Essays in Empirical Monetary Economics

Dissertation

zur Erlangung des Doktorgrades Dr. rer. pol.

der Fakultät für Wirtschaftswissenschaften
der Universität Duisburg-Essen

von

Gabriel Arce Alfaro

aus

San Jose, Costa Rica

Betreuer: Prof. Dr. Christoph Hanck
Lehrstuhl für Ökonometrie

Essen, April 2023

DOI: 10.17185/duepublico/81576
URN: urn:nbn:de:hbz:465-20240301-063901-3

Erstreferent: Prof. Dr. Christoph Hanck
Zweitreferent: Prof. Dr. Ludger Linneman
Tag der mündlichen Prüfung: 11.10.2023

Acknowledgments

This thesis would not have been possible without the support, guidance and encouragement that I received from many people during the past years. To all of you, I want to express my most sincere words of gratitude. To my main supervisor Christoph Hanck for his unconditional guidance and support. To Ludger Linnemann, second supervisor of this thesis, for all his valuable comments and advice. To the memory of Ansgar Belke who sadly left us too soon. I am highly indebted to my colleague and friend Sina Asshoff for all her support throughout this process, many e-mails and countless cups of coffee. I am deeply grateful to my friend and co-author Boris Blagov for his willingness to share his knowledge with me.

Throughout the writing up of this PhD thesis, I benefited greatly from research stays at the Bank of Mexico (Banxico), Deutsche Bundesbank and Bank of England. I am very grateful to Raul Ibarra Ramirez, Miroslava Quiroga Treviño, Patrick Huertgen, PhD interns from the Bank of Mexico and Bank of England, as well as colleagues from the Current Economic Conditions Division (CECD) and Monetary Analysis (MA) from the Bank of England.

I gratefully acknowledge the financial support from the Ruhr Graduate School in Economics and Mercator Scholarship. To my RGS colleagues, many thanks for all the fruitful discussions and comments. To the managing director, Helge Braun, many thanks for all his support during the past years. Special thanks to Torsten Schmidt whose support was invaluable during the writing up of this thesis.

To my family for their unconditional love and encouragement. Lastly, to Sophie whose support and love became a constant source of motivation.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The importance of monetary economics as a field of research lies in the potential monetary policies have in shaping real economic outcomes. Monetary policy decisions implemented by central banks have the capacity to affect aggregate variables such as output, inflation, cost of borrowing, among others; and thus, affect households' and firms' economic decisions. However, these shifts in monetary policy are endogenous responses to the state of the economy, where central banks set their monetary policy in a forward-looking manner in anticipation to future events, e.g. lowering (raising) their policy rate when expecting a recession (boom). Therefore, most of the literature in empirical monetary economics is concerned with identifying how much of the variation in their policy instruments are systematic responses to the state of the economy (endogenous responses) and how much is non-systematic (exogenous).

It is therefore of primary interest to understand and identify the exogenous component of monetary policy as it plays a role as an independent source of influence on aggregate variables. Christiano et al. (1999) formally introduce the concept of a monetary policy shock by relating the "setting" of the policy rate to its systemic and non-systemic components as follows:

$$S_t = f(\Omega_t) + \sigma_t \epsilon_t^s. \tag{1.1}$$

In eq. (1.1), the monetary policy instrument S_t is set as a function of the central bank's information set Ω_t and the exogenous component, $\sigma_s \epsilon_t^s$. In this regard, the latter represents the monetary policy shock.

Since the seminal contribution from Sims (1980), vector autoregressive models (VARs) have become a predominant statistical tool to perform economic forecasts and recover the structure of the macroeconomy from the data. In this regard, monetary economics appears as one of the strands of the literature where VAR models have been widely applied to the analysis of the macroeconomic effects of unexpected shocks. We can therefore characterize the standard structural vector autoregressive model as:

$$B_0 y_t = v + \sum_{i=1}^{p} B_i y_{t-i} + \epsilon_t,\ \epsilon_t \sim N(0, I_K), \tag{1.2}$$

where $y_t = (y_{1t}, ..., y_{Kt})'$ represents the $K \times 1$ vector of endogenous variables, v is a vector of constants of dimension $K \times 1$, B_i is an $K \times K$ matrix of structural coefficients with $i = 1, ..., p$ and ϵ_t represents the structural innovations. Thus, premultiplying eq. (1.2) by B_0^{-1} yields its reduced-form representation counterpart:

$$y_t = c + \sum_{i=1}^{p} A_i y_{t-i} + u_t,\ u_t \sim N(0, \Sigma_u), \tag{1.3}$$

with A_i and u_t representing the reduced-form coefficients and innovations, correspondingly. Thus, the core empirical strategy of this thesis lies on finding a linear mapping between the unorthogonalized reduced-form errors and the structural shocks from equation (1.3) from the identity $u_t = B_0^{-1}\epsilon_t$, satisfying the condition: $B_0^{-1}(B_0^{-1})' = \Sigma_u$, with Σ_u representing the reduced-form variance-covariance matrix.

The impact matrix, B_0, contains the contemporaneous (causal) relationship among the variables of the system. Given the symmetry of Σ_u, we can only uniquely identify $K(K+1)/2$ out of the K^2 parameters in B_0, leaving the model with $K(K-1)/2$ unrestricted parameters.

This condition is commonly known in the literature as the identification problem.

The very nature of the identification problem lies on the assumptions placed on the matrix B_0 to recover the unobserved structural responses from the estimated reduced-form counterpart. Thus, an important amount of literature has focused on motivating innovative ways of recovering the structural responses, i.e. identification strategies.

This thesis consists of four self-contained chapters. Although, each of them addresses independent research questions, they are intertwined by the aim of providing a better understanding of the transmission of economic shocks and the channels through which monetary policy affects the economy. The empirical strategy used throughout this dissertation builds on some extension of the standard VAR model and thus, the underlying identification problem.

Chapter 2 is concerned with the effects of uncertainty shocks on the inflation expectations formation process among economic agents. The seminal contribution from Bloom (2009) gave rise to an active field of research on the effects of uncertainty shocks on the economy (Baker et al., 2016; Leduc and Liu, 2016; Bachmann et al., 2019). This chapter contributes to the literature by obtaining an estimate for monetary policy uncertainty (MPU) and estimating its effects on macroeconomic variables. We proxy monetary policy uncertainty as the second moment of the central bank policy rate. To do so, we follow Mumtaz and Zanetti (2013) and estimate a structural vector autoregressive model with stochastic volatility-in-mean for the United States (U.S.).

Our results suggest that MPU shocks have a contractionary effect on output and a negative response on short-run inflation expectations. However, the magnitude of the response of inflation expectations is lower compared to inflation, suggesting rigidity in the inflation expectations' formation process. Furthermore, we observe that long-run inflation expectations are not affected by changes in MPU.

Considering that important changes have occurred in the U.S. economy during the last decades, from the great moderation to the global financial crisis, we relaxed the modeling

assumption of constant coefficients and extended the methodology to include time-varying parameters and thus, capture the evolving responses of inflation expectations to changes in the MPU. We observe that, despite having an important effect throughout most of the sample, MPU shocks have affected short and long-run inflation expectations in the last decade, even when accounting for the zero lower bound.

Chapter 3 investigates the degree of commonality in the bank lending rates and credit volumes across the euro area. To achieve that, we estimate a multi-level dynamic factor model (DFM) with time-varying factor loadings. The multi-level structure allows us to differentiate between fluctuations explained by common components and those that are driven by country-specific factors while the time-varying component in the model enables us to obtain the share which is explained by each of these components over time. Our results show that the common factor accounts for the highest share in explaining fluctuations in the bank lending rates. However, this degree of commonality has decreased since the sovereign debt crisis and remained below its previous level. On the other hand, the share explained by the common factor in the credit sector has remained stable throughout the sample, increasing only during periods of high financial volatility. Furthermore, our results suggest a difference in the relative importance of the common factor in countries which experienced fiscal or financial stress during the sovereign debt crisis, pointing at an increase in financial fragmentation within the euro area with respect to their members' financing conditions.

In Chapter 4, we introduce a novel identification strategy for structural vector autoregressive models. The method proposed combines information from the forecast error variance (FEV) with classical sign restrictions. Our methodology aims to overcome frequent shortcomings of purely sign-identified models. Namely, restrictions on the sign responses alone lead to accepting a broad range of equally likely models with different shapes and magnitudes of the variables responses. All of this together results in wider error bands compared to alternative identification methods, causing uninformative conclusions about the economic

consequences of unexpected shocks.

Results from a Monte Carlo simulation show that the method proposed successfully recovers the true data generation process (DGP) responses while delivering sharper identification. Furthermore, we illustrate the applicability of the method proposed with two applications on the effects of a monetary policy tightening. For the first application, the monetary policy model for the U.S. economy, we show that uninformative responses of prices to a contractionary monetary policy are explained by retaining draws which are associated with a low contribution of the monetary policy shock to inflation FEV. When identifying the model with the proposed identification scheme we observe that this is no longer the case and thus, resolve retaining puzzling responses of prices to a tightening in the policy rate. The second application, the case of the U.K. economy, shows how failing to exclude draws that have a contribution to the FEV close to zero on the variable of interest can potentially lead to conclude that the persistence of the shock is shorter lasting and of lower magnitude. This can be alleviated when applying the proposed identification method.

Building on the latter, Chapter 5 reevaluates the effects of a shortfall in oil supply exploiting the advantages of the identification method developed in the previous chapter of this thesis. We re-estimate some of the benchmark models in the oil market literature updating the dataset from Kilian (2009). Therefore, we can observe that contrasting magnitudes in the responses of real prices to flow supply shocks can be explained by differences in the contribution of this shock to oil prices FEV. Thus, in order to consider estimation uncertainty, we incorporate bounds on the contribution of oil supply shocks to real oil prices FEV following results from Caldara et al. (2019). Our baseline results suggest that there is a substantial response of oil prices to a disruption in oil supply. These results contrast with models that ad-hoc assume a contribution close to zero of oil supply shocks to oil prices FEV, either via incorporating zero restrictions or by setting an upper bound on the price elasticity of oil supply, and thus observe a very mild response of oil prices to a cut in oil supply flow. We

then consider a medium-scale model for the oil market, similar to Lippi and Nobili (2012) and Känzig (2021), accounting for additional sources of fluctuations in oil prices. When comparing our baseline results with classical sign restrictions we observe that allowing, in the set of retain draws, rotations that are associated with a close to zero contribution of oil supply shocks to oil prices FEV, leads to a higher degree of uncertainty around the median response of oil prices to a shortfall in oil supply. Therefore, our identification method stands as an appealing method to obtain more informative results in terms of the magnitude and persistence of the responses.

Although each of the following four chapters in this dissertation deals with different research questions, their common objective is to provide a better understanding of crucial topics in the realm of empirical macroeconomics.

Chapter 2

Monetary Policy Uncertainty and Inflation Expectations

Abstract

Do inflation expectations react to changes in the volatility of monetary policy? They do, but only until the global financial crisis. This paper investigates whether increasing the dispersion of monetary policy shocks, which is interpreted as elevated uncertainty surrounding monetary policy, affects the inflation expectation formation process. Based on U.S. data since the 1980s and a stochastic volatility-in-mean structural VAR model we find that monetary policy uncertainty reduces both inflation expectations and inflation. However, after the Great Recession this link has disappeared, even when controlling for the Zero Lower Bound.

Published in Oxford Bulletin of Economics and Statistics, DOI: 10.1111/OBES.12516
Co-author: Boris Blagov
Keywords Monetary policy uncertainty, inflation expectations, SVAR volatility-in-mean, time-varying coefficients.

1 Introduction

"Inflation targeting, at least in its best-practice form, consists of two parts: a policy framework of constrained discretion and a communication strategy that attempts to focus expectations and explain the policy framework to the public. Together, these two elements promote both price stability and well-anchored inflation expectations." - Bernanke (2003)

The seminal contribution of Lucas (1976) has highlighted the importance of expectations for the effectiveness of economic policy. Inflation expectations play a key role for wage and price setting (Clarida et al., 2000; Svensson, 2000). As a consequence, over the past 40 years, the conduct of monetary policy has markedly changed with increasing emphasis on communication and transparency (Bernanke and Woodford, 1997; Blinder et al., 2008). For example, in 1994 the Federal Open Market Committee (FOMC) began releasing statements regarding their monetary policy decisions and in 2012 the U.S. Federal Reserve (FED) officially adopted an inflation target of 2%. Since the global financial crisis, the main monetary policy instrument - the federal funds rate (FED funds) - has been constrained by the zero lower bound (ZLB). Recent literature has looked at the relevance of the zero lower bound in light of unconventional monetary policy. While Swanson and Williams (2014) and Debortoli et al. (2019) suggest that the ZLB has not been particularly binding due to the effectiveness of forward guidance, Ikeda et al. (2020) find strong evidence that the constraint has been empirically relevant in the U.S. and Japan. Therefore, managing inflation expectations has become even more important for the transmission mechanism of monetary policy.

The effectiveness of monetary policy relies in part on the ability of economic agents to anticipate monetary policy movements. Consequently a large part of monetary economics has dealt with studying the effects of unanticipated monetary policy shocks. However, a growing literature has also considered a different take on the relationship - how does unpredictability of monetary policy affect the economy? This is referred to monetary policy uncertainty (MPU).

Although MPU has been studied not least due to its relevance for central bank credibility (e.g., Stulz, 1986; Swanson, 2006; Neely, 2005), it has gained considerable emphasis since the global financial crisis, along with other types of uncertainty.[1]

A commonality across the empirical MPU literature findings is that, irrespective of the chosen proxy, an increase in monetary policy uncertainty suppresses economic activity, increases unemployment and leads to a decline in prices. The theoretical underpinnings of these findings suggest that the consumption channel plays an important role - risk averse agents hold back consumption, which creates a decline in demand (Mumtaz and Zanetti, 2013). Essentially, MPU shocks are a materialization of negative demand shocks and hence propagate through the expectations of economic agents. However, the focus of the inflation expectations literature has mostly been on the effects of monetary policy shocks, not on monetary policy uncertainty shocks (e.g., Leduc et al., 2007; Canova and Gambetti, 2009; Leduc and Sill, 2013). Not controlling for the uncertainty component might overlook an important aspect in the inflation expectation formation process. Hence, in this paper we investigate empirically the link between monetary policy uncertainty and inflation expectations.

To do so, an important question is how to measure monetary policy uncertainty. MPU has been typically defined as some function of the ability of economic agents to forecast monetary policy instruments, i.e. interest rates. One example is measuring surprises to agents via options and yield curve movements (e.g., Swanson, 2006; Bauer et al., 2012; Chang and Feunou, 2014). More recently, natural language processing (less formally text analysis) has also been employed to create monetary policy uncertainty proxies, either through newspaper-based articles (Baker et al., 2016; Husted et al., 2020) or the FOMC meetings (Hansen et al., 2018).

A further example, popular in structural models, is to postulate specific distributional

[1]See for example Bauer et al. (2012); Kang et al. (2014); Chang and Feunou (2014); Mumtaz and Zanetti (2013); Istrefi and Piloiu (2014); Mumtaz and Theodoridis (2015); Sinha (2016); Creal and Wu (2017); Kurov and Stan (2018); Hansen et al. (2018); Istrefi and Mouabbi (2018); Husted et al. (2020); Bauer et al. (2019); Alessandri and Mumtaz (2019).

assumptions on key target central bank variables, either the money supply (Stulz, 1986) or, more recently, the interest rates (Mumtaz and Zanetti, 2013; Creal and Wu, 2017; Alessandri and Mumtaz, 2019). In these contributions MPU is modelled in the second order moment, i.e. the variance of the central bank policy instrument. An increase in MPU is then defined as an increase of the variance of monetary policy shocks. The hypothesis is that larger monetary policy shocks worsen forecasts of economic agents and thus make it harder to anticipate the correct movements of the target variable.

We incorporate monetary policy uncertainty following the latter approach and estimate a structural VAR (SVAR) with stochastic volatility-in-mean to study the interaction between monetary policy uncertainty and inflation expectations. Capturing the joint dynamics of inflation, inflation expectations and economic activity is natural in the framework of SVARs. Our work is related to the growing literature which focuses on the effects of economic policy uncertainty on the economy (Bloom, 2009; Baker et al., 2016; Mumtaz and Zanetti, 2013; Bachmann et al., 2019). We focus specifically on monetary policy uncertainty, departing from recent works which look at the macroeconomic effects of changes in broader measures of uncertainty (Istrefi and Piloiu, 2014; Fernández-Villaverde et al., 2015).

In a sample spanning from 1982 to 2019 we find that, on average, in the U.S. short run inflation expectations do indeed decline following a MPU shock, although not in the same magnitude as inflation, suggesting that expectations are rigid. On the other hand, long-run inflation expectations, which are typically found not to react to monetary policy shocks (Canova and Gambetti, 2009), do not seem to be affected by MPU shocks. Furthermore, we show that the relationship between monetary policy uncertainty and inflation expectations has not remained stable over time. Since the great recession, short run inflation expectations have not reacted to MPU shocks, while inflation has. These findings suggest that, while MPU might have been important in the past for the expectation formation process, this has not been the case over the past decade even when we control for the zero lower bound.

The reminder of this paper is structured as follows. The next section lays out the methodology used in this article. Section 3 is devoted to the summary of the data set. Sections 4 and 5 discuss the results and Section 6 concludes.

2 Modeling monetary policy uncertainty

This section presents the adopted empirical strategy for estimating the effects of uncertainty on inflation expectations. We follow the approach of modeling monetary policy uncertainty as a second moment shock to the monetary policy instrument (e.g. Stulz, 1986). We adopt the framework of Mumtaz and Zanetti (2013), estimating a stochastic volatility-in-mean SVAR, which has already been employed for the analysis of monetary policy uncertainty on economic activity (Mumtaz and Zanetti, 2013; Alessandri and Mumtaz, 2019). In this setup, MPU is captured by the time-varying variance of monetary policy shocks, which are extracted from the central bank interest rate series. In the spirit of Swanson (2006) an increase in the variance lowers the ability of economic agents to forecast interest rates.[2] For the remainder of this paper, we will use the terms uncertainty and volatility interchangeably.[3]

An appealing property of the model is that it constitutes a one-step procedure in modelling uncertainty, i.e. both the coefficients and the uncertainty measure are estimated jointly from the data, thus providing an important feedback loop. This is in contrast to a model where the proxy is constructed first from an exogenous source and then put into a model as it avoids potential measurement errors.[4]

[2]This definition of uncertainty is popular in the literature. For example Jurado et al. (2015) frame financial and macroeconomic uncertainty as the variance of a time-varying forecast error from a dynamic factor model.

[3]We do acknowledge that there is a difference between implied and realized volatility. Our results draw on a measure based on realized interest rate volatility. See Bachmann et al. (2019) for a discussion on the subject.

[4]This is not to say that one-step approaches do not come without drawbacks. For example, model misspecification would undermine the generated regressors (the uncertainty indicators). The choice of one versus two-step approaches presents a trade-off between internal consistency and robustness to misspecification. For discussion of one-step versus two-step estimations of uncertainty indices see Bianchi et al. (2018).

2.1 SVAR with stochastic volatility in-mean

The framework of Mumtaz and Zanetti (2013) may be summarized by the following key equations. The VAR is given by:

$$Z_t = c + \sum_{j=1}^{P} \beta_j Z_{t-j} + \sum_{m=0}^{M} \gamma_m \tilde{h}_{t-m} + \Omega_t^{1/2} e_t, \quad e_t \sim N(0, I_K), \tag{2.1}$$

where Z_t is a vector of endogenous variables of dimension $K \times 1$ and the vector $\tilde{h}_t = [h_{1t}, \dots h_{Kt}]'$ contains the log volatilities of the K structural shocks. The shocks are extracted from the reduced form variance-covariance matrix Ω_t using an orthogonal decomposition $\Omega_t = A^{-1} H_t A^{-1'}$ and A^{-1} collects the contemporaneous relationships between the variable. H_t is a diagonal matrix with $exp\{\tilde{h}_t\}$ on its diagonal. The processes for the log-volatilities are assumed to be autoregressive of first order, of the form:

$$\tilde{h}_t = \Theta \tilde{h}_{t-1} + \eta_t, \quad \eta_t \sim N(0, \xi), \quad \mathbb{E}(e_t \eta_t) = 0, \tag{2.2}$$

such that Θ is a $K \times K$ diagonal matrix and $\eta_t = [\eta_{1t}, \dots, \eta_{Kt}]'$. Furthermore we assume that the first and second moment shocks are uncorrelated.[5]

Let r_t denote the monetary policy instrument in the vector of endogenous variables, Z_t. Then h_{rt} is our MPU measure, which affects the levels of the rest of the variables through the γ coefficients. The innovations η_t in eq. 2.2 are the uncertainty shocks in our framework with η_{rt} being the MPU shock. An increase in η_{rt} means that the monetary policy instrument may take a wider range of values as the shock dispersion increases.

[5]This assumption eases the computational burden but it may be relaxed, as for example in Alessandri and Mumtaz (2019).

2.2 Model with time-varying coefficients

This model captures the average response to uncertainty shocks over the full sample. Throughout the last 40 years the U.S. economy has experienced significant events, from the great moderation to the global financial crisis, which could have contributed to changes between the relationships of economic variables. Furthermore, it has been shown that inflation expectations depend on the own experience of the economic agents, i.e. people that lived through periods of high inflation expect, on average, higher inflation (Malmendier and Nagel, 2016; Draeger and Lamla, 2018). Finally, a large part of the literature on uncertainty has shown that the relationship between uncertainty and economic activity might be non-linear (Caggiano et al., 2014; Alessandri and Mumtaz, 2019), and specifically monetary policy shocks (Castelnuovo and Pellegrino, 2018). Therefore, we deem important to investigate whether the inflation expectations formation process has evolved over time and how the link between monetary policy uncertainty and expectations has behaved throughout different periods.

To do so we estimate an extended version of the model by allowing the coefficients to change over time, along the lines of Primiceri (2005) and Mumtaz and Theodoridis (2020). Thus our analysis falls in a broad class of studies analyzing the time-variation of economic shocks (Benati and Surico, 2008; Mumtaz and Zanetti, 2015; Liu et al., 2019).

We introduce time-varying parameters (TVP) for the coefficients of the estimated VAR, β and γ, thus capturing potential changes among the economic variables and how uncertainty is perceived over time. Equation (2.1) is modified to incorporate time-variation in the following way:

$$Z_t = c_t + \sum_{j=1}^{P} \beta_{j,t} Z_{t-j} + \sum_{m=0}^{M} \gamma_{m,t} \tilde{h}_{t-m} + \Omega_t^{1/2} e_t, \quad e_t \sim N(0, I_K), \tag{2.3}$$

We postulate a random walk process for the coefficients:

$$B_t = B_{t-1} + v_t, \qquad v_t \sim N(0, Q) \tag{2.4}$$

where $B_t = [\beta_t, \gamma_t]'$ and Q governs the amount of time-variation in the parameters.

2.3 Estimation strategy

Due to the presence of the volatility terms in eq. (2.3) the conventional maximum likelihood approach is not applicable. The model is estimated via Bayesian methods with Gibbs sampling, i.e., drawing the parameters iteratively from their conditional distributions.

The estimation procedure for the TVP and no-TVP specification is mostly the same. The parameters may be divided into several blocks based on their distributional assumptions. The reduced-form coefficients $B = [\beta, \gamma]$, and $B_t = [\beta_t, \gamma_t]$ in the TVP case, respectively, along with the Q matrix, the stochastic volatility block $\mathcal{H} = \{H_1, ..., H_T\}$, where H is a diagonal matrix containing the $\tilde{h}$, the parameters Θ in the equation, and the contemporaneous responses A.

In order to simplify the exposition we introduce notation $\Psi = \{A, B, Q, \mathcal{H}, \Theta\}$, which collects the different blocks of parameters. Let Ψ^{-i}, denote the exclusion of the i-th block such that $\Psi^{-B} = \{A, Q, \mathcal{H}, \Theta\}$.

To conduct inference we draw the i-th block from the conditional probability distribution $p(i|\Psi^{-i})$, which is derived as a function of a conjugate prior distribution $\overline{p}(i)$. The prior for the reduced-form coefficients B and $\tilde{h}$ is based on a GLS estimation on a training sample. For arbitrary starting values the estimation proceeds in the following iterative procedure:

1. Conditional on $\mathcal{H}$, A may be drawn as a linear regression from $p(A|\Psi^{-A})$, given the form in eq. (4.2) in conjunction with the algorithm of Cogley and Sargent (2005).

2. The reduced-form parameters, B, are obtained from a linear regression with heteroskedasticity and may be estimated via GLS. Following Carter and Kohn (1994a) we introduce stochastic volatility via a matrix Q, employ the Kalman filter for $t = 1, ..., T$ to get $\beta_{T|T}$ and $\gamma_{T|T}$ and draw the parameters from $p(\beta_{T|T}, \gamma_{T|T}|Q, \Psi^{-B,-Q})$.

3. Conditional on the draws for B, Q may be drawn from an inverse Wishart distribution.

4. For the stochastic volatility estimation we follow Cogley and Sargent (2005), where the draws for $\mathcal{H}$ may be obtained using a Metropolis-Hastings algorithm conditional on the model parameters.

The only notable difference when estimating the TVP specification is that the time variation in the model is governed by the process given in (2.4), for which the specification of Q is of importance. The constant parameter version is nested here for extremely low values of Q, while high values permit larger jumps in the parameters. Given that the random walk process introduces potential explosiveness, exceptionally high values could make inference impossible. To deal with this issue, in the seminal contribution of Primiceri (2005), Q is scaled a-priori via an additional parameter, k_Q. This choice is of specific importance. Our goal is to find the appropriate amount of time-variation without imposing overly strong restrictions on the parameters movement nor running into estimation difficulties. In Primiceri (2005) k_Q is chosen using a grid-search with the aim to maximize the marginal likelihood of the model with an optimal value of $k_Q = 0.01$ for U.S. data. However, our data set is monthly, hence there is no guarantee that this is an appropriate choice for our application.

We follow recent advancements in Bayesian computation of TVP models and treat k_Q as a hyperparameter to be estimated from the data as in Amir-Ahmadi et al. (2018) by using a random walk metropolis hasting algorithm with an inverse gamma distribution as a prior. We plot the posterior densities of k_Q in the Appendix, Section B. The corresponding estimates of k_Q lead to a median value of 0.063. We use this median of the posterior density of k_Q to calculate the impulse response functions of the TVP model.

3 Data and estimation

Our data set consists of five macroeconomic variables: the civilian unemployment rate, the long-run and short-run inflation expectations, the consumer price index (CPI) inflation rate and the monetary policy rate. Given that our measure of uncertainty relies on the volatility of the variables, the Zero Lower Bound presents a significant issue. For this reason we choose the shadow short rate (SSR) as a proxy for the main monetary policy instrument. Thus, we also capture the unconventional monetary policy measures which are reflected in the SSR. For robustness we also consider the model with federal funds rate (FED Funds).

The key variables of interest in our study are the inflation expectations. Therefore, we use the proxies developed in Haubrich et al. (2012), where expectations are extracted using a term-structure model of the interest rates. This data has several appealing properties. First and foremost, the financial sector is a crucial channel for the expectations of economic agents, thus proxies based on interest rate data should provide timely adjustment of inflation expectations. Second the expectation measures are available both for short and long-term (one year ahead and five years ahead), starting in 1982. In contrast, the long run expectations of the Survey of Professional Forecasters (SPF) only starts in 2005. Third, the data is at a monthly frequency. This increases the degrees of freedom in our highly parametrized non-linear model. Finally, the term-structure model accounts for liquidity and risk-premia. This is an important difference to other measures of inflation expectations based on interest rate data, like the treasury inflation-protected securities (TIPS) that do not account for such premia. However, this comes at a cost, as the proxy is an outcome of a term-structure model, hence carries model uncertainty, which cannot be incorporated into the analysis. For example, misspecification could lead to an over- or underestimation of our probability intervals - a different inflation expectations measure might exhibit different dynamics.[6] Figure

[6]We have also explored an alternative specification, where the one year ahead SPF data has been interpolated using a Mixed Frequency VAR as in Schorfheide and Song (2015). We find that our conclusions remain the same. Results are available upon request.

2.1 plots the quarterly aggregated measure from Haubrich et al. (2012) versus the SPF.

Figure 2.1: One-year ahead inflation expectations from the Survey of Professional Forecasters and from Haubrich et al. (2012). Quarterly frequency. Source: Federal Reserve Bank of Philadelphia and the Federal Reserve Bank of Cleveland.

The unemployment rate, CPI inflation, and the FED Funds were obtained from the Federal Reserve Bank of St. Louis Database. The shadow short rate (SSR) was obtained from Wu and Xia (2016).[7] Finally, we use the estimations from Haubrich et al. (2012) for the long-run and short-run inflation expectations taken from the Federal Reserve Bank of Cleveland. The sample ranges from January 1982 to June 2019.

3.1 Statistical identification and inference

To determine the optimal number of lags in eq. (5.1) we employed information criteria tests on the reduced-form specification without stochastic volatility terms.The proposed lag lengths

[7]Our findings are robust to the choice of SSR as we have also considered the alternative measure from Krippner (2013).

are 2 and 12 for BIC and AIC, respectively. We choose 2 lags for the baseline specification as to prevent parameter proliferation in the time-varying case. In the constant parameter case we also test 12 lags and find that the results are robust to that specification.[8]

Following the estimation strategy from Section 2.3 we initialize the model estimation using the first 10 years as training sample. Thus, after accounting for the lag length the effective sample begins in March 1992. We use 100 000 iterations of the Gibbs sampler, of which the first 95 000 are discarded as burn-in. We test for the convergence of the chain with standard procedures and discuss the results in Appendix D.

To identify the monetary policy shocks we use Cholesky decomposition with a lower triangular matrix, where following the sticky-information literature (Castelnuovo and Surico, 2010), we order inflation expectations above inflation. Thus, we assume that they do not respond within the same month to new information regarding the interest rates or inflation.[9] Apart from that we follow standard ordering of unemployment, inflation and the monetary policy rate (Christiano et al., 1999). Nevertheless, we test our results to different specifications and find that they are consistent to the baseline. This is due to the fact that we are interested in the impulse responses following a shock to the stochastic volatilities, η. Thus, a major determinant of the effects of uncertainty shocks are the γ coefficients and variable ordering plays a secondary role. This holds as long as the estimated stochastic volatilities h remain the same across orderings. This is an advantage of the framework over setups where the uncertainty measure is an additional series whose level is shocked (e.g. Istrefi and Piloiu, 2014).

[8]The results available in the Appendix.

[9]This has been found to be true for SPF data, see Coibion and Gorodnichenko (2015).

4 Estimated uncertainty measures

Figure 2.2 plots the MPU series implied by the two models, the TVP and the constant parameter specification, respectively as well as an array of other popular uncertainty indicators from the literature. The measures from the two models are highly similar, with the constant parameter specification giving a bit more weight on the MPU at the height of the financial crisis as well as during 2013 - the year of the "taper tantrum". Other than that, monetary policy uncertainty is estimated to have been particularly high during the early '90s and following the events of the dot-com bubble. These dynamics mirror closely the estimated monetary policy volatility of Mumtaz and Zanetti (2013), whose sample ends in 2010Q4.

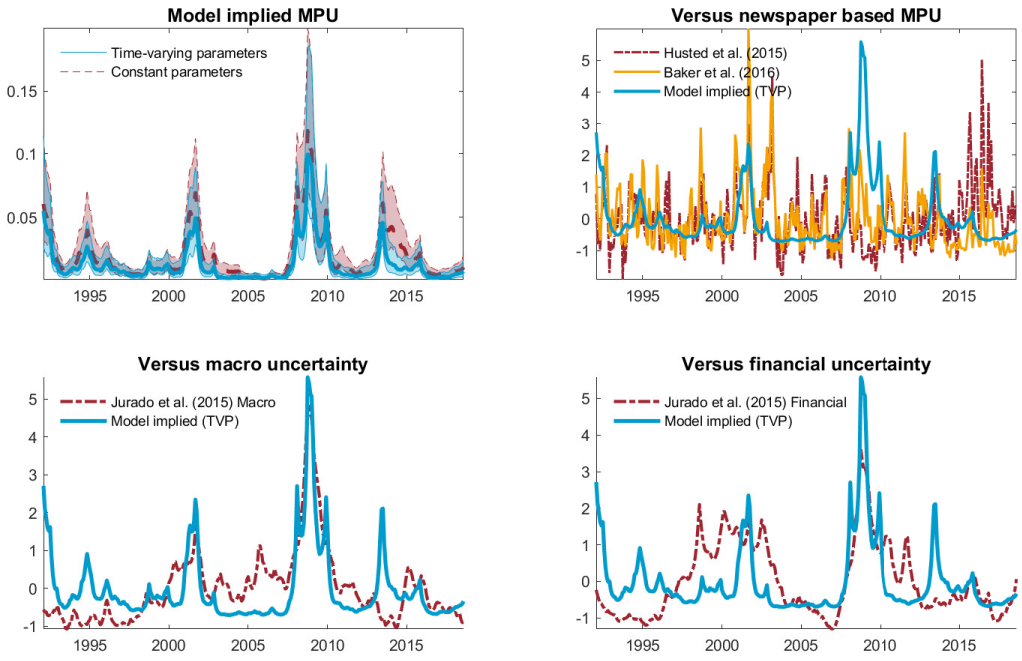

Figure 2.2: Measures of monetary policy uncertainty. Top left: estimates from the time-varying parameter (TVP) and constant parameter models with probability intervals (shaded). Top right: Model implied MPU versus the indices MPU indices from Baker et al. (2016) and Husted et al. (2020), standadrised. Bottom left: Macro uncertainty from Jurado et al. (2015), standardized. Bottom right: financial uncertainty from Jurado et al. (2015), standadrised.

On the top right of Figure 2.2 the model implied measure is plotted against the MPU index

from Baker et al. (2016) and Husted et al. (2020). All time series have been standardized to facilitate comparison. The latter two are based on newspaper data and are much more volatile, compared to the model estimate. Apart from the differences in smoothness, the three series appear to follow similar dynamics throughout the 1990s. However, since the early 2000s there are notable differences. The model implies that the highest peak of monetary policy volatility was surrounding the global financial crisis. The measure of Baker et al. (2016) - during the dot-com bubble and its aftermath, while the index of Husted et al. (2020) suggests that MPU was highest throughout the first years of the Trump administration.[10] The model estimate is by construction smoother; its dynamics are given by the autoregressive process in eq. (2.2). It is also based on the estimated stochastic volatility of the monetary policy shocks identified using our SVAR, hence tightly related to the underlying shadow short rate series and our identification assumptions.

