

CAMP Working Paper Series
No 5/2015

Commodity prices and fiscal policy design: Procyclical despite a rule

Hilde C. Bjørnland and Leif Anders Thorsrud



© Authors 2015

This paper can be downloaded without charge from the CAMP website <http://www.bi.no/camp>

Commodity prices and fiscal policy design: Procyclical despite a rule*

Hilde C. Bjørnland[†] Leif Anders Thorsrud[‡]

August 23, 2018

Abstract

Recent studies offer evidence of reduced fiscal procyclicality to commodity price changes in resource-rich countries, a feature commonly attributed to the adoption of fiscal policy rules. We revisit this issue, and find that by controlling for global activity shocks while allowing for time-varying changes in both fiscal policy and the volatility of shocks, this finding does not hold. To show this we develop a time-varying Dynamic Factor Model, allowing for a multiple of shocks, stochastic volatility and time-varying parameters, and estimate it on data for Norway, whose handling of resource wealth is often cited as exemplary.

JEL-codes: C32, E32, E62, F41, Q43

Keywords: Time-varying Dynamic Factor Model, commodity prices, fiscal policy, sovereign wealth fund

*This Working Paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. We thank three anonymous referees, Martin Eichenbaum, Knut Moum, Rick van der Ploeg, James Stock, Ragnar Torvik and seminar and conference participants at Melbourne University, Norges Bank, the CAMP-ESOP-OxCarre Workshop 2015, the 23rd SNDE Symposium 2015 and the 10th International Conference on Computational and Financial Econometrics (CFE 2016) for valuable comments. This work is part of the research activities at the Centre for Applied Macro and Petroleum economics (CAMP) at the BI Norwegian Business School.

[†]Centre for Applied Macro- and Petroleum economics (CAMP) - BI Norwegian Business School and Norges Bank. Email: hilde.c.bjornland@bi.no

[‡]Norges Bank and Centre for Applied Macro- and Petroleum economics (CAMP) - BI Norwegian Business School. Email: leif.a.thorsrud@bi.no

1 Introduction

In countries where resource revenue constitutes a large component of total government revenues, commodity price fluctuations will have a direct impact on public spending. Many resource-rich countries are therefore advised to adopt some type of fiscal policy framework (i.e., a fiscal spending rule), which, if operated countercyclically, should shelter the economy from commodity price fluctuations and prevent over-spending on the part of the government, see e.g. [Barro \(1979\)](#) for related arguments from the tax and consumption smoothing literature, or [Portes and Wren-Lewis \(2014\)](#) for a recent overview.

The adoption of a fiscal rule, however, does not in itself ensure that fiscal policy works to insulate the domestic economy from commodity price fluctuations: The constructed rule may be too lax over the commodity price cycle, the actual conduct of fiscal policy might not be in accordance with the rule, or both. Hence, what works in theory may not necessarily work in practice.

We examine whether fiscal rules actually work in practice. More specifically, we analyze fiscal policy's response in a resource-rich economy to oil price shocks and the extent to which this response insulated the domestic economy from oil price fluctuations or, indeed, even exacerbated their effect. To account for the changing nature of economic conditions and complexity of fiscal rules, we address this question by developing a time-varying Dynamic Factor Model (DFM), in which we allow the volatility of structural shocks, the systematic fiscal policy responses, and the macroeconomic conditions, to change. Our proposed model is comparable to existing time-varying DFMs, but differs in how the factors are identified in terms of economic quantities, thus permitting the identification of structural shocks. This is where we find the methodological novelty of our approach. From the perspective of its empirical application, this, we believe, is the first time fiscal policy has been evaluated in this way not only for a resource-rich country, but or for any country.

We focus on a particular country, Norway, whose handling of its petroleum wealth has been described as exemplary (see e.g. [OECD \(2005\)](#), [OECD \(2007\)](#) and [Velculescu \(2008\)](#) among many others). Unlike most oil exporters, Norway has adopted a fiscal framework in 2001, with a view to shielding the fiscal budget, and therefore also the domestic economy, from oil price fluctuations. Under the framework, revenue from the sale of oil and gas is deposited into what is known as a Savings Fund. Only the expected real return of the

fund is drawn annually to finance public spending or tax cuts. By comparing how fiscal policy responds to oil price shocks before and after the rule's implementation, our study provides us with a natural experiment allowing us to assess fiscal policy performance over commodity price cycles.

Having said that, analyses of fiscal procyclicality in oil rich economies come with an important caveat: since the price of oil has moved in tandem with global demand changes in the fiscal policy's response to changes in oil price could be due to the role of global demand in the most recent business cycles, not necessarily to an increase in the price of oil. In fact, recent studies tend to emphasize the role of global demand as a driver of oil prices, see, e.g., [Kilian \(2009\)](#).¹

Therefore, and in line with these findings, we control for shocks to global activity in our analysis of fiscal policy responses to oil price shocks. Previous studies addressing the role of fiscal policy in resource rich countries have typically ignored this issue, treating instead oil prices as exogenous, see e.g., [Pieschacon \(2012\)](#) and [Céspedes and Velasco \(2014\)](#) among others. In particular, [Céspedes and Velasco \(2014\)](#) draw their conclusion after estimating the response in government expenditures and revenues to commodity prices using a large panel of commodity exporting countries over two different cycles, while [Pieschacon \(2012\)](#) designs a counterfactual analysis comparing the impulse responses in a series of variables to an exogenous oil price shock in Norway and Mexico in the period 1986-2006. In doing so, both studies provide evidence of reduced fiscal procyclicality to commodity price changes in the recent commodity price boom, something they attribute to improved institutional quality, i.e., adopted fiscal policy rules. Yet if global demand is an important source of variation in commodity prices, in particular in the recent commodity price boom, we would expect fiscal policy to be, exactly, countercyclical. Not necessarily because governments has reduced spending relative to GDP, but simply because domestic GDP has increased with global demand.

And indeed, when we control for global activity, we confirm that the countercyclical

¹[Kilian \(2009\)](#) shows that if oil prices increase in response to spurs in demand (rather than disruptions of supply capacity, see, e.g., [Hamilton \(1983, 2003\)](#)), global economic activity will be positively affected, at least in the short run. Corroborating results are shown in, e.g., [Barsky and Kilian \(2002\)](#), [Lippi and Nobili \(2012\)](#), [Peersman and Van Robays \(2012\)](#), [Kilian and Hicks \(2013\)](#), [Kilian and Murphy \(2014\)](#), [Charnavoki and Dolado \(2014\)](#), [Aastveit et al. \(2015\)](#) and [Bjørnland and Thorsrud \(2016\)](#) for various oil importing and oil exporting countries.

fiscal responses found in the recent oil price boom should be attributed to global activity shocks and their domestic propagation, rather than the adopted fiscal framework. In particular, we find that in the wake of oil price shocks (that are orthogonal to global activity), fiscal policy is procyclical on impact and over response horizons. If anything, fiscal policy has become more (not less) procyclical since the adoption of the fiscal policy rule in 2001, in absolute value and relative to GDP. Hence, taking everything else as given, then the government's spending rule has not effectively insulated the economy from oil price shock. This finding stands in rather stark contrast to the arguments put forward in, e.g., [Pieschacon \(2012\)](#).

