00	15024.64	12692.67	13036.26	13054.99
HQ	101271.8	16432.46	13849.22	16328.00	14944.88	12622.28	12965.88	12984.60

	Tests for identification						**Tests for causal direction**	
H_0	rank $\boldsymbol{S}_T = 1$			rank $\boldsymbol{S}_T = 2$			B symmetric	
Test stat.	LM1	LM2	LM3	LM1	LM2	LM3	LM ($\mathcal{T}_T^{LM2}$)	LR ($\mathcal{T}_T^{LR}$)
	623.48 (1)	863.87 (9)	861.05 (9)	522.95 (1)	522.95 (1)	522.03 (1)	968.32 (3)	3093.37 (3)

	Tests for volatility spillovers								
Direction	None	agg dem. → prec dem.		oil sup. → prec dem.		both shocks → prec dem.	prec dem. → oil sup.		
H_0	G_1, Γ_1 diag.	$G_{21,1} = 0$	$\Gamma_{21,1} = 0$	$G_{23,1} = 0$	$\Gamma_{23,1} = 0$	$G_{2n,1}, \Gamma_{2n,1} = 0,\ \forall n \neq 2$	$G_{32,1} = 0$	$\Gamma_{32,1} = 0$	$G_{32,1}, \Gamma_{32,1} = 0$
$\mathcal{T}_T^{LM1}$	325.23 (12)	142.88 (1)	68.17 (1)	92.94 (1)	97.57 (1)	94.35 (4)	99.11 (1)	100.42 (1)	68.34 (2)
$\mathcal{T}_T^{LM2}$	424.70 (12)	289.23 (1)	111.15 (1)	140.50 (1)	153.32 (1)	163.81 (4)	307.32 (1)	159.36 (1)	222.24 (2)

Table 3.1: *Information criteria for alternative model specifications and results of specification tests. LM tests for rank deficiency of matrix $\boldsymbol{S}_T$ and additional diagnostic analysis of the reduced-form model are detailed in OA E.2, while tests for symmetric loadings of structural shocks and various directions of volatility spillovers (using LR and LM test statistics) are discussed in Section 2.6.*

also incorporated to account for seasonal effects in the conditional mean. Additional details and visual representations of these variables are provided in OA E.1.

The reduced-form model includes an intercept and 24 lags, chosen based on information criteria and the absence of serial correlations in the error terms. As expected, the data provide strong evidence of conditional heteroskedasticity in the error terms, as the null hypothesis of homoskedasticity is rejected at any conventional significance level across both univariate and multivariate ARCH-LM tests on squared residuals under various specifications. Regarding volatility transmission schemes, I compare seven alternative GARCH-SVAR specifications, each with different restrictions on the GARCH parameter matrices G_1 and Γ_1. Specifically, I estimate a model with both G_1 and Γ_1 fully unrestricted (SVAR-f), a model with both matrices restricted to be diagonal (SVAR-d) as in previous studies, two models allowing partial volatility spillovers but without global identification (SVAR-p-a and SVAR-p-b), and three globally identified models as defined in (2.15). Additionally, I consider a BEKK(1,1,1) specification (VAR0-BEKK) to model the conditional covariance of the log-return processes directly. The upper panel of Table 3.1 summarizes the information criteria for these alternative specifications. Models capturing the conditional covariance of structural shocks show substantial improvements in model fit compared to direct

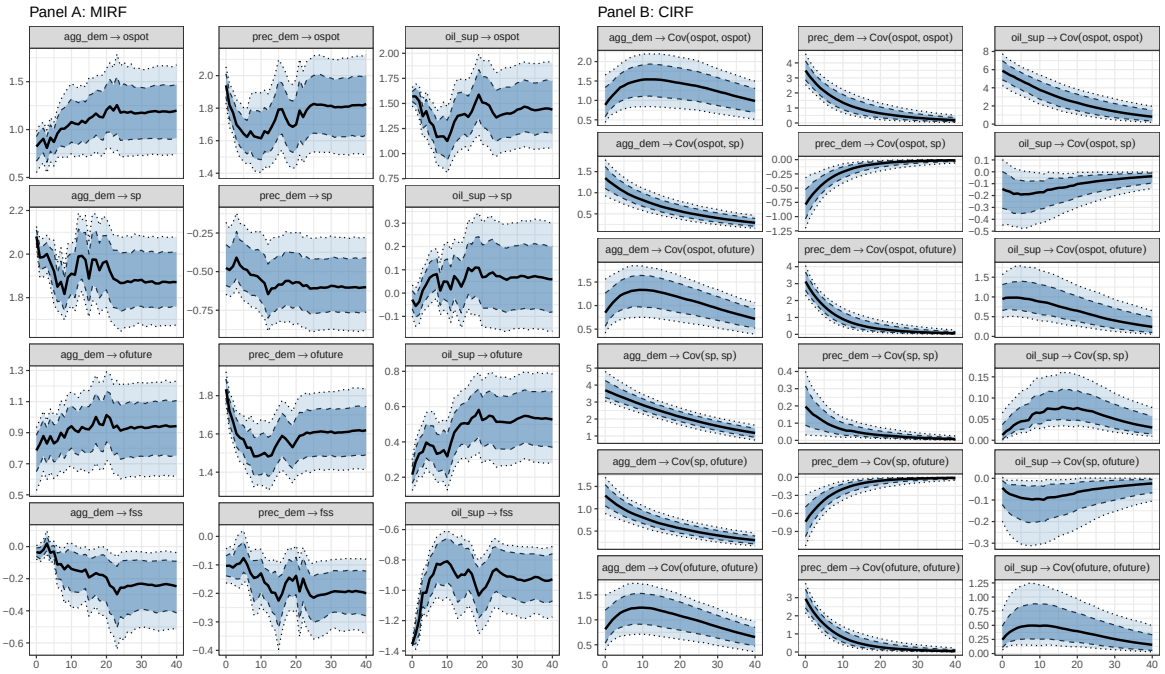

Figure 3.1: *Mean (panel A) and covariance (panel B) impulse response functions to oil price shocks over a 40-day horizon. The MIRF shock dose is set to one standard deviation, while CIRF shocks correspond to the 99%-quantile of each variance innovation's empirical distribution. Solid lines show the median response, with dashed and dotted lines indicating 68% and 90% pointwise confidence bands, respectively, based on 1,000 moving-block bootstrap replications.*

modeling of the data's conditional covariance. The information criteria values from the SVAR approaches are, on average, one-sixth of those for the BEKK model – a notable finding given that SVAR models include an additional 255 slope parameters in $\boldsymbol{\alpha}$. Among the GARCH-SVAR models, the globally identified models (SVAR-g) that relax the diagonal restriction provide a better fit to the data. Of these, the transmission scheme labeled 'SVAR-g-a' minimizes all considered criteria (AIC, BIC, and HQ). As a result, I select SVAR-g-a as the benchmark model, with details and results for alternative specifications available in OA E.9. I further test the identification criterion in Lemma 2.2 using the LM tests proposed by Lanne and Saikkonen (2007) and Lütkepohl and Milunovich (2016) and find strong evidence against the null hypothesis of any rank deficiency in matrix $\boldsymbol{S}_T$.

Panel A in Figure 3.1 displays the MIRFs for model variables, including the future-spot spreads (fss) computed as the difference in responses between spot and futures prices, while panel B presents their CIRFs. In Figure 3.2, panel A, shows the historical decompositions (HDs) of the oil spot price (yellow) and stock price (blue), while panel B plots each identified daily shock. Furthermore, I study the effect of these shocks

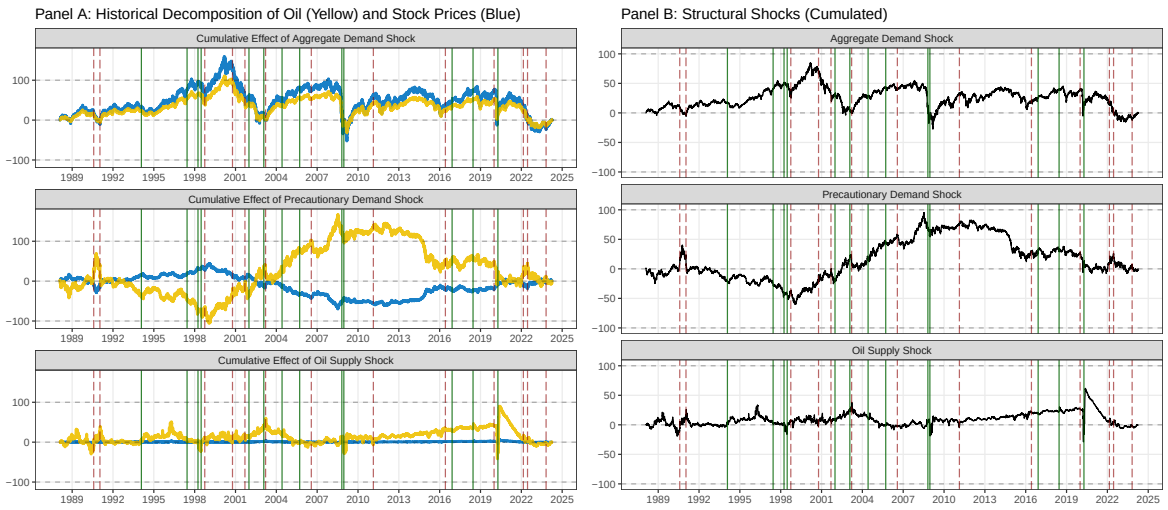

Figure 3.2: *Historical decomposition of oil spot price and stock price (panel A) and cumulative identified oil price shocks (panel B). Red dashed lines denote major geopolitical events, including the Iraqi invasion of Kuwait (1990), U.S. and UN interventions in Iraq (1991, 1998), the 9/11 attacks, the 2003 Iraq invasion, sanctions on Iran (2006, 2022), Russia's invasion of Ukraine (2022), and Israel's invasion of Gaza following the Oct. 7 attacks. Green solid lines indicate significant OPEC production adjustments (see OA E.1 for details).*

(monthly averages) on lower-frequency variables within a classical set of variables for global oil market, which includes oil production (OProd), a measure of global real economic activity (GREA), the real price of oil (ORP), and oil inventory (OInv). Additionally, I assess effects of these shocks on a selection of U.S. macroeconomic variables, including output, price and inflation expectations, interest rate and credit spreads.[10] Figure 3.3 presents the MIRFs for these lower-frequency oil market and macro variables, obtained using local projection (Jorda 2005). Further details on data sources, transformations, and local projection specifications are provided in OA E.1.

3.2 Transmission of oil price shocks and their volatilities

In an ideal setting, established economic theories would guide the choice of variance transmission schemes, allowing us to directly attach economic interpretations

[10]I employ two alternative measures of global real activity: the world industrial production, as used in Baumeister and Hamilton (2019), and the index constructed by Kilian (2009) based on dry cargo bulk freight rates. The results, which are largely similar, are presented in Figure 3.3 and Figure XYZ in OA E.8. The set of monthly U.S. macroeconomic variables includes industrial production (IP), personal consumption expenditures (PCE) inflation, 12-month inflation expectations from the University of Michigan survey (MICH), the yield on 2-year Treasury securities (GS2), and the EBP constructed by Gilchrist and Zakrajšek (2012). As alternatives to the EBP, I also consider additional risk measures, with detailed results provided in Figure 3.4 and Figure XYZ in OA E.8.

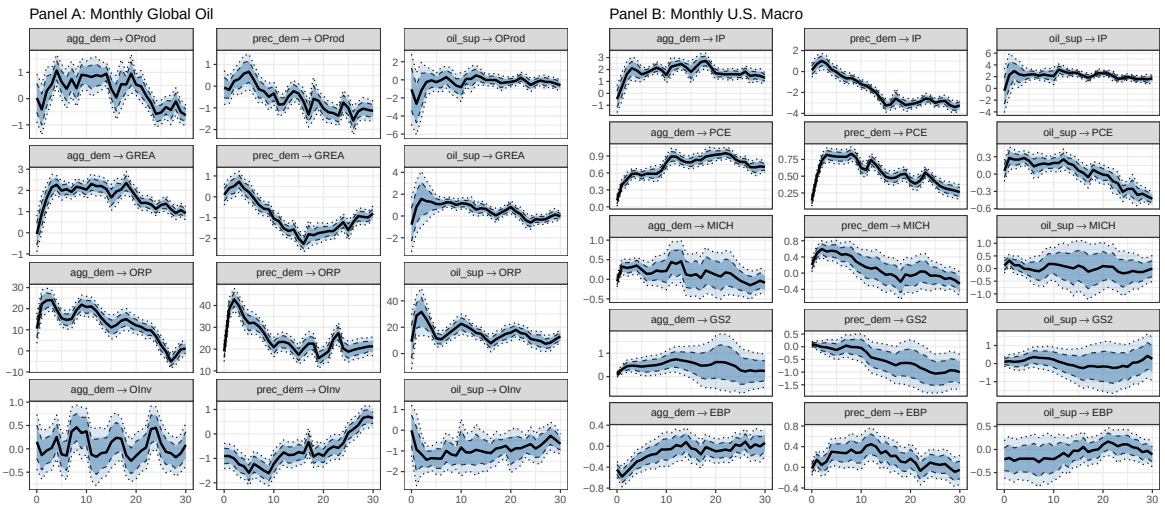

Figure 3.3: *Monthly mean impulse response functions in a global oil market model (panel A) and a U.S. macro mode (panel B), shown over 30 months. Solid lines represent median responses, while dashed and dotted lines indicate 68% and 90% pointwise confidence bands, respectively, based on heteroskedasticity-consistent standard errors. See OA E.3 for further details.*

to the statistically identified shocks. However, since the transmission scheme here was selected based on likelihood criteria, the shocks, despite being globally identified, require careful labeling. The empirical results provide clear guidance for this process: as evident from the MIRFs and CIRFs displayed in Figure 3.1, the three daily oil price shocks exhibit distinct and well-delineated properties in terms of their effects on the conditional mean and variance of asset prices. The first shock has the most pronounced impact on both the conditional mean and variance of stock prices. This shock is labeled as an *aggregate demand shock*, reflecting unexpected shifts in current macroeconomic conditions that drive changes in the demand for the 'flow' of oil. The second structural shock exerts the largest influence on both the mean and variance of oil futures prices – its impact on the mean is almost double (tenfold) and on the variance is nearly fourfold (twelvefold) that of the aggregate demand shock (or the third shock). This shock is labeled as a *precautionary demand shock*, encapsulating heightened uncertainties regarding future macroeconomic conditions as well as anticipated disruptions in oil supply relative to expected demand, independent of current production levels (see Anzuini et al. 2015 for a similar conceptualization). The third identified shock, which emerges as the dominant driver of increases in both the mean and variance of daily oil spot prices – its effects surpassing those induced by other shocks – is labeled as the *oil supply shock*, representing unexpected and sudden disruptions to the current oil production.