Since economic shocks are by assumption unforecastable, the measure of Mumtaz and Zanetti (2013) is in spirit much closer to the approach taken by Jurado et al. (2015) where uncertainty is related to the second moments of a forecast distribution. We plot both macroeconomic and financial uncertainty based on Jurado et al. (2015) in the bottom panels of Figure 2.2. While all three measures display elevated uncertainty surrounding the global financial crisis there appear to be notable differences in the other periods. Macro uncertainty displays lower but more prolonged peaks around the turn of the millennium and has even fallen during the "taper tantrum" periods. Similarly, financial uncertainty also does not fluctuate post the financial crisis and on average appears to follow two long cycles throughout the sample as opposed to MPU, which is estimated to have peaked shortly and only during specific events.[11]

Next, we report the volatility of the shocks to short and long run inflation expectations,

[10] This is valid for our data span, which cuts short right before the index of Husted et al. (2020) reaches its historical maximum in August 2019 of 407 (long run mean 110).

[11] We control for financial uncertainty by extending our model with the returns from Standard and Poors' Index. We compare the estimated financial uncertainty from that model and report that it matches the Jurado et al (2015) measure much more closely. See Appendix C.2.

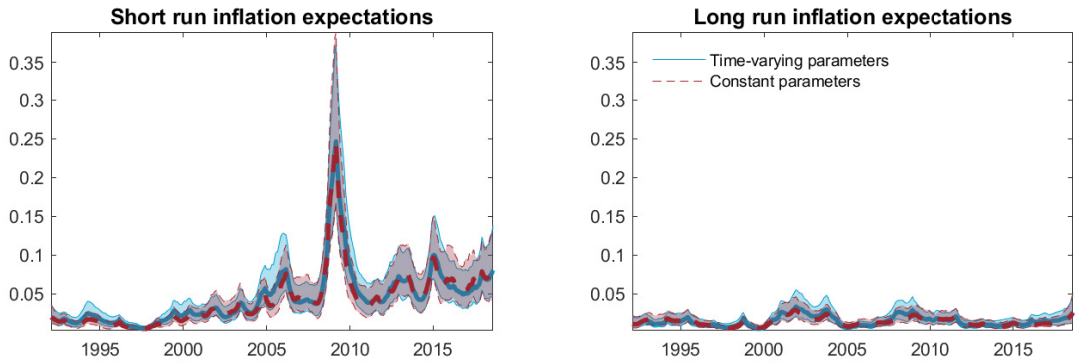

Figure 2.3: Measures of inflation expectations uncertainty from a time-varying parameter model and a constant coefficients specification. Median responses (solid and dashed line) and 68% probability intervals (shaded area).

which could be interpreted as uncertainty surrounding the inflation expectations. They are plotted on Figure 2.3 for each specification. The estimates are overlapping almost perfectly. Notably, long run inflation expectations display a low and stable level of uncertainty, while short run inflation expectations uncertainty increased during the Great Recession and remained elevated compared to the period before the 2008 crisis.[12] Long run inflation expectations appear to be well anchored as their uncertainty level does not fluctuate even through dramatic events such as the global financial crisis. These estimates may also be seen through the lens of central bank credibility. Constant uncertainty surrounding the long run expectations suggests trust in the monetary authority - a crucial condition for an efficient transmission of monetary policy. These findings are in line with the inflation gap literature that establishes low volatility of trend inflation (Cogley and Sbordone, 2008; Cogley et al., 2010; Coibion and Gorodnichenko, 2011). However, long run interest rates have been shown to respond to economic news (Guerkaynak et al., 2005), which suggests that higher volatility of long run inflation expectations could also be expected. The model, however, recovers the large volatility swings in the short run inflation expectations only. Given the similarity between the two series, this raises the question whether the model always recovers the true

[12]The estimated volatilities in the model with FED Funds are almost identical to the ones obtained with the SSR. Estimates are shown in the appendix, Figure B.1.

peaks in the volatilities.

To test this hypothesis we conduct a Monte Carlo exercise using artificially generated data. First, we generate data from a model in which short run inflation expectations volatility $h_{SR,t}$ for $t = 1, \ldots, T$ is extremely low, while long run inflation expectations volatility is high. Then we estimate the model on that data set to see whether it would recover the volatility series correctly. We mirror this setup with artificially generated data where the opposite is true, in the data generating process short run inflation expectations volatility is high and long run is low. We generate the data 200 times, 100 for each setup and estimate correspondingly 200 times the model on the artificial data set. We find that in both cases the model correctly attributes the volatility to the correct scenario. Further details as well as a plot of the estimates may be found in Appendix, Section B.

5 The effects of monetary policy uncertainty shocks

Next we will examine the responses of the variables to a one standard deviation monetary policy uncertainty shock. Overall am MPU shock manifests as a negative demand shock, as documented by Mumtaz and Zanetti (2013) and Alessandri and Mumtaz (2019). Namely, an increase in MPU leads to a contraction of the real sector due to risk aversion and consumption smoothing behaviour and a decline in economic output (Mumtaz and Zanetti, 2013), as well as a fall in prices and the interest rate. With our monthly data set we observe the same effects, unemployment increases while inflation and interest rates decline. However, when we allow for gradual evolution of the relationships between the variables we find interesting developments, especially in relation to inflation expectations.

A concise way to highlight the time-variation in the responses is to show the average responses of the model variables to a monetary policy uncertainty shock during four different periods: i) 1993 as the first full year of impulse responses after accounting for the training

sample and the lag values, ii) 2003 as a boom period after the recovery of the dot-com crash, iii) 2008 to look to the effects during the Great Recession and iv) 2018 the last full year of estimates. We do this on Figures 2.4 through 2.8. For comparison we also plot the responses from the constant parameter model, which by construction would be the identical over the chosen periods.

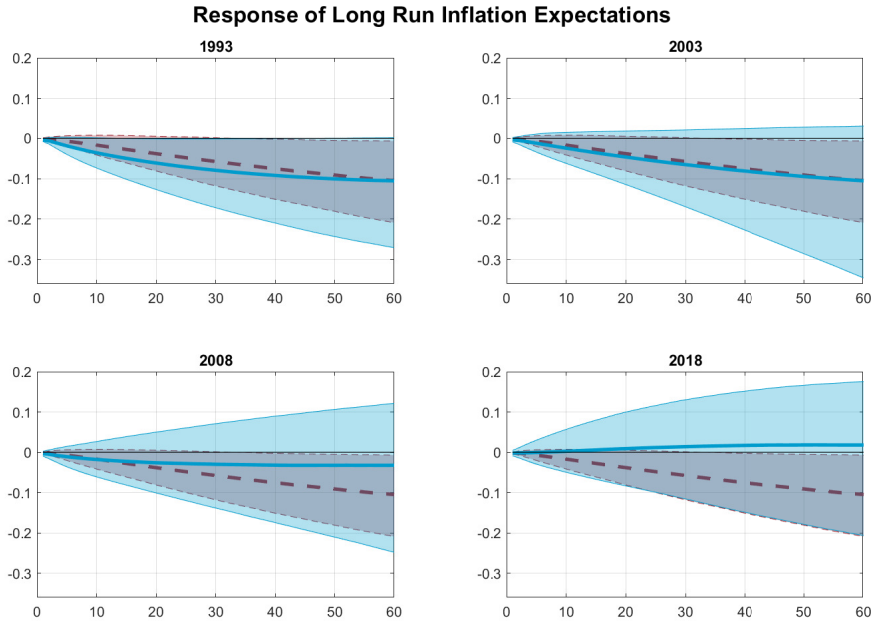

Figure 2.4: Impulse response of long run inflation expectations following a monetary policy uncertainty shock, yearly averages. Solid line denotes the median response and shaded area is the 68% probability intervals for the time-varying parameter model. The dashed line is the response from a constant parameter case.

Figure 2.4 shows the response of long run inflation expectations to a monetary policy uncertainty shock. We find that with the exception of the beginning of the sample, the response of long-run inflation expectations is close to zero. The 68% probability intervals bands suggest that this response is not different from zero for the remaining part of the sample. The median response has evolved from slightly negative towards showing no link to reacting to monetary policy uncertainty. The constant parameter specification also captures a negative median response of the long run inflation expectations, naturally with a tighter

probability intervals due to the higher degrees of freedom. These findings are of particular interest since long run inflation expectations are associated with the level of credibility the central bank maintains. This is a especially important for monetary policy effectiveness in a context where forward guidance and other unconventional measures have been implemented. Along with the estimated low long run volatility these results reinforce the notion that long run inflation expectations are well anchored.

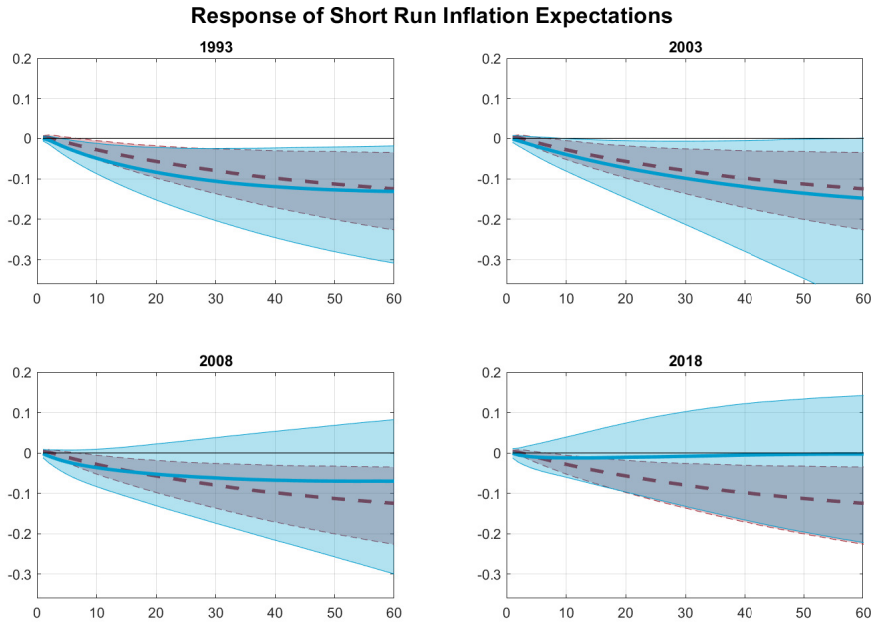

Figure 2.5: Impulse response of short run inflation expectations following a monetary policy uncertainty shock, yearly averages. Solid line denotes the median response and shaded area is the 68% probability intervals for the time-varying parameter model. The dashed line is the response from a constant parameter model.

Short run inflation expectations, depicted in Figure 2.5, respond somewhat differently. The constant parameter model shows a negative response with the zero line falling outside of the probability intervals, in contrast to the long run expectations case. In the TVP version we see that these findings are driven primarily by the response of the short run expectations throughout the first half of the sample. The median response of the TVP case (solid lines) lies below the constant parameter case (dashed lines) in 1993 and 2003 while it

is less negative after the global financial crisis. Correspondingly the 68% error bands show a decline in the short run inflation expectations not only in the beginning of the sample, as long run expectations do, but also before the financial crisis. After the beginning of the Great Recession the negative response to MPU shocks has dissipated, indicating a distinct pre- and post-financial crisis dynamics of inflation expectations.

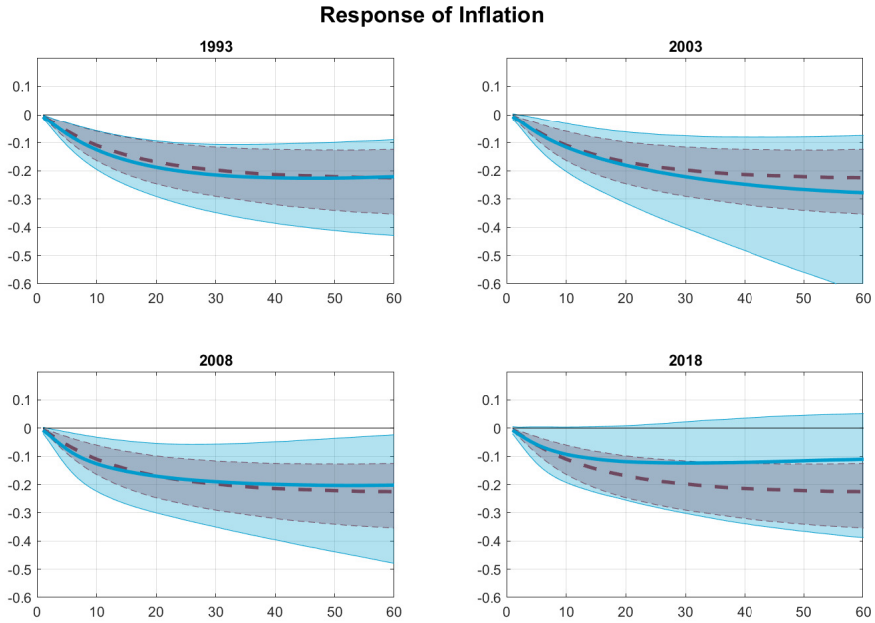

Figure 2.6: Impulse response of inflation following a monetary policy uncertainty shock, yearly averages. Solid line denotes the median response and shaded area is the 68% probability intervals for the time-varying parameter model. The dashed line is the response from a constant parameter model.

Next, we look at the response of inflation, shown in Figure 2.6, which is also negative. It is, however, different than the response of inflation expectations in Figures 2.4 and 2.5 in several key aspects. First, both the TVP and no-TVP model suggest a median estimated (including the error bands) outside of the zero line throughout almost the entire sample with the exception of the very last periods - inflation declines following an MPU shock. Given the lack of a strong response of both expectations series post the financial crisis, it appears that the transmission mechanism is not through the expectations channel. Second, the response

of inflation has always been several magnitudes stronger than that of inflation expectations, indicative of a rigidity in the expectation formation process when it comes to MPU shocks. This is in contrast to the dynamics following an uncertainty shock to the real variable, unemployment, which is also a negative demand shock. In that case all three variables react in a similar magnitude (See Appendix B).

To highlight these points we present an alternative view of the MPU effects in Figure 2.7. It contains a cross section of the median responses of inflation and inflation expectations at specific horizons following an MPU shock - from 12 to 60 months ahead. The figure shows the gradual transition towards the flat response throughout the last decade of low interest rates. Combined with the low estimated volatilities this evidence points toward a strong anchoring of inflation expectations.

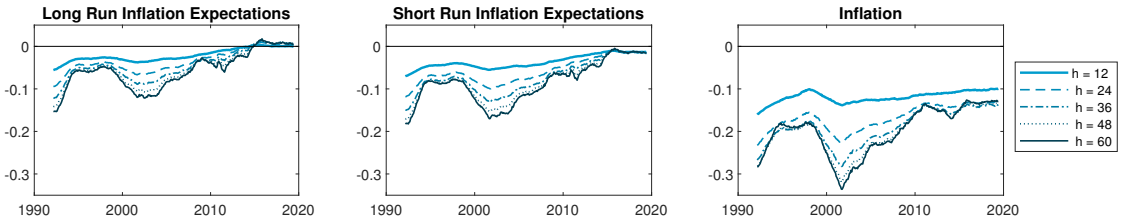

Figure 2.7: Cross section of the median impulse responses at specific horizons after the shock - from $h = 12$ months to $h = 60$ months.

Figure 2.8 depicts the responses of unemployment and the monetary policy rate. Unemployment has always increased following an MPU shock, albeit in different magnitude. By far the strongest response is estimated to be during the financial crisis, a result probably exacerbated through omitted variable bias as additional information is needed in the model to capture the sharp increase in unemployment during the crisis. Nevertheless, even in good times, a standard deviation shock in MPU increases unemployment by about 0.2 percentage points.

Following an MPU shock, the response of the monetary policy rate mimics the dynamics of inflation - it has declined throughout the first half of the sample. However, post the financial

crisis interest rates do not seem to react any more. Thus, it appears that the importance of MPU shocks has declined in recent years. Mumtaz and Theodoridis (2018) also show highlighted that uncertainty shocks, namely macro uncertainty, have become less important over time, specifically towards real and financial variables. However, they do not find this for prices. Our results suggest that, at least monetary policy uncertainty has become less important for inflation, yet only in the recent periods not covered by their sample. A further distinction is that our real variable, unemployment, reacts stronger to an MPU shock during and post the GFC, which is in line with other contributions in the literature, such as Angelini et al. (2019) and Alessandri and Mumtaz (2019).

These findings may also be related to a large strand of the literature on state dependent effects of uncertainty (e.g. Caggiano et al., 2014, 2017; Castelnuovo and Pellegrino, 2018; Liu et al., 2019).[13] These contributions show that uncertainty shocks have had different effects throughout time, nonetheless by highlighting that the impact is more pronounced in recessionary periods. However, we find diminishing effects of MPU shocks over time (as for example also Mumtaz and Theodoridis (2018) do, albeit with respect to macro uncertainty). In part, this is explained by our model choice - we allow for a gradual parameter evolution to capture long run changes of the macroeconomic dynamics. The state dependent literature, on the other hand, utilizes models that allow for large and sudden parameter breaks, such as change-point, threshold and smooth-transition approaches.

Furthermore, the VAR structure imposes that all variables are interconnected - as the monetary policy rate reacts less to MPU shocks over time, the other variables might do so as well. For example, Liu et al. (2019) show that first moment demand shocks suppress prices and interest rates, while unemployment rises.[14] However, we show that second moment shocks matter also on their own - the median response of the monetary policy rate is almost

[13] See Castelnuovo (2019) for an overview of the literature.

[14] The original paper reports the results to a positive demand shock from a change-point VAR model identified via sign restrictions conditioning on the regime. Thus, the model should preserve its signs following a negative demand shock.

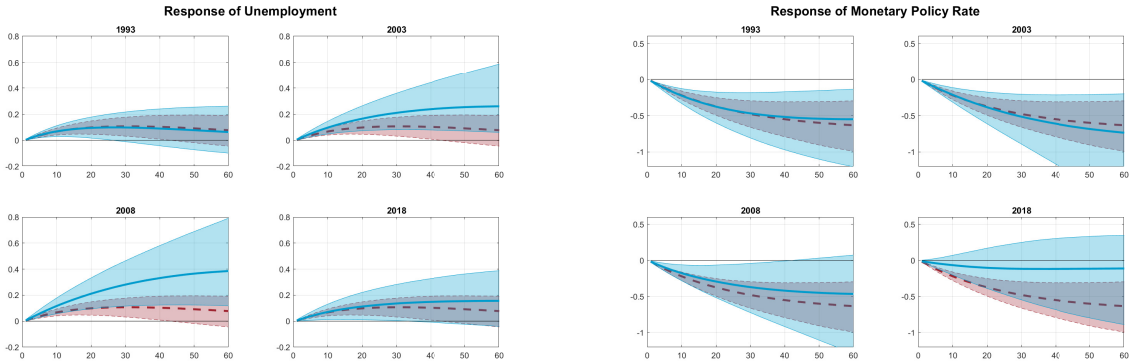

Figure 2.8: Impulse responses unemployment and monetary policy rate following a monetary policy uncertainty shock, yearly averages. Solid line denotes the median response and shaded area is the 68% probability intervals for the time-varying parameter model. The dashed line is the response from a constant parameter model.

flat towards the end of our sample 2018, while inflation and unemployment continue to react to MPU shocks.

Finally, the model is robust with regard to different specifications. We have explored robustness with respect to the number of lags, the variables included and their ordering. Notably, given the importance of the macro-financial nexus we test for a possible omitted variable bias by including a financial variable, namely, the S&P500 returns. We find that controlling for financial uncertainty does not alter the estimated proxy for monetary policy uncertainty shocks, neither it affects the impulse response analysis. Moreover, the model implied financial uncertainty from the S&P500 appears to match rather well the financial uncertainty index by Jurado et al. (2015). These findings are reported in the Appendix, Section B. Furthermore, changing the order of the variables, a prominent issue with SVAR models with Cholesky decomposition and even more so with TVP-VARs, does not lead to different conclusions (Section B in the Appendix).

6 Concluding remarks

In this paper we investigate the evolution of monetary policy uncertainty on the economy, with a special focus on the consequences unexpected increases of MPU on inflation expectations. Using monthly data for the US we estimate a SVAR with stochastic volatility-in-mean. Starting with a constant parameter specification we find further evidence that MPU shocks act as negative demand shocks with rising unemployment and decreasing price level. More importantly, we find that inflation expectations react negatively to an increase in the monetary policy uncertainty with a different magnitude for long and short inflation expectations.

Considering the evidence that inflation expectations differ across age groups and personal experience, we then estimate a model where we allow the parameters among the model variables to change over time. We treat the question of how much parameter evolution is allowed in the model as open and estimate it from the data.

This extension allows us to observe that long run and short run inflation expectations have reacted differently to MPU shocks in the 90s as opposed to today. Long run inflation expectations have remained mostly robust and unperturbed by MPU shocks apart from the beginning of the 90s. On the other hand, short run inflation expectations have had a pronounced negative response up to the Great Recession. After the 2008 crisis, short run inflation expectations have become less and less affected by MPU shocks. This finding is surprising given that most episodes associated with high MPU have occurred post 2008. This distinction between long and short run inflation expectations suggests a level of rigidity surrounding long run inflation expectations associated with a credible central bank. Contrary to short run inflation expectations, inflation reacts negatively throughout the whole sample and more strongly than the inflation expectations in the short run Jurado et al. (2015).

Overall we find that MPU shocks have become less important over time. This is true not only for inflation expectations and inflation but also for unemployment and the interest rate. These findings suggests that monetary policy uncertainty might be less obstructing

for the conduct of monetary policy than perceived. Neither interest rates react nowadays to MPU shocks as they did in the past, nor inflation expectations appear to respond as much as historically observed.

B Appendix

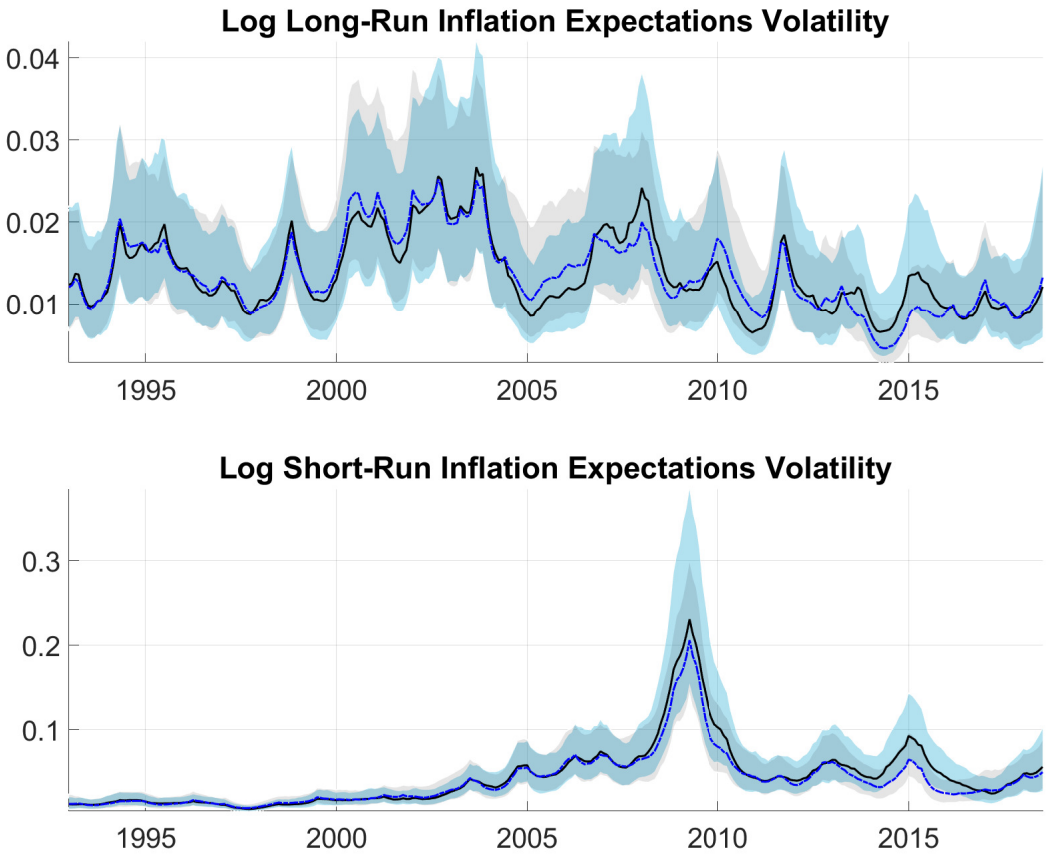

Figure B.1: Log-volatilities of short and long run inflation expectations based on a model with the SSR (solid line) and a FED Funds model (blue dashed line) with 68% probability intervals.

Monte Carlo simulation

The data generating process (DGP) for the Monte Carlo simulations is given by

$$Z_t = \hat{c} + \sum_{j=1}^{P} \hat{\beta}_j Z_{t-j} + \sum_{m=0}^{M} \hat{\gamma_m} \tilde{h}_{t-m} + u_t, \quad u_t \sim N(0, \hat{\Omega}_t), \tag{2.5}$$

with $\hat{\Omega}_t = \hat{A}^{-1} H_t \hat{A}^{-1'}$ where H is specified as in the original setup as being a diagonal matrix with $\tilde{h}$ on the diagonal. The parameters of the DGP are set at the posterior mode from the estimation of the model based on real data, which ensures a stable VAR process. We then proceed to draw a series of shocks for u_t for $t = 1, ..., 500$, which is in line with the sample size we have in practice. To draw from $N(0, \hat{\Omega}_t)$ we need to specify a series for $\tilde{h}_t$. The shapes of the true volatility series are inspired by the estimates with the real data - we clone the first 250 observations from our model twice, thus creating a data set with two large crises. This results in a couple of large peaks in long-run inflation volatility in the first setup (MC1), while we force short run inflation volatility to be extremely low with a mean at 0.005. Correspondingly in the second setup (MC2) we adopt a low long-run volatility series, similar to the one in the real data.

We then proceed to draw 200 data sets, 100 for each setup and estimate the model 200 times. After each estimation we have an uncertainty measure for long and short run inflation expectations and its probability intervals. On Figure B.2 we plot the median volatility estimates for each data set (100 per setup) along with the true DGP from which said data set was created. We find that the model correctly captures the shape of each special case. Note that each gray line on the graph is the median outcome of one data set, for ease of readibility we omit the probability intervals surrounding each run.

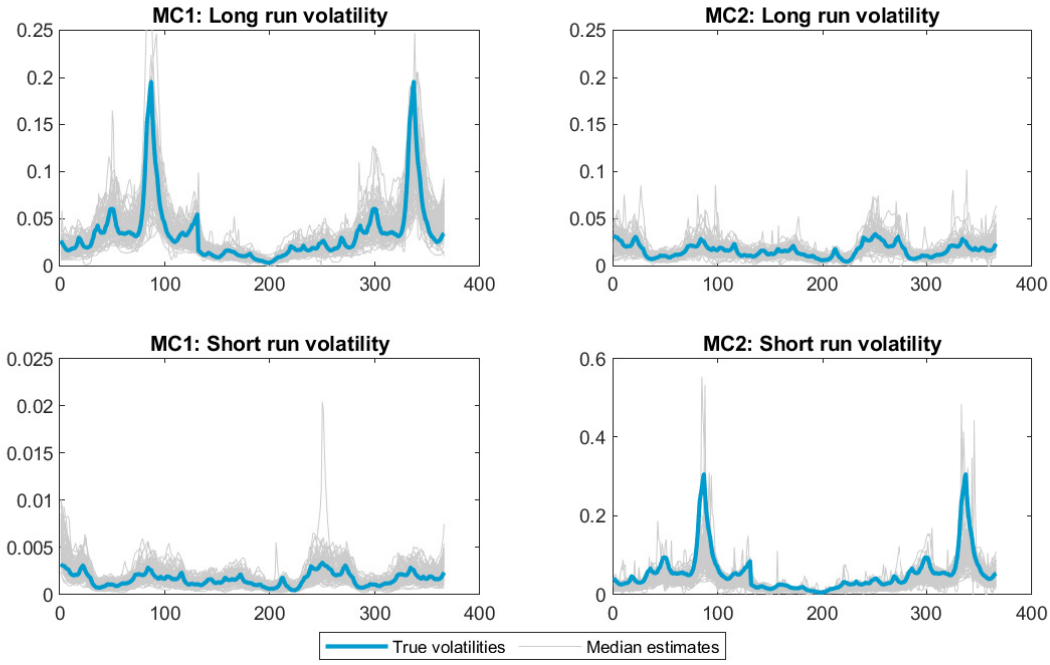

Figure B.2: True versus estimated volatility from 200 Monte Carlo simulations. Left: Monte Carlo setup 1 (MC1) with high long run inflation volatility and low short run inflation volatility. Right: Monte Carlo setup 2 (MC2) with high short run inflation volatility.

Robustness Section

Robustness to the lag specification

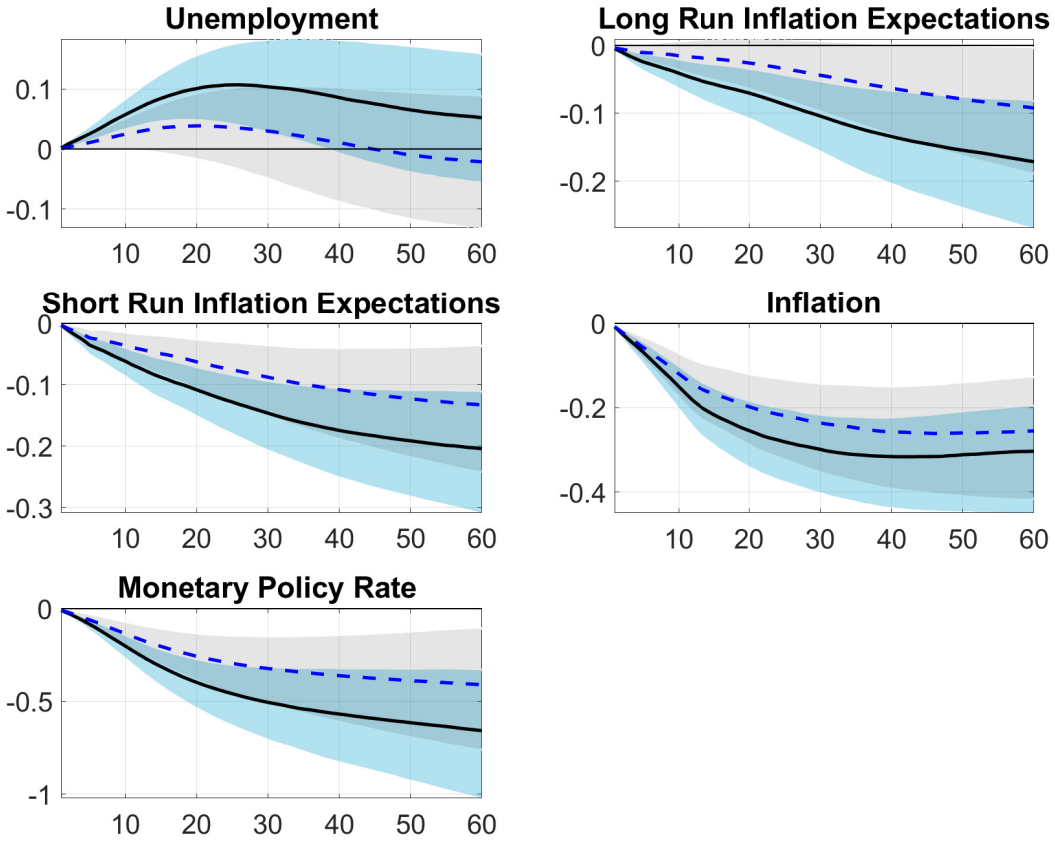

Figure B.3: Median responses to a monetary policy uncertainty shock. Model with shadow short rate (solid line) and model with the federal funds rate (dashed line). The shaded areas show the 68% probability interval.

Figure B.3 plots the impulse responses of the constant parameter models the Federal funds rate with 12 lags. The only quantitative difference is the case of long-run inflation expectations with the SSR, where the response is significant, while the FED Funds model shows identical behavior to the 2 lag SSR specification.

Extending the model to control for financial uncertainty

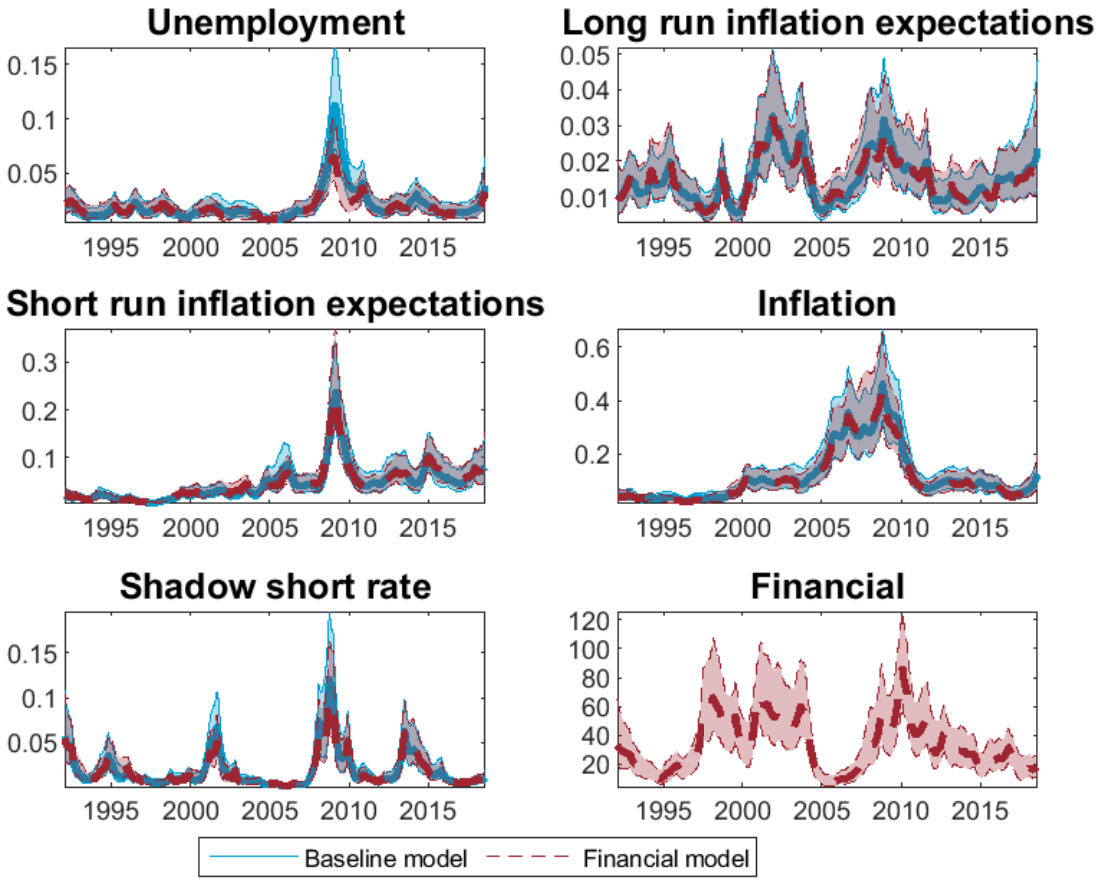

Figure B.4: Solid blue line depicts the median estimate of the stochastic volatilities of the baseline model. Dashed red lines are the median estimates of the stochastic volatilities of the alternative financial specification. Shaded areas show the corresponding 68% probability intervals.