Furthermore, following a global activity shock that increases oil prices, the picture becomes somewhat more nuanced, with some components of public spending alternating from countercyclical to acyclical in the last decade, while others are more procyclical. The main takeaway, then, is that of a tendency for more (not less) fiscal procyclicality since the adoption of the spending rule. At the same time, the non-resource economy responded strongly procyclically following the global activity shock, and this shock also explains an increasing share of the variation in the economy over time. The strong countercyclical fiscal policy responses (relative to GDP) in the last boom, as reported by [Céspedes and Velasco \(2014\)](#), among others, are therefore most likely to be due to global activity shocks and their domestic propagation, rather than to fiscal policy governed by a rule.

Our framework contributes to the literature in two additional ways. First, a comparison of different countries' fiscal policy responses, implicitly assumes that each of the commodity countries is at the same stage of development. This is seldom the case, as countries may have extracted the windfall gain at different points in time. For instance, Mexico was producing oil as early as the turn of the twentieth century. Today it is a mature oil exporter, with oil and gas production accounting for 7-8 percent of GDP. Norway, on the other hand, discovered its oil fields 70 years later and oil and gas production accounts today for close to 25 percent of total GDP. In this sense, a comparison of the effects of oil price shocks on a mature and a new oil producer, as is done by, e.g., [Pieschacon \(2012\)](#), is likely to indicate different responses which may ultimately have nothing to do with the adoption of fiscal rules per se, but simply reflect different stages of development.² Moreover, apart from being commodity exporters, Mexico and Norway in most

²Applying a related argument, [Alexeev and Conrad \(2009\)](#) control for initial endowment when comparing

other areas are highly dissimilar. We argue that it is more informative to compare fiscal responses in one country consistently over time than to compare fiscal responses across countries at a given time.

Second, countries adopt fiscal rules in response to changing economic conditions. Fiscal policy design is often particularly complex insofar as countries combine the objectives of sustainability with the need for flexibility in response to shocks, see [Schaechter et al. \(2012\)](#) for details. Norway is no exception. Under the fiscal framework, the government's non-oil structural deficit is supposed to equal the long-run real return of the sovereign wealth fund (GPF: Government Pension Fund - Global), projected to be 4 per cent. To ensure against inherent procyclical behaviour, the fiscal guidelines allow for temporary deviations from the rule over the business cycle. The GPF is therefore a hybrid of a savings- and a stabilization fund. So when we compare the economic effects of fiscal policy designs before and after their implementation, we need to control for time-varying changes in macroeconomic conditions as much as in the specific policy implementation.

The Dynamic Factor Model we develop permits us to address these shortcomings in the existing literature in a consistent manner. We include stochastic volatility components to allow for changes in the size of the structural shocks, e.g., the Great Moderation effect and the recent financial crisis and Great Recession,³ together with time-varying factor loadings to allow for changes in systematic policy responses across time, e.g., due to institutional changes in the fiscal rule,⁴ or due to persistent deviations from the adopted rule. The time-varying Dynamic Factor Model developed here compares to the models used by [Lopes and Carvalho \(2007\)](#), [Del Negro and Otrok \(2008\)](#), and [Ellis et al. \(2014\)](#), but differs in the way we identify the dynamic factors, and in the way we model the law of motion of the dynamic factors. In our contribution the dynamic factors are all identified in terms of economic quantities. Importantly, this allows us, in contrast to most other factor model studies, to build on the structural VAR literature, [Primiceri \(2005\)](#) in particular, and to identify the structural shocks driving the dynamic factors. For this reason we

growth performance in resource-rich countries, and find natural resources to *enhance* long term growth, findings quite contrary to the traditional resource curse literature.

³See, e.g., [Cogley and Sargent \(2005\)](#), [Primiceri \(2005\)](#), and [Nakov and Pescatori \(2010\)](#), and the references therein, for a broader discussion of these effects and their possible causes.

⁴Norway has had a number of different fiscal guidelines, from a spend as you go policy in the 1980s, to a neutral fiscal stance in the 1990s and a fiscal spending rule after 2001, see [Section 2](#) for details.

are also able to trace out the effect of different shocks, i.e., global demand and oil price shocks, on a number of public and non-public variables.

The remainder of the paper is structured as follows. In Section 2 we briefly describe the historical evolution of Norway’s fiscal framework, paying particular attention to the introduction of the fiscal rule. Section 3 details the model and the estimation procedure. Section 4 discusses the oil price and global activity shocks and analyses their effect on fiscal policy and the domestic economy. Section 5 concludes.

2 The Fiscal framework

Since the mid 1990s, Norway has been transferring the totality of its petroleum cash flow to a sovereign wealth fund. The fund, established in 1990 as the Government Petroleum Fund, was renamed in 2006 as the Government Pension Fund Global (GPF).⁵ As noted in the introduction, the GPF is a hybrid savings and stabilisation fund. As a savings fund, its main purpose is to save and invest petroleum income in international capital markets, the product of which can be put to use in the Norwegian economy at a later date (i.e., for and by future generations). As a stabilization fund, it seeks to protect and stabilize the budget, and the wider economy, from excess volatility in petroleum revenues, see [Johnson-Calari and Rietveld \(2007\)](#) for details on saving and stabilization funds.

The idea of establishing a hybrid fund arose following large budget deficits and a poor economic environment in Norway following the dramatic drop in the price oil in 1986, see [Lie \(2013\)](#). In its first few years, however, the fund failed to generate a surplus, and fiscal policy guidelines at the time suggested leaving the use of petroleum revenue unchanged, i.e., a neutral fiscal stance, as measured by the structural, non-oil budget balance, see [Ministry of Finance \(2001\)](#). High oil prices, large surpluses on the government budget and high allocations to the GPF in the late 1990s, made it difficult to maintain such fiscal neutrality. In 2001, the government therefore devised a strategy for fiscal policy accommodating a prudent increase in the spending of petroleum revenues. According to the policy guidelines, only the expected real return on the Sovereign Wealth Fund (projected to be 4 percent) could be returned to the budget for general spending purposes. Since under the new fiscal rule the expected return from the Sovereign Wealth Fund would

⁵The change highlighted the fund’s role in saving government revenue to finance an expected increase in future public pension costs. Despite its name, the fund has no formal pension liabilities.

be used to finance the non-oil budget deficit, the budget would eventually balance.⁶ The rule was expected to smooth the spending generated by the oil wealth, while maintaining the strength of Norway’s internationally exposed sector, and insulating the economy from Dutch disease (crowding out of the private sector).