Notably, the labeling of these shocks closely adheres to conventions established in the classical literature on low-frequency models. As we carefully examine the MIRFs, CIRFs, and HDs in this section, and turn to narrative evidence based on real-time FEVD in the next section to analyze market perceptions of the origins of forecast uncertainty, the empirical basis for these interpretations becomes even more evident. The MIRFs of these shocks on low-frequency variables, such as oil production, economic activity, and credit spreads, exhibit response profiles that align with classical low-frequency models, while also revealing some notable distinctions. In what follows, I describe each identified oil price shock's characteristics and transmission mechanisms, followed by a discussion of volatility spillovers across these shocks.

Aggregate demand shock A positive aggregate demand shock leads to an increase in both spot and futures prices for crude oil and has the largest impact on the conditional mean and variance of stock prices, as illustrated in both panels of Figure 3.1. As shown in Figure 3.2, this shock has been the primary driver of historical stock price fluctuations while also contributing substantially to variations in oil prices – consistent with previous findings from low-frequency models (e.g., Kilian and Murphy 2014). Monthly MIRFs reveal that global oil production rises by nearly 1%, with significance lasting about 15 months, while both world and U.S. industrial production show sustained increases. The real price of crude oil and the PCE index respond significantly, with oil prices peaking at 25% above baseline within three months, and the PCE index continuing to rise over two years. Contractionary monetary policy is observed as the 2-year Treasury yield peaks at 75 bps after one year. Financial conditions significantly improve initially with EBP dropping by 60 bps, which fades to insignificance after three quarters.

Precautionary demand shock Following a precautionary demand shock, both spot and futures oil prices rise in close alignment, resembling Gazzani et al. (2024)'s 'forward-looking demand shock', which reflects how such shocks, embodying heightened risks of future supply disruptions or increased demand, are typically priced into both spot and futures price simultaneously. In the low-frequency oil model shown in panel A of Figure 3.3, this shock initially leads to a modest increase in oil production and a sharp rise in the real price of oil, which exceeds 40% in the short run (up to six months), thus lending it characteristics of an oil consumption shock in the near term. However, in the medium to long term, this shock predicts

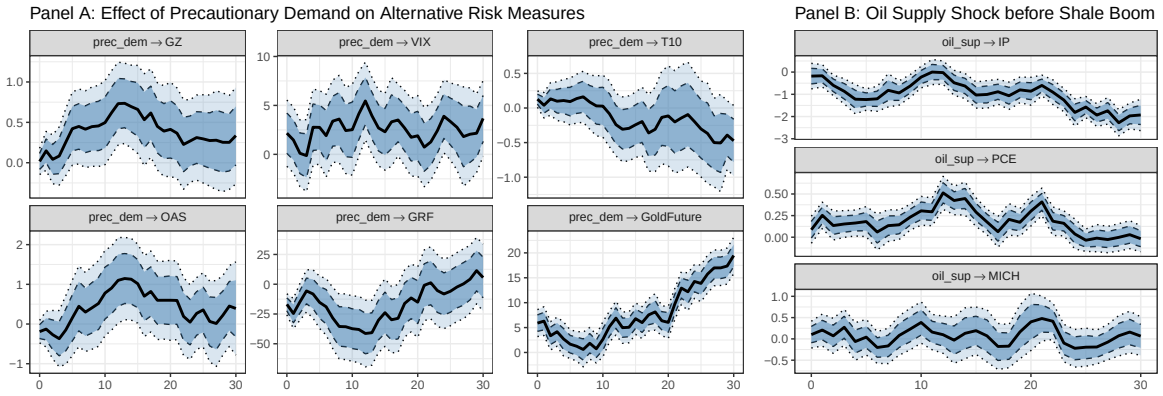

Figure 3.4: *Monthly mean impulse response functions of alternative risk measures in a U.S. macro model, shown over a 30-month horizon, to the precautionary demand shock (panel A) and oil supply shock using pre-shale boom data (panel B). See Figure 3.3 and OA E.8 for additional details.*

a persistent decline in both oil production and economic activity, accompanied by a buildup in oil inventory beginning around two years after the shock.

Stock prices drop sharply in response to this shock, as heightened risk aversion and economic uncertainty prompt investors to shift away from riskier assets, which presents a notable distinction from results in Gazzani et al. (2024). This effect is particularly evident in the U.S. macro model, where the precautionary demand shock uniquely tightens financial conditions measured by the EBP and elevates risk perceptions across various indicators within a 1-1.5-year horizon (see panel A of Figure 3.4). Significant increases in risk aversion over similar horizons are observed in the original GZ credit spread from Gilchrist and Zakrajšek (2012), the global factor in world risky asset prices from Miranda-Agrippino and Rey (2020) (GRF), the VIX index, and option-adjusted spreads between high-yield bonds and Treasury securities (OAS). As investors shift from risky assets to safe havens, as similarly noted in Anzuini et al. (2015), gold futures increase by approximately 15%, while yields on 2-year and 10-year Treasury notes fall by as much as 1% and 0.5%, respectively, over a two-year period, possibly signaling an accommodative monetary stance. Elevated uncertainty also associates with a 3% decline in industrial production at the medium horizon, likely reflecting pauses in investment and consumption amid economic concerns (see, e.g., Bloom 2014), while PCE inflation rises (though less than with an aggregate demand shock) and inflation expectations increase within the first year before returning to baseline (see Anzuini et al. 2015 for a comparable finding). This shock, therefore, can create stagflation over the medium term. Figure 3.2 indicates that the precautionary demand shock has contributed substantially to historical oil

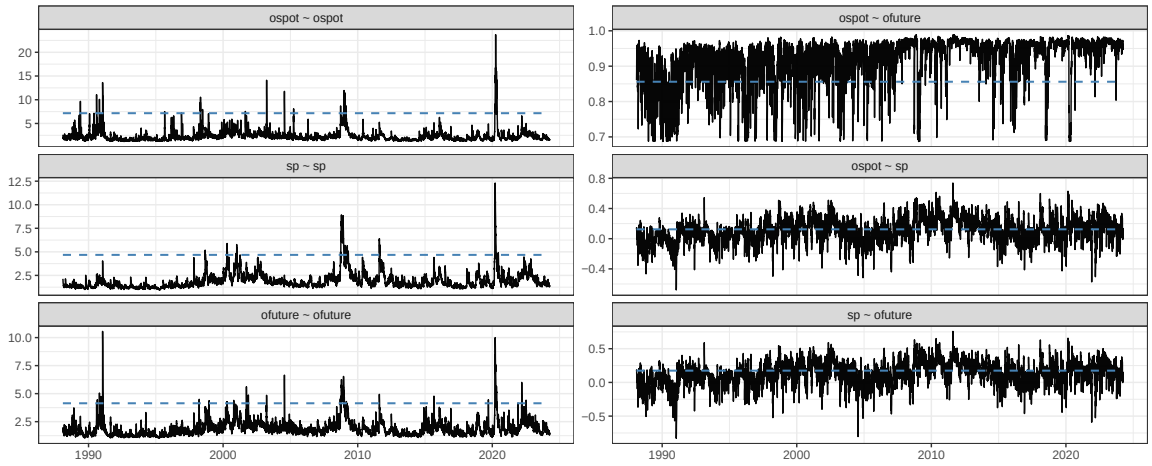

Figure 3.5: *Structural conditional variances for model variables (left column) and their conditional correlations (right column), as implied by the benchmark volatility scheme (SVAR-g-a). Blue dashed lines indicate unconditional variances and correlations.*

price movements but less so to stock market variation. As shown in panel B of the figure, this shock is closely tied to geopolitical events, as also argued in Anzuini et al. (2015).[11] We will explore these associations further in the next section.

Oil supply shock Following an unexpected disruption in the current oil production, spot price rises immediately over 1.5% above baseline. While futures prices also increase, their adjustment is smaller and delayed. Following a temporary production shortfall, stock prices initially decline but quickly recover, rendering the effect insignificant, similar to findings in Kilian and Park (2009). The world and U.S. industrial production exhibit a comparable pattern, with an initial decline that fades, followed by a brief overshoot. Compared to other shocks, the real oil price increase in the low-frequency model less persistent. Figure 3.2 further reveals that oil supply shocks minimally contribute to historical oil price movements and have negligible effects on stock price fluctuations. These results are consistent with previous findings in classic studies on global oil markets using low-frequency models (e.g., Kilian

[11] As discussed earlier, and similar to interpretations in Gazzani et al. (2024) and Anzuini et al. (2015), the precautionary demand shock reflects both economic uncertainties and anticipated future supply disruptions relative to expected oil demand. However, this paper does not take a stance on whether macroeconomic uncertainty or geopolitical instability in oil-producing regions exerts a greater influence on oil prices and domestic economic activity. A notable recent contribution is Kilian et al. (2024), which disentangles level shocks from uncertainty shocks, finding that macroeconomic uncertainty induces recessions, while downside geopolitical risks exert only modest effects on macroeconomic aggregates.

and Murphy 2014).

The divergent responses of spot and futures prices create an immediate 1.4% reduction in the spread. Monthly model results in Figure 3.1 further show that global oil production sharply drops by 2.5%-3% initially and hardly recovers to baseline over a 10-month horizon. Facing an unexpected supply shortfall, inventories are drawn down to meet demand, resulting in persistently low inventory levels. However, this depletion limits investors' ability to smooth demand over time, raising the convenience yield and thus the value of holding physical oil over futures contracts, which contributes to the sharp drop in the spread. Additionally, the muted futures price response may lead to backwardation, where the convenience yield surpasses carrying costs, reflecting market expectations for a subsequent decline in the spot price – consistent with monthly oil price responses in panel A of Figure 3.3. This behavior in spreads aligns with findings in Valenti (2022), who document a significant and lasting drop in spreads following oil supply disruptions.

Interestingly, the responses of U.S. macroeconomic aggregates to oil supply shocks are somewhat puzzling. Given oil's role as a fundamental industrial input and energy source, a supply-driven price surge would conventionally be expected to increase production costs, compress profit margins, and introduce inflationary pressures as higher costs pass on to consumers. However, due to advances in shale oil fracking, which have substantially increased domestic output since 2010, the U.S. has become the world's largest oil producer. Thus, certain segments of the economy may benefit from rising oil prices. To explore this, I re-estimate the model using pre-shale boom data (1988–2010). The subsample results, shown in panel B of Figure 3.4, align more closely with conventional expectations: output declines significantly and persistently, while both PCE inflation and inflation expectations rise for approximately two years, generating stagflationary conditions.

Volatility Spillovers As seen in panel B of Figure 3.1, volatility drivers differ across markets: uncertainty surrounding the aggregate demand dominates volatility in equity markets, while oil spot market volatility is primarily driven by oil supply uncertainty, and futures market volatility is largely influenced by precautionary demand. Unlike the precautionary demand and oil supply shocks – which lower the covariance between oil and stock prices due to heightened risk perception and potential flight-to-safety behaviors – the aggregate demand shock uniquely raises

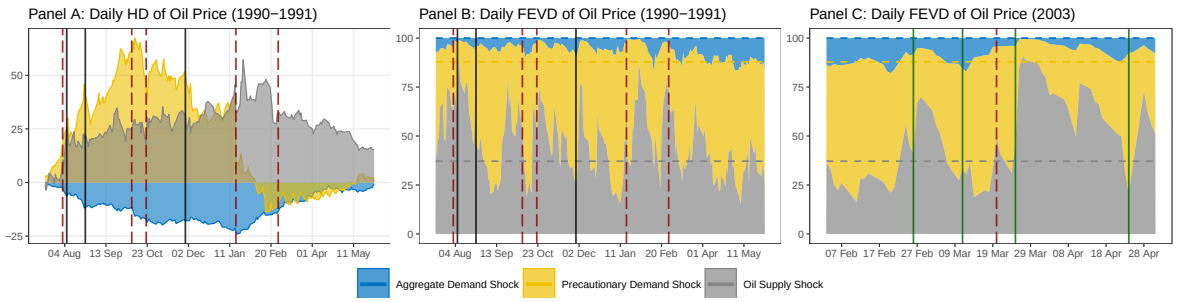

Figure 3.6: *Daily FEVD and HD for oil and stock price dynamics during the Gulf War (1990-1991) and the 2003 Iraq War. Red dashed lines mark key geopolitical developments: in panels A-B, the Iraqi invasion of Kuwait (Aug. 2, 1990), heightened tensions between Iraq and Israel (Oct. 8), and diplomatic efforts from France and Saudi Arabia (Oct. 22). The timeline also includes the start and end of Operation Desert Storm, aerial bombing campaign against Iraq (Jan. 17 - Feb. 28, 1991); in panel C, the start of the 2003 Iraq invasion (Mar. 20). Black solid lines in panel A-B denote dates of UN economic (Resolutions 661 and 665) and military (678) sanctions against Iraq. Green solid lines in panel C highlight oil supply-related events, including the OIC's announcement of a possible oil embargo in anticipation of military action against Iraq (Feb. 26), OPEC's production decisions (maintaining output on Mar. 11 and cutting production by 2 mbpd in an emergency meeting on April 24), and a reported dramatic drop in Iraqi oil exports by the UNOIP (Mar. 25). For additional details, see Figure 1.1.*

covariance between stock prices and both spot and futures oil prices when its variance unexpectedly rises. Figure 3.5 illustrates these dynamics over time, showing positive oil-equity covariance during periods of global economic expansion, such as the early 2000s commodity boom driven by strong global growth and the post-GFC recovery during the first half of 2010s (see Bernanke 2016 for similar evidence). In contrast, during periods of heightened geopolitical tension, such as the Gulf War (1990-1991), oil-equity covariance turns sharply negative.