Figure B.4 indicate that our baseline stochastic volatility estimates, from which our monetary policy uncertainty measure is derived, are unaffected by including the S&P 500 returns as in Mumtaz and Zanetti (2013) to account for financial uncertainty. Additionally, one can observe that despite the fact that the monetary policy uncertainty from the financial model follows very closely the estimates from the baseline one, the estimated volatility of the S&P 500 returns differs significantly from the monetary policy uncertainty measure. Moreover,

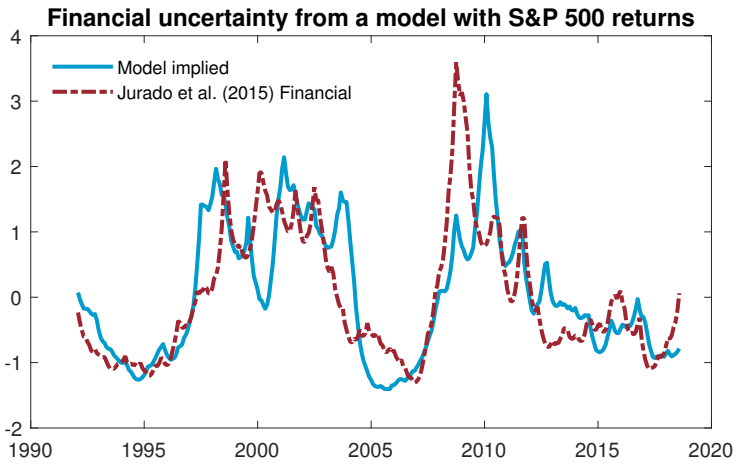

Figure B.5: Comparison between financial uncertainty estimated from the model by extending it with S&P 500 returns and plotting the stochastic volatility to its shocks versus Jurado et al. (2015).

figure B.5 shows the associated financial uncertainty from the model the alternative financial model of figure B.4 appear to coincide well with the financial uncertainty index of Jurado et al. (2015).

Robustness to the variable order

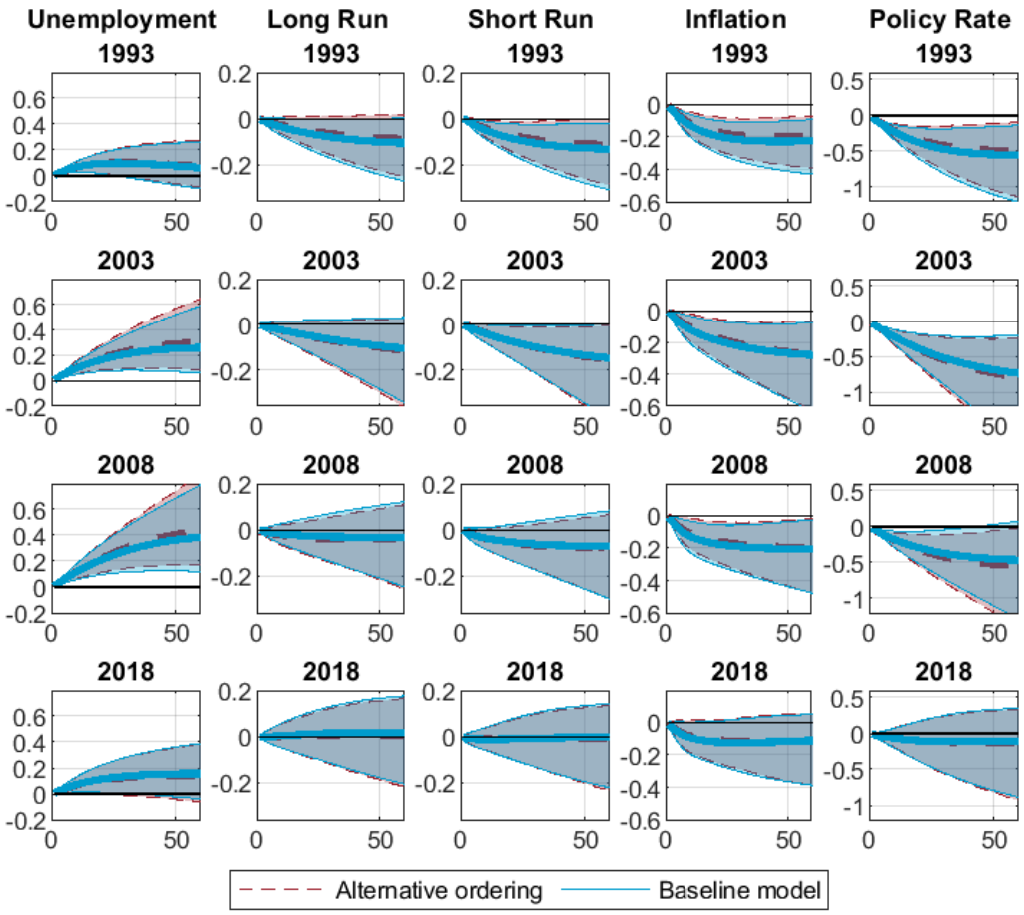

Figure B.6: Solid blue line depicts the median responses to a monetary policy uncertainty shock of the baseline model at different time periods. Dashed red lines are the median responses to a monetary policy uncertainty shock of the alternative ordering specification. Shaded areas show the corresponding 68% probability intervals.

Impulse responses: unemployment uncertainty shock

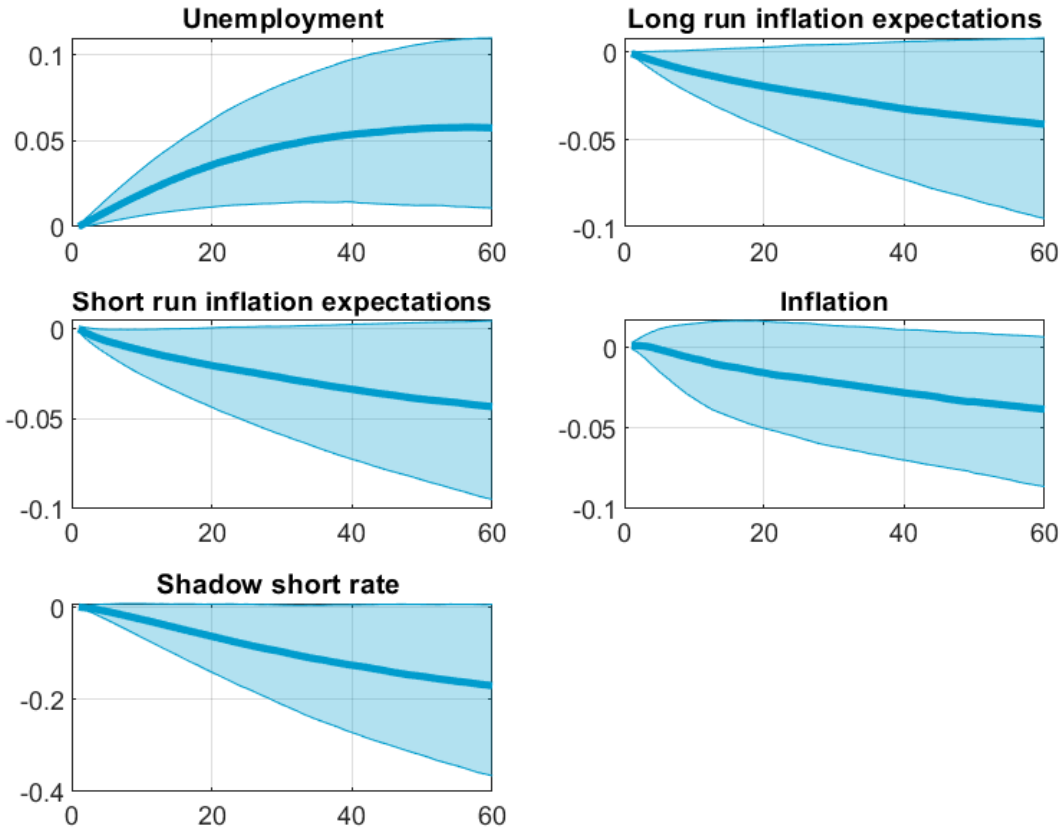

Figure B.7: Solid blue line depicts the median responses to a unemployment uncertainty shock of the baseline model. Shaded areas show the corresponding 68% probability intervals.

Convergence

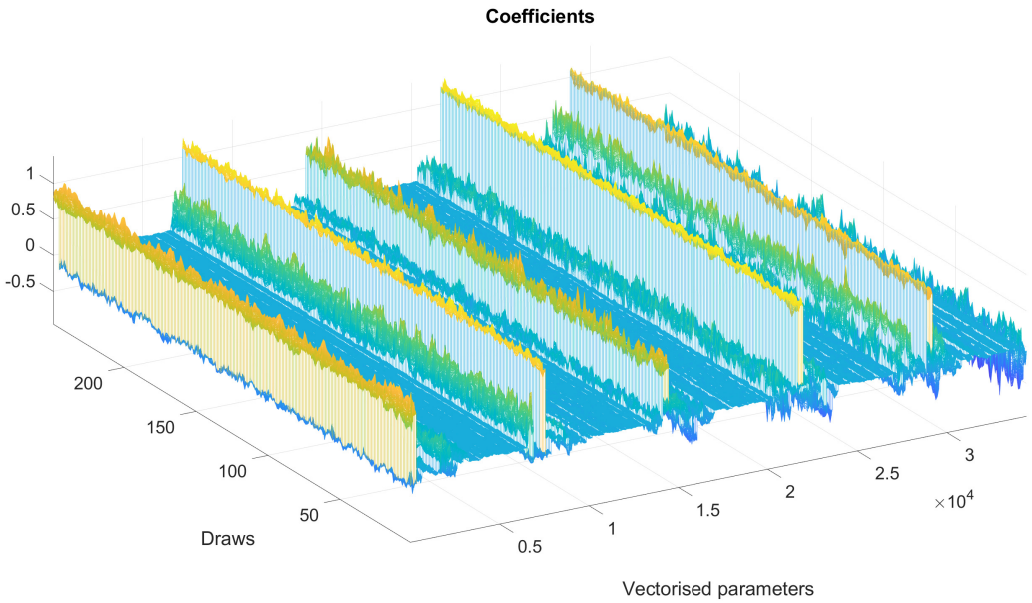

Figure B.8: Recursive means of the VAR coefficients.

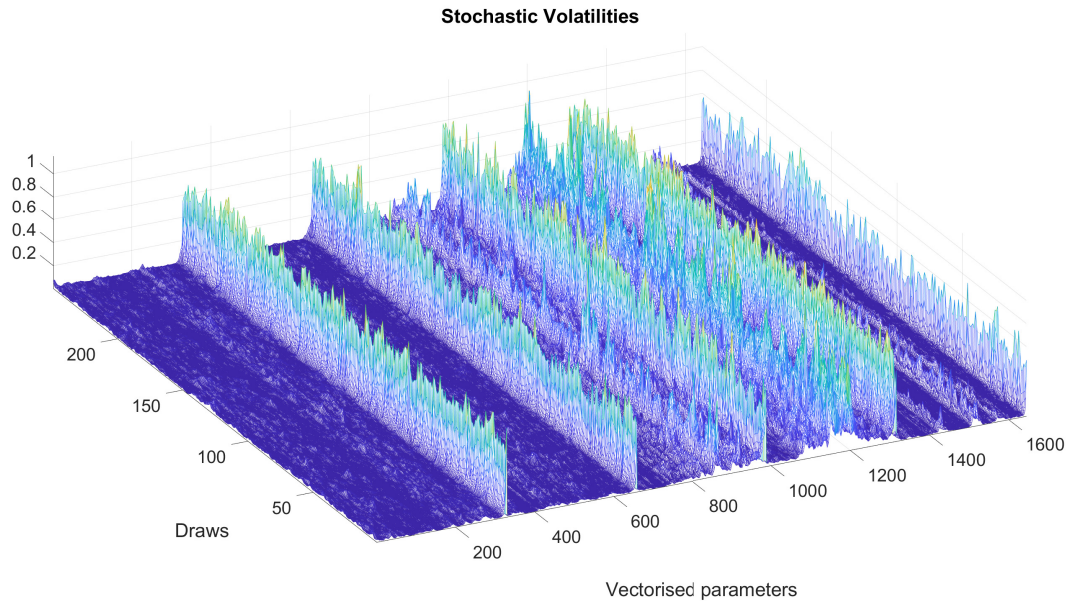

Figure B.9: Recursive means of the volatilities.

To ensure convergence of the Gibbs sampler, we estimate the recursive means of the retained draws for the time-varying parameters and stochastic volatilities of the baseline model. Figure B.8 and B.9 present the recursive means for every 20 draws of the vectorized time-varying parameters and stochastic volatilities. Both graphs suggest convergence in the Gibbs sampler given that along the different draws, the recursive means remain stable and do not depict any sudden increase among the different draws.

TVP model

The posterior density of the k_Q parameter is plotted on Figure 3.2.

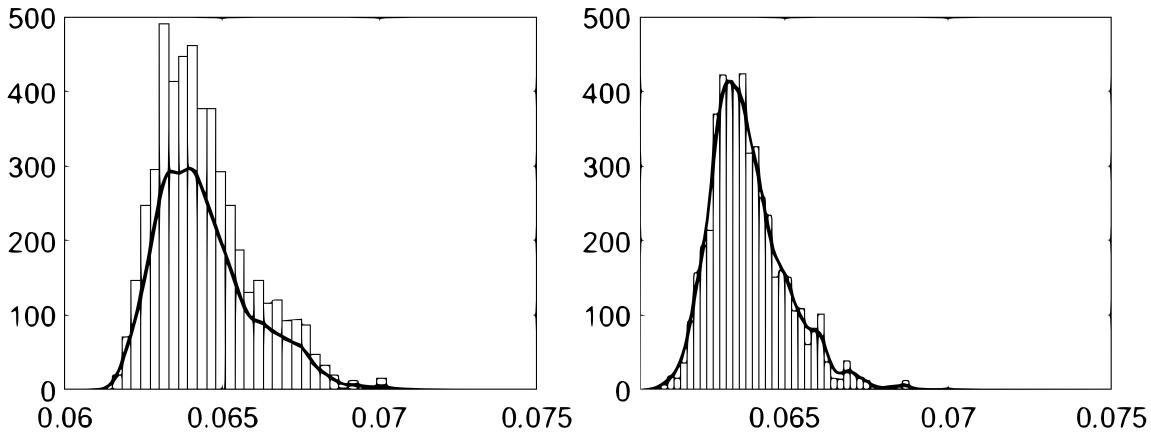

Figure B.10: Posterior densities of the time-variation parameter k_Q. Left: Shadow short rate model. Right: Federal funds rate model.

The strong evidence for time-variation in the data is also supported in the estimated stochastic volatilities. As the relationship between the variables change, the degrees of freedom for the model to fit the data increase, hence the stochastic volatility declines. We find that while the most important events of the estimated MPU remain the same across both specifications, the MPU estimated under the TVP model is markedly lower after the dot-com bubble.

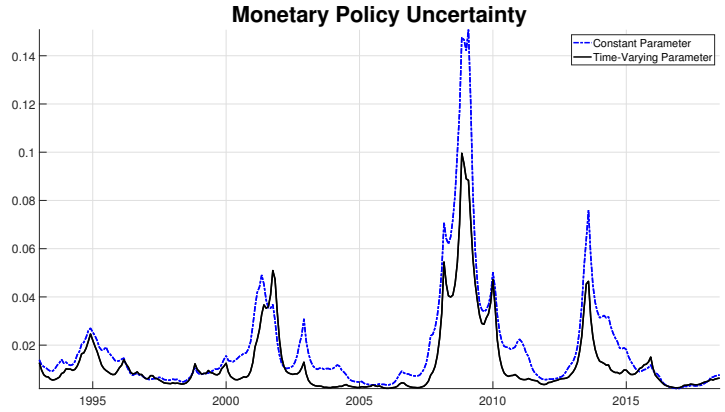

Figure B.11: Median estimates of Monetary Policy Uncertainty

Chapter 3

Financial integration or financial fragmentation? A euro area perspective

Abstract

In this article we address the question of how strongly bank lending rates and credit volumes co-move across the euro area. Following the breakdown in the interest rate pass-through across the euro area, we aim to disentangle the relative importance of country-specific and common components in explaining the variance of the macro and financial variables by using a time-varying two-level dynamic factor model. Our results show that a high share is explained by the common component. However, we find a persistent decline in the importance of the common factor in the bank lending rates, indicating the presence of financial fragmentation. Furthermore, we find persistent heterogeneity across member states, specifically those hit hard by the sovereign-debt crisis.

Published in Economic Modelling, DOI: 10.1016/j.econmod.2022.105902
Co-author: Boris Blagov
Keywords Co-movements, financial fragmentation, dynamic factor model

1 Introduction

The period during and after the Global Financial Crisis was characterized by a profound and synchronized decline in economic activity. The rapid spread of the financial crisis globally highlighted the importance of the analysis of financial cycles and degree of interconnectedness among financial sectors. Financial cycles are broadly referred to as common dynamics of financial variables within or across countries. For example, stock prices, house prices and credit demand typically increase during expansions. The existence of a medium-term global financial cycle has been strongly supported by recent literature (Rey, 2015; Miranda-Agrippino and Rey, 2020), which reinforces the importance of studying the degree of common dynamics (henceforth commonality) within the financial sector across countries. Originally, co-movement was referred to as patterns of positive correlation (Baur, 2003). However, more recent contributions have taken this concept to the analysis of common underlying drivers that explain a high fraction of the changes in the macro-financial variables (Del Negro and Otrok, 2008; Mumtaz and Surico, 2012; Mumtaz and Musso, 2021).

The concept of financial cycles is highly important for the euro area. Higher synchronization, especially in financial variables, could imply efficacy of monetary and macroprudential policies as a centralized action would be sufficient to dampen economic shocks. Cross-country heterogeneity, however, be that due to differences in financial regulation, fiscal policies or country specific shocks, reduces the commonality across macro-financial variables (Corsetti et al. (2020)). Therefore, financial fragmentation increases, leading to businesses and economic agents facing different financing conditions and divergent prices for otherwise similar assets, rendering a one-size-fits-all solution ineffective.[1]

The sovereign debt crisis was a prime example. Financial fragmentation manifested, among other ways, in divergent dynamics of bank lending rates across member states. While

[1] A good discussion on various aspects of the term financial fragmentation can be found in Claessens (2019).

the European Central Bank (ECB) lowered and then kept the policy rate at record low levels, lending rates for businesses and consumers in some countries stagnated and even increased. As a consequence of fiscal and financial stress the interest rate pass-through broke down in Italy, Spain, Ireland, Portugal and Greece, reducing investment demand and stifling the recovery (Ciccarelli et al., 2013; Blagov et al., 2015). These countries have since then recovered (before the world plunged in the recession triggered by the pandemic). However, how has the degree of financial fragmentation in bank lending rates evolved? Is it a cyclical phenomenon? Has the commonality, which declined throughout the sovereign debt crisis as evident by the breakdown in lending rates pass-through, increased again?

To answer these questions, we analyze the evolution of the co-movement in the bank lending rates and credit volumes in the euro area using a large monthly dataset for seven euro area countries that account for more than 70% of the euro area GDP. To this end, we employ a sectoral multi-level dynamic factor model (DFM) with time-varying factor loadings. The sectoral aspect is that we group macro and financial variables by categories which permits attaching an economic interpretation to the factors. The multi-level feature of the DFM allows for a differentiation between fluctuations driven by a common component for all countries and those that have a local origin, i.e. driven by country-specific dynamics. The time-varying factor loadings allow for a dynamic variance decomposition analysis, i.e., we can capture the evolution of the explanatory power of common and country-specific fluctuations over time. The combination of these features allows us to quantify the extent to which the dynamics of the bank lending rates and credit volumes are explained by common drivers and how this has evolved throughout the past 20 years.

We find that overall the common component explains on average a high share of the variance in the bank lending rates and credit volumes, ranging between 34% and 53% for the former, and 26% and 50% for the latter. However, we find a clear downward trend in the relative importance of the common factor with respect to bank lending rates. The

commonality has declined throughout the sovereign debt crisis and never returned to pre-crisis levels. This result holds when we look into the different categories of lending rates. For the credit sector, the common factor explains a high share of the variation in the data, increasing in periods of high financial volatility. In contrast, with the bank lending rates we do not observe a trend in the relative importance of the common component within this category. Finally, we observe a clear distinction with regard to the common factor importance in the countries which experienced fiscal or financial stress throughout the crisis, i.e. Italy, Spain, Portugal, and Ireland compared to Germany, France and the Netherlands. These findings suggest that financial fragmentation in the euro area has persistently increased, at least with respect to firms' financing conditions. Large heterogeneity and a "north-south" divide appear to continue to exist.

The remainder of this article is organized as follows: the next section lays out the empirical method and estimation procedure, as well as the data used in this article. Sections 3 discusses the results. Section 4 concludes.

2 Methodology

In order to capture the degree of commonality and investigate the main drivers of the variations in the macro-financial variables, we estimate a two-level DFM as in Kose et al. (2003) and incorporate the contributions by Del Negro and Otrok (2008) by including time-varying factor loadings and stochastic volatilities. By employing this methodology we capitalize on an extensive data set while reducing the number of explanatory variables to a small set of factors. Another characteristic is that the multi-level structure allows us to disentangle the explanatory characteristics associated with each individual country from the ones that are common to all of them. The following equations characterize the model:

$$Y_{it} = B_{it}^{C} F_{t}^{C} + B_{it}^{E} F_{t}^{E} + u_{it}. \tag{3.1}$$

The observational equation (3.1) relates the matrix of endogenous variables Y_{it} to a set of K^C unobserved country-specific factors F^C, a set of K^E unobserved common factors F^E and the u_{it} idiosyncratic components. Y_{it} contains all the N macroeconomic and financial variables with $i = 1, \ldots, N$ and $t = 1, \ldots, T$. Additionally, indices C and E stand for "country-specific" and common "European", respectively.

We collect the factors in a matrix F_t of size $1 \times (N^C \times K^C + K^E)$, where N^C is the number of countries. Given the orthogonality of the factors, we can describe the VAR(p) process for F_t in the following transition equation describing each of the k^{th} columns of F_t:

$$F_{kt} = c_k + \sum_{j=1}^{P} b_{kj} F_{kt-j} + \sigma_{kt}^{1/2} e_{kt}, \quad e_{kt} \sim N(0,1) \tag{3.2}$$

The idiosyncratic components of the observational equation (3.1) follow an AR(q) process

$$u_{it} = \sum_{j=1}^{q} d_{ij} u_{it-j} + h_{it}^{1/2} v_{it}, \quad v_{it} \sim N(0,1) \tag{3.3}$$

The stochastic volatilities of the factors and idiosyncratic components σ_{kt} and h_{it} are, correspondingly, modeled as AR(1) processes:

$$\ln \sigma_{kt} = \ln \sigma_{kt-1} + g_{k}^{1/2} \epsilon_{kt}, \quad \epsilon_{kt} \sim N(0,1) \tag{3.4}$$

$$\ln h_{it} = \ln h_{it-1} + G_{i}^{1/2} \eta_{it} \quad \eta_{it} \sim N(0,1) \tag{3.5}$$

We model the time-variation in the factor loadings following Del Negro and Otrok (2008). The law of motion characterizing the time-variation in the factor loadings is described by a

random walk

$$B_{it} = B_{it-1} + Q_i^{1/2}\gamma_t, \quad \gamma_t \sim N(0, I_{2K(KP+1)}) \tag{3.6}$$

where the matrix B_{it} collects the corresponding factor loadings associated with the common components, B_{it}^E, and with the country-specific ones, B_{it}^C.

2.1 Priors and starting values

We first set the initial values of the factors using principal components over a training sample T_0, obtaining $\hat{F}_t$ and setting its variance to an identity matrix. Making use of the initial estimate for the factors $\hat{F}_t$, we estimate the equation $Y_{it} = B_i\hat{F}_t + u_{it}$, obtaining the initial values for the factor loadings $B_{i,0}$. The prior for the factor loadings, $P(B_{it}) \sim N(B_{i,0}, V_B)$, is assumed to be normal with mean $B_{i,0}$ and covariance $V_{i,B}$, both set using the OLS estimates from the training sample T_0 period.

The prior for the transition equation of the factor loadings, Q_i, has an inverse Wishart distribution with a scale parameter $Q_{0,i} = var(B_i) \times T_0 \times \kappa$. The hyperparameter κ, which determines the amount of time-variation, is set as in Cogley and Sargent (2005) at $\kappa = 3.5 \times 10^{-4}$ to ensure gradual parameter changes and thus capture long-term structural shifts.

The time varying coefficients of the transition equation of the factors, equation (3.2), have a normal prior distribution $P(\beta_k^F) \sim N(\beta_0^F, V_0^F)$, where $\beta_k^F = (c_k, b_{k1}^F, \ldots, b_{kP}^F)'$ and its prior variance is equal to an identity matrix . Similarly, the coefficients of the transition equation for the idiosyncratic component also follow a normal prior distribution $P(\beta_k^u) \sim N(\beta_0^u, V_0^u)$, with $\beta_k^u = (d_{i1}^u, \ldots, d_{iq}^u)'$.

Finally the prior for the stochastic volatilities of the factors and idiosyncratic component, g_k and G_i correspondingly, follow an Inverse Gamma distribution $P(g_k) \sim IG(g_0, d_o)$ and $P(G_i) \sim IG(G_0, D_o)$, where both scale parameters g_0 and G_0 are set to 0.01. d_o and D_0 are the corresponding degrees of freedom. The initial conditions for the stochastic volatilities, σ_{k0} and h_{i0}, are set following Del Negro and Otrok (2008) to address the issue of unidentified

scale of the factors. The authors highlight that the corresponding sign of the factors and their loadings are not identified independently. This is not an issue here since our analysis relies on the variance decomposition, which, as a product of the two, is invariant to the sign identification.

2.2 Estimation

The estimation is carried out via the Gibbs sampler, the steps can be summarized as follows:

1. Conditional on the draw of the factors and stochastic volatilities, the time-varying factor loadings are drawn from the conditional posterior distribution using the Carter and Kohn (1994b) algorithm.

2. Similarly using the Carter and Kohn (1994b) algorithm, conditional on the factor loadings, the stochastic volatilities $\sigma_{kt}^{1/2}$ and $h_{it}^{1/2}$, and the autoregressive coefficients b_{kj}, we obtain the draws from the conditional posterior distribution of the factors F_{kt}.

3. Finally, conditional on the factors, the stochastic volatilities of the model: $\ln \sigma_{kt}$ and $\ln h_{it}$ are drawn using the particle Gibbs sampler following Lindsten et al. (2014).

We use 35,000 iterations discarding the first 25,000 as burn-in to ensure convergence of the algorithm.

2.3 Data description

The data set is composed of 30 monthly macroeconomic and financial variables spanning from January 2003 until December 2020, for each of the seven countries analyzed: Germany, France, Ireland, Italy, the Netherlands, Portugal and Spain. The selection of countries is mainly motivated by data availability since we aim to consider a balanced panel containing a wide range of macroeconomic and financial variables. Furthermore, another important

consideration in the data selection was to include a representative sample of stressed and non-stressed countries as mentioned in Altavilla et al. (2020). In that regard, we consider the following two groups: Germany, France, and the Netherlands as the non-stressed group and Ireland, Italy, Portugal, and Spain as the group of countries that experienced significant periods of financial and fiscal distress.

We use the Monetary and Financial Institutions (MFI) data set from the European Central Bank (ECB) that collects harmonized financial variables. This allows us to examine different categories of credit volumes and lending rates while ensuring that they are directly comparable. Another advantage, besides direct comparability, is the level of disaggregation of the different components of interest rates and credit volumes; for example, one can observe how much common European factors affect variation in credits to house purchase in contrast with credits to consumption.

To preserve stationarity in the dataset, some variables are transformed. In this regard, we log difference credit volumes, stock prices, imports, exports and industrial production. Furthermore, we look at the first difference of Government Bonds and sentiment indicators.

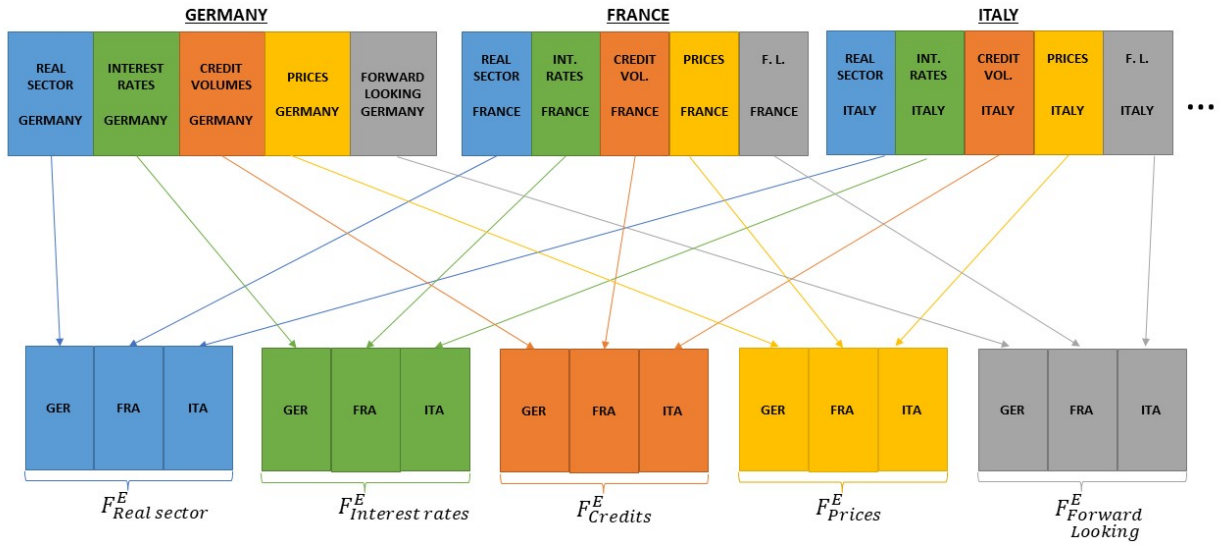

Figure 3.1: An example for Germany, France, and Italy of the categorization and common factor extraction for the *first level* of the Dynamic Factor Model. The full dataset includes also Spain, the Netherlands, Portugal and Ireland.

Figure 3.1 illustrates the corresponding five categories for Germany, France and Italy that are used to extract the common components of the model, i.e. the first level of the DFM (in the model we use all seven countries). At the second level of the DFM we obtain the country-specific dynamics of the categories. We do so by estimating a linear regression on the common components per country and category, thus orthogonalizing the country-specific datasets to the common factors.[2]

Figure 3.2 shows the estimated common factors together with the observed time series of each of the categories. The interest rate factor exhibits a downward trend following the Great Recession, these results are in line with Del Negro et al. (2019) who find a downward trend in interest rates among industrialized countries. The estimated factors for the real sector and forward-looking categories seem to capture mostly the business cycle dynamics,

[2]A visual interpretation with both levels is available in the Appendix, Figure C.1

as evident by their particular movements around the crisis periods, both the global financial crisis and the recent pandemic. Notably, the real sector factor appears to follow primarily the developments of the unemployment rates, which increased sharply in 2008 and even more so in 2020. The common credit volume factor does not exhibit notable volatility around the crises. Finally, the prices common factor fluctuates throughout the entire sample, showing an important decrease during the Great Recession.

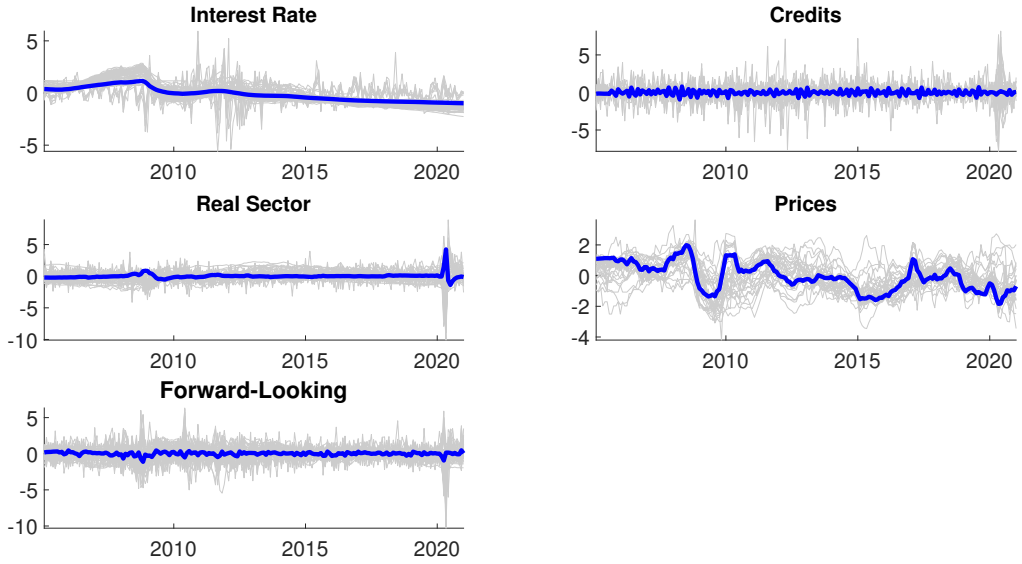

Figure 3.2: Estimated common factors. Blue line shows the median estimate of the common factors. Gray lines are the observed time series of the corresponding categories.