However, fiscal policy also plays an important role in cushioning output fluctuations in two additional ways. First, it stabilizes the fiscal impulse over and above longer term smoothing by allowing deviations from the 4 percent rule to counteract large cyclical variations in economic activity or sharp swings in the value of the Fund. This should give the government manoeuvrability in fiscal policy should oil prices drop or the mainland economy contract. As the government stated in its white paper, “*Fiscal policy should continue to have the main responsibility for stabilising developments in the Norwegian economy.*” (Ministry of Finance (2001), p. 8). Second, to prevent fiscal policy from exacerbating the effect of oil price fluctuations on the Norwegian economy, the rule is expressly defined in terms of the structural non-oil balance; allowing for the full effect of the automatic fiscal stabilizers in contrast to inherently procyclical rules on the actual deficit.

Since the 2001 adoption of the fiscal rule, the GPF has developed rapidly and is today the largest sovereign wealth fund in the world; its value is currently close to 200 percent of Norway’s GDP. This notwithstanding, very little is actually known about how, or indeed if, the rule manages to shield the resource rich economy from oil price fluctuations, as theory predicts.

A mere glimpse at some stylized facts in Figure 1, however, suggests that the Norwegian economy is far from sheltered from oil price fluctuations. The figure displays rolling correlations between oil prices and two key fiscal variables: value added and wages in the public sector. Figure 1 indicates an upward drift in the correlation between either of the fiscal policy variables and oil prices from the early 2000s. While these are unconditional moments, they are nevertheless consistent with an interpretation whereby fiscal policy has tended to respond more procyclically to higher oil prices since 2001. However, the figure also indicates variation in the correlation coefficients across over the sample. In particular, they were just as high in the late 1980s, when politicians pursued a policy of

⁶For this reason, the fiscal rule is defined as a balanced budget rule. Many other countries adopt additional rules restricting spending. For instance, Sweden has both a balanced budget rule and an expenditure rule, see [Schaechter et al. \(2012\)](#) for additional details.




Figure 1. Rolling 30 quarter (backward looking) correlations between public sector spending and oil price growth. Dots indicate when the correlations are significant different from zero (at the 5% significance level).

spend as you go, but fell during the financial crisis, when fiscal policy became counter-cyclical in response to global demand shocks driving both the oil price and GDP in the mainland (non-resource) economy downwards. This suggest that allowing for some form of time variation and controlling for different macroeconomic shocks would seem to be important, and it is what we seek to do in this paper.

3 The model

International fluctuations in the oil market affect macroeconomic developments in oil exporters and oil importers differently. [Bjørnland and Thorsrud \(2016\)](#) proposed a simple theoretical model aimed at capturing the most important transmission channels as seen from an oil (or resource) exporter’s perspective. The model was successfully tested empirically within a Dynamic Factor Model (DFM) framework. Here we extend this work along two important dimensions. First, we focus on how fiscal policy in a resource rich economy responds to international oil market fluctuations. Second, we focus on changes through time by developing a structural time-varying DFM.

Importantly, the model allows a more parsimonious representation of the co-movement of a large cross section of variables than with standard time series techniques; the direct and indirect spillovers between the different sectors of the economy can, moreover, be estimated simultaneously. In particular, the DFM includes four factors with associated shocks with the potential to affect all sectors of the economy. First, we include a measure

of global activity and the real price of oil as two separate factors in the model to capture developments in the oil market as well as international business cycle conditions. This allows us to identify two shocks that can affect oil prices: a global activity shock and an oil price shock, both of which increase the real price of oil, though with potentially very different macroeconomic implications.

Second, we include two separate activity factors for the resource and non-resource (domestic) industries of the economy, thereby allowing the public sector, and the domestic economy in general, to respond differently to a windfall gain due to an activity shock in the resource sector (new discoveries, increased extraction rates etc.) and a windfall gain due to higher oil prices. What's more, we can now also account for spillovers from the petroleum sector to the non-oil sectors due to increased demand for resources, in addition to the spending effect coming via the public sector. Previous studies, such as [Pieschacon \(2012\)](#), typically assume that since resource sector output only provides a source of income from export sales, which the government collects, there will be no spillovers to the rest of the economy while the resources are being extracted. Similar conclusions are drawn by [Husain et al. \(2008\)](#). This is hardly the case for Norway. In particular, as shown theoretically and empirically in [Bjørnland and Thorsrud \(2016\)](#), when the extraction of resources demands complex technical solutions, as it does in Norway, the transfer of learning-by-doing from the resource to the non-resource sectors of the economy can be substantial.⁷

In total, the model identifies four structural shocks: a global activity shock, an oil price shock, a resource activity shock and a domestic activity shock. While controlling for the domestic disturbances and the systematic interaction between the four factors, our main focus is on the fiscal responses caused by the global activity shock and the oil price shock, both of which can affect oil prices. We therefore ask whether fiscal policy has been less procyclical with higher oil prices since the adoption of the fiscal policy rule, and whether the new rule helped shelter the economy from oil price fluctuations?

In the DFM, the factors and shocks will be linearly related to a large panel of domestic variables, including tradable and non-tradable, e.g., public, sectors of the economy. The large panel is needed to account for the sectoral spillovers that exists between the different

⁷Similar findings are also found for mineral-abundant Australia in [Bjørnland and Thorsrud \(2016\)](#), and, albeit using a very different methodology, for a variety of resource-rich countries in [Allcott and Keniston \(2018\)](#) and [Smith \(2014\)](#).

industries of the economy, but also allows us to include a broad range of measures used in the literature to assess the degree of fiscal pro- or countercyclicality. To account for changing policy regimes, due to, e.g., the introduction of the fiscal rule in 2001, we allow for time-varying factor loadings. Finally, to account for changes in the volatility of the structural shocks, due to, e.g., Great Moderation effects, we allow for stochastic volatility.

3.1 A structural time-varying Dynamic Factor Model

Technically, the time-varying DFM we develop relates to the model proposed in [Del Negro and Otrok \(2008\)](#).⁸ We deviate in the way we identify the latent factors, allowing us to model the dynamics of the factors as an endogenous system, and thus to identify structural shocks.

Formally, the observation and transition equations of the time-varying DFM can be written as follows:

$$y_t = z_{0,t}a_t + \dots + z_{s,t}a_{t-s} + e_t \quad (1a)$$

$$a_t = \Phi_1 a_{t-1} + \dots + \Phi_h a_{t-h} + A0_t^{-1} \Sigma_t \epsilon_t \quad (1b)$$

$$e_t = \Phi_1 e_{t-1} + \dots + \Phi_p e_{t-p} + \Upsilon_t u_t \quad (1c)$$

Equation (1a) is the observation equation, and the $N \times 1$ vector y_t represents the observables at time t . $z_{j,t}$ is a $N \times q$ matrix with dynamic factor loadings for $j = 0, 1, \dots, s$, and s denotes the number of lags used for the dynamic factors a_t .⁹ As mentioned above, we set $q = 4$ and identify two foreign factors, global activity and the real oil price; and two domestic factors, one related to the resource sector, the other to the remaining non-resource sectors.