Comparative analyses (in OA E.4) using the VAR0-BEKK model and a SVAR model with diagonal GARCH specification (SVAR-d), as adopted in earlier studies, reveal that these models capture far less of the nuanced second-order dynamics seen here. This underscores the importance of allowing for volatility spillovers when studying uncertainty transmission across oil and stock markets. As shown in the lower section of Table 3.1, instantaneous spillovers from unexpected changes in volatility, particularly from aggregate demand affecting precautionary demand, and from precautionary demand impacting oil supply, are statistically significant across various configurations. Notably, the variance of the precautionary demand shock is particularly sensitive to shifts in the other shocks' volatility, while aggregate demand shock variance remains unaffected. CIRFs displayed in OA E.5 further demonstrate that volatility of the aggregate demand shock does not respond to other shocks' volatility

at higher horizons, though it indirectly transmits to oil shocks over time, with oil supply volatility adjusting gradually with considerable lags.[12]

3.3 Decomposing forecast uncertainty during specific periods

The high-frequency structural model offers a real-time decomposition of market-perceived uncertainty as it digests information to forecast oil and asset prices. Due to the heteroskedastic nature of the structural shocks, the composition of forecast uncertainty is time-varying, often diverging markedly from conventional FEVD results, which are based on unconditional shock variances. In the following, I examine key episodes, including the COVID-19 pandemic in early 2020, Russia's war on Ukraine in 2022, the Gulf War in the early 1990s, and the 2003 Iraq War – as illustrated in Figure 1.1 and 3.6, respectively – to demonstrate that the model captures relevant market information with exceptional granularity.

The Gulf War and the 2003 Iraq War In low-frequency oil models, consensus generally holds that aggregate demand shocks have a minimal effect on oil price fluctuations during the Gulf War, while debate persists on whether oil supply or precautionary demand shocks exert a stronger influence. The daily structural model may offer a more nuanced perspective. As shown in panel A of Figure 3.6, although oil supply shocks contribute substantially to oil price movements, precautionary demand was the primary driver following Iraq's invasion of Kuwait on Aug. 2, 1990. This influence accumulates, pushing WTI spot prices up by as much as 40% amid escalating fears, particularly on Oct. 8, when Saddam Hussein threatened Israel after Israeli forces firing on Palestinian protesters in Jerusalem, intensifying concerns of regional conflict (Jacobs 1991). However, the dynamic changed swiftly with the launch of Operation Desert Storm, a U.S.-led coalition's aerial bombing campaign authorized by UN Security Council Resolution 678 on Jan. 17, 1991. The role of precautionary demand abruptly diminishes, while oil supply shocks become the predominant force driving oil prices, contributing up to 60% of the increase. This shift aligns with widespread physical destruction in Kuwaiti oil production facilities due to fires and spills set by Iraqi forces. Remarkably, the model, using high-frequency financial data, captures these shifts in the market with high resolution.

[12]While no volatility spillover from aggregate demand to oil supply is observed on impact, this does not preclude spillovers at higher horizons. Zero elements in a matrix do not imply zero elements in its higher powers, as often seen in Granger-causality analysis, where effects can emerge over time.

For instance, the model detects a sharp drop in precautionary demand's contribution on Oct. 22, when diplomatic overtures from Saudi Arabia hinted at a potential Iraqi withdrawal. The one-day-ahead FEVD in panel B offers a parallel perspective, revealing that oil supply uncertainty was already perceived by the market as a primary driver of forecast uncertainty as early as August 1990. Notably, just days after Iraq's invasion of Kuwait, the passing of UN Security Council Resolution 661, imposing economic sanctions on Iraq, elevated the oil supply shock's contribution to the one-day-ahead FEV of oil spot prices to 91%, nearly 2.5 times its unconditional level of 37%.

A similar dynamic is observed during the 2003 U.S.-led invasion of Iraq. Almost a month leading up to the invasion, on Feb. 26, leaders of the Organisation of the Islamic Conference (OIC) suggested a potential oil embargo to dissuade the U.S. from military action, immediately raising market participants' perception of oil supply risk as the leading source of forecast uncertainty. A further surge occurred five days into the war, when the UN reported a 75% drop in Iraqi oil exports compared to the previous week. This unexpectedly severe supply disruption heightened oil supply's contribution to forecast uncertainty to approximately 91%, which sustained elevated levels above 70% for the following week.

Covid-19 pandemic and Russian invasion of Ukraine The oil market has experienced extreme turbulence in early 2020 with the onset of the COVID-19 pandemic. While traditional FEVD estimates based on unconditional shock variances suggest that aggregate demand shocks account for roughly 12% of the one-day-ahead FEV of WTI spot prices, this dynamic shifted significantly as the pandemic evolved. On February 25, when the CDC issued its first public warning of a COVID-19 outbreak and as major U.S. cities reported their 'patient zero' cases, the contribution of aggregate demand shocks to forecast variance surged temporarily to approximately 44%, as illustrated in panel A of Figure 1.1. Another peak occurred on March 17 and 18, coinciding with Treasury Secretary Mnuchin's warning that unemployment could reach 20% without substantial federal action and the Senate's approval of an emergency relief aid package. From late March onward, however, as Saudi Arabia increased oil production amidst a price war with Russia, oil supply shocks emerged as the primary driver of forecast uncertainty. On April 2, following pressure from President Trump on Saudi Crown Prince Mohammed bin Salman to initiate produc-

tion cuts, WTI prices jumped 24.7% in anticipation of a supply reduction, with oil supply uncertainty accounting for nearly 90.6% of the forecast variance of WTI spot prices. After OPEC and Russia formalized a historic agreement to cut production by nearly 10 mbpd on April 9, with the cuts taking effect on May 1, oil supply shocks remained the dominant source of forecast uncertainty until mid-May.

In contrast, oil supply shocks played a limited role in driving forecast uncertainty for oil prices during the Russian invasion of Ukraine, with few exceptions. For instance, as illustrated in panels E of Figure 1.1, the forecast uncertainty for oil prices contributed by oil supply spiked on March 8, when the U.S. Senate introduced legislation to ban Russian oil imports. Notably, on dates when FOMC announced interest rate hikes, aggregate demand uncertainty increased its share of forecast variance for oil prices, accounting for about one-third of the one-day-ahead forecast variance of WTI prices – nearly three times its unconditional FEVD contribution.

Panel F further shows that uncertainty regarding current economic conditions consistently drove forecast uncertainty for equity markets. However, on Mar. 11 when President Biden signed an executive order enforcing the ban, the precautionary demand shocks accounted for 40.1% of the forecast variance in stock prices, over six times their typical contribution. Similarly, on July 5, when oil prices dropped by about 9%, marking the steepest daily decline since March, concerns intensified over potential demand reductions as news emerged of Covid-19 mass testing and lockdowns in China. This led to fears of a deepening recession, tripling the precautionary demand shock's share of forecast variance from 7% to 22%.

4 Conclusion

This paper develops a high-frequency SVAR framework to identify oil price shocks and analyze the transmission of their uncertainties. Using daily data on oil spot and futures prices alongside stock indices, the model identifies three structural shocks and demonstrates how each shock influences asset prices and volatilities, as well as how they propagate through U.S. financial and macroeconomic systems. Leveraging the stylized characteristics of high-frequency financial data, such as volatility clustering, effectively captured by a GARCH model, the approach achieves global identification of these shocks while allowing for volatility spillovers across them. For estimation, the paper introduces a QML estimator shown to be consistent and

asymptotically normal under mild conditions. In simulation studies, this estimator, together with LM and LR-type specification tests, exhibits robust finite-sample properties. Empirical findings suggest that aggregate demand uncertainty dominates volatility in equity markets, while oil supply and precautionary demand uncertainty primarily drive the volatility in oil spot and futures markets, respectively. Precautionary demand shocks, reflecting heightened economic uncertainty and risk aversion, sharply increase both spot and futures prices and cause notable declines in stock prices as investors pivot away from risk assets. Conversely, an unexpected oil supply disruption triggers a strong and immediate rise in spot prices, with a more muted response in futures prices, producing a marked reduction in the futures-spot spread. The real-time forecast error variance decomposition further captures evolving market perceptions of uncertainty, showing that oil supply uncertainty was the primary driver of oil price forecast uncertainty between late March and early May 2020, while contributing minimally during the 2022 Russian invasion of Ukraine.

This framework opens avenues for future research into the economic foundations of volatility transmission schemes across distinct shocks, which may offer new perspectives on structural identification and guide multivariate volatility modelling.

References

Aastveit, K. A., H. C. Bjørnland, and J. L. Cross (2023): "Inflation expectations and the pass-through of oil prices," *The Review of Economics and Statistics*, 105, 733–743.

An, S. and F. Schorfheide (2007): "Bayesian analysis of DSGE models," *Econometric Reviews*, 26, 113–172.

Anzuini, A., P. Pagano, and M. Pisani (2015): "Macroeconomic effects of precautionary demand for oil," *Journal of Applied Econometrics*, 30, 968–986.

Baumeister, C. (2023): "Pandemic, war, inflation: Oil markets at a crossroads?" Working Paper 31496, National Bureau of Economic Research.

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

Bernanke, B. S. (2016): "The relationship between stocks and oil prices," https://www.brookings.edu/articles/the-relationship-between-stocks-and-oil-prices, brookings' research and commentary, accessed: 2024-11-09.

Bertsche, D. and R. Braun (2022): "Identification of structural vector autoregressions by stochastic volatility," *Journal of Business & Economic Statistics*, 40, 328–341.

Bloom, N. (2014): "Fluctuations in uncertainty," *Journal of Economic Perspectives*, 28, 153–76.

Bouakez, H. and M. Normandin (2010): "Fluctuations in the foreign exchange market: How important are monetary policy shocks?" *Journal of International Economics*, 81, 139–153.

Bougerol, P. and N. Picard (1992): "Strict Stationarity of Generalized Autoregressive Processes," *The Annals of Probability*, 20, 1714 – 1730.

Braun, R. (2023): "The importance of supply and demand for oil prices: Evidence from non-Gaussianity," *Quantitative Economics*, 14, 1163–1198.

Braun, R. and R. Brüggemann (2023): "Identification of SVAR models by combining sign restrictions with external instruments," *Journal of Business & Economic Statistics*, 41, 1077–1089.

Brüggemann, R., C. Jentsch, and C. Trenkler (2016): "Inference in VARs with conditional heteroskedasticity of unknown form," *Journal of Econometrics*, 191, 69 – 85.

Fengler, M. and J. Polivka (2024): "Structural volatility impulse response analysis," Swiss Finance Institute Research Paper 24-63.

Francq, C. and J.-M. Zakoian (2019): *GARCH models*, John Wiley Sons, Ltd.

Gazzani, A., F. Venditti, and G. Veronese (2024): "Oil price shocks in real time," *Journal of Monetary Economics*, 144, 103547.

Gilchrist, S. and E. Zakrajšek (2012): "Credit spreads and business cycle fluctuations," *American Economic Review*, 102, 1692–1720.

Gouriéroux, C., A. Monfort, and J. Renne (2017): "Statistical inference for independent component analysis: Application to structural VAR models," *Journal of Econometrics*, 196, 111–126.

Hafner, C. M. and H. Herwartz (2006): "Volatility impulse responses for multivariate GARCH models: An exchange rate illustration," *Journal of International Money and Finance*, 25, 719–740.

——— (2023a): "Asymmetric volatility impulse response functions," *Economics Letters*, 222, 110968.

——— (2023b): "Correlation impulse response functions," *Finance Research Letters*, 57, 104176.

Hafner, C. M., H. Herwartz, and S. Maxand (2022): "Identification of structural multivariate GARCH models," *Journal of Econometrics*, 227, 212–227, annals Issue: Time Series Analysis of Higher Moments and Distributions of Financial Data.

Hafner, C. M., H. Herwartz, and S. Wang (2024): "Statistical identification of independent shocks with kernel-based maximum likelihood estimation and an application to the global crude oil market," *Journal of Business & Economic Statistics*, 0, 1–16.

Hafner, C. M. and A. Preminger (2009a): "Asymptotic theory for a factor GARCH model," *Econometric Theory*, 25, 336–363.

——— (2009b): "On asymptotic theory for multivariate GARCH models," *Journal of Multivariate Analysis*, 100, 2044–2054.

Hale, T., N. Angrist, R. Goldszmidt, B. Kira, A. Petherick, T. Phillips, S. Webster, E. Cameron-Blake, L. Hallas, S. Majumdar, and H. Tatlow (2021): "A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker)," *Nature Human Behaviour*, 5, 529–538.

Herrera, A. M. and S. K. Rangaraju (2020): "The effect of oil supply shocks on US economic activity: What have we learned?" *Journal of Applied Econometrics*, 35, 141–159.

Herwartz, H., B. Theilen, and S. Wang (2024): "Unraveling the structural sources of oil production and their impact on CO2 emissions," *Energy Economics*, 132, 107488.

Jacobs, R. H. (1991): "A chronology of the Gulf War," *Arab Studies Quarterly*, 13, 143–165.

Jorda, O. (2005): "Estimation and inference of impulse responses by local projections," *American Economic Review*, 95, 161–182.

Kilian, L. (2008): "The Economic Effects of Energy Price Shocks," *Journal of Economic Literature*, 46, 871–909.

——— (2009): "Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market," *American Economic Review*, 99, 1053–1069.

——— (2024): "How to construct monthly VAR proxies based on daily surprises in futures markets," *Journal of Economic Dynamics and Control*, 168, 104966.

Kilian, L. and L. T. Lewis (2011): "Does the Fed respond to oil price shocks?" *The Economic Journal*, 121, 1047–1072.

Kilian, L. and H. Lütkepohl (2017): *Structural Vector Autoregressive Analysis*, Cambridge University Press.

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

——— (2014): "The role of inventories and speculative trading in the global market for crude oil," *Journal of Applied Econometrics*, 29, 454–478.

Kilian, L. and C. Park (2009): "The impact of oil price shocks on the u.s. stock market," *International Economic Review*, 50, 1267–1287.

Kilian, L., M. D. Plante, and A. W. Richter (2024): "Geopolitical Oil Price Risk and Economic Fluctuations," Working Papers 2403, Federal Reserve Bank of Dallas.

Kilian, L. and X. Zhou (2022): "Oil prices, gasoline prices, and inflation expectations," *Journal of Applied Econometrics*, 37, 867–881.

——— (2023): "The econometrics of oil market VAR models," in *Essays in Honor of Joon Y. Park: Econometric Methodology in Empirical Applications*, Emerald Group Publishing Limited, vol. 45 of *Advances in Econometrics*, 65–95.