2.4 Model specification

To determine the lag order of the factors in equation (3.2), we follow the results from the Bayesian Information Criteria which suggests a lag order of 6. The number of factors in factor models is always a subject for careful consideration. We have allocated the variables in our dataset to five categories with the purpose of being able to attach an economic interpretation to the factors. The hierarchical structure of the two-level DFM induces proliferation in the factor count. For example, assuming one common and one country-specific factor per

category results in 5 common factors (1 per category) and 35 country-specific factors (7 countries in total, each with 5 categories), while a "two-two" factor setup would result in 80 factors. Our dataset consists of 30 variables per country, which dictates our choice for a more parsimonious specification. Therefore we set the number of common and country specific factors to one per category.

Another route is to use formal testing. For example, the test of Bai and Ng (2002), performed on the whole dataset without conditioning on any number of categories, suggests the use of 6 factors. Hence, our choice of 5 factors (while categorizing the data) does not appear disconnected.

3 Results

To answer the question of how important the common latent factors within the euro area are in explaining the volatility in the bank lending rates and credit volumes, we calculate the variance decomposition for each series as:

$$var(Y_{it}) = (B_{it}^{C})\Theta_{t}^{C}(B_{it}^{C})' + (B_{it}^{E})\Theta_{t}^{E}(B_{it}^{E})' + var(u_{it}), \tag{3.7}$$

where $\Theta_t^j = diag(var(F_t^j))$ with $j = [C, E]$. Equation (3.7) captures the fraction of the variance of the observed variables which is explained by the common components $S_t^E = \frac{(B_{it}^{E})\Theta_{t}^{E}(B_{it}^{E})'}{var(Y_{it})}$ and the country-specific component $S_t^C = \frac{(B_{it}^{C})\Theta_{t}^{C}(B_{it}^{C})'}{var(Y_{it})}$. It is worth mentioning that the time variation in the variance decomposition comes from the specification of a law of motion for the idiosyncratic component and not from the time-varying factor loadings. The incorporation of time-varying factor loadings in the model specification is motivated by increases in explanatory power, as our aim is to capture the biggest share of the observed variance in our dataset. As a robustness check and to ensure that the overall results from the variance decomposition are not driven by the specification of the time-varying factor

loadings, we considered an alternative specification with a one-third smaller hyperparameter κ (which accounts for the amount of time variation in the factor loadings), therefore, we suppress the amount of time variation close to being fixed and observe that the results are not sensitive to this specification[3].

We obtain a variance decomposition for each series, i.e. the index i in equation (3.7) runs from 1 to 210. Thus, the model allows for a broad spectrum of analysis since the variance decomposition can be aggregated across various dimensions. We will begin the analysis by first taking averages across all countries per category before diving into more granular data: decomposition per individual time series, such as short or long-term interest rates and decompositions across the country dimension.

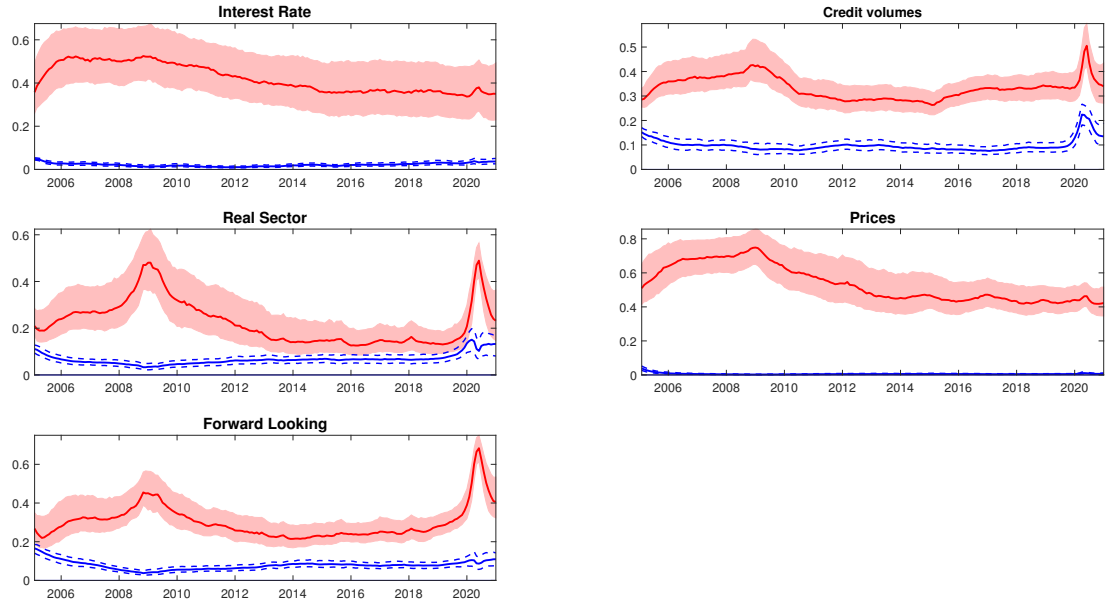

Figure 3.1: Average variance decomposition per category. Solid red line shows the contribution of the common component. Solid blue line represents the contribution from the country-specific component. Median estimates are plotted with their corresponding 68% probability intervals.

Figure 3.1 depicts the relative contribution of each of the factors in explaining the average

[3]Results are shown in Appendix B.

variance of the data in each of the categories for all countries. We present the median estimate for each of the contributions to the variance decomposition, together with the 68% credible sets[4].

For all categories we observe that the common factors are the main component explaining a high share of the variation in the data relative to the country-specific one. The top panel of Figure 3.1 shows how the importance of the common factor has evolved on average over all lending rates in our dataset. There is a clear downward trend which starts around the global financial crisis and continues throughout the sovereign debt crisis up until around 2015. Notably, the average contribution of the common component has not returned to the pre-crises levels. The common factor now explains approximately 35% of the variation of the interest rates compared to the 52% in the pre-crisis period. This finding illustrates that financial fragmentation has on average increased during the sovereign debt crisis, at least when it comes to the financing conditions that businesses face. More importantly, however, it has remained elevated ever since. This suggests that this was not a cyclical phenomenon. While financial fragmentation did increase, it did not decline endogenously as the economies recovered out of the crisis.

In contrast, credit category depicts an important share of the variation explained by the common factor throughout the whole sample analyzed, increasing during the periods of high financial volatility. In addition to the latter, country-specific factor explains on average around 12% of the variation of the credit volumes with an important increase in the recent downturn driven by the pandemic. The real sector exhibits high commonality during recessionary periods, raising its importance up to 50% of the variance explained by the common factor and with the country-specific factor ranging from 5% to 15%. Similarly, the forward-looking factor follows very closely the dynamics displayed in the real sector category, with the distinction that it presented an even higher level of commonality during

[4]We display the average contribution of the country-specific and common components for each member state in Figure C.1 of the Appendix.

the global pandemic, with a common component explaining up to 69% of the variance in this category. The country-specific components exhibit a smaller contribution to the variance decomposition compared to the common component.

For the credit category, the country-specific component reaches the highest level during the current Covid-19 pandemic. Throughout the first half of 2020 the country-specific component explained up to 25% of the variance. For the real sector and forward-looking categories, the country-specific component explains between 10% to 15%. For nominal variables, i.e. the interest rate and price categories, the country factors do not seem to explain a considerable amount of the variation in the data.

We exploit the richness of the harmonized data set from the MFI by analyzing the dynamics within the different types of bank lending rates and credit volumes. Figure 3.2 displays the average contribution of the common factors across different categories of interest rates and credit volumes. Figure 3.2a shows a clear downward trend in all categories of the interest rates. This trend begins around the global financial crisis and continues throughout the sovereign debt crisis with varying slopes for the different rates. The average contribution of the common component has not returned to the pre-crisis level for any of the time series considered, suggesting that this was not a cyclical phenomenon.

In terms of magnitude, we observe that lending rates to house purchases exhibit the highest level of commonality throughout the entire sample analyzed. These results contrast with the findings by Breitung and Eickmeier (2016) who find low levels of co-movements in the housing market. Additionally, we observe that lending rates to non-financial corporations - both short and long term - co-move more in comparison with other categories of composite cost of borrowing. It is worth mentioning that the decline in the relative contribution of the common component in the bank lending rates has not been compensated by an increase in the country-specific component but rather a higher proportion explained by the idiosyncratic component. Therefore, we analyze the overall model fit and present an example in Figure

C.2. Overall, we observe that, while the common factors fit quite well the dynamics of different lending rates (especially for non-stressed countries), common components have failed to follow closely recent developments in lending rates, mainly among countries that suffered from financially stressed periods.

Similarly to the result observed in the lending rates, Figure 3.2b shows that within different categories of credit volumes, small credits (below one million euros) to non-financial corporations present the highest level of commonality. This is not the case for large credits to non-financial corporations which can be explained by the specifics of how large firms obtain funds. To minimize risk, banks may group with other banks to give out a large credit to one corporation. In general, market forces play little role with credits of such magnitudes. Overall, the credit volumes showed a high fraction of the variance explained by the common factors throughout the time period analyzed, with increases in the importance during periods of high financial volatility.

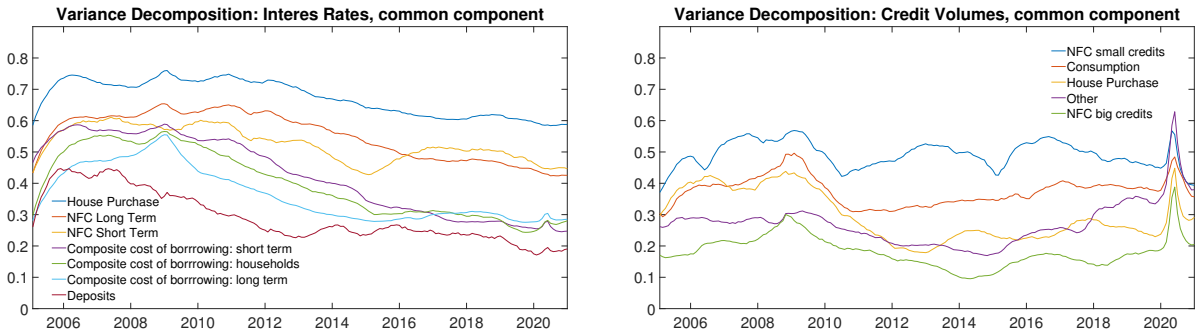

(a) Interest rates category

(b) Credit volumes category

Figure 3.2: Contribution of the common component in the interest rates and credit volumes categories per component. NFC stands for "non-financial corporations". Small credits are below one million euros. Short and long term refer to maturities below and above one year, respectively.

Next, we disaggregate the relative contribution of the common factors among countries and categories. We find a clear distinction between stressed countries (Ireland, Italy, Portugal and Spain) and countries which did not face a financial stress period (Germany, France and

the Netherlands).

This difference is shown in Figure 3.3 where we plot the relative importance of the common factors for the stressed countries in pale blue color and the non-stressed countries in black. We observe a high degree of country heterogeneity in the interest rate, the real sector and the forward-looking categories. The first subfigure decomposes the interest rate category per country. We draw attention to the countries that experienced financial stress, which all exhibit much lower commonality than the rest. In the cases of Portugal and Spain, the dynamics of lending rates can be explained only about 20% to 30% by the variance of the common component. Italy and Ireland also appear detached from the common European dynamics, however, to a smaller extent. Notably, for all countries classified as "non-stressed", namely Germany, France and the Netherlands, the common component explains a much larger fraction of the variance of the interest rates, which goes as high as 75%. However, even for Germany, France and the Netherlands a downward trend in the interest rate category is evident [5].

It is of particular interest, that a pronounced downward trend in the commonality in the interest rates category is visible for the sovereign debt crisis period and how there is almost a kink around 2015 in all countries, where the trend flattens. The economic explanation behind the downward trend may be primarily attributed to the decrease in the interest rate pass-through, which was well documented in that time - an increase in the financial fragmentation (e.g. Ciccarelli et al., 2013). The flattening of the common component is evidence that, despite the reforms following the crisis, financial fragmentation remained.

Similar developments are also present in the price category in Figure 3.3. Naturally, the peak of the explanatory power of the common component dynamics was during the global financial crisis, when many macroeconomic variables, including prices, jointly plunged. However, in the years afterwards, the importance of the common component has gone down

[5]Figure C.4 in the appendix shows how this result is present across different types of interest rates and for all member states considered.

in all countries, similarly to the interest rates category. In the case of prices, however, we do not find a distinction between stressed or non-stressed member states.

While in most cases the contribution of the common component to the variance decreased or remained flat, in few instances, such as the credit category, we see an upward trend. This indicates a convergence in the dynamics of the growth of credit volumes across the countries. Notably the not financially distressed countries, Germany, France and Netherlands, and to a lesser extent Portugal and Ireland, show a gradual increase in the explanatory power of the common component relative to the country and idiosyncratic components since the mid 2010s. One possible economic explanation behind this finding is that growth of the volume of credits was similar because all countries experienced an expansion phase following the sovereign debt crisis.

Another plausible cause could be the changes in banking regulation following the Great Recession. In 2012 the euro area members agreed to allow the European Stability Mechanism to give out loans directly to stressed banks rather than going through each individual member. The ECB also stepped into the role of a bank regulator by forming a deposit insurance program in addition to national programs. This could have fueled common credit volume growth dynamics across countries.

These findings are particularly relevant during times of expansive monetary policy and a low and stable interest rate environment. A prolonged period of unchanged interest rates would ideally translate into stable lending rates for businesses and consumers alike as well as stable price dynamics. Nonetheless, we find that the opposite is the case - the idiosyncratic component has gained much more relevance compared to the late 2000s across all possible dimensions: countries, categories and individual interest rate series. A plausible explanation for this aspect is the composition of the dataset. While we profit from monthly observations which are comparable across countries, we may overlook information from lower-frequency variables such as: government spending, tax policies, among others. If these variables are

relevant for explaining variations in the data, incorporating them will attenuate the relevance of the idiosyncratic component in the model.

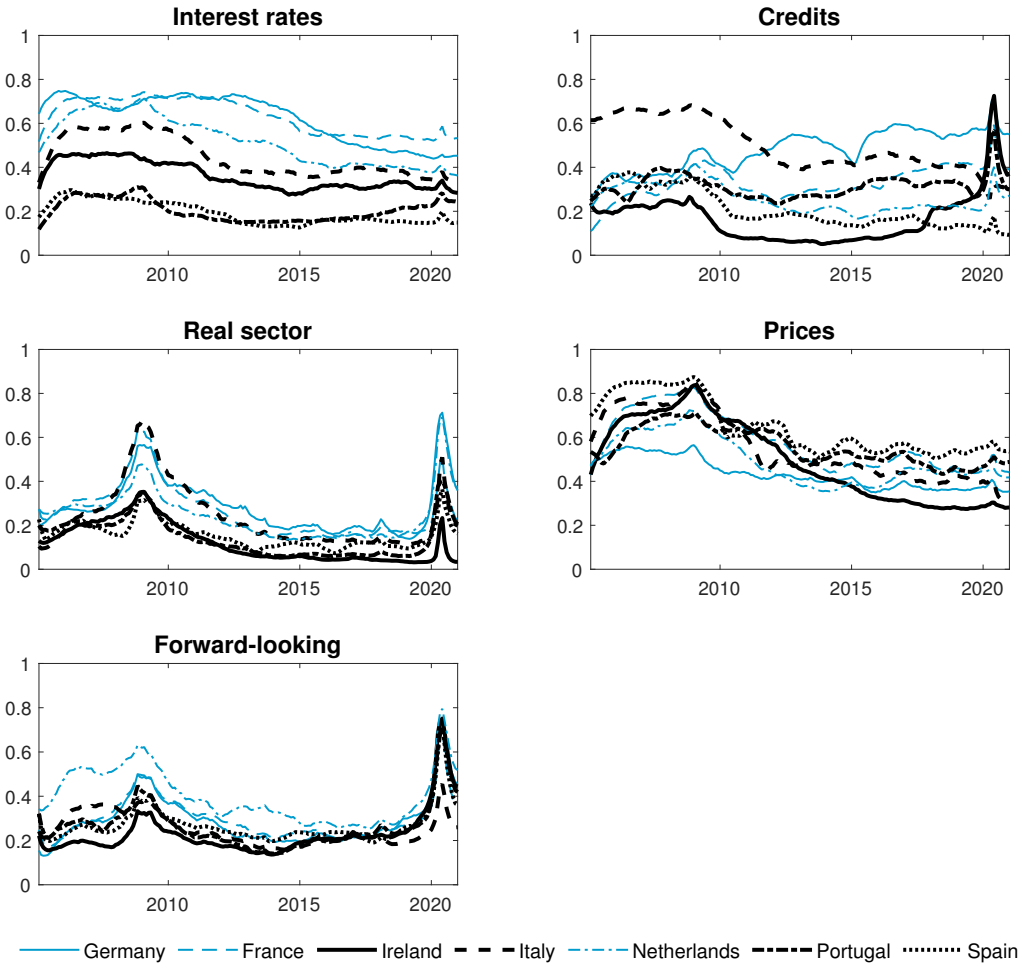

Figure 3.3: Country-wise relative contribution of the common component in each category. The **black** lines group the stressed countries while the light blue lines indicate countries which did not experience severe financial stress during the sovereign debt crisis.

Our results are in line with the findings of Breitung and Eickmeier (2016) who observe high commonality in stock prices and interest rates. However, given the time-varying feature of our approach, we observe that this commonality has decreased in recent years across different types of lending rates. Nonetheless, this is not the case for stock prices, as shown in Figure 3.4. In contrast with Breitung and Eickmeier (2016), we find that common drivers have

a fairly high contribution in explaining the variation in credit volumes, although, it varies across different credit categories. Furthermore, findings of heterogeneity patterns are aligned with the ones from the WGEM Team on Real and Financial Cycles (2018). In this article, the authors find evidence of a "north-south" division, as the one discussed previously. We observe this division primarily in the interest rate, real sector and forward-looking categories. However, in contrast with their findings, we do not rely on aggregate measures of interest rates but rather look into different categories of lending rates. This allows us, not only to observe differences depending on maturities, but also to observe sector specific behaviors. For example, to compare lending rates for house purchase with respect to lending rates for non-financial corporations.

Finally, we find neither a distinction of stressed versus non-stressed member states nor particular trends in the other categories. Perhaps a somewhat unexpected finding is the low commonality across the variables associated with the real sector. One may expect that a more integrated euro area should lead to synchronized business cycles. However, high synchronization seems to arise only during large recessions as all variables exhibit large (mostly downward) movements.

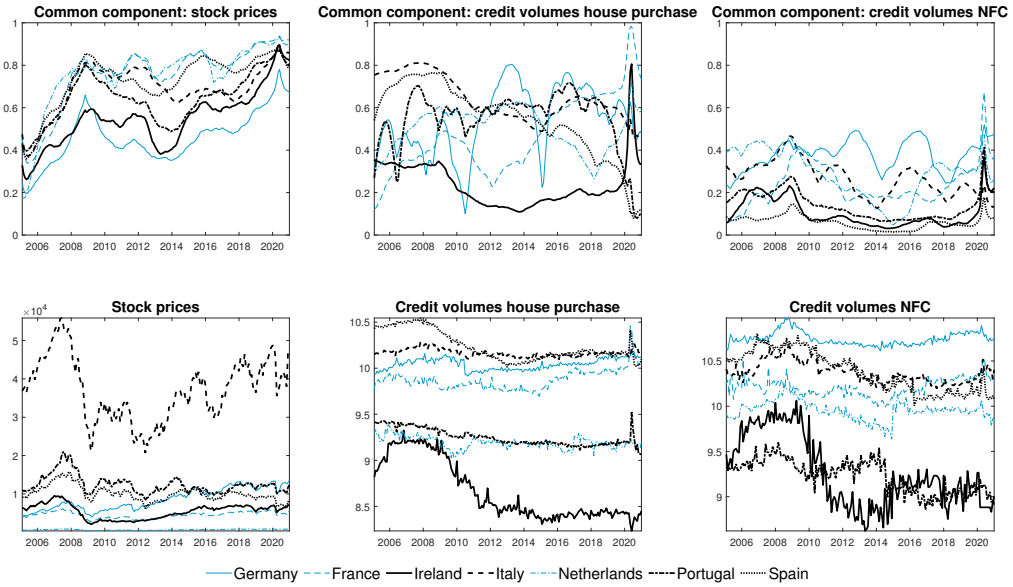

Figure 3.4: Top row: relative contribution of the common component for variables selected. Bottom row: observed time series of the variables selected, raw data. **Black** lines group the stressed countries while light blue lines indicate countries which did not experience severe financial stress during the sovereign debt crisis.

4 Conclusions

We use a two-level dynamic factor model to investigate the relative importance of the common factors in explaining the variations in the bank lending rates and credit volumes for the euro area between the period of January 2003 to December 2020. We observe that the common factors explain a high share of the variations in the data for all five categories analysed. Within the interest category we recognize a downward trend in the relative importance of these common factors in explaining the variation in the data. This decrease in the commonality started after the Great Recession and has not returned to the pre-crisis level, implying an increase in the financial fragmentation among the euro area countries. The trend is evident across all dimensions of the data: across countries, across categories and at an individual level.

When comparing the financially stressed countries with their non-stressed counterparts, we observe a clear heterogeneity in the fraction explained by the common factors in the interest rate, real sector and forward-looking category. The heterogeneity is mostly visible when it comes to the lending rates, where the commonality is extremely low in Spain, Italy, Ireland, and Portugal. The fraction of the variance of the rates explained by the common component is only between 20% to 40%, versus 40% to 60% for Germany, France, and the Netherlands. Such a large divide suggests continued impairment of the monetary policy transmission presenting potential avenues for future research.

We do not find such heterogeneities when it comes to the amount of credits as measured by the credit volumes or in the price category. There does not appear to be a particular divide between the two groups of countries. Finally, the real sector, as well as the forward-looking variables that reflect mostly the economic situation, exhibit an expected increase in the level of commonality during recessionary periods. Otherwise the degree of co-movement appears surprisingly low. During expansions the common business cycle factor explains only about 20% of the variation of the variables.

C Appendix

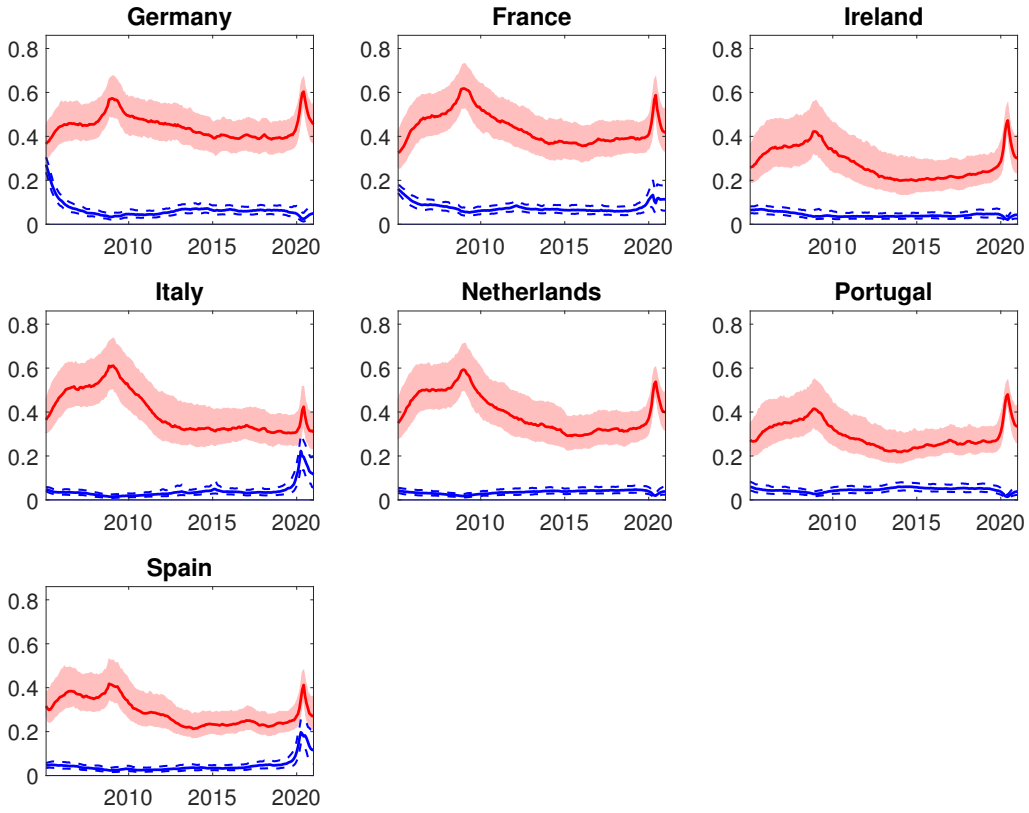

Figure C.1: Average variance decomposition by country. Solid red line shows the contribution of the common component. Solid blue line represents the contribution from the country-specific component.

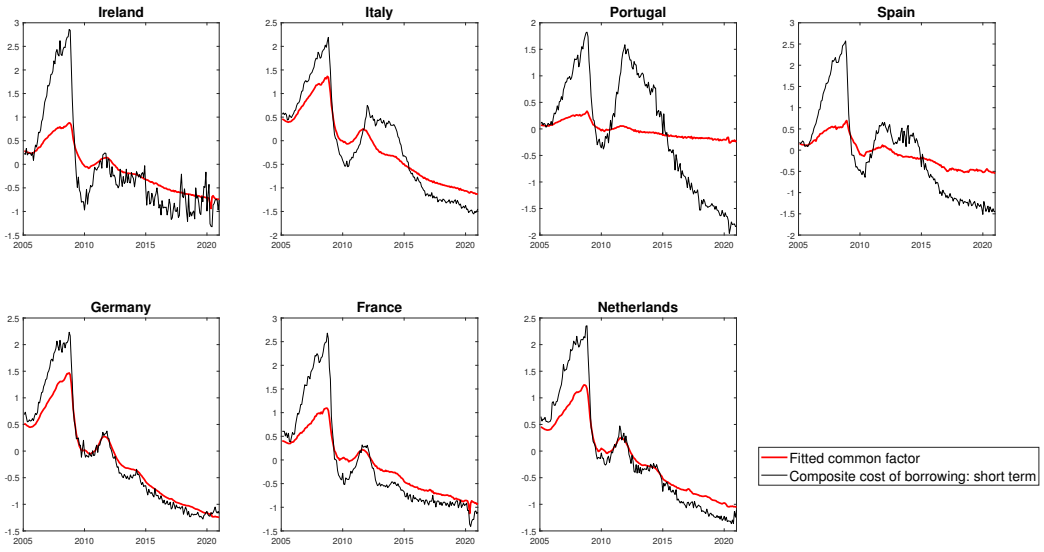

Figure C.2: Fitted values using the common factor. Black line plot the short term composite cost of borrowing and red lines plot the fitted values using the common component.

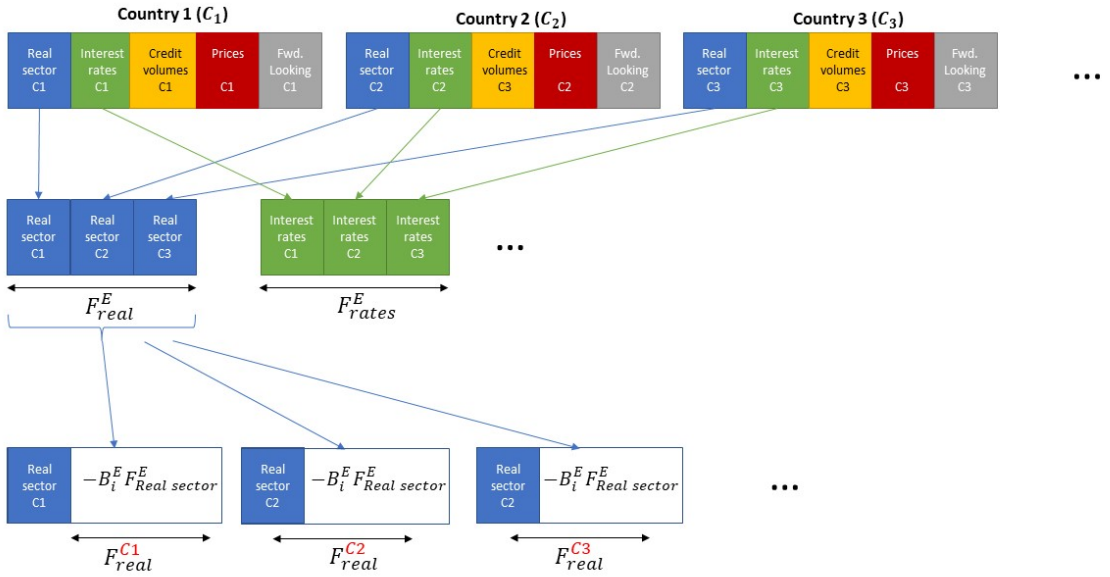

Figure C.3: Summarizes the structure of the model where the country-specific factors are obtained after extracting the effects of the common factor at a country-level.

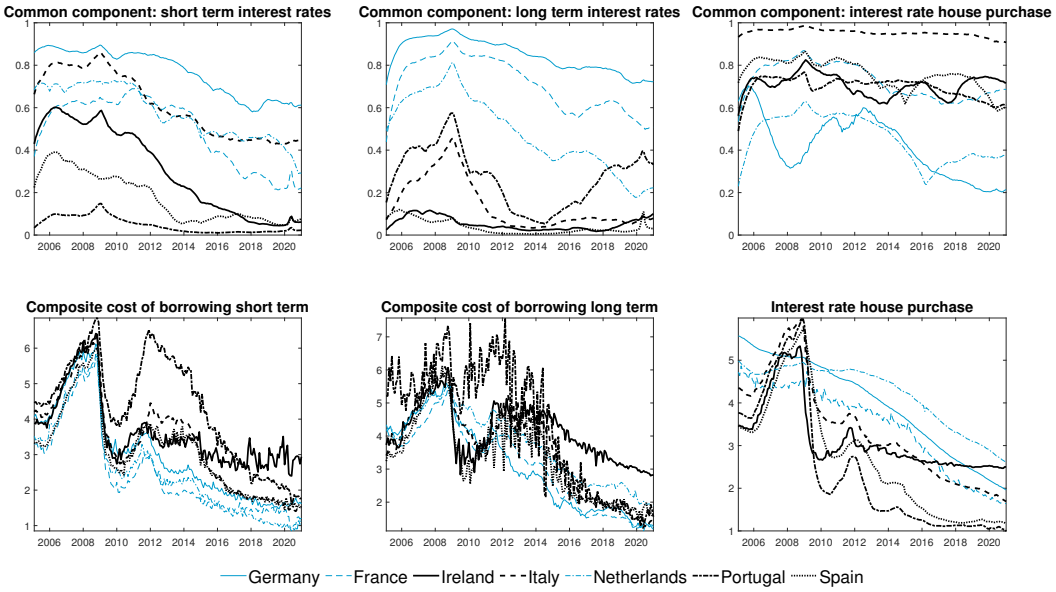

Figure C.4: Country-wise relative contribution of the common component for different variables. The **black** lines group the stressed countries while the light blue lines indicate countries which did not experience severe financial stress during the sovereign debt crisis.

Robustness

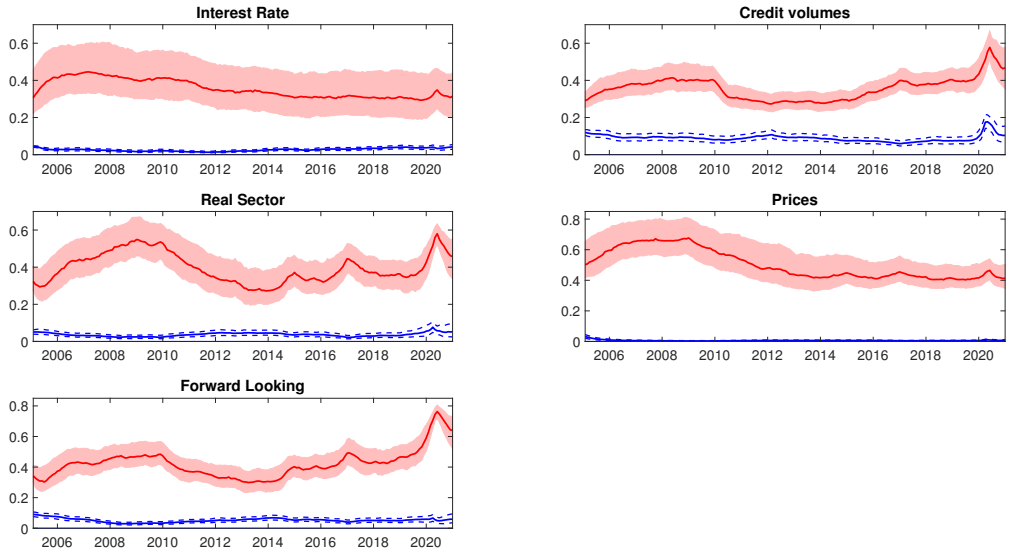

Figure C.5: Average variance decomposition per category under alternative specification. Solid red line shows the contribution of the common component. Solid blue line represents the contribution from the country-specific component. Median estimates are plotted with their corresponding 68% probability intervals.

One plausible concern is that the results derived from equation 3.7 are mainly driven by a priori beliefs on the amount of time variation the factor loadings are allowed to exhibit. To ensure that the results are not sensitive to this specification, we estimate an alternative specification reducing the amount of time variation in the factor loadings of the observational equation 3.1. We do so by reducing to one third the scaling parameter κ from the variance of the transition equation of the factor loadings B_{it}. Figure C.5 shows that the average variance decomposition across categories present very similar patterns as the ones showed in figure 3.1.

In contrast with Mumtaz and Musso (2021), we do not set the priors for the coefficients in equation 3.2 using dummy observations. Instead, we use a training sample to set the initial values. In this regard, we estimated the model with different choices of the training sample

T_0 and observed that the results were not sensitive to this specification.