⁸[Del Negro and Otrok \(2008\)](#) apply a time-varying DFM to analyse international business cycle synchronization. Related models have also been applied in [Aguilar and West \(2000\)](#), [Lopes and Carvalho \(2007\)](#), [Eickmeier et al. \(2011\)](#), [Liu et al. \(2014\)](#), and [Ellis et al. \(2014\)](#).

⁹In the proposed model the observables are a function of time-varying factor loadings and covariances. An alternative assumption would have been to allow instead for time variation in the parameters associated with the law of motion for the factors, as done in, e.g., [Ellis et al. \(2014\)](#) and [Eickmeier et al. \(2011\)](#). We do not follow this route. As described in online Appendix E, the factor loadings in the observation equation of the system can be estimated one equation at the time. The parameters of the law of motion for the factors must be estimated jointly. With four factors and a substantial number of lags in the transition equation, see Section 3.2, it increases the computational burden considerably, and would likely not have resulted in any meaningful estimates.

The dynamic factors follow a VAR(h) process, given by the transition equation in (1b), with $h > s$. We work with the convention that $\epsilon_t \sim i.i.d.N(0, I)$ such that we can write:

$$A0_t \Omega_t A0_t' = \Sigma_t \Sigma_t' \quad (2)$$

where Ω_t is the covariance of the error terms in the transition equation, and $A0_t$ and Σ_t is a lower triangular matrix and a diagonal matrix, respectively:

$$A0_t = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ a_{o_{21,t}} & 1 & \ddots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ a_{o_{q1,t}} & \cdots & a_{o_{qq-1,t}} & 1 \end{bmatrix} \quad \Sigma_t = \begin{bmatrix} \sigma_{1,t} & 0 & \cdots & 0 \\ 0 & \sigma_{2,t} & \ddots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & \sigma_{q,t} \end{bmatrix} \quad (3)$$

This decomposition of the covariance matrix Ω_t builds on the work of [Primiceri \(2005\)](#), and facilitates identification of the model's structural shocks, ϵ_t , and their associated time-varying volatility, captured by Σ_t . In particular, the lower triangular structure of $A0_t$ implies that we identify the structural shocks using a recursive identification scheme. The economic rationale for this choice is set out in [Section 3.3](#).

Lastly, equation (1c) describes the time series process for the $N \times 1$ vector of idiosyncratic errors e_t . We will assume these evolve as independent AR(p) processes with stochastic volatility. The parameter matrix Φ_k for $1 \leq k \leq p$ is therefore:

$$\Phi_k = \begin{bmatrix} \Phi_{1,k} & 0 & \cdots & 0 \\ 0 & \Phi_{2,k} & \ddots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & \Phi_{N,k} \end{bmatrix} \quad (4)$$

As above, denoting the covariance matrix of the heteroscedastic errors in (1c) as H_t , and under the assumption that $u_t \sim i.i.d.N(0, I)$, we have that:

$$H_t = \Upsilon_t \Upsilon_t' \quad (5)$$

where Υ_t is the diagonal matrix:

$$\Upsilon_t = \begin{bmatrix} \eta_{1,t} & 0 & \cdots & 0 \\ 0 & \eta_{2,t} & \ddots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & \eta_{N,t} \end{bmatrix} \quad (6)$$

The model's time-varying parameters and stochastic volatilities are assumed to follow random walk processes. Let $Z_t = [z_{0,t}, \dots, z_{s,t}]$ and $z_t = \text{vec}(Z_t')$ (the matrix Z_t stacked by rows) be a vector of the factor loadings at time t , ao_t the vector on non-zero and non-one elements of the matrix $A0_t$ (stacked by rows), and σ_t and η_t the vectors of diagonal elements of the matrices Σ_t and Υ_t , respectively. The dynamics of the model's time-varying parameters are consequently specified as follows:

$$z_t = z_{t-1} + w_t \quad (7a)$$

$$ao_t = ao_{t-1} + s_t \quad (7b)$$

$$h_t^\sigma = h_{t-1}^\sigma + b_t \quad (7c)$$

$$h_t^\eta = h_{t-1}^\eta + v_t \quad (7d)$$

where $h_t^\sigma = \log(\sigma_t)$ and $h_t^\eta = \log(\eta_t)$

The time-varying factor loadings capture potential changes in how the variables in the domestic economy relate to the factors across time, and, in particular, the changing nature of fiscal policy in Norway in recent decades, see Section 2. The random walk assumptions for h_t^σ and h_t^η simplify estimation of the model, as no autoregressive parameters need to be estimated, and is common in the macroeconomic literature entertaining stochastic volatility. Given that numerous studies have found evidence of time variation in elasticities in the oil market in recent decades, see, e.g., [Blanchard and Gali \(2007\)](#), [Dargay and Gately \(2010\)](#), and [Baumeister and Peersman \(2013\)](#), and given the increase in the resource industry's share of GDP in Norway from the early 1980s to today, we also allow ao_t to vary across time, implying that the contemporaneous spillovers between the factors in the model are time varying as well.

All the errors in the model are assumed to be jointly normally distributed, and we work with the following assumptions on the covariance matrix of the errors:

$$\text{var} \left(\begin{bmatrix} u_t \\ \epsilon_t \\ w_t \\ s_t \\ b_t \\ v_t \end{bmatrix} \right) = \begin{bmatrix} I_N & 0 & 0 & 0 & 0 & 0 \\ 0 & I_q & 0 & 0 & 0 & 0 \\ 0 & 0 & W & 0 & 0 & 0 \\ 0 & 0 & 0 & S & 0 & 0 \\ 0 & 0 & 0 & 0 & B & 0 \\ 0 & 0 & 0 & 0 & 0 & V \end{bmatrix} \quad (8)$$

Here, as already indicated above, I_N and I_q are identity matrices of dimension $N \times N$

and $q \times q$. W and S are assumed to be block diagonal matrices:

$$W = \begin{bmatrix} W_1 & 0 & \cdots & 0 \\ 0 & W_2 & \ddots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & W_N \end{bmatrix} \quad S = \begin{bmatrix} S_1 & 0 & \cdots & 0 \\ 0 & S_2 & \ddots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & S_{q-1} \end{bmatrix} \quad (9)$$

where W_i for $i = 1, \dots, N$ is a $m \times m$ matrix, with $m = q(s+1)$, and S_1 is a 1×1 matrix, S_2 is a 2×2 matrix, and so on.¹⁰ B is a $q \times q$ matrix, while V is a diagonal $N \times N$ matrix.