King, M., E. Sentana, and S. Wadhwani (1994): "Volatility and links between national stock markets," *Econometrica*, 62, 901–933.

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

Lanne, M., M. Meitz, and P. Saikkonen (2017): "Identification and estimation of non-Gaussian structural vector autoregressions," *Journal of Econometrics*, 196, 288–304.

Lanne, M. and P. Saikkonen (2007): "A multivariate generalized orthogonal factor GARCH model," *Journal of Business & Economic Statistics*, 25, 61–75.

Lenza, M. and G. E. Primiceri (2022): "How to estimate a vector autoregression after March 2020," *Journal of Applied Econometrics*, 37, 688–699.

Lewis, D. J. (2021): "Identifying shocks via time-varying volatility," *Review of Economic Studies*, 88, 3086–3124.

Lütkepohl, H. (2005): *New Introduction to Multiple Time Series Analysis*, Springer.

Lütkepohl, H. and G. Milunovich (2016): "Testing for identification in SVAR-GARCH models," *Journal of Economic Dynamics and Control*, 73, 241–258.

Lütkepohl, H. and A. Netšunajev (2014): "Disentangling demand and supply shocks in the crude oil market: How to check sign restrictions in structural VARs," *Journal of Applied Econometrics*, 29, 479–496.

Meitz, M. and P. Saikkonen (2008): "Ergodicity, mixing, and existence of moments of a class of markov models with applications to GARCH and ACD models," *Econometric Theory*, 24, 1291–1320.

Miranda-Agrippino, S. and H. Rey (2020): "U.S. monetary policy and the global financial cycle," *The Review of Economic Studies*, 87, 2754–2776.

Newey, W. K. and D. McFadden (1994): "Chapter 36 Large sample estimation and hypothesis testing," in *Handbook of Econometrics*, Elsevier, 2111–2245.

Normadin, M. and L. Phaneuf (2004): "Monetary policy shocks: Testing identification conditions under time-varying conditional volatility," *Journal of Monetary Economics*, 1217–1243.

Sentana, E. and G. Fiorentini (2001): "Identification, estimation and testing of conditionally heteroskedastic factor models," *Journal of Econometrics*, 102, 143–164.

Straumann, D. and T. Mikosch (2006): "Quasi-maximum-likelihood estimation in conditionally heteroscedastic time series: A stochastic recurrence equations approach," *The Annals of Statistics*, 34, 2449 – 2495.

Valenti, D. (2022): "Modelling the global price of oil:Is there any role for the oil futures-spot spread?" *The Energy Journal*, 43, 41–66.

Valenti, D., A. Bastianin, and M. Manera (2023): "A weekly structural VAR model of the US crude oil market," *Energy Economics*, 121, 106656.

A Proofs of theorems

Proof of Theorem 2.4: At $h = 1$, $\mathbb{E}[\sigma_{t+1}|\mathcal{F}_{t-1}, \eta_t = \eta^*] - \mathbb{E}[\sigma_{t+1}|\mathcal{F}_{t-1}] = G_1\Sigma_t[(\eta^* \odot \eta^*) - \mathbf{1}_N]$, while for $h > 1$, $\mathbb{E}[\eta_{t+h} \odot \eta_{+h}|\mathcal{F}_{t-1}, \eta_t = \eta^*] = \mathbb{E}[\eta_{t+h} \odot \eta_{+h}|\mathcal{F}_{t-1}] = \mathbf{1}_N$ and thus

$$\begin{aligned}\mathbb{E}[\sigma_{t+h}|\mathcal{F}_{t-1}, \eta_t = \eta^*] - \mathbb{E}[\sigma_{t+h}|\mathcal{F}_{t-1}] &= (G_1 + \Gamma_1)^{h-1}(\mathbb{E}[\sigma_{t+1}|\mathcal{F}_{t-1}, \eta_t = \eta^*] - \mathbb{E}[\sigma_{t+1}|\mathcal{F}_{t-1}]) \\ &= (G_1 + \Gamma_1)^{h-1}G_1\Sigma_t[(\eta^* \odot \eta^*) - \mathbf{1}_N].\end{aligned}$$

Hence, $\mathcal{V}^{\xi}_{t+h}(\eta^*|\mathcal{F}_{t-1}) = \left((G_1 + \Gamma_1)^{h-1}G_1\Sigma_t[(\eta^* \odot \eta^*) - \mathbf{1}_N]\mathbf{1}_N'\right) \odot I_N$. To show (2.9), note that $\xi_t|\mathcal{F}_{t-1}, \eta_t = \eta^* = \Sigma_t^{1/2}\eta^*$ a.s., while $\mathbb{E}[\xi_t\xi_t'|\mathcal{F}_{t-1}] = \Sigma_t$, thus

$$\mathcal{V}^{y}_{t+h}(\eta^*|\mathcal{F}_{t-1}) = \sum_{i=0}^{h}\Theta_i(\mathrm{Cov}[\xi_{t+h-i}|\mathcal{F}_{t-1}, \eta_t = \eta^*] - \mathrm{Cov}[\xi_{t+h-i}|\mathcal{F}_{t-1}])\Theta_i' = \sum_{i=0}^{h-1}\Theta_i\mathcal{V}^{\xi}_{t+h-i}(\eta^*|\mathcal{F}_{t-1})\Theta_i'.\square$$

Proof of Theorem 2.5: First, note that the h-step forecast error variance can be written as $\mathbb{E}[u_{t,h}u_{t,h}'|\mathcal{F}_t] = \sum_{i=0}^{h-1}\sum_{j=1}^{N}\mathbb{E}[\sigma_{j,t+h-i}|\mathcal{F}_t]\Theta_i e_j e_j'\Theta_i'$. Since the Sequence $\{\eta_t\}$ is iid s.t. $\mathbb{E}[(\eta_{t+h-1} \odot \eta_{t+h-1})|\mathcal{F}_t] = \mathbf{1}_N$, we have for $h \geq 2$,

$$\begin{aligned}\mathbb{E}[\sigma_{t+h}|\mathcal{F}_t] &= \gamma_0 + G_1\mathbb{E}[\Sigma_{t+h-1}|\mathcal{F}_t]\mathbb{E}[(\eta_{t+h-1} \odot \eta_{t+h-1})|\mathcal{F}_t] + \Gamma_1\mathbb{E}[\Sigma_{t+h-1}|\mathcal{F}_t]\mathbf{1}_N \\ &= \gamma_0 + (G_1 + \Gamma_1)\mathbb{E}[\sigma_{t+h-1}|\mathcal{F}_t].\end{aligned}$$

Recursive substitutions lead to

$$\begin{aligned}\mathbb{E}[\sigma_{t+h}|\mathcal{F}_t] &= \left(I_N + (G_1 + \Gamma_1) + (G_1 + \Gamma_1)^2 + \ldots (G_1 + \Gamma_1)^{h-2}\right)\gamma_0 + (G_1 + \Gamma_1)^{h-1}\sigma_{t+1} \\ &= \left(I_N - (G_1 + \Gamma_1)^{h-1}\right)(I_N - (G_1 + \Gamma_1))^{-1}\gamma_0 + (G_1 + \Gamma_1)^{h-1}\sigma_{t+1} \\ &= (I_N - (G_1 + \Gamma_1))^{-1}\gamma_0 + (G_1 + \Gamma_1)^{h-1}\left[\sigma_{t+1} - (I_N - (G_1 + \Gamma_1))^{-1}\gamma_0\right],\end{aligned}$$

where the $N \times N$ matrix inverse $(I_N - (G_1 + \Gamma_1))^{-1}$ is well-defined since $\rho(G_1 + \Gamma_1) < 1$. Note that the second term is a positive coefficient times the difference between

the conditional and unconditional variances $(I_N - (G_1 + \Gamma_1))^{-1}\gamma_0 = \mathbb{E}[\sigma_t^2] = \mathbf{1}_N$. Another way to see this is by plugging in the parameter restrictions imposed $\gamma_0 = (I_N - (G_1 + \Gamma_1))\mathbf{1}_N$, and hence $\mathbb{E}[\sigma_{t+h}|\mathcal{F}_t] = \mathbf{1}_N + (G_1 + \Gamma_1)^{h-1}[\sigma_{t+1} - \mathbf{1}_N]$, for all $h \geq 1$, which completes the proof. □

Proof of Lemma 2.2: First, given the conditional second-order moment of the reduced-form model residual

$$\mathbb{E}[u_t u_t'|\mathcal{F}_{t-1}] = B\mathbb{E}[\xi_t \xi_t'|\mathcal{F}_{t-1}]B' = B\Sigma_t B',$$

both the structural impact multiplier B and the variances of the structural shocks Σ_t are identified up to an orthogonal transformation. Let $B^* = BQ'$, $\Sigma_t^* = Q\Sigma_t Q'$ with $Q \in O(N)$, then (B^*, Σ_t^*) and (B, Σ_t) are observationally equivalent, i.e., $B^*\Sigma_t^* B^{*\prime} = BQ'Q\Sigma_t Q'QB' = B\Sigma_t B'$. The associated vector of non-structural orthogonalized shocks is given by $\xi_t^* = B_t^{*-1}$ with conditional covariance matrix given by

$$\Sigma_t^* = Q\Sigma_t Q' = \begin{bmatrix} \sum_{k=1}^N \sigma_{kt}^2 q_{1k}q_{1k} & \sum_{k=1}^N \sigma_{kt}^2 q_{1k}q_{2k} & \dots & \sum_{k=1}^N \sigma_{kt}^2 q_{1k}q_{Nk} \\ \sum_{k=1}^N \sigma_{kt}^2 q_{2k}q_{1k} & \sum_{k=1}^N \sigma_{kt}^2 q_{2k}q_{2k} & \dots & \sum_{k=1}^N \sigma_{kt}^2 q_{2k}q_{Nk} \\ \vdots & \vdots & \ddots & \vdots \\ \sum_{k=1}^N \sigma_{kt}^2 q_{Nk}q_{1k} & \sum_{k=1}^N \sigma_{kt}^2 q_{Nk}q_{2k} & \dots & \sum_{k=1}^N \sigma_{kt}^2 q_{Nk}q_{Nk} \end{bmatrix}.$$

Since structural shocks in ξ_t^* are uncorrelated, Σ_t^* must be a diagonal matrix and thus $\sum_{k=1}^N \sigma_{kt}^2 q_{ik}q_{jk} = 0$ for all $i \neq j$ for all t. This condition can be expressed as

$$\underbrace{\begin{bmatrix} \sigma_{11}^2 & \sigma_{21}^2 & \dots & \sigma_{N1}^2 \\ \sigma_{12}^2 & \sigma_{22}^2 & \dots & \sigma_{N2}^2 \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{1T}^2 & \sigma_{2T}^2 & \dots & \sigma_{NT}^2 \end{bmatrix}}_{\boldsymbol{S}_T} \begin{bmatrix} q_{i1}q_{j1} \\ q_{i2}q_{j2} \\ \dots \\ q_{iN}q_{jN} \end{bmatrix} = \mathbf{0}, \ \forall i, j \in \{1, \dots, N\}, i \neq j.$$

Since rank $\boldsymbol{S}_T = N$, the only solution to the system above is $q_{ik}q_{jk} = 0$ for all $i \neq j$ and all $k = 1, \dots, N$. However, since $Q \in O(N)$, $\sum_{j=1}^N q_{jk}^2 = 1$ for all k, each column in Q can not have two (or more) non-zero elements and the only non-zero element is either +1 or −1. Consequently, Q must take the form $Q = \mathcal{P}\mathcal{D}$, B is identified up to column permutations and sign flips and the structural shocks in ξ_t are identified up to a reordering. The GARCH model (2.6) can be uniquely and compactly written as

$$\sigma_t = \gamma_0 + \sum_{i=1}^q G_i(\xi_{t-i} \odot \xi_{t-i}) + \sum_{i=1}^p \Gamma_i \sigma_{t-i}.$$

As the variances do not exhibit equivalent representations, the conditional variance vector of the permuted shocks $\xi_t^* = \mathcal{P}'\xi_t$, reads as

$$\begin{aligned}\sigma_t^* = \mathcal{P}'\sigma_t &= \mathcal{P}'\gamma_0 + \sum_{i=1}^{q} \mathcal{P}'G_i\mathcal{P}\mathcal{P}'(\xi_{t-i} \odot \xi_{t-i}) + \sum_{i=1}^{p} \mathcal{P}'\Gamma_i\mathcal{P}\mathcal{P}'\sigma_{t-i} \\ &= \gamma_0^* + \sum_{i=1}^{q} (\mathcal{P}'G_i\mathcal{P})(\mathcal{P}'\xi_{t-i} \odot \mathcal{P}'\xi_{t-i}) + \sum_{i=1}^{p} (\mathcal{P}'\Gamma_i\mathcal{P})\sigma_{t-i}^* \\ &= \gamma_0^* + \sum_{i=1}^{q} G_i^*(\xi_{t-i}^* \odot \xi_{t-i}^*) + \sum_{i=1}^{p} \Gamma_i^*\sigma_{t-i}^*,\end{aligned}$$

with $G_i^* = \mathcal{P}'G_i\mathcal{P}$ and $\Gamma_i^* = \mathcal{P}'\Gamma_i\mathcal{P}$. This completes the proof. □

Proof of Theorem 2.6: For the consistency of standard estimators of VAR reduced-form parameters $(\hat{\alpha}_T', \hat{\omega}_T')' \xrightarrow{p} (\alpha_0', \omega_0')'$, since $\{u_t\}$ is a MDS and the process is stationary and ergodic, as we show next, by Lemma 3.1 of Lütkepohl (2005) there exists non-singular matrix $M_Z \in \mathbb{R}^{NP\times NP}$ s.t., $ZZ'/T = M_Z + o_p(1)$ and $T^{-1/2}\sum_{t=1}^{T} \text{vec}\left(u_t Z_{t-1}'\right)$ is uniformly tight. By a continuous mapping Theorem, we have $\hat{\varepsilon}_t := \varepsilon_t(\hat{\vartheta}_{r,T}) \xrightarrow{p} \varepsilon_t$ with $\varepsilon_t = \varepsilon_t(\vartheta_{r,0})$. The conditional variances can be embedded in a stochastic recurrent equation (see, e.g., Straumann and Mikosch 2006):

$$\sigma_t = \gamma_0 + \left[((\eta_t \odot \eta_t)' \otimes G_1(\phi))\,\Delta_\sigma' + \Gamma_1(\phi)\right]\sigma_{t-1} = \gamma_0 + \mathbf{\Gamma}_{t-1}\sigma_{t-1}.$$

It is well-known in the literature on stochastic dynamic systems that the process has a non-anticipative strictly stationary solution if the associated top Lyapunov exponent is strictly negative (Bougerol and Picard 1992):

$$\boldsymbol{\gamma} = \inf\left\{\mathbb{E}\left[\frac{1}{h+1}\log\|\mathbf{\Gamma}_0\mathbf{\Gamma}_{-1}\ldots\mathbf{\Gamma}_{-h}\|\right],\ h \in \mathbb{N}\right\} < 0.$$

Since $\boldsymbol{\gamma} \leq \mathbb{E}[\log\|\mathbf{\Gamma}_t\|]$ with equality in case $N = 1$, Assumption (A3) implies that the top Lyapunov exponent is strictly negative on parameter space Φ. Thus, $\{\sigma_t\}$ is geometric ergodic and has a non-anticipative strictly stationary solution, which extends to the sequence $\{u_t\}$ by Proposition 1 of Meitz and Saikkonen (2008). The stationarity of the sequence $\{y_t\}$ follows from Assumption (A3).