Variable	Source	Trans	Category
Unemployment	Eurostat	1	Real sector
Imports	Eurostat	5	Real sector
Exports	Eurostat	5	Real sector
International Trade	Eurostat	1	Real sector
Real Exchange Rate	IMF	1	Prices
Harmonized Index of Consumer Prices	Eurostat	1	Prices
Producer Price Index	Eurostat	1	Prices
Commodity Import Price Index	IMF	1	Prices
Core Inflation	Eurostat	1	Prices
Stock Prices Index	MacroBond	5	Forward-Looking
Government Bonds: 10 Years	MacroBond	2	Interest rates
Yield Curve	MacroBond	1	Forward-Looking
Composite Cost of Borrowing: Households	ECB, MFI statistics	1	Interest rates
Composite Cost of Borrowing: Long Term	ECB, MFI statistics	1	Interest rates
Composite Cost of Borrowing: Short Term	ECB, MFI statistics	1	Interest rates
Interest Rate for House Purchase	ECB, MFI statistics	1	Interest rates
Interest Rate to Non-Financial Corporations: short term	ECB, MFI statistics	1	Interest rates
Interest Rate to Non-Financial Corporations: long term	ECB, MFI statistics	1	Interest rates
Interest rate on deposits	ECB, MFI statistics	1	Interest rates
Sentiment Indicators: Construction Confidence Indicator	Eurostat	2	Forward-Looking
Sentiment Indicators: Consumer Confidence Indicator	Eurostat	2	Forward-Looking
Sentiment Indicators: Industrial Confidence Indicator	Eurostat	2	Forward-Looking
Index of Financial Stress	ECB	2	Forward-Looking
Sovereign Systemic Stress Composite Indicator	ECB	2	Forward-Looking
Industrial Production	Eurostat	5	Real sector
Credit volumes to non-financial corporations: small (below one million)	ECB, MFI statistics	5	Credits
Credit volumes to non-financial corporations: big (above one million)	ECB, MFI statistics	5	Credits
Credit volumes for consumption	ECB, MFI statistics	5	Credits
Credit volumes for house purchase	ECB, MFI statistics	5	Credits
Credit volumes other credits	ECB, MFI statistics	5	Credits

Table 3.1: This table summarizes the time series used in the model. Transformation code (TC): 1-level; 2-first difference; 3-second difference; 4-log-level; 5-first difference of logarithm

Chapter 4

Improving structural identification in sign-identified VAR models via forecast error variance

Abstract

This article discusses the benefits of using information from the forecast error variance as an additional restriction in sign-identified models. We show with a Monte Carlo simulation that our method successfully recovers the true Data Generating Process while narrowing down the set of admissible models and thus provides sharper identification. We exemplify the usefulness of this approach for the case of a contractionary monetary policy shock applied to the U.S. and U.K. economy. For the former, we show that uninformative responses of prices to a monetary policy shock are explained by retained draws with a low share of the forecast error variance explained by this shock. We provide evidence that the uninformative responses are no longer present when our approach is applied. The latter application illustrates how sharper identification lead to more informative results for the response of inflation to a monetary policy tightening.

Keywords Structural vector autoregressive model, identification problem, sign restrictions, forecast error variance.

1 Introduction

Sign restrictions has become one of the predominant approaches to identify shocks in structural vector autoregressive models (SVAR). First introduced by Faust (1998), Canova and De Nicolo (2002), and Uhlig (2005) sign-identified models offer the advantage of placing a set of commonly accepted signs, motivated by economic theory or previous empirical evidence, which reflect conventional views on the variables responses to exogenous shocks.

On these lines, sign restrictions appear as a less restrictive approach than exclusion restrictions on the impact effects. However, this identification method entails different drawbacks. The identification scheme suffers from the risk of allowing different types of models in the set of retained draws. This is translated into including potentially contradicting models that imply impulse responses which are inconsistent with each other, different timing of the strongest response and which deliver different estimates in the historical decomposition (HD) and forecast error variance decomposition (FEVD). Thus, allowing potentially implausible or problematic responses in the retained set will not only result in imprecise estimations but additionally hinder determining the true responses after an exogenous shock.

Recent literature has proposed different approaches to circumvent the deficiencies of purely sign-identified models. Kilian and Murphy (2012) point out that they will deliver an implausibly high oil-supply elasticity. Thus, the authors exploit the idea that the elasticity of oil supply with respect to the real price of oil should be rather small and hence, by setting an upper bound on the elasticity value, narrow down the set of admissible rotations. Arias et al. (2019) combine sign and zero restrictions to identify the effects of a monetary policy shock observing that, in contrast with the results from Uhlig (2005), a monetary policy tightening has contractionary effects on output. The authors observe that, while keeping the sign of output response to a monetary policy shock unrestricted, adding exclusion restrictions reduces the set of structural parameters retained and therefore exhibits a drop on output.

Antolín-Díaz and Rubio-Ramírez (2018) propose an alternative identification to the pure

sign-identified SVAR by imposing restrictions on the historical decomposition. They place restrictions on shock contributions to specific historical events, leaving out models in which shocks challenge the conventional narrative in those periods.

This article proposes an alternative identification scheme combining sign restrictions with information from the forecast error variance (FEV). We suggest that in most applications, the researcher is interested in investigating the effects of an exogenous shock j on variable i. Therefore, employing information about the contribution of shock j on variable i regarding its share to the FEV as an additional restriction improves the identification of the shock of interest. In practice, information from the FEV has been applied in the context of sign-identified models (Faust, 1998; Uhlig, 2003; Barsky and Sims, 2011; Francis et al., 2014). This approach, also known as Max-Share identification, identifies a shock that maximizes its contribution to the FEV on the target variable. Commonly applied to identify technology shocks, the Max-Share identification suffers from confounding shocks (Dieppe et al., 2021). More recently, Volpicella (2021) imposes bounds on the FEV to identify the structural shocks and therefore avoids the shock labeling problem present in the Max-Share identification scheme.

Our approach is closer in spirit to Antolín-Díaz and Rubio-Ramírez (2018) in aiming to reduce the set of admissible rotations by incorporating an additional restriction to the sign-identified SVAR. However, in contrast with Antolín-Díaz and Rubio-Ramírez (2018), we narrow down the set of admissible models by evaluating the condition that the shock of interest should be informative in terms of its contribution to the FEV to the target variable(s). In this sense, our approach still possesses the advantages of sign-identified models, while overcoming limitations like the large and disperse responses. Furthermore, our approach does not suffer from the shock labeling limitation of the Max-Share approach, as we do not identify a unique shock that maximizes the FEV but rather set the restriction such that the shock of interest on the target variable should be the most informative in relative terms.

We exemplify the advantages of this approach with a classical application in the literature, namely the effects of a monetary policy tightening. We investigate the effects of a monetary shock for the United States (U.S.) and for the United Kingdom (U.K.) economy. For the first example, we impose the restriction that the monetary policy shock should be informative, in terms of the FEV, for the response of inflation, avoiding ruling out a-priori the possibility of a price puzzle, i.e. a positive response of prices after a contractionary monetary policy. In our second application we add the restriction that the contribution of the monetary policy shock on prices should be at least bigger than 5% during the first quarter and illustrate how we can narrow the responses and thus, provide more informative results in terms of the magnitude and persistence of the responses.

The benefits of our alternative approach are supported by a simulation exercise where we compare the results of the proposed identification method to commonly employed identification schemes in recovering the impulse responses from a Data Generating Process (DGP). The Monte Carlo study allows us to observe that the proposed identification method successfully recovers the structural responses generated by the DGP. The simulation exercise shows the improvements with respect to conventional identification methods.

The remainder of this article is structured as follows. Section 2 describes the methodology. Section 3 present the results of the Monte Carlo simulation. Section 4 illustrates the usefulness of this approach with two applications to the analysis of a monetary policy shock. Finally section 5 concludes.

2 Methodology

Consider the following structural vector autoregressive model (SVAR) with p lags given by:

$$B_0 y_t = v + \sum_{i=1}^{p} B_i y_{t-i} + \epsilon_t,\ \epsilon_t \sim N(0, I_K), \tag{4.1}$$

where $y_t = (y_{1t}, ..., y_{Kt})'$ is a $K \times 1$ vector of endogenous variables. v is the $K \times 1$ vector of constants. B_i is a $K \times K$ matrix of structural coefficients with $i = 1, ..., p$. The reduced-form counterpart of equation (4.1) is represented by:

$$y_t = c + \sum_{i=1}^{p} A_i y_{t-i} + u_t,\ u_t \sim N(0, \Sigma_u), \tag{4.2}$$

with $A_i = B_0^{-1} B_i$ and $u_t = B_0^{-1} \epsilon_t$. Thus we aim to find a linear mapping between the reduced-form and structural errors such that $B_0^{-1}(B_0^{-1})' = \Sigma_u$. Under stationary conditions, there exists a moving average representation MA(∞) of equation (4.2) such that $y_t = A(L)u_t$, where $A(L)$ is a matrix polynomial in the lag operator L. The MA(∞) representation of equation (4.1) is given by $y_t = B(L)\epsilon_t$; therefore, we aim to find a linear mapping between the reduced-form and its structural counterpart such that: $u_t = B_0^{-1}\epsilon_t$ and the impact matrix B_0 satisfies the condition $B_0^{-1}(B_0^{-1})' = \Sigma_u$. Under full invertibility one plausible solution for the decomposition of the covariance matrix Σ_u is to obtain $\tilde{B}_0$ by using the Cholesky decomposition of Σ_u.

2.1 Sign Restrictions

One commonly applied identification strategy is to impose sign restrictions on the IRFs according to previous knowledge coming from either economic theory or previous empirical evidence. Following Uhlig (2005), this is achieved by first obtaining an orthonormal matrix, also referred as rotation matrix, Q such that $QQ' = I_K$. Q is obtained with a probability ϕ, which does not contain information about Q. The set of drawn rotations belong to the space of orthonormal matrices $Q \in \Theta(K)$. Therefore, the identification strategy will only imply set-identification, where the set of sign restrictions of shock j at horizon h can be expressed by S_{hj}.

2.2 Max-Share

A strand of the literature has proposed identifying the rotation that is associated with the highest share of the FEV on the target variable (Uhlig, 2003; Barsky and Sims, 2011; Francis et al., 2014). Defining the h-step forecast error of equation (4.1) is as:

$$y_{t+h} - E_{t-1}y_{t+h} = \sum_{l=0}^{h} B_l \bar{B}_0 Q \epsilon_{t+h-l} \tag{4.3}$$

where $\bar{B}_0 = \tilde{B}_0 Q$ and Q corresponds to the orthonormal rotation matrix. Hence, one can make use of a $K \times 1$ selection vector e_i and thus, obtain the shock that maximizes the share to the FEV of variable i at horizon h by solving the following maximization problem:

$$\Omega_{i,j}(h) = \frac{e_i'(\sum_{l=0}^{h} B_{i,l}\bar{B}_0\gamma\gamma'\bar{B}_0'B_{i,l}')e_i}{e_i'(\sum_{l=0}^{h} B_{i,l}\Sigma_u B_{i,l}')e_i},\ s.t.\ \gamma'\gamma = 1 \tag{4.4}$$

where γ denotes the selection vector choosing the jth column of Q. Hence, the impulse vector associated to the jth column of the orthogonalization is contained in the $K \times 1$ vector $\tilde{B}_0\gamma$. The selection vector e_i chooses the ith row of the matrix of MA coefficients $B_{i,l}$.

This approach has been implemented for the analysis of technology shocks (Barsky and Sims, 2011; Francis et al., 2014). In these examples, one can think of the maximization problem in equation (4.4) as identifying a structural shock that maximizes the share to the FEV of total factor productivity (TFP) and that it has no contemporaneous effect on TFP (orthogonal). Ordering the target variable first (TFP for this example), the maximization problem can be expressed as follows:

$$\max_{\gamma} \sum_{h=0}^{H} \Omega_{1,2}(h)\ s.t.\ \gamma\gamma' = 1\ and\ \gamma(1,1) = 0 \tag{4.5}$$

2.3 Sign Restrictions + FEV

We exploit the property that the forecast error variance contains information about the magnitude in which economically meaningful shocks explain variations in the variables of interest. In this sense, it is therefore intuitive to focus on the set of models in which the shock of interest is informative for the variables whose responses we want to analyze.

It is often the case that sign-identified VARs exhibit wide error bands leading to uninformative results. As highlighted in Kilian and Murphy (2012), pure sign-identified models allow the risk of accepting, in the set of admissible rotations, implausible or potentially misleading responses. Our approach aims to avoid that pitfall by narrowing down the set of admitted rotations. To discriminate between different responses we not only restrict them to those which satisfy the sign restrictions, but additionally are informative regarding the contribution of, for example, shock j to variable i to the FEV. The algorithm proceeds in the following way:

Algorithm: sign restrictions + FEV

---

1. Obtain the reduced-form estimates of $y_t = A(L)u_t$ and compute the Cholesky factor $\tilde{B_0} = chol(\Sigma_u)$.

2. Following Rubio-Ramirez et al. (2010), obtain M draws of an orthonormal rotation matrix Q and estimate $\bar{B_0} = \tilde{B_0}Q$, where $\bar{B_0}$ satisfies $\bar{B_0}\bar{B_0}' = \Sigma_u$.

3. Evaluate whether the responses associated with the draw of $\bar{B_0}$ satisfy the S_{hj} signs imposed.

4. Examine the FEVD corresponding to the draw of $\bar{B_0}$ and check if the shock j is the most informative in terms of its contribution to the FEV for variable i, formally: $\Omega_{i,j}(h) >$

$\Omega_{i,s}(h),\ \forall s \neq j$.

5. If $\bar{B}_0$ satisfies both restrictions then retain this draw, otherwise discard it.

6. Repeat steps 2-5 M times.

Notice that if we exclude Step 4, the model reduces to a purely sign-identified model.

3 Monte Carlo simulation

3.1 Bivariate case

The Monte Carlo simulation in this section compares the performance of different identification strategies to recover the true Data-Generating Process. The DGP comes from the following structural model:

$$B_0 y_t = \sum_{i=1}^{p} B_i y_{t-1} + \epsilon_t,\ \epsilon_t \sim N(0, I_K)\ . \tag{4.6}$$

We set the lag order to $p = 1$ and generate 100 observations, repeating this process 500 times from the following DGP:

$$\begin{pmatrix} 1.01 & -2.12 \\ 2.19 & 1.05 \end{pmatrix} \begin{pmatrix} y_{1,t} \\ y_{2,t} \end{pmatrix} = \begin{pmatrix} -1.01 & -2.02 \\ 1.52 & -0.3 \end{pmatrix} \begin{pmatrix} y_{1,t-1} \\ y_{2,t-1} \end{pmatrix} + \begin{pmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \end{pmatrix}. \tag{4.7}$$

Table 4.1 describes the identification methods here considered to identify the structural shocks from their reduced-form counterpart. We study four identification schemes: recursive identification, pure sign restrictions, sign restrictions + FEV and Max-Share identification. We assume the correct signs of the responses for the sign-identified and the sign restrictions

+ FEV model. We do not explore the misspecification case in sign-restrictions that would arise if imposing the incorrect signs on the contemporaneous response.

Table 4.1: Identification schemes

Identification I: Cholesky

	Shock 1	Shock 2
Variable 1	1	0
Variable 2	1	1

Identification II: Pure Sign Restriction

	Shock 1	Shock 2
Variable 1	**	+
Variable 2	**	+

Identification III: Sign Restriction + FEV

	Shock 1	Shock 2
Variable 1	**	+
Variable 2	**	+
FEV		$\Omega_{1,2}(h) > \Omega_{1,1}(h)$

Identification IV: Max-Share

	Shock 1	Shock 2
Variable 1	**	**
Variable 2	**	**
Max-FEV	$\max_\gamma \sum_{h=0}^{H} \Omega_{1,2}(h)$ *s.t.* $\gamma\gamma' = 1$ *and* $\gamma(1,1) = 0$	

We plot the median responses to a shock to the second variable in Figure 4.1. The solid black line represents the true responses from the DGP while the solid lines with markers depict the responses of the different identification schemes considered. All models, with

the exception of the Cholesky model, correctly recover the true response of variable 1 to an exogenous shock in variable 2. The recursive structure in the Cholesky identification implies a negative response of variable 1 for the first 3 periods. In addition, the Cholesky identification exhibits a response for the second variable more than twice as high than the true response. Sign restrictions present a considerably higher response than the true IRF but only on impact. On the other hand, the median estimate of the Max-Share identification and sign restrictions + FEV depict a response fairly close to the true IRF.

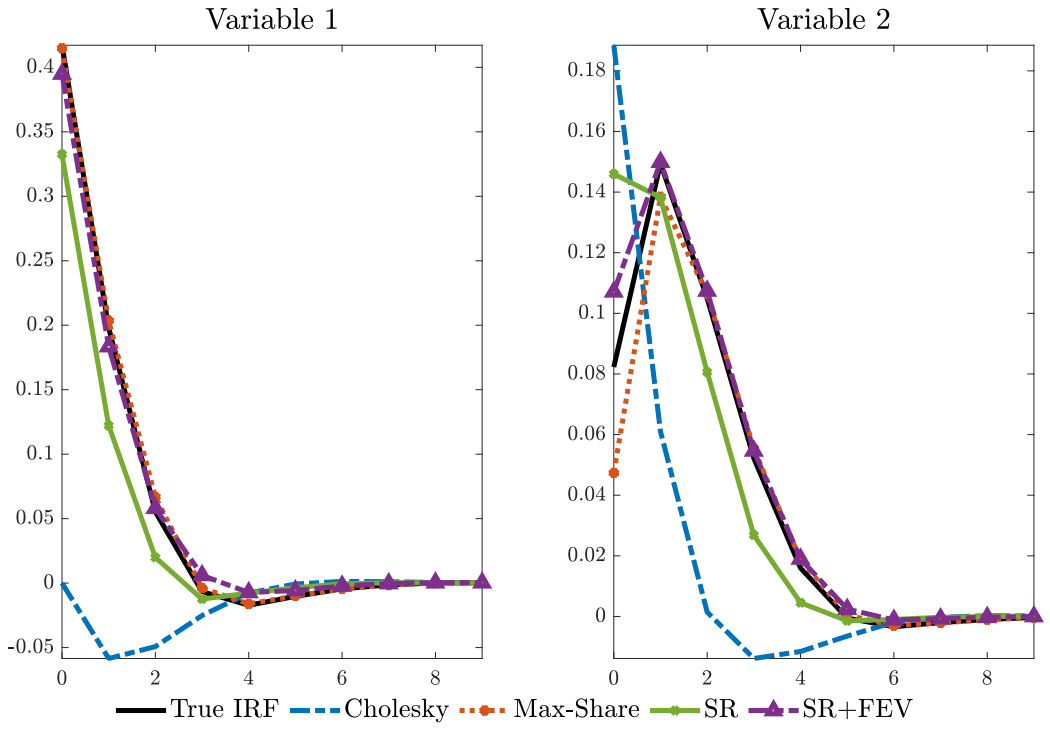

Figure 4.1: Median impulse responses to a shock to ϵ_2.

A more detailed analysis of the response of variable 2 to an exogenous shock to itself is shown in Figure 4.2[1]. When considering the 68% error bands we can then observe that, although with higher uncertainty, the true response is contained within the error bands for the sign-restricted model. In contrast, the credible interval of the Max-Share identification scheme excludes the true response on impact. Narrowing down the set of admissible rotations

[1]The response of variable 1 with its corresponding error bands is plotted in Appendix D.1.

by including an additional restriction from the FEV not only brings the median response closer to the true IRF, but also avoids excluding the true DGP value.

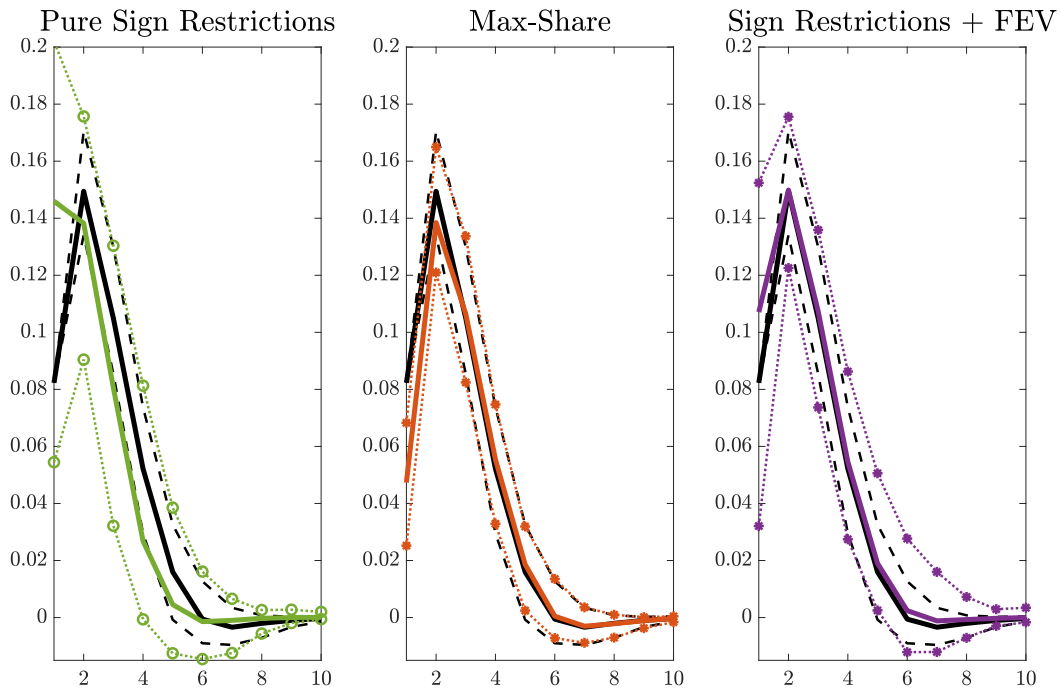

Figure 4.2: Response of variable 2 to a shock in ϵ_2. Solid lines represent the median estimates, dashed lines correspond to the 68% error bands. Solid black line represents the true DGP response.

3.2 Three-variable case

In this subsection we now consider a Monte Carlo simulation for the 3-variable case. Following the previous bi-variate case, we consider a VAR(1) as the true DGP. We obtain the impact matrix B_0 and the matrix of structural parameters B_1 randomly from a normal distribution with the only requirement that it satisfies the stability condition for a VAR. Thus, we generate

M = 500 from the following DGP:

$$\begin{pmatrix} 1.09 & -1.66 & -0.61 \\ 0.96 & -0.64 & -0.57 \\ -0.57 & 0.32 & 0.034 \end{pmatrix} \begin{pmatrix} y_{1,t} \\ y_{2,t} \\ y_{3,t} \end{pmatrix} = \begin{pmatrix} -0.705 & -0.080 & 0.076 \\ 0.407 & -1.318 & -0.884 \\ 0.066 & 0.679 & 0.41 \end{pmatrix} \begin{pmatrix} y_{1,t-1} \\ y_{2,t-1} \\ y_{3,t-1} \end{pmatrix} + \begin{pmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \\ \epsilon_{3,t} \end{pmatrix}. \quad (4.8)$$

Median responses of the four identification strategies considered here are depicted in Figure 4.3. As in the previous section, we consider four identification schemes along the true DGP.[2] Our results suggest that pure sign restrictions and sign restrictions with the additional restriction from the FEV successfully recover the true response after an exogenous shock in the third variable. It is worth noting that for the 3-variable case, Cholesky identification and the Max-Share approach deliver similar responses. For the response of the first two variables, Cholesky identification is bounded to be zero on impact by construction (recursive assumption), however, the responses under the Max-Share identification present a similar behavior without being constrained to zero on impact. A plausible explanation for this behavior is the $\gamma(1,1) = 0$ assumption of the Max-Share approach. In this sense, this assumption seems to have a significant effect on the responses of the rest of the variables, mirroring to some extent the dynamics of the Cholesky identification.

[2]The identification strategies are described in Table 4.5 of the appendix.

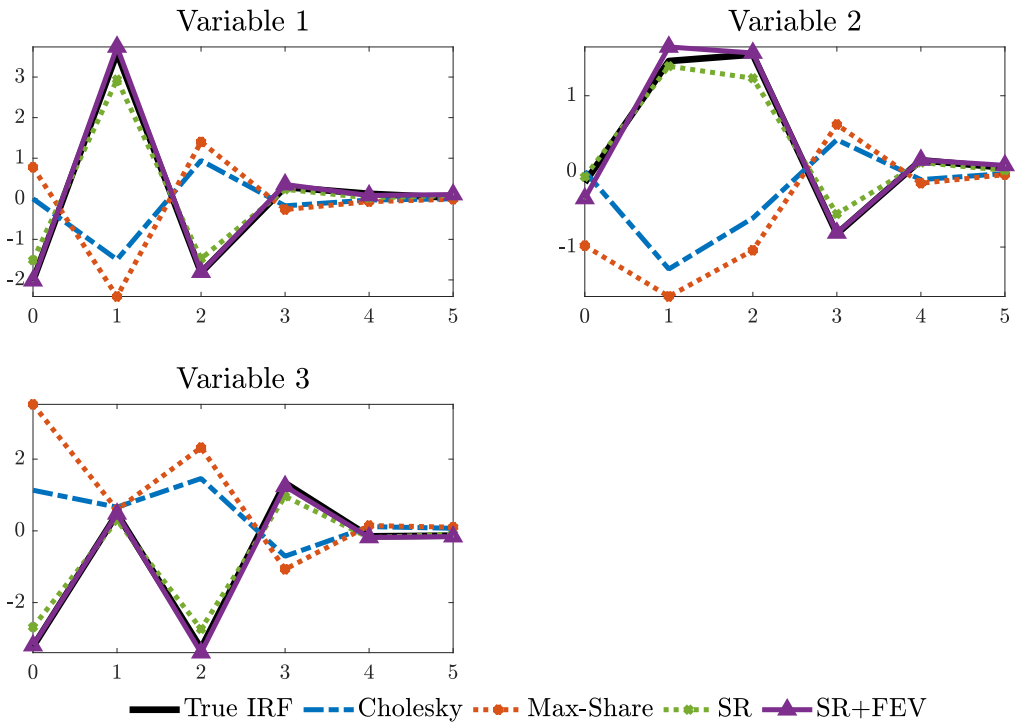

Figure 4.3: Median impulse responses to a shock to ϵ_3.

4 Empirical applications

We illustrate the practicality of our methodology with two empirical applications on the effects of a contractionary monetary policy shock. The macroeconomic implications of a monetary policy shock (MP shock) have been largely investigated using SVARs with several identification methods. Early contributions implied a recursive structure on the shock propagation, for example Christiano et al. (1999). Nonetheless, to avoid the risk of placing questionable timing restrictions on the contemporaneous responses, sign restrictions have become a popular alternative strategy to map reduced-form shocks with their structural counterparts.

4.1 Monetary policy shock: U.S. application

We start this application by showing the results of a purely sign-identified model for the United States. Our quarterly data span from 1982Q1 to 2019Q4. The set of endogenous variables contains the one-year ahead inflation expectations measured by the survey of professional forecasters, the growth rate of real gross domestic product (GDP), the consumer price index (CPI) inflation rate and the Federal Funds rate. We set the VAR order to 4 following results from the Hannan-Quinn (HQ) and the Bayesian (BIC) information criteria. The sign restrictions on the impact responses are shown on the first row of Table 4.2. Following related studies, we restrict the sign of output to depict a negative response to a monetary policy tightening. Furthermore, we do not want to rule out a price puzzle by imposing a sign on the response of prices to a monetary policy shock. However, we aim to capture the effects of an unanticipated change in the policy rate. Therefore, we impose a negative response of inflation expectations to a tightening in the monetary policy in an attempt to avoid capturing information shocks[3].

Table 4.2: Identification restrictions

MP Shock	Inflation Expectations	GDP	Inflation	FED Funds
Sign Restriction	-	-	**	+
FEV Restriction	**	**	$\Omega_{i,j}(h) > \Omega_{i,s}(h),\ \forall s \neq j$	**

Note: ** indicates the unrestricted responses

The model estimation is carried out using Bayesian methods with uninformative priors. Details on the estimation and inference are provided in section D of the appendix. The effects of a contractionary monetary policy shock under a sign-identified model are depicted in Figure 4.1. In the first place we can observe that an unexpected MP shock leads to a drop in inflation expectations by -0.067%. The negative response is certainly restricted by the signs

[3]Campbell et al. (2012) highlight that monetary policy announcements convey two types of shocks: one related to future commitment of monetary policy actions (Odyssean) and another shock which updates agents' information about future economic outlook (Delphic).

imposed in Table 4.2, nonetheless, the expectations response is characterized by exhibiting a considerable level of uncertainty. GDP declines for the first two quarters, showing an impact response of -0.27%. As indicated previously, we do not rule out the presence of a price puzzle by imposing a sign on the impact response of prices to a MP shock, proceeding in that way leads to an uninformative result of the response of prices.

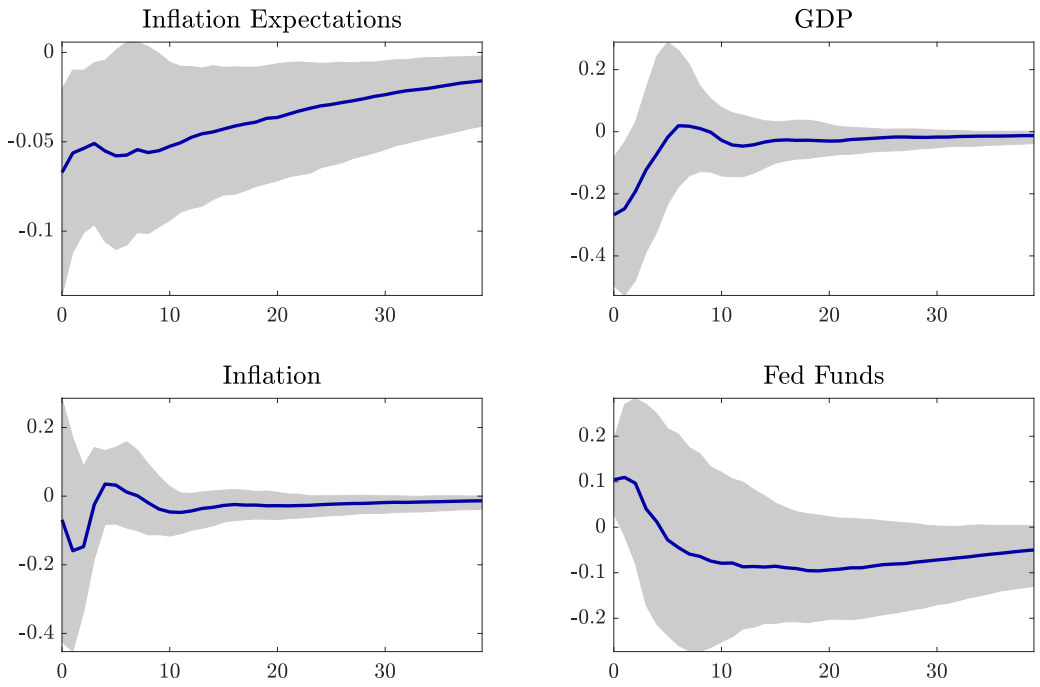

Figure 4.1: Sign-identified model. Impulse responses to a monetary policy shock. Solid blue line depicts the median responses. Shaded areas show the corresponding 68% probability intervals.

We have a thorough look into the different responses of prices to an MP shock in Figure 4.2. Randomly, we select 5 responses associated to different draws, noticing that those response which are associated to a positive response of prices (price puzzle) are also connected to a low share of the FEV explained by the policy shock. On the other side, negative responses of inflation to a contractionary monetary policy are linked to an informative MP shock in terms of its contribution to the FEV.

Figure D.2 of the Appendix shows that the biggest concentration of retained draws associated with uninformative monetary policy shocks for inflation, where the biggest mass of the smoothed kernel densities for the posterior distribution of FEV is concentrated on lower bound with a share of the MP shock explaining less than 10% after a year ($h = 4$).

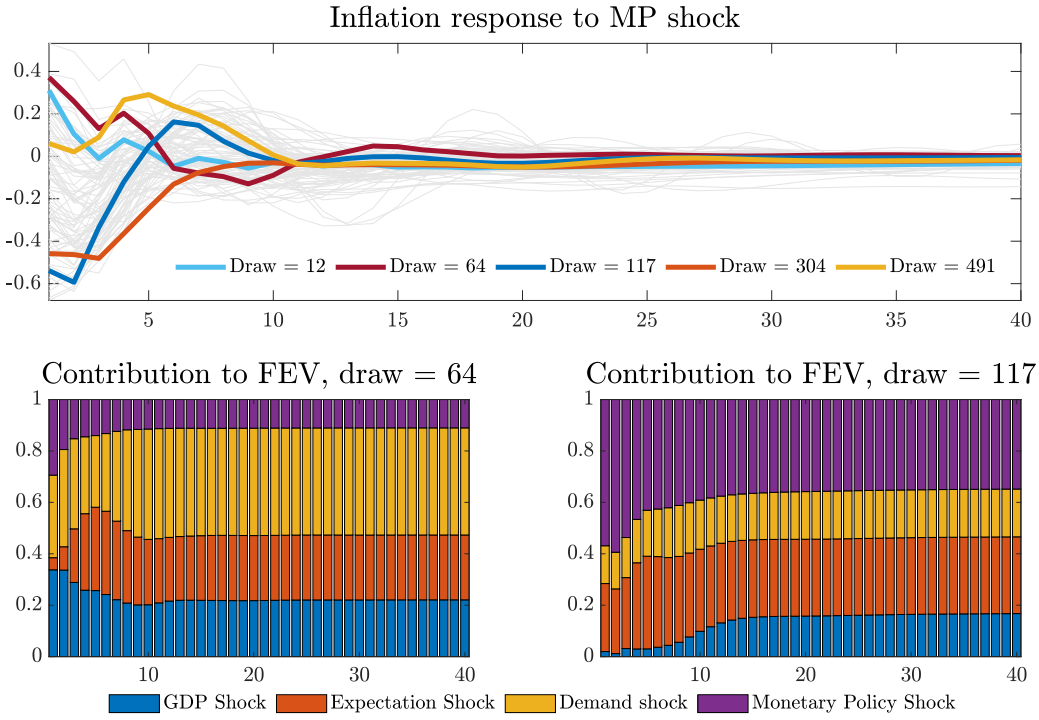

Figure 4.2: Top row: median response of the retained draws of inflation response to a monetary policy shock. Bottom row: contribution to the FEV for selected draws.

Sign restrictions + FEV

As we are interested in exploring the effects of MP shocks on macroeconomic variables, among them inflation, we consider undesirable to retain draws for which the contribution of the MP shock to the FEV of inflation is low. In this regard, we now proceed to include an additional restriction on top of the already imposed signs on the impact responses. In this sense, we add the restriction that the monetary policy shock has to be the most informative shock for inflation in terms of the FEV. This restriction is reflected on the second row of Table 4.2. We

plot the responses to a MP shock under the sign restrictions + FEV identification strategy in Figure 4.3. In contrast with the sign-identified model, a sharper identification now delivers informative results with regards to the response of prices to a monetary policy tightening. Inflation now drops by -0.53% on impact and remains negative for four quarters, thereafter, the effect of an unanticipated increase in the policy rate dies out. Inflation expectations now show a slightly stronger decrease, dropping to -0.08% after the first quarter. However, the expectations response is now less persistent and only showing a negative effect for two quarters. The weaker response on impact of the policy rate shapes a weaker response of output to a contractionary monetary policy.