The restrictions put on (8) imply that the structural shocks driving the dynamics of transition equation, ϵ_t , are independent of the shocks driving the evolution of the time-varying factor loadings, w_t (and all other disturbances in the model). This allows us to utilize the standard SVAR machinery to analyse impulse responses and variance decompositions following ϵ_t shocks, while at the same time identifying changes in, e.g., systematic fiscal policy, captured by the w_t shocks. However, although less restrictive assumptions regarding (8) can be justified, see [Primiceri \(2005\)](#) for a broader discussion, they would increase the computational burden and complexity. As the proposed model is already heavily parametrized, it would not, in our opinion, be a feasible option in the current setting.

3.2 Model specification and data

We build on extensive out-of-sample tests in [Bjørnland and Thorsrud \(2016\)](#) and earlier findings in the oil market literature when we specify the lag orders in the baseline model. Accordingly, we allow for one lag of the dynamic factors in the observation equation (1a) of the system, i.e., $s = 1$. This is somewhat more restrictive than what was found to fit the data best in the constant parameter model of [Bjørnland and Thorsrud \(2016\)](#). On the other hand, allowing for time-varying parameters increases the potential for good model fit, and therefore also the need for many lags in the observation equation of the system. Further, as shown in, e.g., [Hamilton and Herrera \(2004\)](#) among others, a large number of lags is needed to capture the dynamics in the oil-macro relationship. For this reason we allow for up to eight lags in equation (1b) describing the law of motion of the factors, implying that $h = 8$. Finally, to capture autocorrelation in the observables not explained by the common factors we set $p = 1$ in the autoregressive processes for the idiosyncratic

¹⁰That is, S_1 is associated with $ao_{21,t}$ in (3), S_2 is associated with $ao_{31,t}$ and $ao_{32,t}$ in (3), etc.

errors. In Section 4.4 we show that our results are robust to more parsimonious model specifications.

To accommodate resource movement and spending effects within the petroleum exporting economy, as well the potential for learning spillovers between the resource sector and the rest of the economy, the observable y_t vector includes a broad range of sectoral employment, production and wage series, see also Bjørnland and Thorsrud (2016). Turning to the fiscal variables, as described in, e.g., Kaminsky et al. (2004), many indicators can be used to assess the degree of pro- or countercyclical fiscal policy. One of the advantages of the factor model methodology is that we can look at many variables at the same time, possibly allowing for more robust conclusions. For this reason we include value added, wages, and employment in the public sector from the quarterly national account statistics. From the central government's fiscal account we utilize fiscal revenues, expenditures, transfers to municipalities, and operating costs. Naturally, we also include the real exchange rate, a core variable in the Dutch disease literature. A full description of the data is given in online Appendix A.

The two variables meant to capture the developments in the international commodity market are the real price of oil and a world economic activity indicator. The real price of oil is constructed on the basis of Brent Crude oil prices (U.S. dollars), deflated using the U.S. CPI. Our main consideration when constructing the global (or world) activity indicator was to include countries whose economic activity is most likely to affect the global oil market. In addition, to capture possible direct trade linkages, we include the most important trading partners. So, for Norway, we construct global activity as the simple mean of four-quarter logarithmic changes in real GDP in Denmark, Germany, the Netherlands, Sweden, the UK, Japan, China, and the U.S.

In sum, this gives a panel of roughly 50 international and domestic data series, covering a sample period from 1981:Q1 to 2012:Q4. Correcting for the number of lags imposed on the model leaves us with 124 observations which we have used for estimation, covering the sample 1983:Q1-2012:Q4.¹¹ To capture the economic fluctuations of interest, we transform all variables to year-on-year growth. Lastly, we remove the local mean (of the

¹¹The sample periods reflect the longest possible time for which a full panel of observables is available. The vintage of quarterly national account statistics we use was generously provided to us by Statistics Norway. In the official statistics, these numbers (for employment) do not cover the earlier part of our sample. The prolonged vintage of data ends in 2012:Q4.

growth rates) and then standardize the resulting data before estimation.¹²

3.3 Identification and estimation

As is common for all factor models, the factors and factor loadings in (1) are not identified without restrictions. To separately identify the factors and the loadings, and to be able to provide an economic interpretation of the factors, we enforce the following identification restrictions on $z_{0,t}$ in (1a):

$$z_{0,t} = \begin{bmatrix} \tilde{z}_{0,t} \\ \hat{z}_{0,t} \end{bmatrix}, \quad \text{for } t = 0, 1, \dots, T \quad (10)$$

where $\tilde{z}_{0,t}$ is a $q \times q$ identity matrix for all t , and $\hat{z}_{0,t}$ is left unrestricted. As shown in Bai and Ng (2013) and Bai and Wang (2015), these restrictions uniquely identify the dynamic factors and the loadings, but leave the VAR(h) dynamics for the factors completely unrestricted. Accordingly, the innovations to the factors, ϵ_t , can be linked to structural shocks that are implied by economic theory.

As the first two factors in the system - world activity and the real price of oil - are treated as observables, they naturally load with one on the corresponding element in the y_t vector, and we set the two first elements of the idiosyncratic errors e_t equal to 0 for all time periods t . The latent domestic factors - resource and non-resource activity - must be inferred from the data. To ensure unique identification we require the domestic resource factor to load with one on value added in the petroleum sector, and the domestic non-resource factor to load with one on total value added excluding petroleum. Note that while these restrictions identify the latent factors, the factors and the observables are generally not identical due to the influence of the idiosyncratic errors.

Formally, Bai and Ng (2013) and Bai and Wang (2015) do not show in their work that the proposed factor identification scheme applies in a time-varying parameter setting. Although we do not provide a formal proof, we conjecture that it does. Like in Del Negro and Otrok (2008), one could, for example, envision rescaling the factor loadings $z_{0,t}$ for all i by say c_t , $z_{0,t-1}$ could be rescaled by c_{t-1} , and so on. However, doing so would result in $\bar{a}_t = a_t/c_t$ and $\bar{z}_{0,t} = z_{0,t}c_t$, and therefore be in conflict with the dynamic processes for a_t

¹²We standardize the data to simplify the prior specification. Likewise, we remove the local mean of the series to simplify the high dimensional estimation problem. See online Appendix A for further details about the data.

and z_t given in (1b) and (7a). By estimating the model under the normalizing restrictions in (10), such conflicts are avoided. Relatedly, one could be concerned about whether the model is able to separately identify “common” shifts in the factor loadings from changes in the stochastic volatility of the factors. Such “common” shifts are ruled out from the start by assuming that the covariances W in (9) are independent across i .