Let $\tilde{\Sigma}_t(\phi)$ be the conditional variance process where the initial values are drawn from the stationary distribution. Let $\tilde{l}_t$ and $\tilde{\mathcal{L}}_T$ be defined analogously. Denote $\xi_t(\varrho, \vartheta_r) = Q(\varrho)'\varepsilon_t(\vartheta_r)$ and $\xi_t(\varrho) = Q(\varrho)'\varepsilon_t(\vartheta_{r,0})$, $\tilde{\mathcal{L}}_T(\vartheta_s) = \tilde{\mathcal{L}}_T(\vartheta_s|\vartheta_{r,0})$ and $\mathcal{L}_T(\vartheta_s) = \mathcal{L}_T(\vartheta_s|\vartheta_{r,0})$. We next show that any fixed initialization of the conditional variances

is asymptotically irrelevant. By a triangle inequality,

$$\sup_{\vartheta_s\in\Theta_s}|\tilde{\mathcal{L}}_T(\vartheta_s)-\mathcal{L}_T(\vartheta_s)| \leq \underbrace{\frac{1}{T}\sum_{t=1}^{T}\sup_{\vartheta_s\in\Theta_s}\left|\xi_t(\varrho)'\left(\Sigma_t^{-1}(\phi)-\tilde{\Sigma}_t^{-1}(\phi)\right)\xi_t(\varrho)\right|}_{\tilde{\Delta}^{(I)}_{\mathcal{L}_T}}$$
$$+\underbrace{\frac{1}{T}\sum_{t=1}^{T}\max\left\{\sup_{\phi\in\Phi}\log\left|\Sigma_t(\phi)\tilde{\Sigma}_t^{-1}(\phi)\right|,\sup_{\phi\in\Phi}\log\left|\tilde{\Sigma}_t(\phi)\Sigma_t^{-1}(\phi)\right|\right\}}_{\tilde{\Delta}^{(II)}_{\mathcal{L}_T}}$$

For the first part, note that

$$\begin{aligned}\tilde{\Delta}^{(I)}_{\mathcal{L}_T} &= \frac{1}{T}\sum_{t=1}^{T}\sup_{\vartheta_s\in\Theta_s}\left|\mathrm{tr}\left(\tilde{\Sigma}_t^{-1}(\phi)\left(\Sigma_t(\phi)-\tilde{\Sigma}_t(\phi)\right)\Sigma_t^{-1}(\phi)\xi_t(\varrho)\xi_t(\varrho)'\right)\right| \\ &\leq \frac{1}{T}\sum_{t=1}^{T}N\sup_{\vartheta_s\in\Theta_s}\|\Sigma_t^{-1}(\phi)\|\|\Sigma_t(\phi)-\tilde{\Sigma}_t(\phi)\|\|\tilde{\Sigma}_t^{-1}(\phi)\|\|\xi_t(\varrho)\xi_t(\varrho)'\| \\ &\leq T^{-1}NM_\sigma M_{\tilde{\sigma}}M_{\sigma'}\sum_{t=1}^{T}\psi_\sigma^t\|\xi_t(\varrho)\xi_t(\varrho)'\|,\end{aligned}$$

where the first inequality holds since for any $A\in\mathbb{R}^{m\times n}$, $B\in\mathbb{R}^{n\times m}$, $|\mathrm{tr}\,(AB)|\leq(mn)^{1/2}\|A\|\|B\|$ due to Cauchy-Schwarz inequality, and the last inequality follows from Lemma A.1 with $\psi_\sigma\in[0,1)$. Moreover, since $\|Q(\varrho)\|=1$ for all $\varrho\in\mathcal{R}$, for some $0<s<1$, $\mathbb{E}\|\xi_t\|^{2s}\leq\mathbb{E}\|\eta_t\|^{2s}<\infty$, thus,

$$\mathbb{E}\left|\sum_{t=1}^{T}\psi_\sigma^t\|\xi_t(\varrho)\xi_t(\varrho)'\|\right|^s\leq\sum_{t=1}^{T}\psi_\sigma^{st}\|Q(\varrho)\|\mathbb{E}\|\xi_t\xi_t'\|^s\|Q(\varrho)\|\leq\sum_{t=1}^{T}\psi_\sigma^{st}\mathbb{E}\|\xi_t\|^{2s}<\infty,$$

by Assumption (A2) and Lemma A.1.Thus, $\tilde{\Delta}^{(I)}_{\mathcal{L}_T}=O(T^{-1})$. For the second part, since $|A|\leq\|A\|^N$ for $A\in R^{N\times N}$ and by Minkowski inequality and Lemma A.1,

$$\begin{aligned}\frac{1}{T}\sum_{t=1}^{T}\sup_{\phi\in\Phi}\log\left|\Sigma_t(\phi)\tilde{\Sigma}_t^{-1}(\phi)\right| &= \frac{1}{T}\sum_{t=1}^{T}\sup_{\phi\in\Phi}\log\left|I_N+(\Sigma_t(\phi)-\tilde{\Sigma}_t(\phi))\tilde{\Sigma}_t^{-1}(\phi)\right| \\ &\leq \frac{1}{T}\sum_{t=1}^{T}\sup_{\phi\in\Phi}N\log\left(\|I_N\|+\|\Sigma_t(\phi)-\tilde{\Sigma}_t(\phi)\|\|\tilde{\Sigma}_t^{-1}(\phi)\|\right) \\ &\leq \frac{1}{T}\sum_{t=1}^{T}\sup_{\phi\in\Phi}N\|\Sigma_t(\phi)-\tilde{\Sigma}_t(\phi)\|\|\tilde{\Sigma}_t^{-1}(\phi)\| \\ &\leq T^{-1}NM_{\sigma'}M_{\tilde{\sigma}}\sum_{t=1}^{T}\psi_\sigma^t\end{aligned}$$

and similarly, $\frac{1}{T}\sum_{t=1}^{T}\sup_{\phi\in\Phi}\log\left|\tilde{\Sigma}_t(\phi)\Sigma_t^{-1}(\phi)\right| \leq T^{-1}NM_{\sigma'}M_{\sigma}\sum_{t=1}^{T}\psi_{\sigma}^{t}$, thus,

$$\tilde{\Delta}_{\mathcal{L}_T}^{(II)} \leq T^{-1}NM_{\sigma'}\max\{M_{\sigma},M_{\tilde{\sigma}}\}\sum_{t=1}^{T}\psi_{\sigma}^{t} \to 0, \text{ as } T\to\infty.$$

Furthermore,

$$\sup_{\vartheta_s\in\Theta_s}|\mathcal{L}_T(\vartheta_s|\hat{\vartheta}_{r,T})-\mathcal{L}_T(\vartheta_s)| = \frac{1}{T}\sum_{t=1}^{T}\sup_{\vartheta_s\in\Theta_s}|\text{tr}\left(Q(\varrho)\Sigma_t^{-1}(\phi)Q(\varrho)'(\hat{\varepsilon}_t\hat{\varepsilon}_t'-\varepsilon_t\varepsilon_t')\right)| \leq \frac{NM_{\tilde{\sigma}}}{T}\sum_{t=1}^{T}(\hat{\varepsilon}_t\hat{\varepsilon}_t'-\varepsilon_t\varepsilon_t').$$

Therefore, we have uniform convergence of the log-likelihood by a triangle inequality $\sup_{\vartheta_s\in\Theta_s}|\mathcal{L}_T(\vartheta_s|\hat{\vartheta}_{r,T})-\tilde{\mathcal{L}}_T(\vartheta_s)| \to 0$, a.s. Finally, since both $l_t(\vartheta_s|\hat{\vartheta}_{r,T})$ and $\tilde{l}_t(\vartheta_s)$ are continuous and thus measurable functions of an ergodic strictly stationary sequence, they are also ergodic strictly stationary. Applying a uniform ergodic theorem entails that $\sup_{\vartheta_s\in\Theta_s}|\mathcal{L}_T(\vartheta_s|\hat{\vartheta}_{r,T})-\mathbb{E}[-\tilde{l}_t(\vartheta_s)]| = 0$, a.s., as $T\to\infty$. By Lemma A.2, the population objective function $\mathbb{E}[\tilde{l}_t(\vartheta_s)]$ has a clear and unique minimizer at the true value $\vartheta_{s,0}$. This proves the consistency of the QML estimator $\hat{\vartheta}_{s,T} = \text{argmax}_{\vartheta_s\in\Theta_s}\mathcal{L}_T(\vartheta_s|\hat{\vartheta}_{r,T})$. □

Proof of Theorem 2.7: Conditional on the reduced-form parameters and initial values generated from the stationary distribution, the corresponding (negative) score is given by $\nabla_{\vartheta_s}\tilde{l}_t(\vartheta_s) := \frac{\partial\tilde{l}_t(\vartheta_s|\vartheta_{r,0})}{\partial\vartheta_s} = (\nabla_{\varrho}\tilde{l}_t(\vartheta_s)', \nabla_{\phi}\tilde{l}_t(\vartheta_s)')'$. For $i=1,\dots,d_{\varrho}$,

$$\nabla_{\varrho_i}\tilde{l}_t(\vartheta_s) = 2\xi_t(\varrho)'\tilde{\Sigma}_t^{-1}(\phi)\frac{\partial\xi_t(\varrho)}{\partial\varrho_i} = 2\varepsilon_t'Q(\varrho)\tilde{\Sigma}_t^{-1}(\phi)(\varepsilon_t'\otimes I_N)K_{NN}\text{vec}\left(\nabla_iQ(\varrho)\right),$$

where $\nabla_iQ(\varrho) = \frac{\partial Q(\varrho)}{\partial\varrho_i}$ is the $N\times N$ matrix of partial derivatives and K_{NN} is the commutation matrix. Moreover, for $i=1,\dots,d_{\phi}$,

$$\begin{aligned}\nabla_{\phi_i}\tilde{l}_t(\vartheta_s) &= \frac{\partial\log|\Sigma_t(\phi)|}{\partial\phi_i}+\frac{\partial}{\partial\phi_i}\xi_t(\varrho)'\tilde{\Sigma}_t^{-1}(\phi)\xi_t(\varrho)\\ &= \text{tr}\left(\tilde{\Sigma}_t^{-1}(\phi)\frac{\partial\tilde{\Sigma}_t(\phi)}{\partial\phi_i}-\tilde{\Sigma}_t^{-1}(\phi)\xi_t(\varrho)\xi_t(\varrho)'\tilde{\Sigma}_t^{-1}(\phi)\frac{\partial\tilde{\Sigma}_t(\phi)}{\partial\phi_i}\right).\end{aligned}$$

Without loss of generality, suppose $Q(\varrho_0)=I_N$, then $\varepsilon_t = Q(\varrho_0)\xi_t = \xi_t$. Evaluating at the $\vartheta_{s,0}$ and using the elementary relation $\text{vec}(A)'\text{vec}(B) = \text{tr}(A'B)$, the scores are

$$\nabla_{\varrho_i}\tilde{l}_t(\vartheta_{s,0}) = 2((\xi_t'\otimes I_N)\tilde{\Sigma}_t^{-1}\xi_t)'K_{NN}\text{vec}\left(\nabla_iQ\right) = 2\text{tr}\left(\xi_t\xi_t'\tilde{\Sigma}_t^{-1}\nabla_iQ'\right),$$

where the parameter dependence at the true value is suppressed, and

$$\nabla_{\phi_i}\tilde{l}_t(\vartheta_{s,0}) = \text{tr}\left(\tilde{\Sigma}_t^{-1}\frac{\partial\tilde{\Sigma}_t(\phi_0)}{\partial\phi_i}-\tilde{\Sigma}_t^{-1}\tilde{\Sigma}_t^{1/2}\eta_t\eta_t'\tilde{\Sigma}_t^{1/2'}\tilde{\Sigma}_t^{-1}\frac{\partial\tilde{\Sigma}_t(\phi_0)}{\partial\phi_i}\right)$$

$$= \operatorname{tr}\left((I_N - \eta_t\eta_t')\tilde{\Sigma}_t^{-1/2}\frac{\partial\tilde{\Sigma}_t(\phi_0)}{\partial\phi_i}\tilde{\Sigma}_t^{-1/2'}(\phi_0)\right)$$
$$= \operatorname{vec}\left(I_N - \eta_t\eta_t'\right)'(\tilde{\Sigma}_t^{-1/2}\otimes\tilde{\Sigma}_t^{-1/2})\Delta_\sigma^+\frac{\partial\sigma_t(\phi_0)}{\partial\phi_i}.$$