Figure D.2 from the Appendix shows that this additional restriction induces a more uniform distribution of the share of the FEV explained by a monetary policy shock, where in contrast with the sign-identified model, the contribution is now centered between 50%-60% and, thus, avoiding accepting a majority of the rotations which are associated with a low share of the FEV explained by the MP shock.

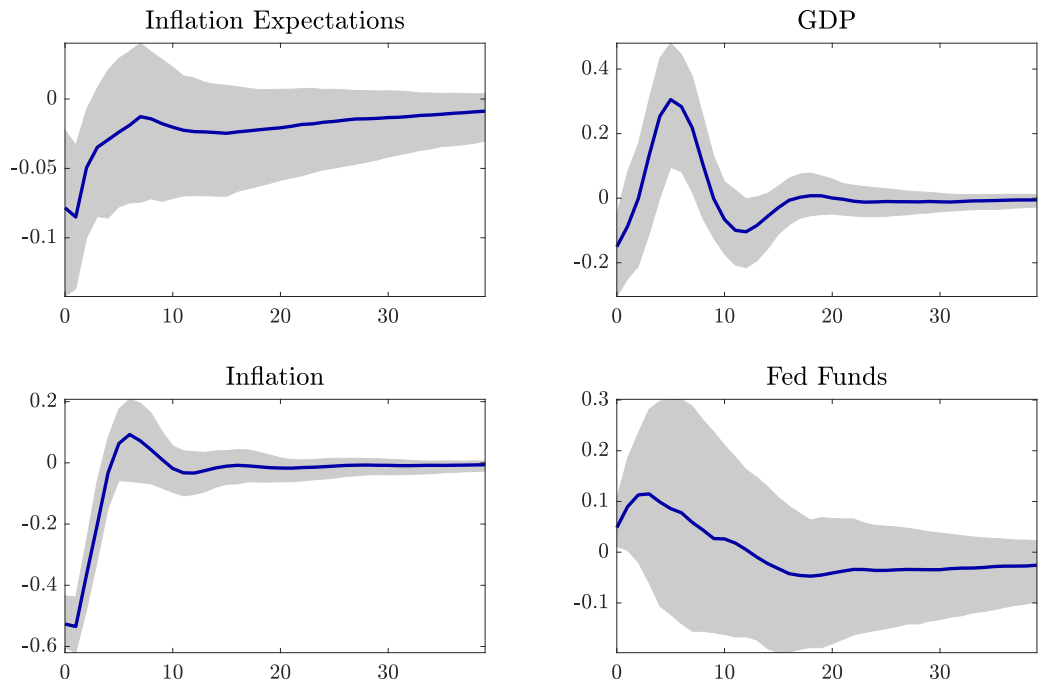

Figure 4.3: Impulse responses to a monetary policy shock. Solid blue line depicts the median responses to a monetary policy shock. Shaded areas show the corresponding 68% probability intervals.

Sign restrictions + FEV: alternative specification

One advantage of using the FEV as a sharper identification mechanism is that it allows to be applied to several variables when the focus of interest is on the response of more than one variable. We illustrate this example by considering the case in which we aim to understand the effects of a monetary policy tightening on inflation and inflation expectations. As mentioned before, it is undesirable to allow rotations which present a low share of the FEV explained by the MP shock. Table 4.3 exemplifies how this additional restriction is imposed on the second row.

We restrict the set of structural shocks to fulfill the sign restrictions and for the monetary policy shock to be informative (in terms of the FEV) to inflation expectations and inflation. Similarly as in the case of the share of the FEV explained by the MP on inflation under the

sign restrictions model, the MP shock also shows that a high share of the retained draws are associated with a low contribution of the MP shock to the inflation expectations FEV. Therefore, we then proceed to incorporate the restriction that the MP shock must be the most informative shock. Figure C.3 of the appendix illustrates how this additional restriction is reflected in a more uniform distribution of the smoothed kernel density of the posterior distribution of the MP shock share of the FEV on inflation expectations. The sign-restricted model groups most of this share on the lower bound whereas the sign-restricted + FEV model allocates the contribution around 50%. The set of restrictions are described in Table 4.3.

Table 4.3: Alternative identification: sign restrictions + FEV

MP Shock	Inflation Expectations	GDP	Inflation	FED Funds
Sign Restriction	-	+	**	+
FEV Restriction	$\Omega_{1,4}(h) > \Omega_{1,s}(h),\ \forall s \neq 4$	**	$\Omega_{3,4}(h) > \Omega_{3,s}(h),\ \forall s \neq 4$	**

We plot the responses after a monetary policy tightening under the alternative specification of sign restriction + FEV in Figure 4.4. Including the restriction that the MP shock is the most informative in terms of the FEV for inflation expectations and inflation now provide a sharper response in both variables while keeping the rest substantially unaffected. We now observe that a MP shock triggers a stronger drop in inflation expectations. Moreover, the effects of a contractionary monetary policy are longer lasting for inflation expectations that under the purely sign-identified model. Interestingly, the inclusion of the FEV restriction on inflation expectations do not alter the responses of inflation, as we observe they respond in the same magnitude as in Figure 4.3.

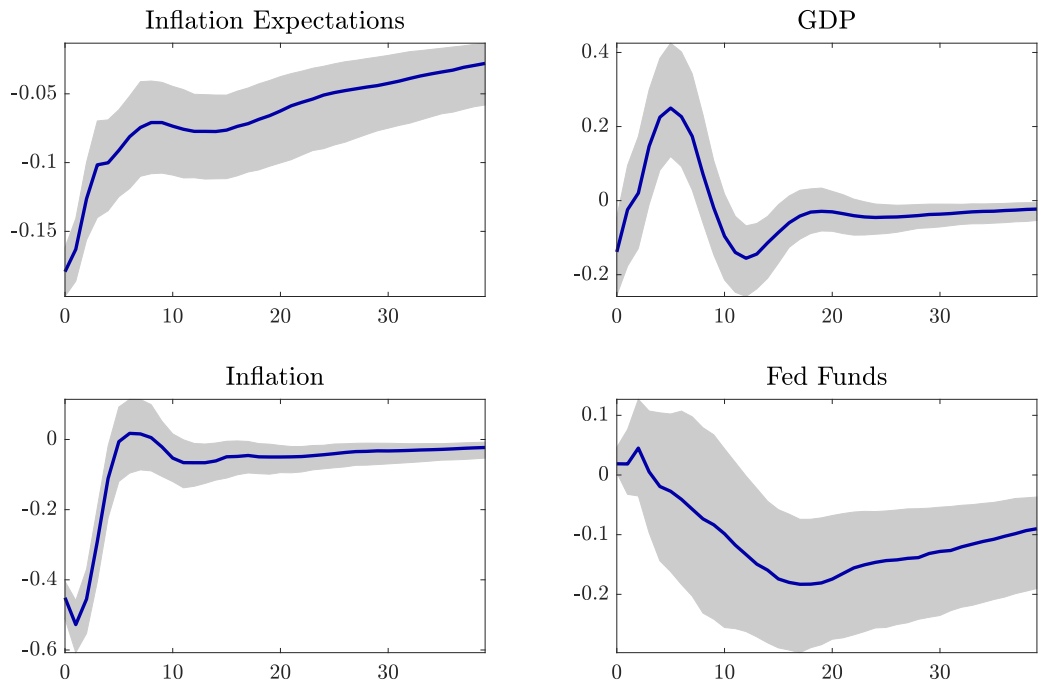

Figure 4.4: Impulse responses to a monetary policy shock. Solid blue line depicts the median responses. Shaded areas show the corresponding 68% probability intervals.

4.2 Monetary policy shock: U.K. application

This section is concerned with the effects of a contractionary monetary shock for the U.K. economy. We estimate a monthly VAR for the period 1992m01 to 2019m12. The set of endogenous variables contain the one-year Gilt rate, CPI inflation rate, industrial production, effective exchange rate, corporate bond spread and the mortgage spread. The one-year Gilt rate is used as a risk-free alternative to the Bank Rate to account for the lack of variability in the latter since 2009. We calculate the corporate bond spread as difference between the Moody's BAA corporate bond index and the U.S. 10-year Government bonds. Similarly, the mortgage spread is calculated using the difference between the 30-year fixed rate mortgage average in the U.S. and the U.S. long-term Government bonds.

In our baseline specification, we se the lag order following results from Information Criteria

and select 3 lags as the VAR order. For this empirical application, we illustrate the flexibility of the methodology proposed by incorporating the FEV restrictions in form of a threshold. In this regard, we set the restriction that the monetary policy shock should account for at least 5% of CPI inflation's FEV for the first 3 months ($h = 3$). This restriction, in addition to the sign responses, are shown in Table 4.4.

Table 4.4: Identification restriction: monetary policy shock

MP Shock	One-year Gilt	CPI	IP	EER	Corporate BS	Mortage Spread
Sign Res.	+	-	**	-	-	**
FEV Res.	**	$\Omega_{2,1}(h) > 0.05$	**	**	**	**

Note: ** indicates the unrestricted responses.

Our starting point in capturing the effects of a tightening in monetary policy for the U.K. economy is the classical sign restriction. When analyzing the set of retained draws one can observe that majority of the accepted rotations have very little contribution of the policy shock to inflation's FEV. As shown in Figure D.4 of the Appendix, the biggest share of the mass of the posterior density of the FEV contribution of the policy shock on inflation, is concentrated close to the zero value. This motivates the inclusion of an additional restriction on the share a contractionary monetary policy shock contributes to inflation's FEV.

We compare the results from the proposed identification strategy with the classical sign restrictions model in Figure 4.5. One can observe that the additional restriction on the FEV yields more informative results in terms of the magnitude and persistence of the response of CPI inflation to a tightening in monetary policy in comparison with the purely sign-identified model. Under classical sign restrictions, one would conclude that effect of the initial shock will die out after three months. In contrast, the sign restrictions + FEV model shows that a contractionary monetary policy shock is more persistent, having an effect on prices up to a year. Moreover, the magnitude of the response sits on the lower bound of the possible responses under sign restrictions, concluding a stronger response of inflation under

the proposed identification mechanism compared to classical sign restrictions.

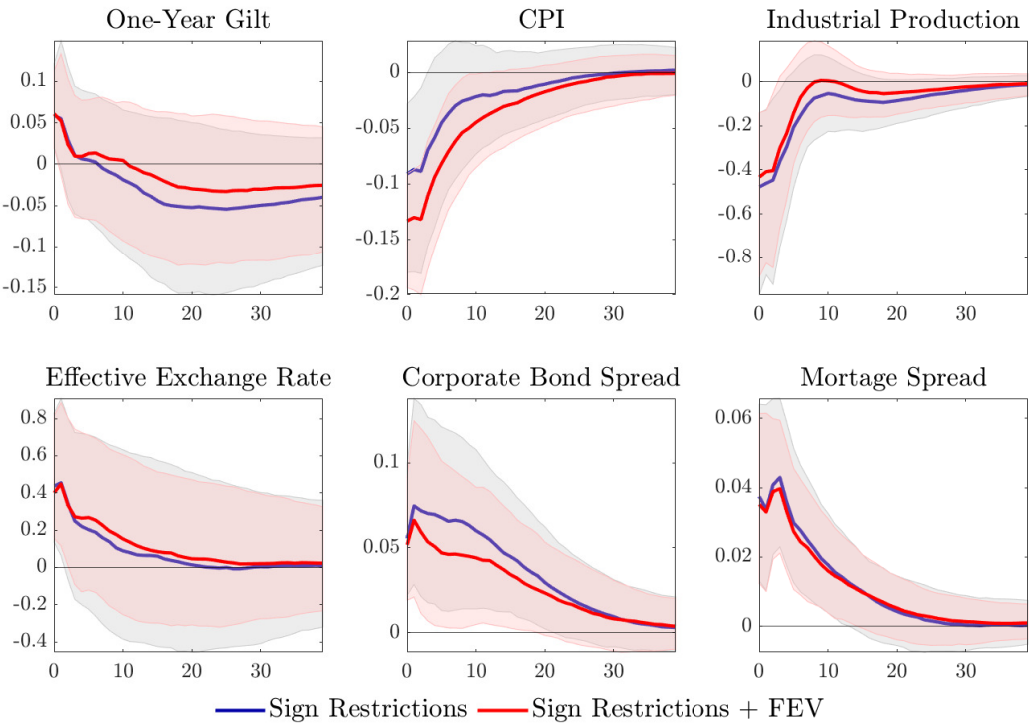

Figure 4.5: Impulse responses to a monetary policy shock. Solid blue line depicts the median responses. Shaded areas show the corresponding 68% probability intervals.

5 Conclusions

In this article we propose an alternative identification strategy to overcome frequent limitations of purely sign-identified models. More specifically, we intend to overcome the conjecture that sign restrictions alone lead to accepting different responses across the retained draws which ultimately lead to uninformative conclusions about the economic responses. This issue can be alleviated by incorporating an additional restriction coming from the forecast error of the shock of interest on the target variable(s). The usefulness of this additional restriction, in comparison with other commonly applied identification methods is first introduced in a Monte Carlo simulation where the gains of pursuing the alternative identification strategy

here proposed, are rather noticeable. The Monte Carlo study shows that the proposed identification method successfully recovers the true response for the two-variable and three-variable case. Furthermore, it demonstrates that the sharper identification achieved by the additional restriction does not induce a credible interval that excludes the true DGP response.

We illustrate the usefulness of this method with two empirical applications for the case of a tightening in monetary policy in the U.S. and U.K. economy. For the former, we observe that a sign-identified model, without imposing a sign on the impact response of prices, leads to an uninformative response of prices due to the inclusion, in the set of retained draws, of responses which are associated with a low contribution of the MP shock to the prices FEV. Restricting the model to retain only draws in which the MP shock is the most informative in terms of the FEV for inflation results in a sharper response of prices. Moreover, the set of retain draws now show a unique type of response in contrast with the purely sign-identified model.

Finally, the application of a contractionary monetary policy shock for the U.K. economy shows how this identification strategy can provide a sharper identification and therefore, more informative responses in larger models. We observe how including an additional restriction of this shock on the contribution to inflation FEV reduces the uncertainty over the response of inflation and delivers a sharper response. Thus, the identification method here proposed yields more informative results in terms of the peak effect and persistence of the response.

D Appendix

Estimation and Inference

Let's consider the reduced-form VAR in (4.2):

$$y_t = c + \sum_{i=1}^{p} A_i y_{t-i} + u_t,\ u_t \sim N(0, \Sigma_u)$$

We obtain the VAR coefficients using Gibbs sampling to draw from its posterior density. Following Koop and Korobilis (2009), we re-write equation (4.2) in the following way:

$$y = (I_K \otimes X)\alpha + U, \tag{4.9}$$

with $y = vec(y_t), X = vec(c_i, y_{i,t-1}, y_{i,t-2}, ..., y_{i,t-p}), \alpha = vec(c, A_1, A_2, ..., A_p)$ and $U = vec(u_t)$. The priors for the VAR coefficients are assumed to follow a normal distribution:

$$\alpha \mid \Sigma, y \sim N(\underline{\alpha}, \Sigma \otimes \underline{V}), \tag{4.10}$$

where $\underline{\alpha}$ denotes the prior mean and $\underline{V}$ states the variance of prior of the coefficients. We set a conjugate prior for the covariance matrix Σ following an inverse Wishart distribution with $\underline{S}$ prior scale matrix and $\underline{v}$ degrees of freedom:

$$\Sigma \mid y \sim IW(\underline{S}, \underline{v}) \tag{4.11}$$

Hence, the posterior density of the VAR coefficients and the covariance matrix have the following form:

$$\alpha \mid \Sigma, y \sim N(\overline{\alpha}, \Sigma \otimes \overline{V}) \tag{4.12}$$

and

$$\Sigma \mid y \sim IW(\overline{S}, \overline{v}), \tag{4.13}$$

where

$$\overline{A} = \overline{V}[\underline{V}^{-1}\underline{A} + X'X\hat{A}],$$
$$\overline{V} = [\underline{V}^{-1} + X'X]^{-1}$$

and

$$\overline{S} = S + \underline{S} + \hat{A}'X'X\hat{A} + \underline{A}'\underline{V}^{-1}\underline{A} - \overline{A}'(\underline{V}^{-1} + X'X)\overline{A},$$

with $\overline{a} = vec(\overline{A})$ and $\hat{A} = (X'X)^{-1}X'Y$ denoting the OLS coefficient estimator. Hence, the Gibbs algorithm can be summarized as follows:

1. Set the priors for the VAR coefficients and covariance matrix Σ. As mentioned above, we set the prior for the coefficients to follow a normal distribution with mean $\underline{a}$ and variance $\Sigma \otimes \underline{V}$. The prior for the covariance matrix follow an inverse Wishart distribution with prior hyperparameters: $\overline{S}$ and $\overline{v}$.

2. Sample the VAR coefficients α from its conditional posterior density: $\alpha|\Sigma, y \sim N(\overline{\alpha}, \Sigma \otimes \overline{V})$.

3. Sample the covariance matrix Σ from its conditional posterior distribution: $\Sigma|y \sim IW(\overline{S}, \overline{v})$.

4. Repeat steps 2-3 M times.

Figures

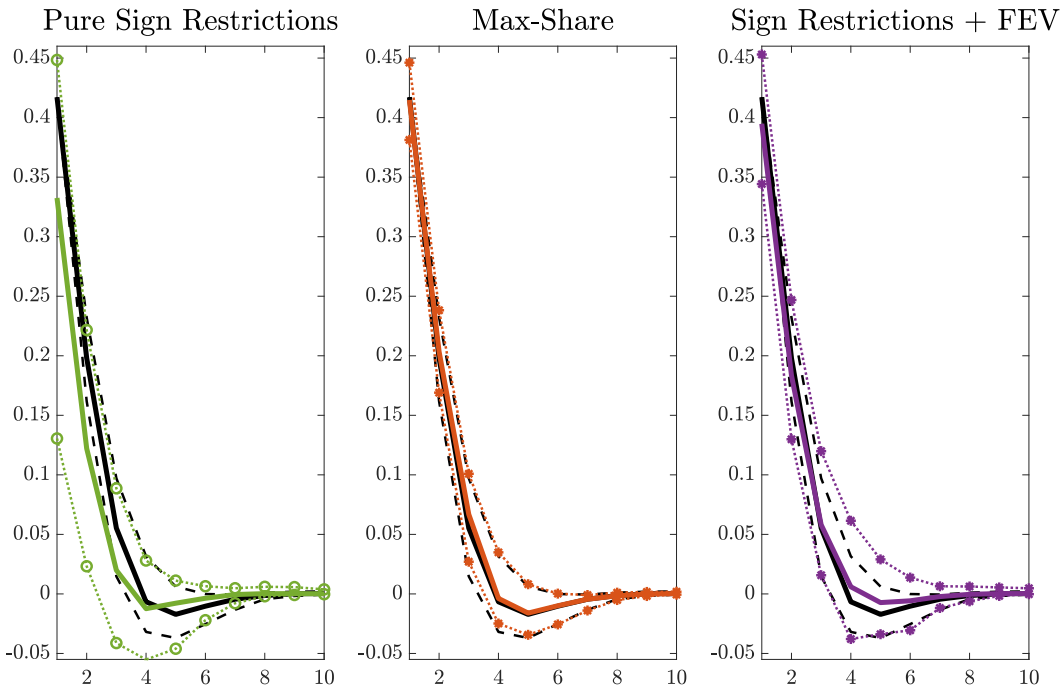

Figure D.1: Response of variable 1 to a shock in ϵ_2. Solid lines represent the median estimates, dashed lines correspond to the 68% error bands

Table 4.5: Identification schemes

Identification I: Cholesky

	Shock 1	Shock 2	Shock 3
Variable 1	1	0	0
Variable 2	1	1	0
Variable 3	1	1	1

Identification II: Pure SR

	Shock 1	Shock 2	Shock 3
Variable 1	**	+	**
Variable 2	**	+	**
Variable 3			

Identification III: SR + FEV

	Shock 1	Shock 2	Shock 3
Variable 1	**	+	**
Variable 2	**	+	**
Variable 3	**	+	**
FEV	$\Omega_{1,3}(h) > \Omega_{1,1}(h), \Omega_{1,2}(h)$		

Identification IV: Max FEV

	Shock 1	Shock 2	Shock 3
Variable 1	**	**	**
Variable 2	**	**	**
Variable 3	**	**	**
Max-FEV	$\max_\gamma \sum_{h=0}^{H} \Omega_{1,2}(h)$ *s.t.* $\gamma\gamma' = 1$ *and* $\gamma(1,1) = 0$		

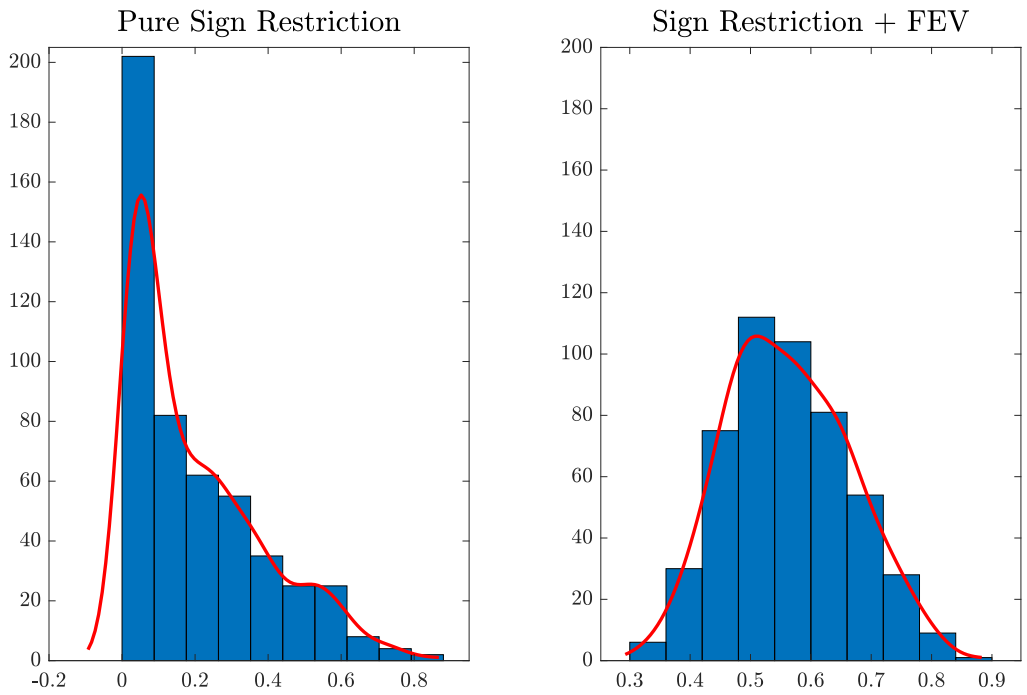

Figure D.2: Posterior density of the FEV contribution of a MP shock on inflation at horizon: $h = 6$.

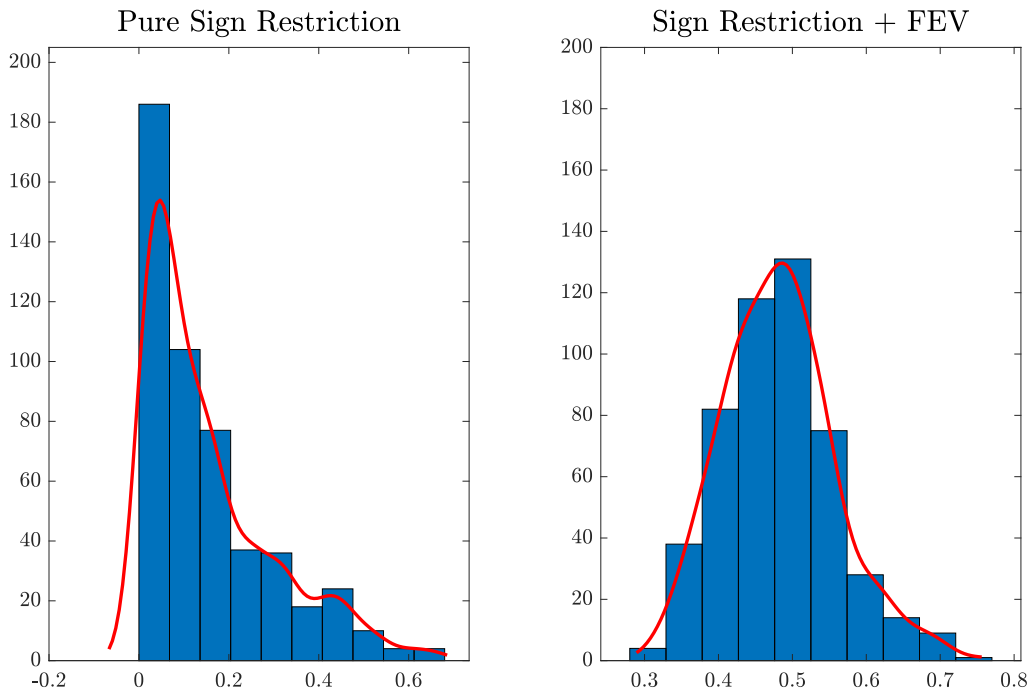

Figure D.3: Posterior density of the FEV contribution of a MP shock on inflation expectations at horizon: $h = 6$.

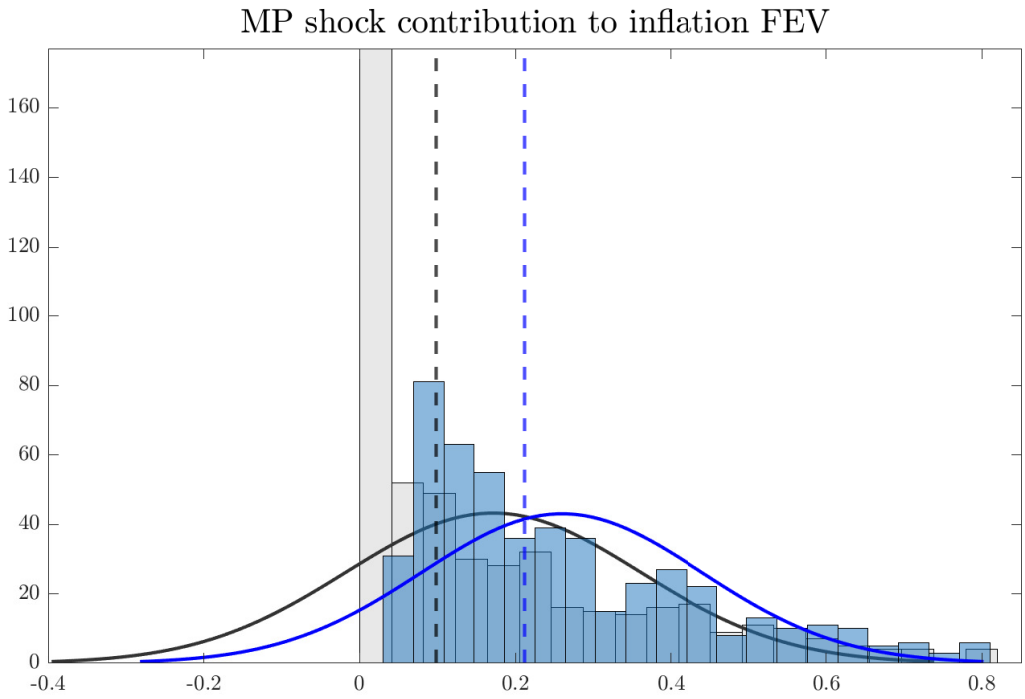

Figure D.4: Posterior density of the MP shock contribution to inflation FEV at $h = 3$. Grey bars represent the model under sign restriction. Blue bars correspond to the model with additional restrictions to the FEV. Dashed lines are the corresponding median estimates

Chapter 5

Reassessing the effects of oil supply shocks: an alternative identification strategy

Abstract

In this article we propose an alternative identification strategy to evaluate the effects of a shortfall in global oil supply. Our baseline model combines information from the forecast error variance (FEV) with classical sign restrictions. We compare the results obtained from the baseline model with some of the benchmark models in the oil market literature. Overall, we observe that differences in the magnitude of the response of oil prices to a drop in oil supply can be explained by variations in the contributions of this shock to oil prices FEV. Our results stand in the upper bound of how responsive oil prices are to oil supply shocks.

Keywords Oil market, SVAR model, identification problem, sign restrictions, FEV.

1 Introduction

The Russian invasion of Ukraine in February 2022 has brought renewed attention to assessing the potential impact of an unexpected reduction in global oil supply, highlighting the importance of understanding the consequences of a sudden disruption in global production. Accounting for approximately 10% of global oil production, an abrupt cut in Russia's oil supply raises concerns regarding the magnitude of the impact on real oil prices and global economic activity. Figure 5.1 shows the behavior of the West Texas Intermediate (WTI) price of oil in real terms in comparison with the evolution of oil supply. The magnitude of the cutback in Russian oil production can be observed in panel (b), where there is an evident break-point in the amount of crude barrels produced after the first quarter of 2022.

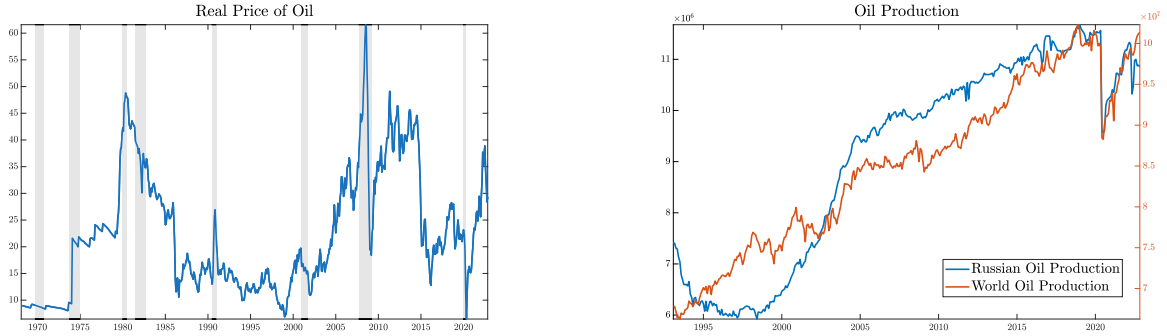

(a) (b)

Figure 5.1: Panel (a) depicts the evolution of the WTI real price of oil. Panel (b) shows the World and Russia's total number of oil barrels produced. Shaded areas correspond to the National Bureau of Economic Research (NBER) recession periods.

In the analysis of the macroeconomic consequences of fluctuations in the oil market, vector autoregressive models (VARs) have prevailed as a commonly used technique. Starting with models that exploited short run (exclusion) restrictions (Kilian, 2008, 2009), the literature has progressively moved to employing restrictions on the signs of the contemporaneous responses (Baumeister and Peersman, 2013; Lippi and Nobili, 2012). In this regard, sign restrictions stand as an appealing alternative to overcome limitations present under recursive

identification. For example, using exclusion restrictions in the Kilian (2009) model implies an expansionary effect after a negative oil-specific demand shock. Thus, it appears intuitive to exclude positive responses of economic activity after an increase in oil prices. However, recent literature has emphasized the risks of accepting potentially unlikely responses in the set of retained draws in classical sign restrictions models (Kilian and Murphy, 2012, 2014; Antolín-Díaz and Rubio-Ramírez, 2018).

We propose an alternative identification strategy that combines information from the forecast error variance (FEV) with classical sign restrictions. To ease comparability, our starting point is the classical 3-variable model from Kilian (2009). When comparing the results of different identification strategies we therefore observe that differences in the real oil price responses are importantly influenced by the underlying contribution the oil supply shock has on real oil prices FEV. Instead of imposing a-priori restrictions on the price elasticity of oil supply (as in Kilian and Murphy (2012)), we take into account estimation uncertainty and incorporate bounds on the contribution of oil supply shocks to real oil prices FEV. We set these bounds following results from Caldara et al. (2019) and observe that abrupt reductions in the oil supply tend to have an important impact on the real prices of oil.

The contribution of this article is twofold. First, we shed some light on the underlying differences in the magnitude of the responses of oil prices to cuts in oil supply. The magnitude of the responses in all models considered ranges from 0.44% - 2.18% with the peak effect usually occurring after 3 to 4 months from the initial shock. Our baseline model exhibits a peak response of 1.62% after 3 months. Second, we propose a different method to identify the shocks that combine conventional sign restrictions with an additional uncontroversial restriction on the contribution to the FEV. In the second part of our results, we consider a medium-scale oil market model as in Lippi and Nobili (2012) and Känzig (2021). We thus exemplify how draws that exhibit a close to zero contribution to oil prices FEV, notably influence the median response of oil prices to a sudden shortfall in oil supply.

The remainder of this article is structured as follows. Section 2 reviews the methodology here considered. Section 3 revisits the effects of a negative oil supply shock and finally, section 4 contain the concluding remarks.

2 Methodology

The core empirical model used among different specifications is a structural vector autoregressive model (SVAR) with p lags given by:

$$B_0 y_t = v + \sum_{i=1}^{p} B_i y_{t-i} + \epsilon_t,\ \epsilon_t \sim N(0, I_K), \tag{5.1}$$

where the vector of endogenous variables $y_t = (y_{1t}, ..., y_{Kt})'$ is of dimension $K \times 1$, v is a $K \times 1$ vector of constants and B_i is an $K \times K$ matrix of structural coefficients with $i = 1, ..., p$. Thus, the researcher is interested in finding a linear mapping between the unorthogonalized reduced-form errors and the structural shocks from equation (5.1) from the identity $u_t = B_0^{-1}\epsilon_t$, where u_t corresponds to the reduced-form errors and B_0 represents the impact matrix which satisfies the condition: $B_0^{-1}(B_0^{-1})' = \Sigma_u$, with Σ_u representing the reduced-form variance-covariance matrix. Therefore, the identification strategy to recover the structural responses from the reduced-form estimates depends on the assumptions placed on the matrix B_0.

In the following subsections we will discuss the identifying assumptions used in the recent oil market literature.