Letting $a_t = [a_t^{gact}, a_t^{oil}, a_t^r, a_t^{nr}]'$, where the superscripts denote global activity (*gact*), oil price (*oil*), resource activity (*r*) and non-resource activity (*nr*), the model identifies four structural disturbances:

$$\epsilon_t = [\epsilon_t^{gact}, \epsilon_t^{oil}, \epsilon_t^r, \epsilon_t^{nr}]' \quad (11)$$

The recursive structure of $A0_t$ in (3) implies that we follow the usual assumption made by both theoretical and empirical models of the commodity market, and restrict global activity to respond to oil price disturbances with a lag. This restriction is consistent with the sluggish behaviour of global economic activity after each of the major oil price hikes in recent decades, see e.g., [Hamilton \(2009\)](#). Furthermore, we do not treat oil prices as exogenous to the rest of the global macro economy. Any unexpected news regarding global activity is assumed to affect real oil prices contemporaneously. This is consistent with recent work in the oil market literature, see, e.g., [Kilian \(2009\)](#), [Lippi and Nobili \(2012\)](#), and [Aastveit et al. \(2015\)](#). In contrast to these papers, and to keep our empirical model as parsimonious as possible, we do not explicitly identify a global oil supply shock.¹³ Turning to the domestic factors, in the very short run, disturbances originating in the Norwegian economy can not affect global activity and the real price of oil. These are plausible assumptions insofar as Norway is a small, open economy. However, both of the domestic factors respond to unexpected disturbances in global activity and the real oil price on impact. In small, open economies such as Norway’s, news regarding global activity will affect variables such as the exchange rate, the interest rate, asset prices, and consumer sentiment contemporaneously, and in consequence overall demand in the economy. Norway is also an oil exporter, and any disturbances to the real price of oil will most likely rapidly affect both the demand and supply side of the economy.

¹³However, as shown in [Kilian \(2009\)](#), and a range of subsequent papers, such supply shocks explain a trivial fraction of the total variance in the price of oil, and do not account for a large fraction of the variation in real activity either (at least during the sample covered here). Also, [Bjørnland and Thorsrud \(2016\)](#) have shown that the effect of an oil price shock on the Norwegian economy is robust to the exclusion of oil supply shocks from the model.

The restrictions suggested here are motivated by the Dutch disease theory model presented in [Bjørnland and Thorsrud \(2016\)](#). As in that study, and as argued above, the identification scheme employed is needed to correctly quantify the domestic spillovers from unexpected windfall gains and changing international business cycle conditions in a small and open resource-rich economy. However, in contrast to the [Bjørnland and Thorsrud \(2016\)](#) study, the domestic shocks and their spillovers are not in focus here and we do not discuss them other than to emphasize that by including the ϵ_t^r shock we are able to control for the fact that the domestic economy, and the public sector in particular, might respond differently to a windfall gain due to an activity shock in the resource sector and a windfall gain due to higher oil prices, see the discussion in [Section 3](#). We note, however, that all observable variables in the model, apart from the ones used to identify the factors, are likely to respond to all shocks on impact inasmuch as they are contemporaneously related to the factors through the unrestricted part of the loading matrix (i.e., the $\hat{z}_{0,t}$ matrix in [equation \(10\)](#)). The recursive structure is therefore only applied to identify the shocks.

Unlike in the theoretical monetary policy literature, there does not exist a common “Taylor rule” principle for modeling the systematic part of fiscal authorities response function. However, like in the monetary policy literature, see, e.g., [Bernanke et al. \(2005\)](#) and [Baumeister et al. \(2013\)](#), the DFM framework allows us to empirically approximate the systematic response of fiscal policy to exogenous changes in the commodity market. To highlight our primary object of interest, let $y_{i,t+k}^{G,\cdot}$ denote the impulse response of a fiscal policy variable at horizon $t+k$. We are then interested in analysing

$$y_{i,t+k}^{G,qact} = Z_{i,t} f_{t+k}(\epsilon_t^{qact}) \text{ and } y_{i,t+k}^{G,oil} = Z_{i,t} f_{t+k}(\epsilon_t^{oil}) \quad (12)$$

where $Z_{i,t}$ is the i 'th row of the factor loadings matrix at time t , $f_{t+k}(\epsilon_t^i)$ is the impulse response function implied by the transition equation of the system (the SVAR), and ϵ_t^{qact} and ϵ_t^{oil} are respectively the structural orthogonal global activity and oil prices shocks, as defined above. The systematic policy part, the $Z_{i,t}$'s, are allowed to change through time as described by [equation \(7a\)](#) above. We can then compare the responses in the fiscal policy variables over time (say, before and after implementing the fiscal policy rule), due to either the structural oil price shock, or due to an innovation in global activity, controlling for changing volatility in all shocks.¹⁴

Finally, we estimate the time-varying DFM using Bayesian estimation, decomposing

¹⁴Note here that when computing the impulse response functions at any point in time we assume that all

the problem of drawing from the joint posterior of the parameters of interest into a set of much simpler ones using Gibbs simulations. Gibbs simulations are a particular variant of Markov Chain Monte Carlo (MCMC) methods which samples a high dimensional joint posterior by drawing from a set of lower dimensional conditional posteriors. In short, the sampling algorithm involves sequentially drawing the model’s unobserved state variables, a_t , e_t , z_t , ao_t , h_t^σ , and h_t^η , and hyper-parameters, Φ , Φ , W , S , B , and V , utilizing 7 blocks until convergence is achieved. In essence, each block involves exploiting the state space nature of the model using the Kalman filter and the simulation smoother suggested by [Carter and Kohn \(1994\)](#). A detailed description of the Gibbs algorithm is given in online Appendix E.

To generate the posterior draws, the simulations in this paper are all based on 20000 iterations of the Gibbs sampler. The first 14000 are discarded and only every sixth of the remaining are used for inference. The estimated factors are reported in Figure 7 in online Appendix B.1. As shown in online Appendix C, the convergence checks seem satisfactory. In online Appendix D we describe the priors used for the initial state variables a_0 , z_0 , ao_0 , h_0^σ , and h_0^η , and for the hyper-parameters Φ , S , B , W , Φ and V . In the appendix we also report various sensitivity analyses, showing that our main results are robust to a set of alternative assumptions regarding the prior specifications.

4 Oil price shocks and systematic fiscal policy

In the following we examine the estimated responses to a set of fiscal and macroeconomic variables from the oil price and the global activity shocks. Our aim is to analyze the response of fiscal policy over time, and in so doing, examine to what extent fiscal policy has contributed to *insulate* the domestic economy from the effects of the oil price and the global activity shocks, or, conversely, to *exacerbate* those effects. To organize the discussion, we examine in particular examine whether we can observe changes in the response patterns prior to and after the introduction of the fiscal rule in 2001.