Let $\mathcal{I} = \begin{bmatrix} I_{\varrho\varrho} & \mathcal{I}_{\varrho\phi} \\ \mathcal{I}_{\varrho\phi}' & \mathcal{I}_{\phi\phi} \end{bmatrix}$ with $\mathcal{I}_{\varrho\varrho} = \mathbb{E}[\nabla_\varrho\tilde{l}_t(\vartheta_{s,0})\nabla_\varrho\tilde{l}_t(\vartheta_{s,0})']$, $\mathcal{I}_{\varrho\phi} = \mathbb{E}[\nabla_\varrho\tilde{l}_t(\vartheta_{s,0})\nabla_\phi\tilde{l}_t(\vartheta_{s,0})']$, and $\mathcal{I}_{\phi\phi} = \mathbb{E}[\nabla_\phi\tilde{l}_t(\vartheta_{s,0})\nabla_\phi\tilde{l}_t(\vartheta_{s,0})']$. We next show that the covariance matrix $\mathcal{I}$ is well-defined. First, note that

$$\mathbb{E}\left|\nabla_{\varrho_i}\tilde{l}_t(\vartheta_{s,0})\nabla_{\varrho_j}\tilde{l}_t(\vartheta_{s,0})\right| \le 4N\mathbb{E}\left[\|\xi_t\xi_t'\tilde{\Sigma}_t^{-1}\|^2\|\nabla_i Q\|\|\nabla_j Q\|\right] = 4N\mathbb{E}\left\|\tilde{\Sigma}_t^{-1/2}\eta_t\eta_t'\tilde{\Sigma}_t^{-1/2}\right\|^2 < \infty,$$

since by Assumption (A5), $\mathbb{E}\|\eta_t\|^4 = \mathbb{E}\left[\operatorname{tr}\left(\xi_t'\tilde{\Sigma}_t^{-1}\xi_t\right)^2\right] \le N^2M_{\tilde{\sigma}}^2\mathbb{E}\|\xi_t\|^4 < \infty$, where $M_{\tilde{\sigma}}^2$ is defined analogously as in Lemma A.1 and $M_{\tilde{\sigma}}^2 < \infty$ due to the compactness of the parameter space. Since $\mathbb{E}\|\eta_t\|^4 < \infty$, $\mathbb{E}\left[\operatorname{vec}\left(I_N - \eta_t\eta_t'\right)\operatorname{vec}\left(I_N - \eta_t\eta_t'\right)'\right]$ is finite. By Lemma A.3, $\mathbb{E}\left[\sup_{\phi\in\mathcal{B}(\phi_0,\delta)}\left\|\tilde{\Sigma}_t^{-1/2}(\phi)\frac{\partial\tilde{\Sigma}_t(\phi)}{\partial\phi_i}\tilde{\Sigma}_t^{-1/2'}(\phi)\right\|^{2s}\right] < \infty$ for some $s > 1$,

$$\mathbb{E}\left\|(\tilde{\Sigma}_t^{-1/2}(\phi_0)\otimes\tilde{\Sigma}_t^{-1/2}(\phi_0))\Delta_\sigma^+\frac{\partial\sigma_t(\phi_0)}{\partial\phi_i}\right\|^2 < \infty,$$

and thus $\mathbb{E}\left|\nabla_{\phi_i}\tilde{l}_t(\vartheta_{s,0})\nabla_{\phi_j}\tilde{l}_t(\vartheta_{s,0})\right| < \infty$. Furthermore,

$$\mathbb{E}\left|\nabla_{\varrho_i}\tilde{l}_t(\vartheta_{s,0})\nabla_{\phi_j}\tilde{l}_t(\vartheta_{s,0})\right| \le 2N\mathbb{E}[\|\xi_t\xi_t'\tilde{\Sigma}_t^{-1}\|\|I_N - \xi_t\xi_t\tilde{\Sigma}_t^{-1}\|\|\nabla_i Q\|\|\partial\tilde{\Sigma}_t(\phi_0)/\partial\phi_i\tilde{\Sigma}_t^{-1}\|] < \infty,$$

by Lemma A.1 and Lemma A.3. Notably, $\mathbb{E}\left[\nabla_{\vartheta_s}\tilde{l}_t(\vartheta_{s,0})|\mathcal{F}_{t-1}\right] = 0$ and thus, $\left\{\nabla_{\vartheta_s}\tilde{l}_t(\vartheta_{s,0})\right\}$ is a square-integrable stationary ergodic martingale difference sequence. By a corresponding CLT and Cramér-Wold device, as $T \to \infty$,

$$\sqrt{T}\nabla_{\vartheta_s}\tilde{\mathcal{L}}_T(\vartheta_{s,0}) = \frac{1}{\sqrt{T}}\sum_{t=1}^{T}\nabla_{\vartheta_s}\tilde{l}_t(\vartheta_{s,0}) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \mathcal{I}).$$

By Lemma 4(i) in Hafner and Preminger (2009a) and Lemma A.3, by a similar argument in the proof of Theorem 2.6 we have $\sup_{\phi\in B(\phi_0,\delta)}\sqrt{T}\left\|\nabla_{\vartheta_s}\tilde{\mathcal{L}}_T(\vartheta_{s,0}) - \nabla_{\vartheta_s}\mathcal{L}_T(\vartheta_{s,0})\right\| \to 0$, a.s. as $T \to \infty$. For the Hessian matrix, denote $\nabla_{\vartheta_{s,i}\vartheta_{s,j}}\tilde{l}_t(\vartheta_s) := \frac{\partial^2\tilde{l}_t(\vartheta_s|\vartheta_{r,0})}{\partial\vartheta_{s,i}\partial\vartheta_{s,j}}$, $\mathcal{J}_{\varrho\varrho} = \mathbb{E}[\nabla_{\varrho\varrho}\tilde{l}_t(\vartheta_{s,0})]$, $\mathcal{J}_{\varrho\phi} = \mathbb{E}[\nabla_{\varrho\phi}\tilde{l}_t(\vartheta_{s,0})]$, $\mathcal{J}_{\phi\phi} = \mathbb{E}[\nabla_{\phi\phi}\tilde{l}_t(\vartheta_{s,0})]$ and $\mathcal{J} = \begin{bmatrix} \mathcal{J}_{\varrho\varrho} & \mathcal{J}_{\varrho\phi} \\ \mathcal{J}_{\varrho\phi}' & \mathcal{J}_{\phi\phi} \end{bmatrix}$, with typical elements given by

$$\nabla_{\varrho_i\varrho_j}\tilde{l}_t(\vartheta_{s,0}) = \operatorname{tr}\left(\tilde{\Sigma}_t^{-1}(\xi_t'\otimes I_N)K_{NN}\operatorname{vec}(\nabla_i Q)\operatorname{vec}\left(\nabla_j Q\right)'K_{NN}'(\xi_t\otimes I_N)\right)$$
$$= \operatorname{vec}\left(\nabla_j Q'\right)'(\xi_t\xi_t'\otimes\tilde{\Sigma}_t^{-1})\operatorname{vec}(\nabla_i Q')$$

$$
\begin{aligned}
&= \operatorname{tr}\left(\nabla_j Q \tilde{\Sigma}_t^{-1} \nabla_i Q' \xi_t \xi_t'\right). \\
\nabla_{\phi_i\phi_j} \tilde{l}_t(\vartheta_{s,0}) &= \frac{\partial}{\partial \phi_j} \operatorname{tr}\left(\tilde{\Sigma}_t^{-1} \frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_i} - \tilde{\Sigma}_t^{-1} \tilde{\Sigma}_t^{1/2} \eta_t \eta_t' \tilde{\Sigma}_t^{1/2\prime} \tilde{\Sigma}_t^{-1} \frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_i}\right) \\
&= 2\xi_t' \tilde{\Sigma}_t^{-1} \frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_i} \tilde{\Sigma}_t^{-1} \frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_j} \tilde{\Sigma}_t^{-1} \xi_t - \xi_t' \tilde{\Sigma}_t^{-1} \frac{\partial^2 \tilde{\Sigma}_t(\phi_0)}{\partial \phi_i \partial \phi_j} \tilde{\Sigma}_t^{-1} \xi_t \\
&\quad + \operatorname{tr}\left(\tilde{\Sigma}_t^{-1} \frac{\partial^2 \tilde{\Sigma}_t(\phi_0)}{\partial \phi_i \partial \phi_j}\right) - \operatorname{tr}\left(\frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_i} \tilde{\Sigma}_t^{-1} \frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_j} \tilde{\Sigma}_t^{-1}\right) \\
&= R_t^I + R_t^{II} + R_t^{III} + R_t^{IV}. \\
\nabla_{\varrho_i\phi_j} \tilde{l}_t(\vartheta_{s,0}) &= -2\operatorname{tr}\left(\tilde{\Sigma}_t^{-1} (\xi_t' \otimes I_N) K_{NN} \operatorname{vec}(\nabla_i Q)\, \xi_t' \tilde{\Sigma}_t^{-1} \frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_j}\right) \\
&= -2\xi_t' \tilde{\Sigma}_t^{-1} \frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_j} \tilde{\Sigma}_t^{-1} (\xi_t' \otimes I_N) K_{NN} \operatorname{vec}(\nabla_i Q).
\end{aligned}
$$

Note that $\mathbb{E}[\nabla_{\varrho_i\varrho_j} \tilde{l}_t(\vartheta_{s,0})] = \mathbb{E}[\operatorname{tr}\left(\nabla_j Q \tilde{\Sigma}_t^{-1} \nabla_i Q' \tilde{\Sigma}_t\right)] = \operatorname{tr}\left(\nabla_j Q \nabla_i Q'\right)$ is well-defined by Assumption (A5) and Lemma A.1, since

$$
\mathbb{E} \sup_{\vartheta_s \in \mathcal{B}(\vartheta_{s,0},\delta)} |\operatorname{tr}\left(\nabla_j Q(\varrho) \tilde{\Sigma}_t^{-1}(\phi) \nabla_i Q'(\varrho) \tilde{\Sigma}_t(\phi)\right)| \leq N M_{\tilde{\sigma}} \sup_{\phi \in \mathcal{B}(\phi_0,\delta)} \mathbb{E}[\overline{\sigma_t}(\phi)] < \infty.
$$

Moreover,

$$
\begin{aligned}
\mathbb{E}[R_t^I] &= 2\mathbb{E}\left[\operatorname{tr}\left(\tilde{\Sigma}_t^{-1} \frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_i} \tilde{\Sigma}_t^{-1} \frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_j}\right)\right] = -2\mathbb{E}[R_t^{IV}] \\
\mathbb{E}[R_t^{II}] &= -\mathbb{E}\left[\operatorname{tr}\left(\tilde{\Sigma}_t^{-1} \frac{\partial^2 \tilde{\Sigma}_t(\phi_0)}{\partial \phi_i \partial \phi_j}\right)\right] = -\mathbb{E}[R_t^{III}],
\end{aligned}
$$

and thus

$$
\begin{aligned}
\mathbb{E}[\nabla_{\phi_i\phi_j} \tilde{l}_t(\vartheta_{s,0})] &= \frac{1}{2}\mathbb{E}[R_t^I] = \mathbb{E}\left[\operatorname{vec}\left(\frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_i}\right)' \operatorname{vec}\left(\tilde{\Sigma}_t^{-1} \frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_j} \tilde{\Sigma}_t^{-1}\right)\right] \\
&= \mathbb{E}\left[\frac{\partial \sigma_t(\phi_0)'}{\partial \phi_i} \Delta_\sigma \left(\tilde{\Sigma}_t^{-1} \otimes \tilde{\Sigma}_t^{-1}\right) \Delta_\sigma^+ \frac{\partial \sigma_t(\phi_0)}{\partial \phi_j}\right] = \frac{\partial \sigma_t(\phi_0)'}{\partial \phi_i} \frac{\partial \sigma_t(\phi_0)}{\partial \phi_j}.
\end{aligned}
$$

By Lemma A.3, the dominating function is integrable within a neighborhood about ϕ_0. Furthermore, $\mathbb{E}[\nabla_{\varrho_i\phi_j} \tilde{l}_t(\vartheta_{s,0})] = -2\operatorname{tr}\left(\frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_j} \nabla_i Q'\right) = \frac{\partial \sigma_t(\phi_0)'}{\partial \phi_i} \Delta_\sigma \operatorname{vec}(\nabla_i Q')$ and by Lemma A.3,

$$
\mathbb{E} \sup_{\vartheta_s \in \mathcal{B}(\vartheta_{s,0},\delta)} \left|\operatorname{tr}\left(\frac{\partial \tilde{\Sigma}_t(\phi)}{\partial \phi_j} \tilde{\Sigma}_t^{-1}(\phi) \nabla_i Q'(\varrho)\right)\right| \leq N M_{\tilde{\sigma}} M_{\dot{\sigma}}^{1/3} < \infty.
$$

Therefore, the limit is well-defined and by ergodic theorem, for any $i, j \in \{1, \ldots, d_{\vartheta_s}\}$,

$$
\sup_{\vartheta_s \in \mathcal{B}(\vartheta_{s,0},\delta)} \left|\frac{1}{T} \sum_{t=1}^{T} \nabla_{\vartheta_{s,i}\vartheta_{s,j}} \tilde{l}_t(\vartheta_s) - \mathcal{J}_{i,j}\right| \to 0,
$$

a.s. By Lemma A.3 and using a similar argument for the score, it can be shown that as $T \to \infty$, $\sup_{\vartheta_s \in \mathcal{B}(\vartheta_{s,0},\delta)} \left\| \nabla_{\vartheta_{s,i}\vartheta_{s,j}} \tilde{l}_t(\vartheta_s) - \nabla_{\vartheta_{s,i}\vartheta_{s,j}} l_t(\vartheta_s) \right\| \to 0$, a.s. By definition of $\hat{\vartheta}_{s,T}$, the first order condition $\frac{1}{\sqrt{T}} \sum_{t=1}^{T} \nabla_{\vartheta_s} l_t(\hat{\vartheta}_{s,T}) = 0$ is satisfied with probability approaching one, by a mean value theorem

$$0 = \frac{1}{\sqrt{T}} \sum_{t=1}^{T} \nabla_{\vartheta_s} l_t(\hat{\vartheta}_{s,T}) + \frac{1}{T} \sum_{t=1}^{T} \nabla_{\vartheta_s\vartheta_s} l_t(\vartheta_s^*) \sqrt{T}(\hat{\vartheta}_{s,T} - \vartheta_{s,0}),$$

where ϑ_s^* is a line segment joining $\hat{\vartheta}_{s,T}$ and $\vartheta_{s,0}$. By the results above and Theorem 2.6, $\frac{1}{T} \sum_{t=1}^{T} \nabla_{\vartheta_s\vartheta_s} l_t(\vartheta_s^*) \to \mathcal{J}$ a.s. Thus, given the reduced-form parameter $\vartheta_{r,0}$, as $T \to \infty$,

$$\sqrt{T}(\hat{\vartheta}_{s,T} - \vartheta_{s,0}) | \vartheta_{r,0} \xrightarrow{d} \mathcal{N}(\mathbf{0}, \mathcal{J}^{-1}\mathcal{I}\mathcal{J}^{-1}).$$