2.1 Recursive identification

Highlighting the importance of separating supply and demand shocks on the real price of oil, Kilian (2009) relies on a recursive identification imposing a zero contemporaneous response from oil production to shocks in economic activity and oil prices. This identification strategy

is obtained by estimating the Cholesky decomposition of variance-covariance matrix Σ_u. This identification method implies the following structure:

$$\begin{bmatrix} u_t^{\text{Oil Prod}} \\ u_t^{\text{Econ Act}} \\ u_t^{\text{Oil Price}} \end{bmatrix} = \begin{bmatrix} b_{11} & 0 & 0 \\ b_{21} & b_{22} & 0 \\ b_{31} & b_{32} & b_{33} \end{bmatrix} \begin{bmatrix} \varepsilon_t^{\text{oil supply shock}} \\ \varepsilon_t^{\text{aggregate demand shock}} \\ \varepsilon_t^{\text{oil-specific demand shock}} \end{bmatrix} = B_0^{-1}\varepsilon_t \tag{5.2}$$

As it is the case for all identification strategies achieved by a recursive structure, the variable ordering is crucial for the responses the system exhibits after an exogenous shock. Kilian (2009) motivates the ordering by assuming a vertical short-run oil supply curve which implies placing the oil supply first, followed by a measurement for economic activity and real oil prices placed last.

2.2 Sign Restrictions

Potential misspecifications derived from imposing incorrect exclusion restrictions can be alleviated by achieving identification based on the sign of the responses. The sign of the responses are commonly motivated by economic theory. Following Uhlig (2005), identification by sign restrictions can be achieved by employing an orthonormal matrix Q, such that $QQ' = I_K$. Q. One commonly applied method for generating the rotation matrix Q was introduced by Rubio-Ramirez et al. (2010), the authors show how this matrix can be obtained from a Haar-uniform draw from O_K, treating all draws of Q equally likely[1]. The set of drawn rotations belong to the space of orthonormal matrices $Q \in \Theta(K)$. Thus, the model will only imply set-identification, where the set of sign restrictions of shock j at horizon h can be expressed by S_{hj}.

Following the tri-variate case from Kilian (2009), the set of sign restrictions imposed to

[1]Recently, Baumeister and Hamilton (2019) and Giacomini and Kitagawa (2021) have question the uninformativeness of this uniformly distributed prior on the reduced-form coefficients. Instead, Baumeister and Hamilton (2019) impose explicit Bayesian priors on the elements of the impact matrix B_0.

the oil market model can be summarized by the following table:

Table 5.1: Identification restrictions on the sign responses

	Oil Production	Economic Activity	Real Price of Oil
Negative oil supply shock	-	-	+
Positive aggregate demand shock	+	+	+
Positive oil-specific demand shock	+	-	+

2.3 Oil supply elasticity bounds

Kilian and Murphy (2012) argue that agnostic identification based on the sign responses are insufficient to deduce how real oil prices will react to oil supply disruptions. Instead, they impose an upper bound on the price elasticity of oil supply claiming that the set of admissible models should be narrowed down only to draws that exhibit an oil supply price elasticity inferior to 0.0258. This upper bound on the fraction b_{13}/b_{33} of the elements of the impact matrix B_0, reduces the set of retained rotations from initially 30,860 draws that satisfied the sign restrictions to 80 which satisfied both restrictions.

2.4 Sign Restrictions + FEV bounds

The baseline approach developed in this paper involves utilizing the property that the FEV holds valuable information regarding the extent to which meaningful economic shocks account for fluctuations in the variables of concern. From this perspective, it is intuitive to concentrate on the set of models in which the shock of interest explains a plausible share of the variation of the variable under analysis.

It is often the case that sign-identified VARs exhibit wide error bands leading to uninformative results. As highlighted in Kilian and Murphy (2012), pure sign-identified models allow the risk of accepting, in the set of admissible rotations, implausible or potentially misleading responses. Our approach aims to avoid that pitfall by narrowing down the set of

admitted rotations.

We first define the h-step forecast error of equation (5.1) is as:

$$\begin{aligned} y_{t+h} - y_t(h) &= \sum_{i=0}^{h-1} B_i B_0 \epsilon_{t+h-i} \\ &= \sum_{i=0}^{h-1} \Theta_i w_{t+h-i}, \end{aligned} \tag{5.3}$$

where B_0 corresponds to the impact matrix. Thus, we can express h-step forecast error for the j^{th} element of y_t as:

$$y_{j,t+h} - y_{j,t}(h) = \sum_{k=1}^{K} \left(\theta_{jk,0} w_{k,t+h} + \cdots + \theta_{jk,h-1} w_{k,t+1}\right) \tag{5.4}$$

Therefore, the contribution of shock k to the forecast error variance of variable j is:

$$\theta_{jk,0}^2 + \theta_{jk,1}^2 + \cdots + \theta_{jk,h-1}^2 = \sum_{i=0}^{h-1} \left(e_j' \Theta_i e_k\right)^2. \tag{5.5}$$

Finally, one can obtained the share explained by innovation k of the FEV of variable j, dividing equation (5.5) by the Mean Square Error (MSE) of $y_{j,t}(h)$:

$$\Omega_{i,j}(h) = \frac{\theta_{jk,0}^2 + \theta_{jk,1}^2 + \cdots + \theta_{jk,h-1}^2}{\sum_{k=1}^{K} \left(\theta_{jk,0}^2 + \cdots + \theta_{jk,h-1}^2\right)} \tag{5.6}$$

We do not just limit the retained set to those draws that meet the sign-response criteria but rather reduce the set of retained draws by imposing bounds on the contribution to FEV of oil supply shocks to real oil prices. The logic behind these bounds is to reflect estimation uncertainty about the contribution of oil supply shocks to fluctuations in the real price of oil and thus, avoiding placing exceptionally restrictive assumptions.

We use the estimation bounds from Caldara et al. (2019) model. In their approach, the authors minimize the distance between the estimated cross-country equation elasticity of oil

supply with the estimates derived from a VAR model. Thus, we set the lower bound of the contribution of oil production to real oil price FEV to 24% and an upper bound to 46%.

The algorithm to obtain the baseline estimations proceeds in the following way:

Algorithm: sign restrictions + FEV bounds

---

1. Obtain the reduced-form estimates of $y_t = A(L)u_t$ and compute the Cholesky factor $\tilde{B_0} = chol(\Sigma_u)$.

2. Following Rubio-Ramirez et al. (2010), obtain M draws of an orthonormal rotation matrix Q and estimate $\bar{B_0} = \tilde{B_0}Q$, where $\bar{B_0}$ satisfies $\bar{B_0}\bar{B_0}' = {\Sigma_u}^2$.

3. Evaluate if the responses of the $\bar{B_0}$ draw satisfy the S_{hj} signs imposed and if the associated FEVD is within the bounds: $\Omega_{Lower}(h) < \Omega_{i,j}(h) < \Omega_{Upper}(h)$.

4. If $\bar{B_0}$ satisfies both restrictions then retain this draw, otherwise discard it.

5. Repeat steps 2-4 M times.

---

3 The effects of a negative oil supply shock

To examine the impact of an abrupt disruption in the oil supply, and with the aim of making the baseline results from the proposed identification strategy comparable to the established literature, our data set consists of three commonly used monthly variables $y_t = [OilProd_t, EconAct_t, RPO_t]$ as in Kilian (2009), where $OilProd_t$ is the global supply of oil

[2]Despite recent attention on the use of uninformative Haar priors typically utilized in SVARs identified using sign restrictions (Baumeister and Hamilton, 2019; Giacomini and Kitagawa, 2021), we follow Inoue and Kilian, 2021 who argue that the effect of such prior is negligible in cases where sign restrictions are employed in combination with additional restrictions.

obtained from the U.S. Energy Information Administration (EIA), $EconAct_t$ is the Index of Global Real Economic Activity proposed by Kilian (2009) and RPO_t is the (WTI) Spot Crude Oil Price deflated by U.S. consumer price index (CPI). Data spans from January 1968 until October 2022.

Our starting point in the analysis of the effects of a negative oil supply shock on real oil prices is the recursive identification proposed in Kilian (2009)[3]. As it is well known, a Cholesky identification implies strong assumptions on the contemporaneous responses of the system. In this regard, the ordering of the variables, although motivated by economic theory, will imply strong assumptions for the dynamics of the model. Namely, the model will assume a vertical short-run oil supply. Moreover, this assumption conveys a price elasticity of oil supply near zero.

In this respect, sign restrictions appear as an appealing alternative to exclusion (zero) restrictions as they are not sensitive to the variable order. In Figure 5.1 we compare the results from a sign-identified model based on the signs from 5.1 with respect to the model based on exclusion restrictions.

One can observe that the response of real oil prices after an unexpected shortfall in oil supply is markedly higher under the sign-identified model, by construction the response of oil prices under the recursive identification is zero in the contemporaneous, however, significantly lower and less persistent than under the sign restrictions model. For the former, the initial effect of a negative oil supply shock dies out after 6 months in contrast with a longer lasting effect of up to 16 months for the sign restrictions model. Another important difference between these two models is the direction of the effect an oil-specific demand shock has on economic activity. Under the recursive identification model a shock to oil prices has an expansionary effect on real activity, as opposed to the sign-identified model which excludes this possibility by construction. Finally, it is possible to notice contrasting drivers in the changes

[3]Figure of the Appendix depicts the impulse response functions of the model identified by a recursive structure.

in real oil prices between these models. Under the model based on exclusion restrictions, real oil prices are majorly driven by oil-specific demand shocks in contrast with a more important contribution of the oil supply shock under sign restrictions.

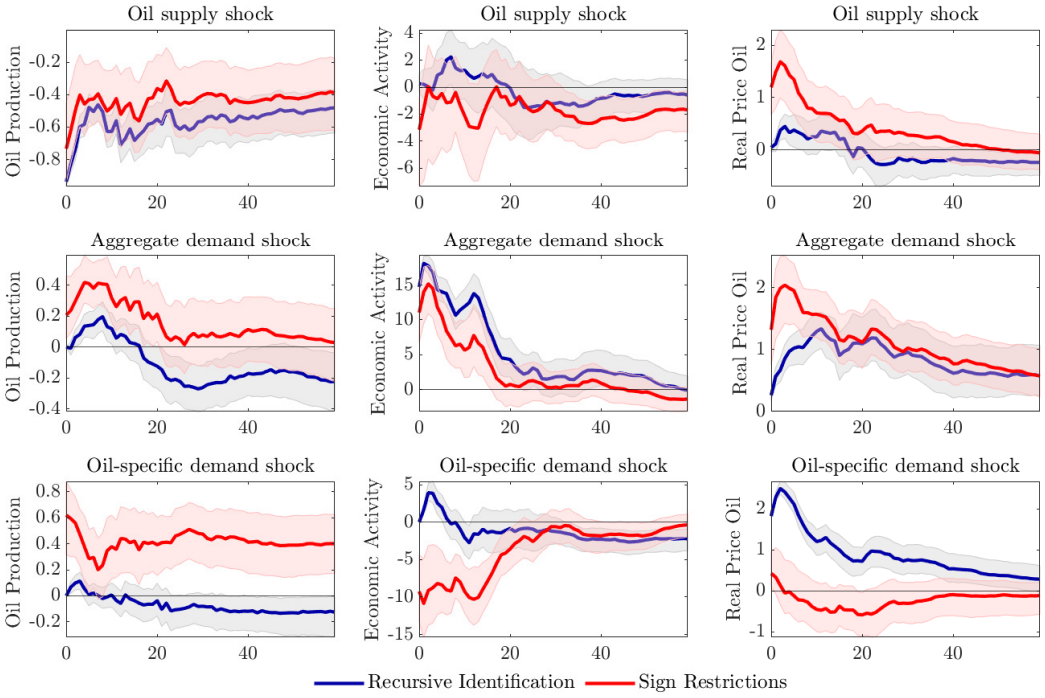

Figure 5.1: Impulse response: recursive identification and sign restrictions. Solid lines are the median responses. Shaded areas represent the 68% error bands.

A prevalent finding in sign identified models is that being too agnostic in the identification strategy can lead to retain draws that exhibit potentially implausible responses. In Figure 5.1 we can observe that the peak effect of oil prices to a cut in oil production spans from 1 percentage points to 2.29. When looking in more detail to the set of retained rotations, it is possible to observe that different responses exhibit a substantially different share of the FEV explained by oil supply shocks. Figure 5.2 shows how mild responses of real oil prices are associated with a low contribution to the FEV explained by oil supply shocks. In contrast, stronger responses of oil prices are linked to a higher share of the FEV explained by oil supply shocks.

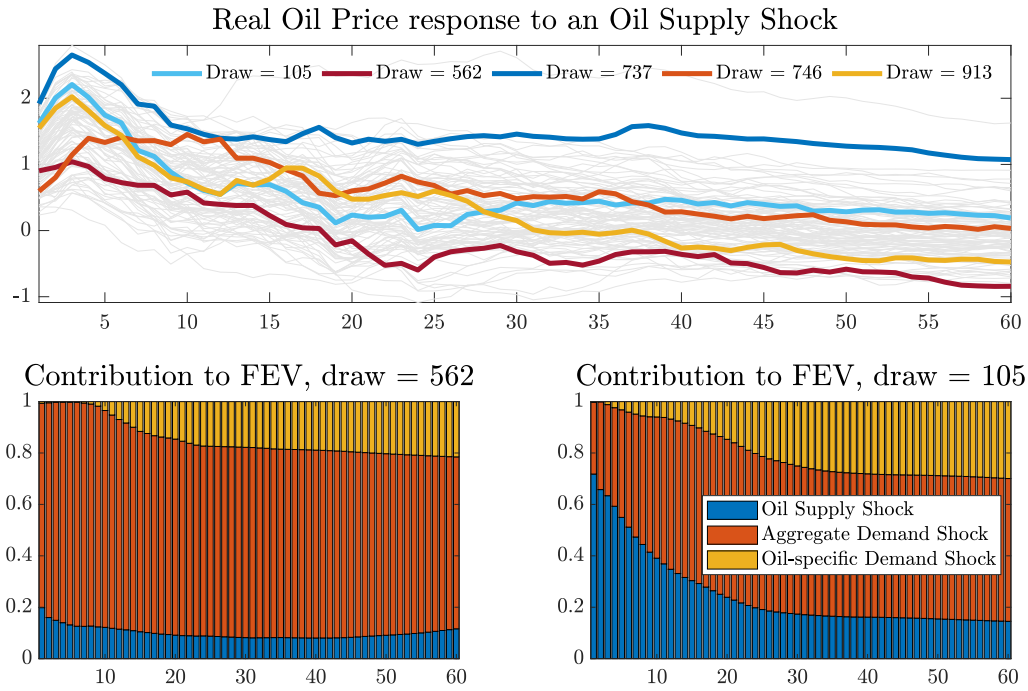

Figure 5.2: Top row: retained draw responses of real oil price to an oil supply shock. Bottom row: contribution to the FEV for selected draws.

Kilian and Murphy (2012) take into account these concerns regarding the admission of responses that exhibit an implausible response of oil prices to an oil supply shock by reducing the set of retained draws to those that show a price elasticity of oil supply smaller than 0.0258. This reduces the initial number of draws from 1.5 million to 80 rotations that fulfill the sign and price elasticity restrictions. As shown in Figure 5.3 this strong assumption on the value of the price elasticity of oil supply leads to a response of real prices which is nearly indistinguishable. Moreover, the contribution to the FEV of oil prices explained by the oil supply shock after 6 months is 3.4% for the elasticity bounds approach compared to 1.9% under the recursive identification. Both results reflect a significantly smaller contribution compared to the sign-identified model which exhibit a contribution to the FEV of 34.6% in the same horizon. As pointed out in Braun (2021), strong assumptions about the price elasticity of oil supply leads to choosing from a set of equally likely models, predetermining

the responses of oil prices that the identified model will exhibit.

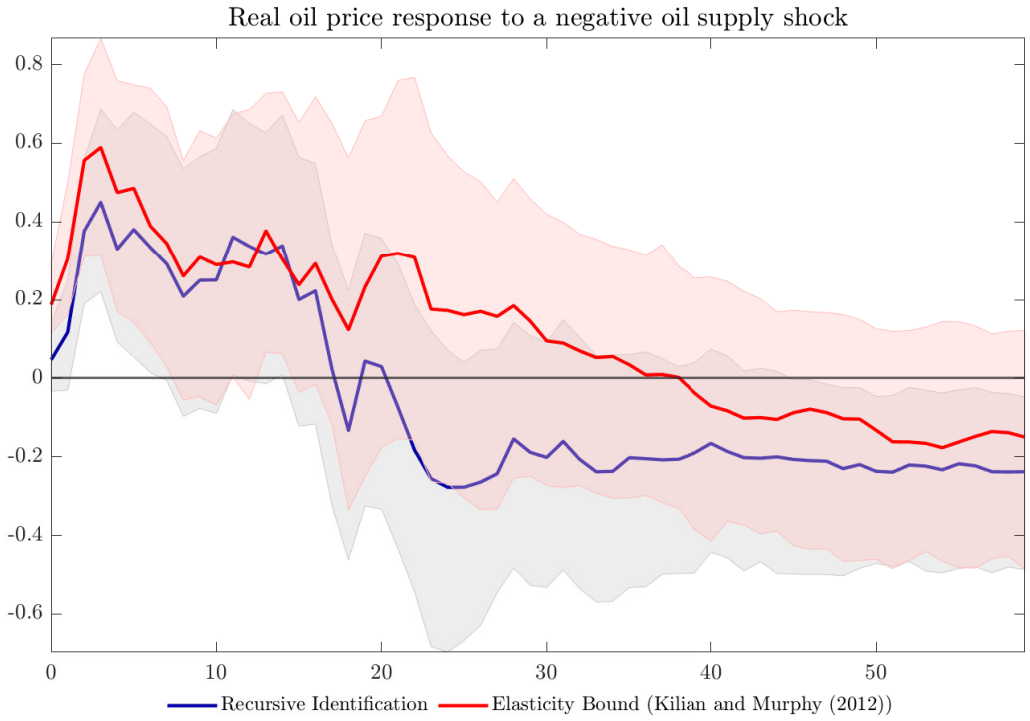

Figure 5.3: Solid lines are the median responses. Shaded areas represent the 68% error bands.

In order to address both type of concerns: admitting potentially implausible responses or restricting the set of retained draws a-priori, we propose an alternative identification strategy combining restrictions on the sign of the impulse response functions with bounds on the FEV contribution of the shock of interest to the variable under examination. To accomplish that, we introduce an upper and lower bound on the oil supply shock contribution to the FEV of real oil prices, derived from the estimates of Caldara et al. (2019). By incorporating an upper and lower bound, we therefore account for estimation uncertainty while reducing the set of admitted rotations.

Our baseline results are shown in Figure 5.4. On panel **(a)** we plot the response of real oil prices to a negative oil supply shock in comparison with the purely sign-identified model[4]. It can be observed that both responses are quantitative equivalent in terms of their median

[4]Figure E.2 of the appendix shows the complete set of responses under the baseline identification strategy.

response, nonetheless, our baseline identification strategy presents an important reduction in the error bands which narrows down the dispersion among the retained set of responses and brings down the uncertainty around the magnitude of the real oil prices responses in comparison with classical sign restrictions. Panel **(b)** shows that including an additional restriction in the form of FEV bounds, reshapes the contribution an oil supply shock has on real oil prices FEV, centering the mass of the contribution around the median estimates of this share and reducing the dispersion on the tails of the distribution.

These results suggest that incorporating an additional restriction to the sign responses in the shape of restricting the bounds of the contribution to the FEV, lead to more informative results as it reduces the dispersion in the magnitude of the responses without tilting the responses beforehand by incorporating potentially over-restrictive assumptions.

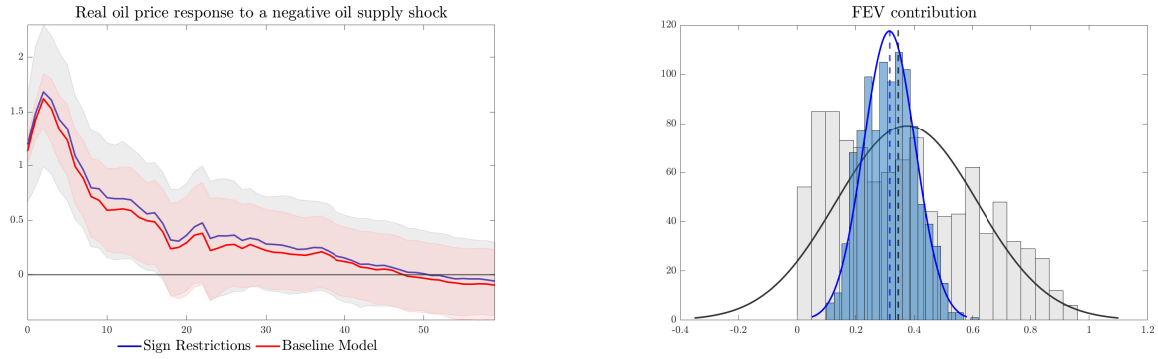

(a) (b)

Figure 5.4: Panel **(a)** Solid lines are the median responses. Shaded areas represent the 68% error bands. Panel **(b)** shows the posterior density of the FEV contribution of an oil supply shock on real oil prices at $h = 6$.

We now consider the case where the researcher knows in advance the driving forces behind the movements in real oil prices. In other words, he/she has a-priori information about the importance of oil supply and oil-specific demand shocks on real oil prices. We examine the alternative scenario in which restrictions on the contribution of these two shocks are placed as an additional restriction to conventional sign restrictions. Conceptually this implies a

slight modification to the step 3 of the algorithm of the baseline identification. In this regard, instead of setting bounds on the FEV $\Omega_{i,j}(h)$, we now place the restriction that the oil supply shock represents the biggest share of real oil price FEV. This restriction is represented as: $\Omega_{i,j}(h) > \Omega_{i,s}(h),\ \forall s \neq j$, where j represents the oil supply shock and s are the other shocks excluding j and h is set for the first 6 months. Conversely, j will represent the oil-specific demand shock for the oil demand driven scenario.

We compare both alternative scenarios in Figure 5.5. As expected the magnitude of the real price of oil response to an unexpected shortfall in oil supply is larger under the supply driven model. However, even under the model where the oil demand shock is the most important contributor to the variance of the real price of oil, the response is substantially higher compared to the models identified using a recursive identification and applying elasticity bounds.

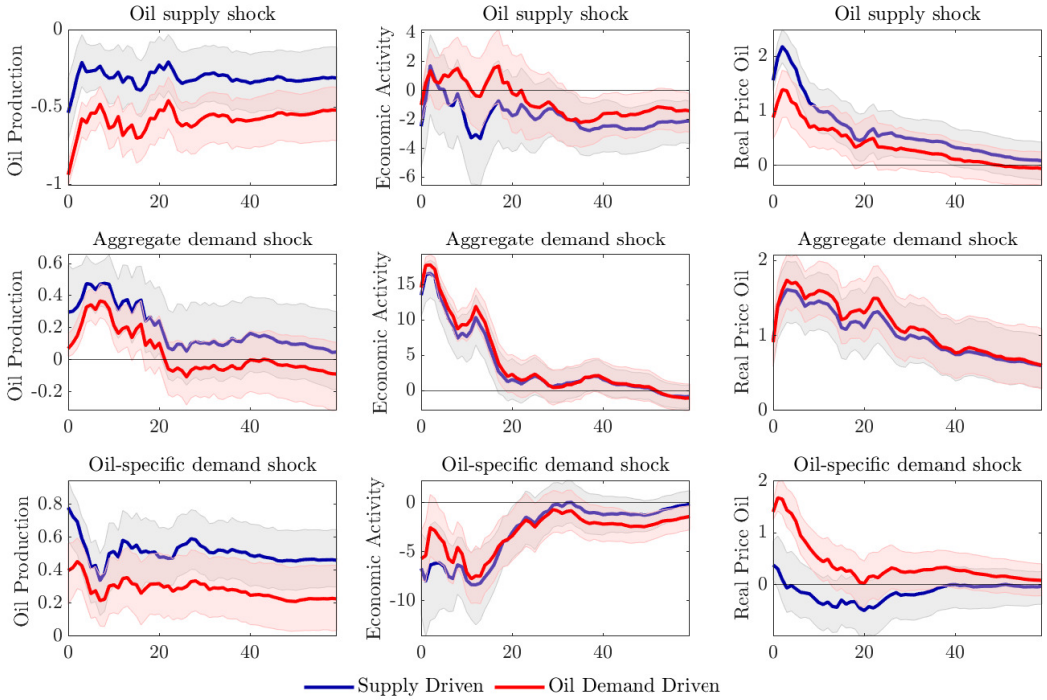

Figure 5.5: Impulse response: supply driven and oil demand driven models. Solid lines are the median responses. Shaded areas represent the 68% error bands.

It is possible to observe that the size of the response of real oil prices to a cut in oil supply is directly associated with the corresponding share of the oil supply shock to the FEV of oil prices. Table 5.2 presents the share of real oil price FEV attributed to an oil supply shock under the different identification methods. It is worth noticing that the similarities in the responses between the exclusion restrictions model and the elasticity bounds by Kilian and Murphy (2012) are mainly explained by the resembling contribution to the FEV. In contrast, the remainder identification strategies present a significantly higher contribution to the FEV and thus, depict a stronger response of real oil prices. Furthermore, it is possible to observe in the last column of 5.2 that despite adding an additional restriction that the oil-specific demand shock is the main driver of the variations in oil prices for the first 6 months, the contribution of the oil supply shock is substantially higher than under recursive identification and elasticity bounds, hence, the difference in the size of the response.

Table 5.2: Oil supply shock contribution to real oil price forecast error variance

	Recursive	KM	Sign Restrictions	SR Bounds	Supply Driven	Demand Driven
$h = 1$	0.001	0.010	0.380	0.350	0.661	0.211
$h = 6$	0.019	0.032	0.346	0.316	0.592	0.234
$h = 12$	0.024	0.035	0.290	0.258	0.505	0.204
$h = 24$	0.036	0.045	0.230	0.202	0.412	0.181

Although neither the supply driven nor the demand driven models are as restrictive as the elasticity bounds from Kilian and Murphy (2012), it is worth highlighting that they reduce the set of admissible models substantially, where the acceptance rate (step 3 of the algorithm) is 0.3% for the demand driven model and 5.6% for the supply one.

3.1 Medium-scale oil market model

In the previous section we investigated the effects shocks in the oil market employing the three-variable system introduced by Kilian (2009). Making use of this small-scale model facilitates the comparability of the baseline results with the existing literature. However, due

to its simplicity, it potentially misses on important sources of fluctuations in the oil market. In this regard, we now evaluate the effects of an unexpected shortfall in oil supply from the perspective of a medium-scale model of six variables. To do so, we utilize the monthly dataset from Känzig (2021), spanning from January 1974 until December 2017. The six variables contained in the SVAR are: the WTI real price of oil, global oil production, global oil inventories, world industrial production (IP), United States (U.S.) IP and U.S. consumer price index (CPI) inflation rate. The lag order is set to 12 following Känzig (2021).

We aim to highlight the benefits of combining restrictions on the contribution to the FEV with classical sign restrictions. To do so, we start by obtaining the results of a purely sign-identified model to a sudden cut in oil supply. The signs of the responses are summarized in Table 5.3.

Table 5.3: Identification: oil supply shock

Oil Supply Shock	Real Oil Price	Oil Prod	Oil Inven	IP	U.S. IP	U.S. CPI
Sign Restriction	+	-	*	-	-	*

Note: $*$ indicates the unrestricted responses

Our objective is to examine the consequences of an abrupt disruption in oil supply, we set the sign of oil production, real oil prices, world IP and U.S. IP comparable to Table 5.1. On the other hand, we leave the signs of world oil inventories and U.S. inflation rate unrestricted as it is not clear the direction of the responses after an unexpected shortfall in oil supply.

The top row of Figure 5.6 presents the results derived from the purely sign-identified model. In Panel **(a)** one can observe that the mass of the contribution of a negative oil supply shock on the FEV of real oil price after 6 months is strongly concentrated in values close to zero. Although, one could argue that oil disruptions contribute little to explaining the fluctuations in real oil prices in the short-run (Kilian, 2009; Kilian and Murphy, 2012, 2014), models with an associated oil supply shock that contribute zero to variations in oil prices in the short-run appear rather implausible. In this regard, we set the additional restriction

that the oil supply shock should explain at least 5% of real oil prices FEV within the first 6 months, $\Omega_{i,j}(h) > 0.05$.

The results of this alternative identification strategy are shown on the bottom row of Figure 5.6. From Panel **(c)**, it is possible to observe that this additional restriction leads to the mass of the posterior density contribution to the FEV to appear more evenly distributed, in contrast with Panel **(a)** where the mass was concentrated around zero. However, it is important to highlight that this additional restriction is not heavily changing the median values of the contribution of the oil supply shock to the FEV of real oil prices. Under the sign restrictions model the median contribution after 6 months is 7.7%, when including the additional restriction the median contribution is 12% in the same horizon. When comparing the responses of oil prices to a negative shock in oil supply from the sign-identified model, Panel **(a)**, with the alternative identification strategy with additional restrictions to the FEV, Panel **(b)**, we can observe that the common conjecture that sign-identified models exhibit an important dispersion in terms of the magnitude of the response (wider error bands) and thus, occasionally lead to uninformative results, can be alleviated by introducing a rather weak additional restriction on the contribution to the FEV[5]. Retaining draws associated with a low contribution from the oil supply shock to real oil price FEV, lead to a rather weak and short lasting response of oil prices. Under the alternative identification strategy, the response is over a percentage point higher than under sign restrictions. Furthermore, the initial shock is more persistent, having an effect for 6 months in contrast with only 3 months under the sign-identified model.

[5]The complete set of responses to a negative oil supply shock is depicted in Figure E.5 of the Appendix.

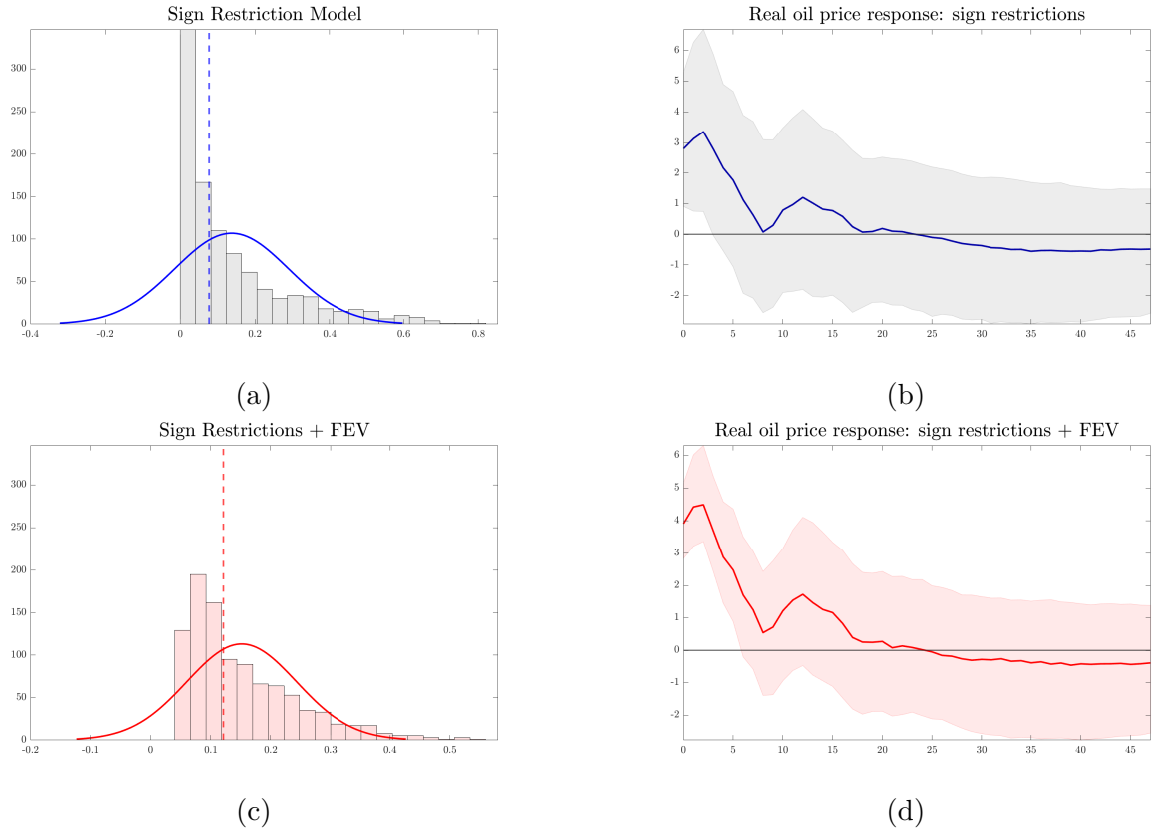

Figure 5.6: Left column shows the posterior density of the oil supply shock contribution to real oil prices FEV at $h = 6$ for the sign-identified model, Panel (**a**), and the model with additional restrictions to the FEV, Panel (**c**). Right column depicts the response of real oil price to a shock in oil supply under sign restrictions, Panel (**b**), and sign restrictions + FEV restrictions, Panel (**d**).

4 Conclusions

In this paper we have reevaluated the effects of an unexpected shortfall in global oil supply using an alternative identification strategy for SVARs which combines information from the FEV with classical sign restrictions. Our baseline identification method aims to overcome the common finding that being too agnostic in the restrictions to the rotation matrix (i.e. imposing only signs on the contemporaneous responses) lead to accepting a wide range of potentially implausible models in the set of retained draws. In the least concerning case,

this would be reflected in a broad set of different responses which exhibit contrasting magnitudes. Ultimately, the sign identified model presents wider error bands which often lead to uninformative results about the magnitude and persistence of the responses to an exogenous shock.

We alleviate this limitation, present in classical sign restrictions, by incorporating an additional restriction in terms of the contribution to the FEV. In the first part of the results we embody this restriction within the 3-equation model from Kilian (2009). In order to account for estimation uncertainty, we use the bounds of the FEV derived from Caldara et al. (2019). Our results suggest that mild responses of oil prices to a cut in oil supply are associated with a low contribution of oil supply shocks to oil prices FEV. Draws that exhibit a share of the FEV close to zero depict an oil price response similar to the one present under recursive identification. Thus, the results in Kilian and Murphy (2012) produce analogous responses to the recursive identification model by narrowing down the set of retained draws to those which satisfies the price elasticity restriction. This additional restriction is analogous to reducing the set to those rotations that present a contribution to the FEV close to zero. When introducing restrictions in the form of bounds to the contribution to the FEV, we can observe that the median response of real prices is almost identical to the purely sign-identified model one. However, the baseline results exhibit an important reduction in the error bands, leading to more informative results in terms of the magnitude of the response.