A complicating factor is the use in the literature of different measures of fiscal policy shocks equal 0. This does not account for slow moving drift in the parameters during the impulse response horizons. As we are mainly focusing on short to medium horizon impulse responses, we do not believe that this is problematic. Likewise, our interest centers on comparing fiscal policy responses across time, and not the long-run responses implied by the time-varying parameter specification.

icy to gauge the degree of pro- or countercyclical fiscal policy,¹⁵ which makes it difficult to compare results across studies. In the following we define fiscal policy as procyclical (countercyclical) if public value added, public wages, public employment, government spending, government operating costs, or transfers increase (decrease) following a shock that increase the real price of oil. We first examine the impulse responses in the *level* of the variables (as in Pieschacon (2012)) and then *relative* to GDP (as Céspedes and Velasco (2014) and others analyse). Finally, using the data from the central government accounts we define the primary balance as income (non-oil tax revenues) minus spending. A procyclical (countercyclical) fiscal policy implies that the primary balance responds negatively (positively) to positive to a shock that increases oil prices, i.e, increasing (decreasing) spending ahead of income. In sum, these definitions are comparable to the largely similar usage in Kaminsky et al. (2004).

4.1 The Great Moderation, Recession and the global oil market

We start by examining the impulse responses in Figure 2 of the two oil market shocks on world activity (upper row) and the real price of oil (lower row) across time, before taking a more detailed look at the responses in the fiscal variables. Note that in so doing, we control for changing volatility over time.¹⁶ In the figure we display the responses for three different time periods: early (1983); intermediate (1997); and late (2012) in the sample. As will be discussed in more detail below, these dates are chosen to reflect periods when

¹⁵For instance, while Pieschacon (2012) analyses impulse responses in government purchases and transfers to an exogenous oil price shock, Céspedes and Velasco (2014) estimate the effect of a change in commodity price on government expenditures relative to GDP. Some do not control for shocks at all, but simply compare the fiscal impulse as a percentage of GDP relative to, say, the change in the output gap, see e.g. Lopez-Murphy and Villafuerte (2010) and Takáts (2012) among many others.

¹⁶As shown in Figure 8 in online Appendix B.2, the volatility of the structural oil price and global activity shocks have indeed varied considerable over the sample. There is a marked decline in the volatility of the global activity shock during the 1980s and 1990s, with a subsequent pick up at the end of the sample. These “facts” are well known and commonly attributed to the Great Moderation and the Great Recession. The structural oil price shock also shows evidence of declining volatility in periods (i.e., during the 1990s), but with marked spikes of heightened volatility, in the early 1990s (the first Gulf War) and during the Great Recession. Similar patterns have also been reported in Baumeister and Peersman (2013), who explain the decline by a fall in oil supply elasticity.

fiscal policy was governed by different fiscal regimes. Details for other time periods can be obtained on request.

The results confirm that our identified shocks are in line with the results found in the oil market literature. After an unexpected one standard deviation increase in global activity, the real price of oil rises substantially on impact, reflecting that the real price of oil is not exogenous to the macro economy, see, e.g., [Kilian \(2009\)](#), [Lippi and Nobili \(2012\)](#). Moreover, after a one standard deviation shock to the real price of oil, world activity falls, although with a lag. This is consistent with the fact that it takes time before the higher production costs associated with the higher oil price work their way through to actual output, see, e.g., [Hamilton \(2009\)](#).

Note also the differences in the response path of world activity and the real oil price to shocks over time. The differences in impact responses reflect the changes in volatility of the structural shocks, as already mentioned. However, we also observe some changes in the response path that relate to the changes in the overall covariance structure of the oil market, see Section [3.1](#) and equation (3). In particular, a world activity shock has stronger impact on the oil price at the end of the sample (2012) than it has in the earlier part of the sample (1983). This is consistent with studies documenting the important role of global demand as a driver for the real price of oil over the past decade, see, [Aastveit et al. \(2015\)](#). For the oil price shock, the changing effects on world activity across time are minor, with the middle 1990s displaying slightly fewer volatile oil price shocks, and subsequently also a milder downturn in the world economy.

4.2 Procyclical or countercyclical fiscal policy?

We now discuss the fiscal responses to the oil price and the global activity shocks, both normalized to increase oil prices. First, [Figure 3](#) compares the evolution of the responses of some key variables in the public sector (value added, wages, employment and spending) to an oil price shock that increases the real price of oil. In each row, we first graph impulse responses for three specific periods in time: 1983, 1997, and 2012. The dates are chosen to reflect three comparable periods: the initial discovery period during which spending increased rapidly; the period just after the GPF started to generate some revenue (fiscal policy was yet to be governed by a rule, and was intended to remain neutral over the business cycle); and 10 years after the adoption of the fiscal rule. The two subsequent

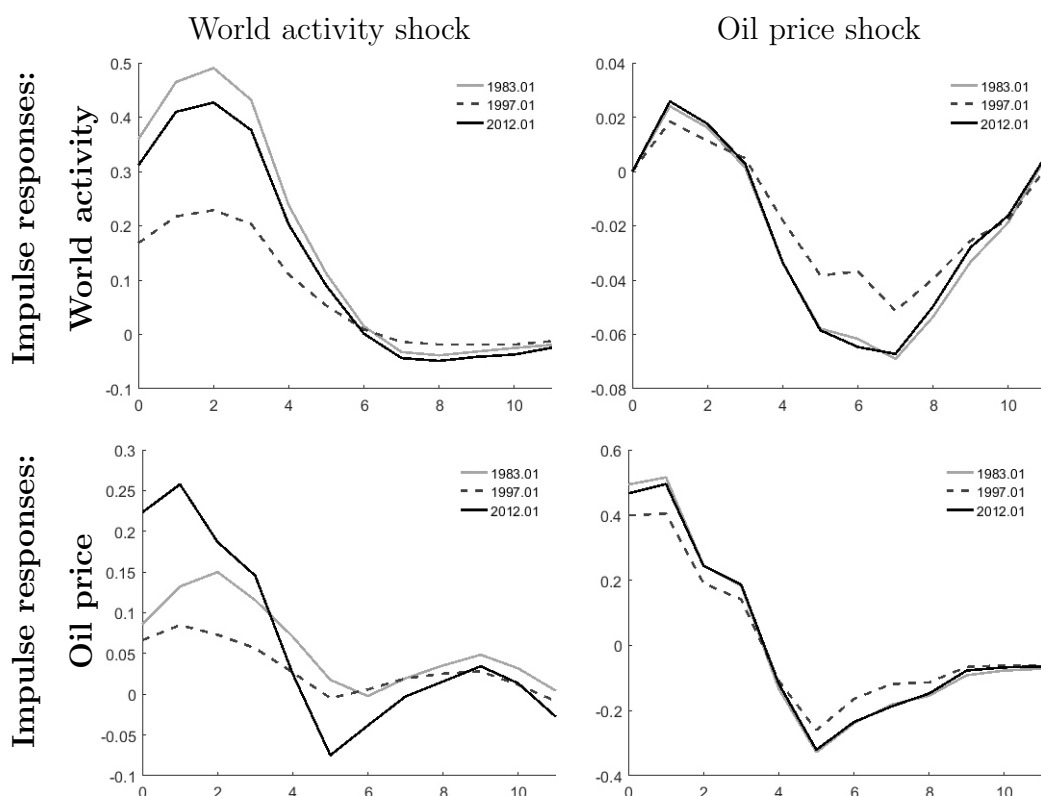


Figure 2. The figure reports the estimated impulse responses at three different periods of time. The initial shock corresponds to a one standard deviation innovation (of the normalized data). All responses are reported in levels (of the normalized data).

graphs offer more detail on the time-varying responses after the first and fourth quarters, but now measured over the whole sample, and with 68 percent posterior probability bands.¹⁷ Figure 4 displays similar responses, but now due to a positive global activity shock.¹⁸

Starting with the oil price shock, i.e., an increase in the oil price that is not due to increased global activity, a few results stand out. First, fiscal policy responds procyclically to the oil price shock over the sample, even more so after the 2001 adoption of the fiscal framework. In particular, the positive effects of an oil price shock on value added, real wages, employment, and spending in the public sector are more pronounced today than in the decade preceding the rule, and for value added and spending, also more procyclical than in the 1980s, see Figure 3.