Denote the $d_{\vartheta_r} \times d_{\vartheta_s}$ expected Jacobian matrix $\mathcal{H} = \begin{bmatrix} \mathcal{H}_{\varrho\alpha} & \mathcal{H}_{\varrho\omega} \\ \mathcal{H}_{\phi\alpha} & \mathcal{H}_{\varrho\omega} \end{bmatrix}$, where the typical elements in the first d_α rows are

$$\begin{aligned}
\nabla_{\varrho_i\alpha} \tilde{l}_t(\vartheta_{s,0}) &= \frac{\partial}{\partial \alpha'} \mathrm{tr}\left(\xi_t \xi_t' \tilde{\Sigma}_t^{-1} \nabla_i Q' \right) = \frac{\partial}{\partial \alpha'} u_t' \Omega^{-1/2\prime} \tilde{\Sigma}_t^{-1} \nabla_i Q' \Omega^{-1/2} u_t \\
&= -(y_t - (Z_{t-1}' \otimes I_N)\boldsymbol{\alpha}) \Omega^{-1/2\prime} (\nabla_i Q \tilde{\Sigma}_t^{-1} + \tilde{\Sigma}_t^{-1} \nabla_i Q') \Omega^{-1/2} (Z_{t-1}' \otimes I_N) \\
\nabla_{\phi_i\alpha} \tilde{l}_t(\vartheta_{s,0}) &= \frac{\partial}{\partial \alpha'} \mathrm{tr}\left(-\tilde{\Sigma}_t^{-1} \xi_t \xi_t' \tilde{\Sigma}_t^{-1} \frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_i} \right) = -\frac{\partial}{\partial \alpha'} u_t' \Omega^{-1/2\prime} \tilde{\Sigma}_t^{-1} \frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_i} \tilde{\Sigma}_t^{-1} \Omega^{-1/2} u_t \\
&= 2(y_t - (Z_{t-1}' \otimes I_N)\boldsymbol{\alpha}) \Omega^{-1/2\prime} \tilde{\Sigma}_t^{-1} \frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_i} \tilde{\Sigma}_t^{-1} \Omega^{-1/2} (Z_{t-1}' \otimes I_N)
\end{aligned}$$

and in the last d_ω rows are

$$\begin{aligned}
\nabla_{\varrho_i\omega} \tilde{l}_t(\vartheta_{s,0}) &= \frac{\partial}{\partial \omega'} \mathrm{vec}\left(\xi_t \xi_t'\right)' \mathrm{vec}\left(\tilde{\Sigma}_t^{-1} \nabla_i Q' \right) = \mathrm{vec}\left(\tilde{\Sigma}_t^{-1} \nabla_i Q' \right)' \frac{\partial \mathrm{vec}\left(\xi_t \xi_t'\right)}{\partial \mathrm{vech}\left(\Omega_t\right)'} \\
\nabla_{\phi_i\omega} \tilde{l}_t(\vartheta_{s,0}) &= \frac{\partial}{\partial \omega'} \mathrm{vec}\left(\xi_t \xi_t'\right)' \mathrm{vec}\left(-\tilde{\Sigma}_t^{-1} \frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_i} \tilde{\Sigma}_t^{-1} \right) = -\mathrm{vec}\left(\tilde{\Sigma}_t^{-1} \frac{\partial \tilde{\Sigma}_t(\phi_0)}{\partial \phi_i} \tilde{\Sigma}_t^{-1} \right)' \frac{\partial \mathrm{vec}\left(\xi_t \xi_t'\right)}{\partial \mathrm{vech}\left(\Omega_t\right)'},
\end{aligned}$$

where

$$\begin{aligned}
\frac{\partial \mathrm{vec}\left(\xi_t \xi_t'\right)}{\partial \mathrm{vech}\left(\Omega_t\right)'} &= \frac{\partial \mathrm{vec}\left(\Omega^{-1/2} u_t u_t' \Omega^{-1/2\prime} \right)}{\partial \mathrm{vech}\left(\Omega_t\right)'} = (I_{N^2} \otimes K_{NN})(\Omega^{-1/2} u_t \otimes I_N)(u_t' \otimes I_N) \frac{\partial \mathrm{vec}\left(\Omega^{-1/2}\right)}{\partial \mathrm{vech}\left(\Omega_t\right)'} \\
&= -4 N_N (\Omega^{-1/2} u_t \otimes I_N)(u_t' \otimes I_N) N_N (\Omega^{-1/2} \otimes I_N)(\Omega^{-1} \otimes \Omega^{-1}) D_N,
\end{aligned}$$

with N_N being a symmetric idempotent matrix such that $N_N = \frac{1}{2}(I_{N^2} + K_{NN})$. If the matrix-root $\Omega^{-1/2}$ is chosen to be symmetric, then $N_N = K_{NN} = I_{N^2}$ and the terms simplify to $-4(\Omega^{-1/2} u_t u_t' \Omega^{-1/2} \otimes I_N)(\Omega^{-1} \otimes \Omega^{-1}) D_N$. Thus, $\mathcal{H}_{\varrho\alpha} = \mathbf{0}_{d_\varrho \times d_\alpha}$, $\mathcal{H}_{\phi\alpha} = \mathbf{0}_{d_\phi \times d_\alpha}$ and typical elements in $\mathcal{H}_{\varrho\omega}$ and $\mathcal{H}_{\phi\omega}$ are $-4\mathrm{vec}\left(\nabla_i Q'\right)' (\Omega^{-1} \otimes \Omega^{-1}) D_N$ and $4 \frac{\partial \sigma_t(\phi_0)'}{\partial \phi_i} \Delta_\sigma (\Omega^{-1} \otimes \Omega^{-1}) D_N$, respectively.

The score functions for the reduced-form estimator $\hat{\vartheta}_{r,T}$ are

$$\nabla_\alpha g_t(\vartheta_r) = 2(y_t - (Z'_{t-1} \otimes I_N)\alpha)'\Omega^{-1}(Z'_{t-1} \otimes I_N)$$
$$\nabla_\omega g_t(\vartheta_r) = \text{vec}\left(\Omega^{-1}u_t u_t'\Omega^{-1}\right)' D_N - \text{vec}\left(\Omega^{-1}\right)' D_N,$$

and the expected Hessian matrix at $\vartheta_{r,0}$ is $\mathcal{G} = \begin{bmatrix} \mathcal{G}_{\alpha\alpha} & \mathcal{G}_{\alpha\omega} \\ \mathcal{G}'_{\alpha\omega} & \mathcal{G}_{\omega\omega} \end{bmatrix}$, where

$$\mathcal{G}_{\alpha\alpha} = \mathbb{E}[\nabla_{\alpha\alpha} g_t(\vartheta_{r,0})] = -2\mathbb{E}[(Z_{t-1} \otimes I_N)\Omega^{-1}[(Z'_{t-1} \otimes I_N)] = -2\mathbb{E}[Z_{t-1}Z'_{t-1}] \otimes \Omega^{-1}$$
$$\mathcal{G}_{\omega\omega} = \mathbb{E}[\nabla_{\omega\omega} g_t(\vartheta_{r,0})] = D'_N \mathbb{E}\left[\left((I_N \otimes \Omega^{-1}u_t u_t') + (\Omega^{-1}u_t u_t' \otimes I_N) - I_{N^2}\right)\frac{\partial \text{vec}\left(\Omega^{-1}\right)}{\partial \text{vec}\left(\Omega_t\right)'}\right] D_N$$
$$= D'_N \frac{\partial \text{vec}\left(\Omega^{-1}\right)}{\partial \text{vec}\left(\Omega_t\right)'} D_N = -D'_N(\Omega^{-1} \otimes \Omega^{-1})D_N$$
$$\mathcal{G}_{\alpha\omega} = \mathbb{E}[\nabla_{\alpha\omega} g_t(\vartheta_{r,0})] = \mathbf{0}_{d_\alpha \times d_\omega}.$$

Let $\mathcal{W}_{\varrho\omega} = \mathbb{E}[\nabla_\varrho \tilde{l}_t(\vartheta_{s,0})\nabla_\omega g_t(\vartheta_{r,0})']$, $\mathcal{W}_{\phi\omega} = \mathbb{E}[\nabla_\phi \tilde{l}_t(\vartheta_{s,0})\nabla_\omega g_t(\vartheta_{r,0})']$ and define $d_{\vartheta_s} \times d_\omega$ matrix $\mathcal{W}_\omega = [\mathcal{W}'_{\varrho\omega} \quad \mathcal{W}'_{\phi\omega}]'$. Therefore,

$$\sqrt{T}(\hat{\vartheta}_{s,T} - \vartheta_{s,0}) \xrightarrow{d} \mathcal{N}(\mathbf{0}, V),$$

with the asymptotic covariance matrix given by (see also Theorem 6.1 in Newey and McFadden 1994)

$$V = \mathcal{J}^{-1}\mathbb{E}\left[\left(\nabla_{\vartheta_s}\tilde{l}_t(\vartheta_{s,0}) - \mathcal{H}\mathcal{G}^{-1}\nabla_{\vartheta_r} g_t(\vartheta_{r,0})\right)\left(\nabla_{\vartheta_s}\tilde{l}_t(\vartheta_{s,0}) - \mathcal{H}\mathcal{G}^{-1}\nabla_{\vartheta_r} g_t(\vartheta_{r,0})\right)'\right]\mathcal{J}^{-1},$$

where

$$\nabla_{\vartheta_s}\tilde{l}_t(\vartheta_{s,0}) - \mathcal{H}\mathcal{G}^{-1}\nabla_{\vartheta_r} g_t(\vartheta_{r,0}) = \begin{bmatrix} \nabla_\varrho \tilde{l}_t(\vartheta_{s,0}) \\ \nabla_\phi \tilde{l}_t(\vartheta_{s,0}) \end{bmatrix} - \begin{bmatrix} \mathbf{0}_{d_\varrho \times d_\alpha} & \mathcal{H}_{\varrho\omega} \\ \mathbf{0}_{d_\phi \times d_\alpha} & \mathcal{H}_{\phi\omega} \end{bmatrix}\begin{bmatrix} \mathcal{G}^{-1}_{\alpha\alpha}\nabla_\alpha g_t(\vartheta_{r,0}) \\ \mathcal{G}^{-1}_{\omega\omega}\nabla_\omega g_t(\vartheta_{r,0}) \end{bmatrix}$$
$$= \begin{bmatrix} \nabla_\varrho \tilde{l}_t(\vartheta_{s,0}) - \mathcal{H}_{\varrho\omega}\mathcal{G}^{-1}_{\omega\omega}\nabla_\omega g_t(\vartheta_{r,0}) \\ \nabla_\phi \tilde{l}_t(\vartheta_{s,0}) - \mathcal{H}_{\phi\omega}\mathcal{G}^{-1}_{\omega\omega}\nabla_\omega g_t(\vartheta_{r,0}) \end{bmatrix}.$$

Denote the last d_ω columns of matrix $\mathcal{H}$ by $\mathcal{H}$, we have

$$V = \mathcal{J}^{-1}\left(\mathcal{I} - \mathcal{W}_\omega G^{-1}_{\omega\omega}\mathcal{H}'_\omega - \mathcal{H}_\omega G^{-1}_{\omega\omega}\mathcal{W}'_\omega + \mathcal{W}_\omega V_{\omega\omega}\mathcal{W}'_\omega\right)\mathcal{J}^{-1},$$

where the asymptotic covariance matrix of $\hat{\omega}_T$ is given by (see Thm.2.1 in Brüggemann et al. 2016)

$$V_{\omega\omega} = G^{-1}_{\omega\omega}\sum_{j=-\infty}^{\infty}\text{Cov}\left(D'_N\text{vec}\left(\Omega^{-1}u_t u_t'\Omega^{-1}\right) - D'_N\text{vec}\left(\Omega^{-1}\right),\right.$$

$$\operatorname{vec}\left(\Omega^{-1}u_{t-j}u_{t-j}'\Omega^{-1}\right)' D_N' - \operatorname{vec}\left(\Omega^{-1}\right)' D_N' \Big) G_{\omega\omega}^{-1}$$
$$= D_N^+ \left(\sum_{j=-\infty}^{\infty} \mathbb{E}\left[\operatorname{vec}\left(u_t u_t'\right) \operatorname{vec}\left(u_{t-j}u_{t-j}'\right)' \right] - \operatorname{vec}(\Omega)\operatorname{vec}(\Omega)' \right) D_N^{+\prime}.$$

This completes the proof. □

Lemma A.1. There exist $0 < M_\sigma, M_{\tilde{\sigma}}, M_{\sigma'} < \infty$, $\psi_\sigma \in [0,1)$, s.t.,

$$\sup_{\phi\in\Phi} \|\Sigma_t^{-1}(\phi)\| \le M_\sigma, \ \sup_{\phi\in\Phi} \|\tilde{\Sigma}_t^{-1}(\phi)\| \le M_{\tilde{\sigma}}, \text{ and } \sup_{\phi\in\Phi} \|\tilde{\Sigma}_t(\phi) - \Sigma_t(\phi)\| \le \psi_\sigma^t M_{\sigma'}, \text{ a.s.}$$

Proof. Note that $\|\Sigma_t^{-1}(\phi)\| = \rho^{1/2}(\Sigma_t^{-1}(\phi)\Sigma_t^{-1}(\phi)) = \max_{j\in\{1,\dots,N\}} \sigma_{jt}^{-1}(\phi) = \underline{\sigma_t}^{-1}(\phi)$ and similarly, $\|\tilde{\Sigma}_t^{-1}(\phi)\| = \underline{\tilde{\sigma}_t}^{-1}(\phi)$, where $\underline{\tilde{\sigma}_t}(\phi)$ is defined analogously. Since the parameter space Φ is compact and σ_t, $\tilde{\sigma}_t$ are positive definite and continuous in ϕ, there exists $M_\sigma^{-1}, M_{\tilde{\sigma}}^{-1} \in \mathbb{R}^+$, s.t. for all t, $0 < M_\sigma^{-1} \le \inf_{\phi\in\Phi} \underline{\sigma_t}(\phi) < \infty$ and $0 < M_{\tilde{\sigma}}^{-1} \le \inf_{\phi\in\Phi} \underline{\tilde{\sigma}_t}(\phi) < \infty$. To show $\sup_{\phi\in\Phi} \|\tilde{\Sigma}_t(\phi) - \Sigma_t(\phi)\| \le \psi_\sigma^t M_{\sigma'}$, note that by recursive substitutions of model (2.7), we have

$$\sigma_t = \Gamma_1^t \sigma_0 + \sum_{i=0}^{t-1} \left[\Gamma_1^i (\gamma_0 + G_1(\xi_{t-i-1} \odot \xi_{t-i-1}) \right].$$