For the medium-scale model, we show how results can be heavily influenced by rotations which exhibit an implausible contribution to the FEV. Essentially, draws that have no contribution in terms of the FEV to real oil price have important implications in terms of the median response of real oil prices to sudden oil disruptions. In this regard, we show how combining sign restrictions with mild restrictions in terms of the contribution to the FEV, lead to a significant decrease in the range of responses real oil prices exhibit to a negative oil supply shock and thus provide more informative results.

E Appendix

E.1 Figures

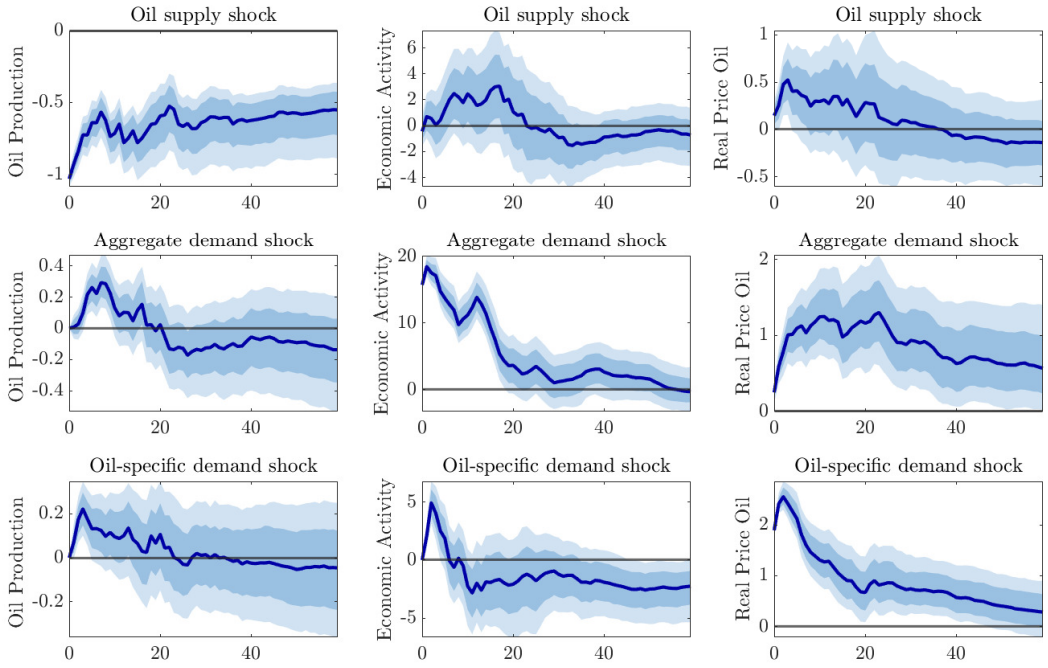

Figure E.1: Solid lines are the median responses identified using a recursive structure. Shaded areas represent the 68% and 90% confidence bands.

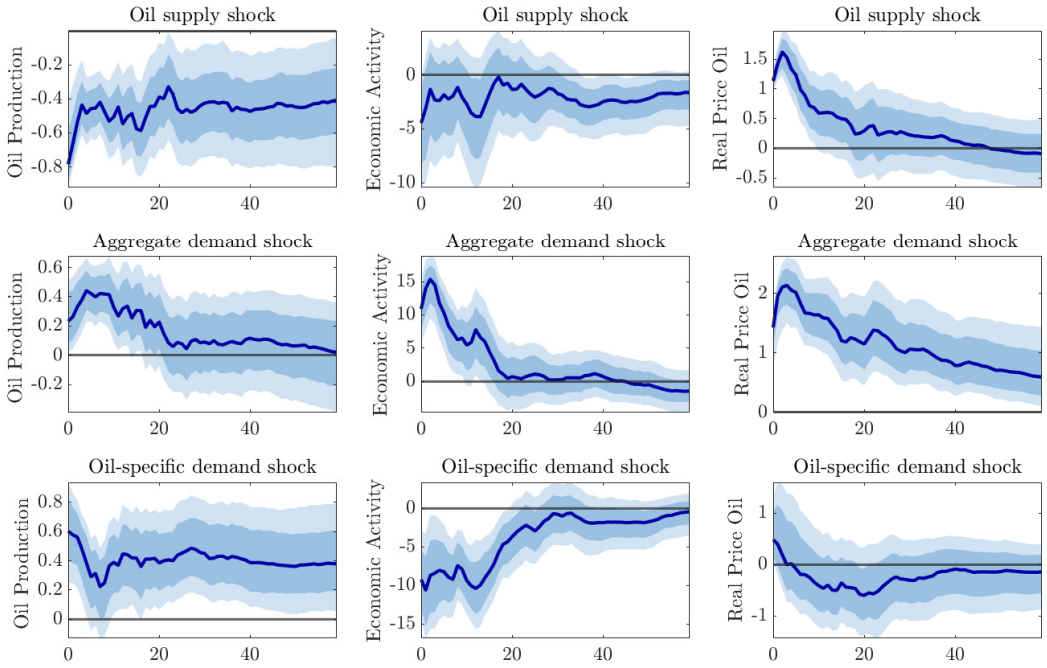

Figure E.2: Solid lines are the median responses identified using sign-restrictions + FEV bounds. Shaded areas represent the 68% and 90% confidence bands.

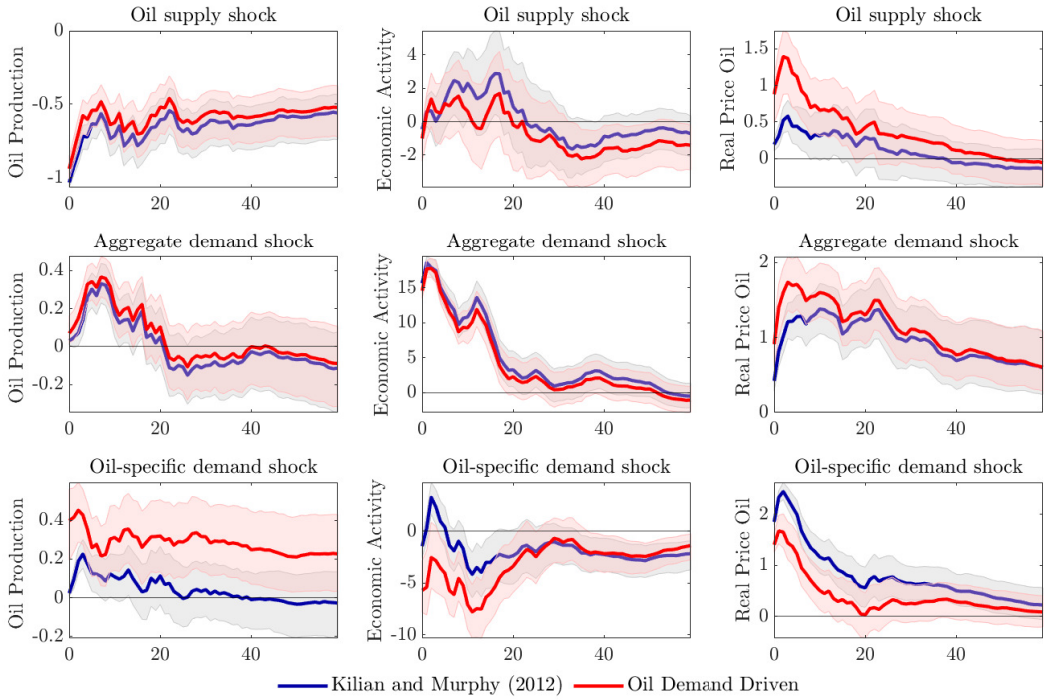

Figure E.3: Solid lines are the median responses. Shaded areas represent the 68% and 90% error bands.

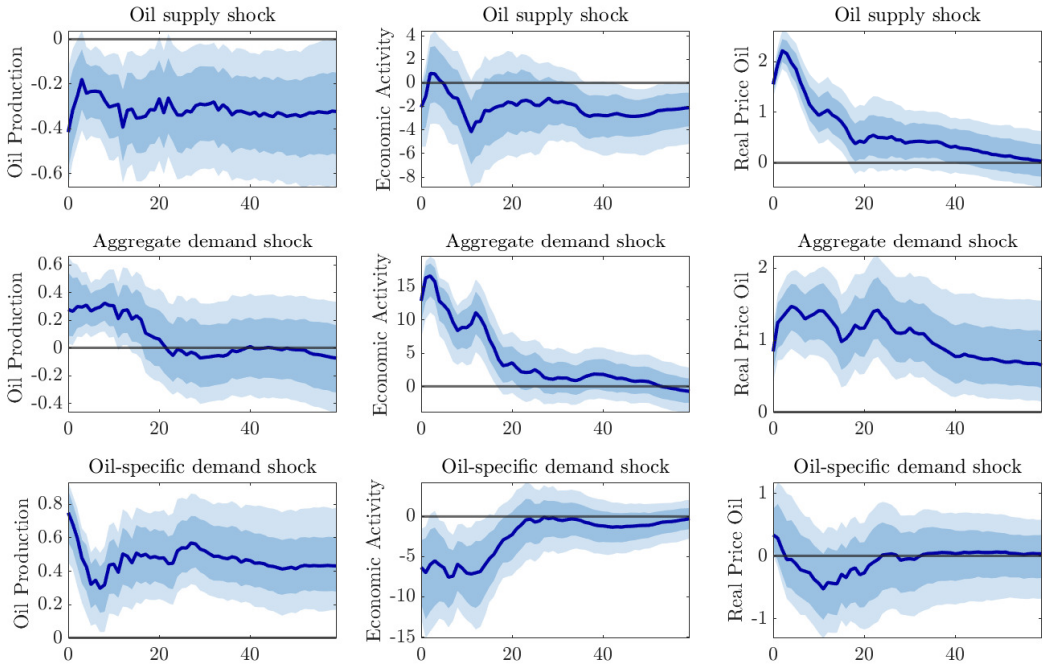

Figure E.4: Solid lines are the median responses identified using sign-restrictions + FEV. Shaded areas represent the 68% and 90% confidence bands.

Table 5.4: Identification assumptions: oil supply shock

Oil Supply Shock	Real Oil Price	Oil Prod	Oil Inven	IP	U.S. IP	U.S. CPI
Sign Restriction	+	-	**	-	-	**
FEV Restriction	$\Omega_{1,2}(h) > 0.05$	$\Omega_{2,2}(h) > 0.05$	**	$\Omega_{4,2}(h) > 0.05$	$\Omega_{5,2}(h) > 0.05$	**

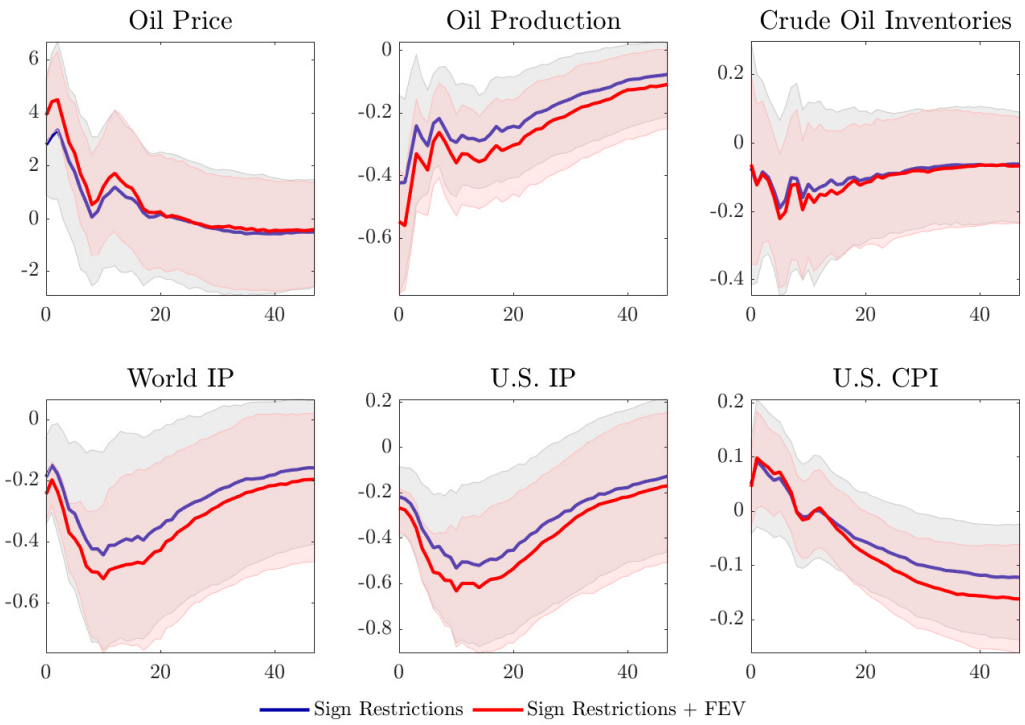

Figure E.5: Impulse response to a negative oil supply shock

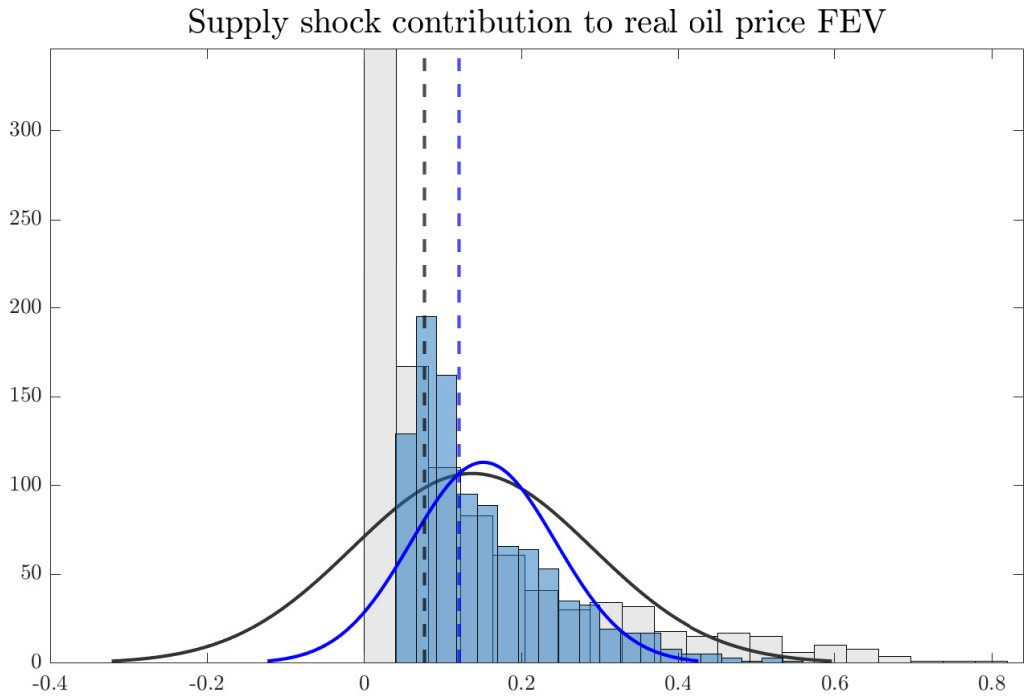

Figure E.6: Posterior mean of the oil supply shock contribution to real oil price FEV at $h = 6$. Grey bars represent the model under sign restriction. Blue bars correspond to the model with additional restrictions to the FEV. Dashed lines are the corresponding median estimates.

Bibliography

Alessandri, P. and Mumtaz, H. (2019). Financial regimes and uncertainty shocks. *Journal of Monetary Economics*, 101:31–46.

Altavilla, C., Canova, F., and Ciccarelli, M. (2020). Mending the broken link: Heterogeneous bank lending rates and monetary policy pass-through. *Journal of Monetary Economics*, 110:81–98.

Amir-Ahmadi, P., Matthes, C., and Wang, M.-C. (2018). Choosing prior hyperparameters: with applications to time-varying parameter models. *Journal of Business & Economic Statistics*, pages 1–13.

Angelini, G., Bacchiocchi, E., Caggiano, G., and Fanelli, L. (2019). Uncertainty across volatility regimes. *Journal of Applied Econometrics*, 34(3):437–455.

Antolín-Díaz, J. and Rubio-Ramírez, J. F. (2018). Narrative sign restrictions for svars. *American Economic Review*, 108(10):2802–29.

Arias, J. E., Caldara, D., and Rubio-Ramirez, J. F. (2019). The systematic component of monetary policy in svars: An agnostic identification procedure. *Journal of Monetary Economics*, 101:1–13.

Bachmann, R., Born, B., Elstner, S., and Grimme, C. (2019). Time-varying business volatility and the price setting of firms. *Journal of Monetary Economics*, 101:82–99.

Bai, J. and Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, 70(1):191–221.

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636.

Barsky, R. B. and Sims, E. R. (2011). News shocks and business cycles. *Journal of monetary Economics*, 58(3):273–289.

Bauer, M., Lakdawala, A., and Mueller, P. (2019). Market-based monetary policy uncertainty. *Available at SSRN 3371160*.

Bauer, M. D. et al. (2012). Monetary policy and interest rate uncertainty. *FRBSF Economic Letter*, 38:1–5.

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

Baumeister, C. and Peersman, G. (2013). The role of time-varying price elasticities in accounting for volatility changes in the crude oil market. *Journal of Applied econometrics*, 28(7):1087–1109.

Baur, D. G. (2003). What is co-movement? *EUR Working Paper No. 20759*.

Benati, L. and Surico, P. (2008). Evolving u.s. monetary policy and the decline of inflation predictability. *Journal of the European Economic Association*, 6(2-3):634–646.

Bernanke, B. S. (2003). A Perspective on Inflation Targeting.

Bernanke, B. S. and Woodford, M. (1997). Inflation Forecasts and Monetary Policy. *Journal of Money, Credit and Banking*, 29(4):653.

Bianchi, F., Kung, H., and Tirskikh, M. (2018). The Origins and Effects of Macroeconomic Uncertainty. Working paper 25386, NBER.

Blagov, B., Funke, M., and Moessner, R. (2015). Modelling the time-variation in euro area lending spreads. BIS working paper No 526, Bank for International Settlements.

Blinder, A. S., Ehrmann, M., Fratzscher, M., De Haan, J., and Jansen, D. J. (2008). Central bank communication and monetary policy: A survey of theory and evidence. *Journal of Economic Literature*.

Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica*, 77(3):623–685. Publisher: National Bureau of Economic Research Cambridge, Mass., USA.

Boivin, J., Kiley, M. T., and Mishkin, F. S. (2010). How has the monetary transmission mechanism evolved over time? In *Handbook of monetary economics*, volume 3, pages 369–422. Elsevier.

Braun, R. (2021). The importance of supply and demand for oil prices: evidence from non-gaussianity. *Bank of England Staff Working Paper No. 957*.

Breitung, J. and Eickmeier, S. (2016). Analyzing international business and financial cycles using multi-level factor models: A comparison of alternative approaches. In *Dynamic Factor Models (Advances in Econometrics, Vol. 35)*. Emerald Group Publishing Limited.

Caggiano, G., Castelnuovo, E., and Groshenny, N. (2014). Uncertainty shocks and unemployment dynamics in U.S. recessions. *Journal of Monetary Economics*, 67:78–92.

Caggiano, G., Castelnuovo, E., and Pellegrino, G. (2017). Estimating the real effects of uncertainty shocks at the zero lower bound. *European Economic Review*, 100:257–272.

Caldara, D., Cavallo, M., and Iacoviello, M. (2019). Oil price elasticities and oil price fluctuations. *Journal of Monetary Economics*, 103:1–20.

Campbell, J. R., Evans, C. L., Fisher, J. D., Justiniano, A., Calomiris, C. W., and Woodford, M. (2012). Macroeconomic effects of federal reserve forward guidance [with comments and discussion]. *Brookings papers on economic activity*, pages 1–80.

Canova, F. and De Nicolo, G. (2002). Monetary disturbances matter for business fluctuations in the g-7. *Journal of Monetary Economics*, 49(6):1131–1159.

Canova, F. and Gambetti, L. (2009). Structural changes in the us economy: Is there a role for monetary policy? *Journal of Economic dynamics and control*, 33(2):477–490.

Carter, C. K. and Kohn, R. (1994a). On Gibbs sampling for state space models. *Biometrika*, 81(3):541–553.

Carter, C. K. and Kohn, R. (1994b). On Gibbs sampling for state space models. *Biometrika*, 81(3):541–553.

Cascaldi-Garcia, D., Ferreira, T. R., Giannone, D., and Modugno, M. (2021). Back to the Present: Learning about the Euro Area through a Now-casting Model. *International Finance Discussion Paper*, 2021(1312):1–46.

Castelnuovo, E. (2019). Domestic and global uncertainty: A survey and some new results. *CAMA Working Paper*, 75/2019.

Castelnuovo, E. and Pellegrino, G. (2018). Uncertainty-dependent effects of monetary policy shocks: A new-Keynesian interpretation. *Journal of Economic Dynamics and Control*, 93:277–296.

Castelnuovo, E. and Surico, P. (2010). Monetary policy, inflation expectations and the price puzzle. *The Economic Journal*, 120(549):1262–1283.

Chan, J. C. (2017). The stochastic volatility in mean model with time-varying parameters: An application to inflation modeling. *Journal of Business & Economic Statistics*, 35(1):17–28.

Chang, B. Y. and Feunou, B. (2014). Measuring Uncertainty in Monetary Policy Using Implied Volatility and Realized Volatility. *Bank of Canada Review*, 1:32–41.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (1999). Monetary policy shocks: What have we learned and to what end? *Handbook of macroeconomics*, 1:65–148.

Ciccarelli, M., Maddaloni, A., and Peydró, J.-L. (2013). Heterogeneous transmission mechanism: monetary policy and financial fragility in the eurozone. *Economic Policy*, 28(75):459–512.

Claessens, S. (2019). Fragmentation in global financial markets: good or bad for financial stability? BIS Working Papers 815, Bank for International Settlements.

Clarida, R., Gali, J., and Gertler, M. (2000). Monetary policy rules and macroeconomic stability: evidence and some theory. *The Quarterly journal of economics*, 115(1):147–180.

Cogley, T., Primiceri, G. E., and Sargent, T. J. (2010). Inflation-gap persistence in the us. *American Economic Journal: Macroeconomics*, 2(1):43–69.

Cogley, T. and Sargent, T. J. (2005). Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics*, 8(2):262–302.

Cogley, T. and Sbordone, A. M. (2008). Trend inflation, indexation, and inflation persistence in the new keynesian phillips curve. *American Economic Review*, 98(5):2101–26.

Coibion, O. and Gorodnichenko, Y. (2011). Monetary policy, trend inflation, and the great moderation: An alternative interpretation. *American Economic Review*, 101(1):341–70.

Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *The American Economic Review*, 105(8):2644–2678.

Coibion, O., Gorodnichenko, Y., Kumar, S., and Pedemonte, M. (2020). Inflation expectations as a policy tool? *Journal of International Economics*, 124:103297. NBER International Seminar on Macroeconomics 2019.

Corsetti, G., Duarte, J., and Mann, S. (2020). Modelling the time-variation in euro area lending spreads. CEPR discussion paper 14968, Centre for Economic Policy Research.

Creal, D. D. and Wu, J. C. (2017). Monetary Policy Uncertainty and Economic Fluctuations. *International Economic Review*, 58(4):1317–1354.

Dahlhaus, T. and Sekhposyan, T. (2018). Monetary policy uncertainty: A tale of two tails. Working paper 2018-50, Bank of Canada.

Debortoli, D., Galí, J., and Gambetti, L. (2019). *On the Empirical (Ir)relevance of the Zero Lower Bound Constraint*, pages 141–170. Volume 34 of Eichenbaum et al. (2020).

Dées, S. and Zorell, N. (2012). Business cycle synchronisation: disentangling trade and financial linkages. *Open Economies Review*, 23(4):623–643.

Del Negro, M., Giannone, D., Giannoni, M. P., and Tambalotti, A. (2019). Global trends in interest rates. *Journal of International Economics*, 118:248–262.

Del Negro, M. and Otrok, C. (2008). Dynamic factor models with time-varying parameters: measuring changes in international business cycles. *FRB of New York Staff Report*, (326).

Dieppe, A., Francis, N., and Kindberg-Hanlon, G. (2021). The identification of dominant macroeconomic drivers: coping with confounding shocks.

Draeger, L. and Lamla, M. J. (2018). Is the Anchoring of Consumers' Inflation Expectations Shaped by Inflational Experience? Working paper series 7042, CESifo.

Eichenbaum, M. S., Hurst, E., and Parker, J. A. (2020). *NBER Macroeconomics Annual 2019, volume 34*, volume 34. University of Chicago Press.

Faust, J. (1998). The robustness of identified var conclusions about money. *Carnegie-Rochester Conference Series on Public Policy*, 49:207–244.

Fernández-Villaverde, J., Guerrón-Quintana, P., Kuester, K., and Rubio-Ramírez, J. (2015). Fiscal volatility shocks and economic activity. *American Economic Review*, 105(11):3352–84.

Francis, N., Owyang, M. T., Roush, J. E., and DiCecio, R. (2014). A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks. *Review of Economics and Statistics*, 96(4):638–647.

Giacomini, R. and Kitagawa, T. (2021). Robust bayesian inference for set-identified models. *Econometrica*, 89(4):1519–1556.

Guerkaynak, R. S., Sack, B., and Swanson, E. (2005). The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models. *American Economic Review*, 95(1):425–436.

Hansen, S., McMahon, M., and Prat, A. (2018). Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach. *The Quarterly Journal of Economics*, 133(2):801–870.

Haubrich, J., Pennacchi, G., and Ritchken, P. (2012). Inflation expectations, real rates, and risk premia: Evidence from inflation swaps. *The Review of Financial Studies*, 25(5):1588–1629.

Husted, L., Rogers, J., and Sun, B. (2020). Monetary policy uncertainty. *Journal of Monetary Economics*, 115:20–36.

Ikeda, D., Li, S., Mavroeidis, S., and Zanetti, F. (2020). Testing the Effectiveness of Unconventional Monetary Policy in Japan and the United States. IMES Discussion Paper Series 20-E-10, Institute for Monetary and Economic Studies, Bank of Japan.

Inoue, A. and Kilian, L. (2021). The role of the prior in estimating var models with sign restrictions. *Center for Financial Studies Working Paper*, (660).

Istrefi, K. and Mouabbi, S. (2018). Subjective interest rate uncertainty and the macroeconomy: A cross-country analysis. *Journal of International Money and Finance*, 88:296–313.

Istrefi, K. and Piloiu, A. (2014). Economic policy uncertainty and inflation expectations. Technical report, Bank de France.

Jurado, K., Ludvigson, S. C., and Ng, S. (2015). Measuring uncertainty. *American Economic Review*, 105(3):1177–1216.

Kang, W., Lee, K., and Ratti, R. A. (2014). Economic policy uncertainty and firm-level investment. *Journal of Macroeconomics*, 39(PA):42–53.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *American Economic Review*, 111(4):1092–1125.

Kilian, L. (2008). Exogenous oil supply shocks: how big are they and how much do they matter for the us economy? *The Review of Economics and Statistics*, 90(2):216–240.

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

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

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

Koop, G. and Korobilis, D. (2009). *Bayesian multivariate time series methods for empirical macroeconomics*. Foundations and Trends in Econometrics.

Kose, M. A., Otrok, C., and Whiteman, C. H. (2003). International business cycles: World, region, and country-specific factors. *American Economic Review*, 93(4):1216–1239.

Krippner, L. (2013). Measuring the stance of monetary policy in zero lower bound environments. *Economics Letters*, 118(1):135–138.

Kurov, A. and Stan, R. (2018). Monetary policy uncertainty and the market reaction to macroeconomic news. *Journal of Banking & Finance*, 86:127–142.

Leduc, S. and Liu, Z. (2016). Uncertainty shocks are aggregate demand shocks. *Journal of Monetary Economics*, 82:20–35.

Leduc, S. and Sill, K. (2013). Expectations and economic fluctuations: an analysis using survey data. *Review of Economics and Statistics*, 95(4):1352–1367.

Leduc, S., Sill, K., and Stark, T. (2007). Self-fulfilling expectations and the inflation of the 1970s: Evidence from the livingston survey. *Journal of Monetary economics*, 54(2):433–459.

Lindsten, F., Jordan, M. I., and Schon, T. B. (2014). Particle Gibbs with ancestor sampling. *Journal of Machine Learning Research*, 15:2145–2184.

Lippi, F. and Nobili, A. (2012). Oil and the macroeconomy: a quantitative structural analysis. *Journal of the European Economic Association*, 10(5):1059–1083.

Liu, P., Theodoridis, K., Mumtaz, H., and Zanetti, F. (2019). Changing macroeconomic dynamics at the zero lower bound. *Journal of Business and Economic Statistics*, 37(3):391–404.

Lucas, R. E. (1976). Econometric policy evaluation: A critique. *Carnegie-Rochester Confer. Series on Public Policy*.

Malmendier, U. and Nagel, S. (2016). Learning from Inflation Experiences. *The Quarterly Journal of Economics*, 131(1):53–87. Publisher: Oxford Academic.

Miranda-Agrippino, S. and Rey, H. (2020). U.S. Monetary Policy and the Global Financial Cycle. *The Review of Economic Studies*, 87(6):2754–2776.

Mumtaz, H. and Musso, A. (2021). The evolving impact of global, region-specific, and country-specific uncertainty. *Journal of Business & Economic Statistics*, 39(2):466–481.

Mumtaz, H. and Surico, P. (2012). Evolving international inflation dynamics: world and country-specific factors. *Journal of the European Economic Association*, 10(4):716–734.

Mumtaz, H. and Theodoridis, K. (2015). The international transmission of volatility shocks: an empirical analysis. *Journal of the European Economic Association*, 13(3):512–533.

Mumtaz, H. and Theodoridis, K. (2018). The changing transmission of uncertainty shocks in the us. *Journal of Business & Economic Statistics*, 36(2):239–252.

Mumtaz, H. and Theodoridis, K. (2020). Dynamic effects of monetary policy shocks on macroeconomic volatility. *Journal of Monetary Economics*, 114:262–282.

Mumtaz, H. and Zanetti, F. (2013). The impact of the volatility of monetary policy shocks. *Journal of Money, Credit and Banking*, 45(4):535–558.

Mumtaz, H. and Zanetti, F. (2015). Labor market dynamics: A time-varying analysis. *Oxford Bulletin of Economics and Statistics*, 77(3):319–338.

Neely, C. J. (2005). Using implied volatility to measure uncertainty about interest rates. *Federal Reserve Bank of St Louis Review*, 87(3):407–426.

Onatski, A. (2009). Testing hypotheses about the number of factors in large factor models. *Econometrica*, 77(5):1447–1479.

Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies*, 72(3):821–852.

Rey, H. (2015). Dilemma not trilemma: The global financial cycle and monetary policy independence. *NBER Working Paper No. 21162*.

Rubio-Ramirez, J. F., Waggoner, D. F., and Zha, T. (2010). Structural vector autoregressions: Theory of identification and algorithms for inference. *The Review of Economic Studies*, 77(2):665–696.

Schorfheide, F. and Song, D. (2015). Real-time forecasting with a mixed-frequency var. *Journal of Business & Economic Statistics*, 33(3):366–380.

Sims, C. A. (1980). Macroeconomics and reality. *Econometrica: journal of the Econometric Society*, pages 1–48.

Sinha, A. (2016). Monetary policy uncertainty and investor expectations. *Journal of Macroeconomics*, 47:188–199.

Stulz, R. (1986). Interest rates and monetary policy uncertainty. *Journal of Monetary Economics*, 17(3):331–347.

Svensson, L. E. (2000). Open-economy inflation targeting. *Journal of international economics*, 50(1):155–183.

Swanson, E. T. (2006). Have increases in federal reserve transparency improved private sector interest rate forecasts? *Journal of Money, Credit and Banking*, pages 791–819.

Swanson, E. T. and Williams, J. C. (2014). Measuring the effect of the zero lower bound on medium- and longer-term interest rates. *American Economic Review*, 104(10):3154–85.

Uhlig, H. (2003). What moves real GNP? *Working Paper, Humbold University*.

Uhlig, H. (2005). What are the effects of monetary policy on output? results from an agnostic identification procedure. *Journal of Monetary Economics*, 52(2):381–419.

Volpicella, A. (2021). Svars identification through bounds on the forecast error variance. *Journal of Business & Economic Statistics*, pages 1–11.

WGEM Team on Real and Financial Cycles (2018). Real and financial cycles in eu countries-stylised facts and modelling implications. *ECB Ocassional Paper Series*, (205).

Wu, J. C. and Xia, F. D. (2016). Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking*, 48(2-3):253–291.

Statement of co-authorship

Chapter 2 "Monetary policy uncertainty and inflation expectations" was created in collaboration with Boris Blagov. We thereby declare that Gabriel Arce Alfaro participated proportionally in all parts of the work.

Gabriel Arce Alfaro Boris Blagov

Chapter 3 "Financial integration or financial fragmentation? A euro area perspective" was created in collaboration with Boris Blagov. We thereby declare that Gabriel Arce Alfaro participated proportionally in all parts of the work.

Gabriel Arce Alfaro Boris Blagov

Eidesstattliche Erklärung

Ich erkläre hiermit, dass ich in keinem laufenden oder früheren Promotionsverfahren zum Erwerb desselben Grades Dr. rer. pol endgültig gescheitert bin.

Ich erkläre hiermit, dass ich die vorliegende Arbeit selbständig ohne unzulässige Hilfe Dritter verfasst, keine anderen als die angegebenen Quellen und Hilfsmittel benutzt und alle wörtlich oder inhaltlich übernommenen Stellen unter der Angabe der Quelle als solche gekennzeichnet habe. Die Grundsätze für die Sicherung guter wissenschaftlicher Praxis an der Universität Duisburg-Essen sind beachtet worden.

Ich habe die Arbeit keiner anderen Stelle zu Prüfungszwecken vorgelegt. Die Gelegenheit zum vorliegenden Promotionsverfahren ist mir nicht kommerziell vermittelt worden. Insbesondere habe ich keine Organisation eingeschaltet, die gegen Entgelt Betreuerinnen und Betreuer für die Anfertigung von Dissertationen sucht oder die mir obliegenden Pflichten hinsichtlich der Prüfungsleistungen für mich ganz oder teilweise erledigt. Hilfe Dritter wurde bis jetzt und wird auch künftig nur in wissenschaftlich vertretbarem und prüfungsrechtlich zulässigem Ausmaß in Anspruch genommen. Mir ist bekannt, dass Unwahrheiten hinsichtlich der vorstehenden Erklärung die Zulassung zur Promotion ausschließen bzw. später zum Verfahrensabbruch oder zur Rücknahme des Titels führen können.

Ort, Datum	Gabriel Arce Alfaro