Thus,

$$\|\tilde{\Sigma}_t(\phi) - \Sigma_t(\phi)\| \le \|\Delta_\sigma^+\| \|\tilde{\sigma}_t(\phi) - \sigma_t(\phi)\| \le \|\Gamma_1^t(\phi)\| \|\tilde{\sigma}_0 - \sigma_0\|,$$

since $\|\Delta_\sigma^+\| = 1$. Choose $\psi_\sigma := \sup_{\phi\in\Phi} \rho(\Gamma_1(\phi)) \in [0,1)$ and define constant $M_{\sigma'} := \|\tilde{\sigma}_0 - \sigma_0\|$, which completes the proof. □

Lemma A.2. For any $\delta > 0$, there exists an open ball $\mathcal{B}(\vartheta_{s,0}, \delta)$, s.t.,

$$\inf_{\vartheta_s \in \mathcal{B}(\vartheta_{s,0},\delta)^c \cap \Theta_s} \mathbb{E}\left[\tilde{l}_t(\vartheta_s)\right] - \mathbb{E}\left[\tilde{l}_t(\vartheta_{s,0})\right] > 0$$

Proof. We start by showing that the population objective function is well-defined at the true-parameter $\vartheta_{s,0}$. Note that $\varepsilon_t' Q(\varrho_0)\Sigma_t(\phi_0)^{-1}Q(\varrho_0)'\varepsilon_t = \eta_t'\eta_t$, by Jensen's inequality and Assumption (A2),

$$\mathbb{E}\left[\tilde{l}_t(\vartheta_{s,0})\right] = \mathbb{E}[\log|\Sigma_t(\phi_0)|] + \mathbb{E}[\eta_t'\eta_t] \le \frac{1}{s}\mathbb{E}\left[\log \prod_{j=1}^{N} \sigma_{jt}^s\right] + \mathbb{E}\|\eta_t\|^2$$
$$\le \frac{1}{s}\log \mathbb{E}[(\overline{\sigma_t}^s)^N] + \mathbb{E}\|\eta_t\|^2 < \infty.$$

Since in case $\mathbb{E}\left[\tilde{l}_t(\vartheta_s)\right] = +\infty$ the inequality holds trivially, we assume that the population objective function is finite and show $\mathbb{E}\left[\tilde{l}_t(\vartheta_s)\right] > -\infty$ everywhere on Θ_s. Since $|\tilde{\Sigma}_t(\phi)|^{-1} = \left(\prod_{j=1}^N \tilde{\sigma}_{jt}(\phi)\right)^{-1} \leq \underline{\tilde{\sigma}_t}^{-1}(\phi) \leq M_{\tilde{\sigma}}^N < \infty$ by Lemma A.1 and due to the elementary relation $(f_1 + f_2)^- \leq f_1^-$ if $f_2 \geq 0$, we have

$$\mathbb{E}\left[\tilde{l}_t^-(\vartheta_s)\right] = \mathbb{E}\left[\left((Q(\varrho)'\varepsilon_t)'\Sigma_t(\phi)^{-1}(Q(\varrho)'\varepsilon_t) + \log|\tilde{\Sigma}_t(\phi)|\right)^-\right] \leq \mathbb{E}\left[\log^-|\tilde{\Sigma}_t(\phi)|\right] < \infty.$$

Furthermore, for any $\vartheta_s \in \Theta_s$,

$$\begin{aligned}
\mathbb{E}\left[\tilde{l}_t(\vartheta_s)\right] - \mathbb{E}\left[\tilde{l}_t(\vartheta_{s,0})\right] &= \mathbb{E}[\log|\tilde{\Sigma}_t(\phi)| - \log|\tilde{\Sigma}_t|] + \mathbb{E}[\xi_t'\tilde{\Sigma}_t^{-1}(\phi)\xi_t - \xi_t'\tilde{\Sigma}_t^{-1}\xi_t] \\
&\quad + \mathbb{E}[\xi_t(\varrho)'\tilde{\Sigma}_t^{-1}(\phi)\xi_t(\varrho) - \xi_t'\tilde{\Sigma}_t^{-1}(\phi)\xi_t] \\
&= \mathbb{E}[\log|\tilde{\Sigma}_t(\phi)\tilde{\Sigma}_t^{-1}|] + \mathbb{E}[\xi_t'(\tilde{\Sigma}_t^{-1}(\phi) - \tilde{\Sigma}_t^{-1})\xi_t] + \mathbb{E}\left[\operatorname{tr}\left(\tilde{\Sigma}_t^{-1}(\phi)(\xi_t(\varrho)\xi_t(\varrho)' - \xi_t\xi_t')\right)\right] \\
&:= (I) + (II) + (III),
\end{aligned}$$

where the parameter dependence of Σ_t and ξ_t at the true value $\vartheta_{s,0}$ is suppressed to simplify the notation. For all t and $i, j \in \{1, \ldots, N\}$, define $\lambda_{ij,t}(\vartheta_s) = \tilde{\sigma}_{it}/\tilde{\sigma}_{jt}(\vartheta_s)$ with $\inf_{\vartheta_s \in \Theta_s} \lambda_{ij,t}(\vartheta_s) > 0$. Then,

$$(I) + (II) = \mathbb{E}\left[\operatorname{tr}\left(\tilde{\Sigma}_t^{-1}(\phi)\tilde{\Sigma}_t - I_N\right) - \log|\tilde{\Sigma}_t\tilde{\Sigma}_t^{-1}(\phi)|\right] = \mathbb{E}\left[\sum_{j=1}^N \left(\lambda_{jj,t}(\vartheta_s) - 1 - \log\lambda_{jj,t}(\vartheta_s)\right)\right].$$

Furthermore, $(III) := \mathbb{E}\left[\sum_{j=1}^N \left(\lambda_{jj,t}(\vartheta_s)q_{jj}^2(\varrho) + \sum_{i\neq j}\lambda_{ij,t}(\vartheta_s)q_{ij}^2(\varrho) - \lambda_{jj,t}(\vartheta_s)\right)\right]$ with $q_{ij}(\varrho) = e_i'Q(\varrho)e_j$. Therefore,

$$\mathbb{E}\left[\tilde{l}_t(\vartheta_s)\right] - \mathbb{E}\left[\tilde{l}_t(\vartheta_{s,0})\right] = \mathbb{E}\left[\sum_{j=1}^N \left(\lambda_{jj,t}(\vartheta_s)q_{jj}^2(\varrho) + \sum_{i\neq j}\lambda_{ij,t}(\vartheta_s)q_{ij}^2(\varrho) - 1 - \log\lambda_{jj,t}(\vartheta_s)\right)\right].$$

where $\lambda_{jj,t}(\vartheta_s)q_{jj}^2(\varrho) + \sum_{i\neq j}\lambda_{ij,t}(\vartheta_s)q_{ij}^2$ is minimized if each column in $Q(\varrho)$ contains only one non-zero element ± 1, i.e., $Q(\varrho) = Q\mathcal{P}\mathcal{D}$ and $\varrho = \varrho_0$ given the identifiability assumption. Without loss of generality, choose the permutation $q_{jj}^2 = 1$ and $q_{ij}^2 = 0$ for $i \neq j$, then $\lambda_{jj,t}(\vartheta_s)q_{jj}^2(\varrho) + \sum_{i\neq j}\lambda_{ij,t}(\vartheta_s)q_{ij}^2(\varrho) = \lambda_{jj,t}(\vartheta_s)$ for all j, and

$$\mathbb{E}\left[\tilde{l}_t(\vartheta_s)\right] - \mathbb{E}\left[\tilde{l}_t(\vartheta_{s,0})\right] = \mathbb{E}\left[\sum_{j=1}^N \left(\lambda_{jj,t}(\vartheta_s) - 1 - \log\lambda_{jj,t}(\vartheta_s)\right)\right] \geq 0$$

by the elementary relation $x - 1 \geq \log x$ on $\mathbb{R}^+$ with the equality applicable if and only if $\lambda_{jj,t}(\vartheta_s) - 1 = 0$ for all j and all t, a.s., which in turn implies $\tilde{\Sigma}_t^{-1}(\phi)\tilde{\Sigma}_t = I_N$ and $\tilde{\Sigma}_t = \tilde{\Sigma}_t(\phi)$, $\forall t$ and thus $\phi = \phi_0$ given the identifiability. Therefore, the inequality holds strictly for all $\vartheta_s \in \mathcal{B}(\vartheta_{s,0}, \delta)^c \cap \Theta_s$ with $\delta > 0$. □

Lemma A.3. For some $s > 1$, there exists an open neighborhood $\mathcal{B}(\phi_0, \delta)$, such that

$$\mathbb{E}\left[\sup_{\phi\in\mathcal{B}(\phi_0,\delta)} \|\tilde{\Sigma}_t^{-1/2}(\phi)\tilde{\Sigma}_t^{1/2}(\phi_0)\|^{2s}\right] < \infty,$$

and for any $i, j \in \{1, \ldots, d_\phi\}$,

$$\mathbb{E}\left[\sup_{\phi\in\mathcal{B}(\phi_0,\delta)} \left\|\tilde{\Sigma}_t^{-1/2}(\phi)\frac{\partial\tilde{\Sigma}_t(\phi)}{\partial\phi_i}\tilde{\Sigma}_t^{-1/2'}(\phi)\right\|^{2s}\right] < \infty, \quad \mathbb{E}\left[\sup_{\phi\in\mathcal{B}(\phi_0,\delta)} \left\|\tilde{\Sigma}_t^{-1/2}(\phi)\frac{\partial^2\tilde{\Sigma}_t(\phi)}{\partial\phi_i\partial\phi_j}\tilde{\Sigma}_t^{-1/2'}(\phi)\right\|^{2s}\right] < \infty.$$

Proof. First, by Assumption (A5),

$$\begin{aligned}
\mathbb{E}\left[\sup_{\phi\in\mathcal{B}(\phi_0,\delta)} \|\tilde{\Sigma}_t^{-1/2}(\phi)\tilde{\Sigma}_t^{1/2}(\phi_0)\|^{2s}\right] &\leq \mathbb{E}\left[\sup_{\phi\in\mathcal{B}(\phi_0,\delta)} \rho(\tilde{\Sigma}_t^{-1}(\phi))^s \rho(\tilde{\Sigma}_t(\phi_0))^s\right] \\
&\leq \mathbb{E}\left[\sup_{\phi\in\mathcal{B}(\phi_0,\delta)} \operatorname{tr}\left(\tilde{\Sigma}_t^{-1}(\phi)\right)^s \operatorname{tr}\left(\tilde{\Sigma}_t(\phi_0)\right)^s\right] \\
&\leq \mathbb{E}\left[\sup_{\phi\in\mathcal{B}(\phi_0,\delta)} N^{2s}\|\tilde{\Sigma}_t^{-1}(\phi)\|^s\|\tilde{\Sigma}_t(\phi_0)\|^s\right] \leq N^{2s} M_{\tilde{\sigma}}^s \mathbb{E}[\overline{\sigma}_t^s(\phi_0)] < \infty,
\end{aligned}$$

with constant $M_{\tilde{\sigma}}^s$ defined analogously as in Lemma A.1 and $M_{\tilde{\sigma}}^s < \infty$ due to the compactness of the parameter space. Moreover, under Assumption (A5), there exists $0 < M_{\dot{\sigma}}, M_{\ddot{\sigma}} < \infty$, s.t., for any $i, j \in \{1, \ldots, d_\phi\}$,

$$\mathbb{E}\left[\sup_{\phi\in\mathcal{B}(\phi_0,\delta)} \left\|\frac{\partial\tilde{\Sigma}_t(\phi)}{\partial\phi_i}\right\|^3\right] \leq M_{\dot{\sigma}} \quad \text{and} \quad \mathbb{E}\left[\sup_{\phi\in\mathcal{B}(\phi_0,\delta)} \left\|\frac{\partial^2\tilde{\Sigma}_t(\phi)}{\partial\phi_i\partial\phi_j}\right\|^3\right] \leq M_{\ddot{\sigma}}, \tag{A.1}$$

by Lemma 3 in Hafner and Preminger (2009b). Thus, for any $1 < s \leq 3/2$,

$$\begin{aligned}
\mathbb{E}\left[\sup_{\phi\in\mathcal{B}(\phi_0,\delta)} \left\|\tilde{\Sigma}_t^{-1/2}(\phi)\frac{\partial\tilde{\Sigma}_t(\phi)}{\partial\phi_i}\tilde{\Sigma}_t^{-1/2'}(\phi)\right\|^{2s}\right] &\leq \mathbb{E}\left[\sup_{\phi\in\mathcal{B}(\phi_0,\delta)} 2N^s\|\tilde{\Sigma}_t^{-1}(\phi)\|^s \left(\left\|\frac{\partial\tilde{\Sigma}_t(\phi)}{\partial\phi_i}\right\|^3\right)^{2s/3}\right] \\
&\leq 2N^s M_{\tilde{\sigma}}^s M_{\dot{\sigma}}^{2s/3} < \infty,
\end{aligned}$$

by Lemma A.1 and Jensen's inequality and

$$\begin{aligned}
\mathbb{E}\left[\sup_{\phi\in\mathcal{B}(\phi_0,\delta)} \left\|\tilde{\Sigma}_t^{-1/2}(\phi)\frac{\partial^2\tilde{\Sigma}_t(\phi)}{\partial\phi_i\partial\phi_j}\tilde{\Sigma}_t^{-1/2'}(\phi)\right\|^{2s}\right] &\leq \mathbb{E}\left[\sup_{\phi\in\mathcal{B}(\phi_0,\delta)} 2N^s\|\tilde{\Sigma}_t^{-1}(\phi)\|^s \left(\left\|\frac{\partial^2\tilde{\Sigma}_t(\phi)}{\partial\phi_i\partial\phi_j}\right\|^3\right)^{2s/3}\right] \\
&\leq 2N^s M_{\tilde{\sigma}}^s M_{\ddot{\sigma}}^{2s/3} < \infty,
\end{aligned}$$

which completes the proof. □