¹⁷We report 68 percent posterior probability bands as the overall posterior uncertainty is large due to the high dimensionality of the model.

¹⁸Note that from now on we normalize the oil price and global activity shocks such that we compare similarly sized innovations across time (see the discussion in Section 4.1).

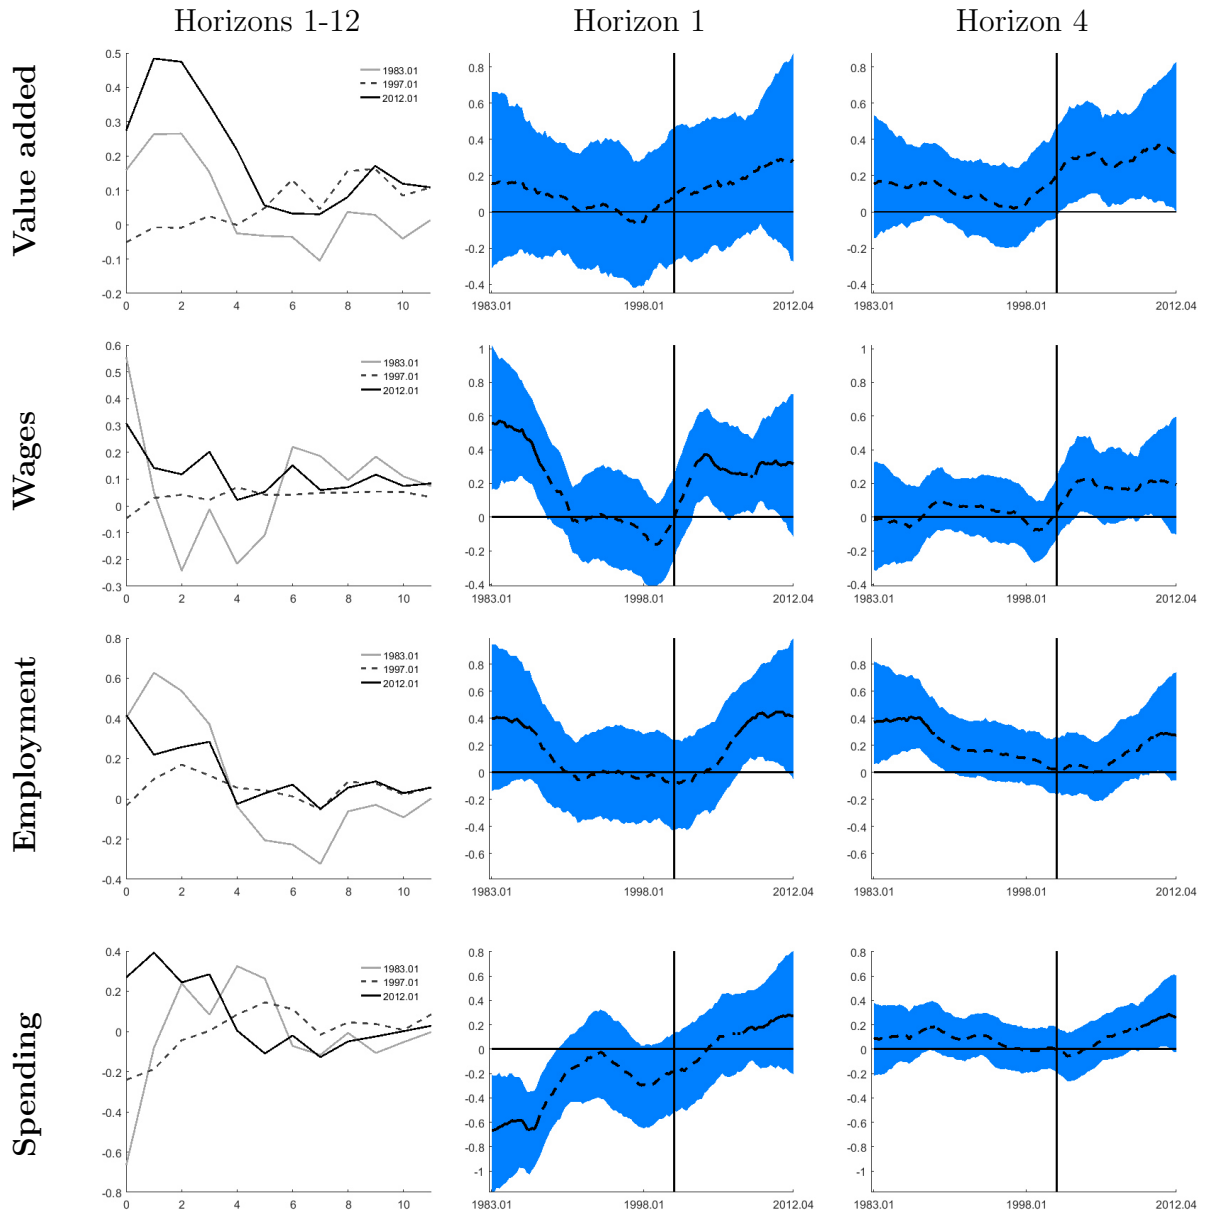


Figure 3. Oil price shock and the public sector: Time-varying responses. The first column displays estimated impulse responses at three different periods of time. The initial shock is normalized to one percent (of the normalized data). All responses are reported in levels (of the normalized data). The subsequent two columns report a snapshot of the responses across the whole sample for two specific response horizons. The color shadings represent 68 percent posterior probability bands. The black line is the median estimate. The line is solid (dotted) whenever the median estimate is outside (inside) the 68 percent area in 2001:Q1. Finally, we plot a vertical line in 2001:Q1 to indicate the introduction of the fiscal rule.

Turning to the global activity shock that spurred a rise in oil prices, c.f. Figure 2, the picture is somewhat more nuanced, with some components of public sector (value added, spending and wages) shifting from a clear countercyclical to more acyclical pattern during

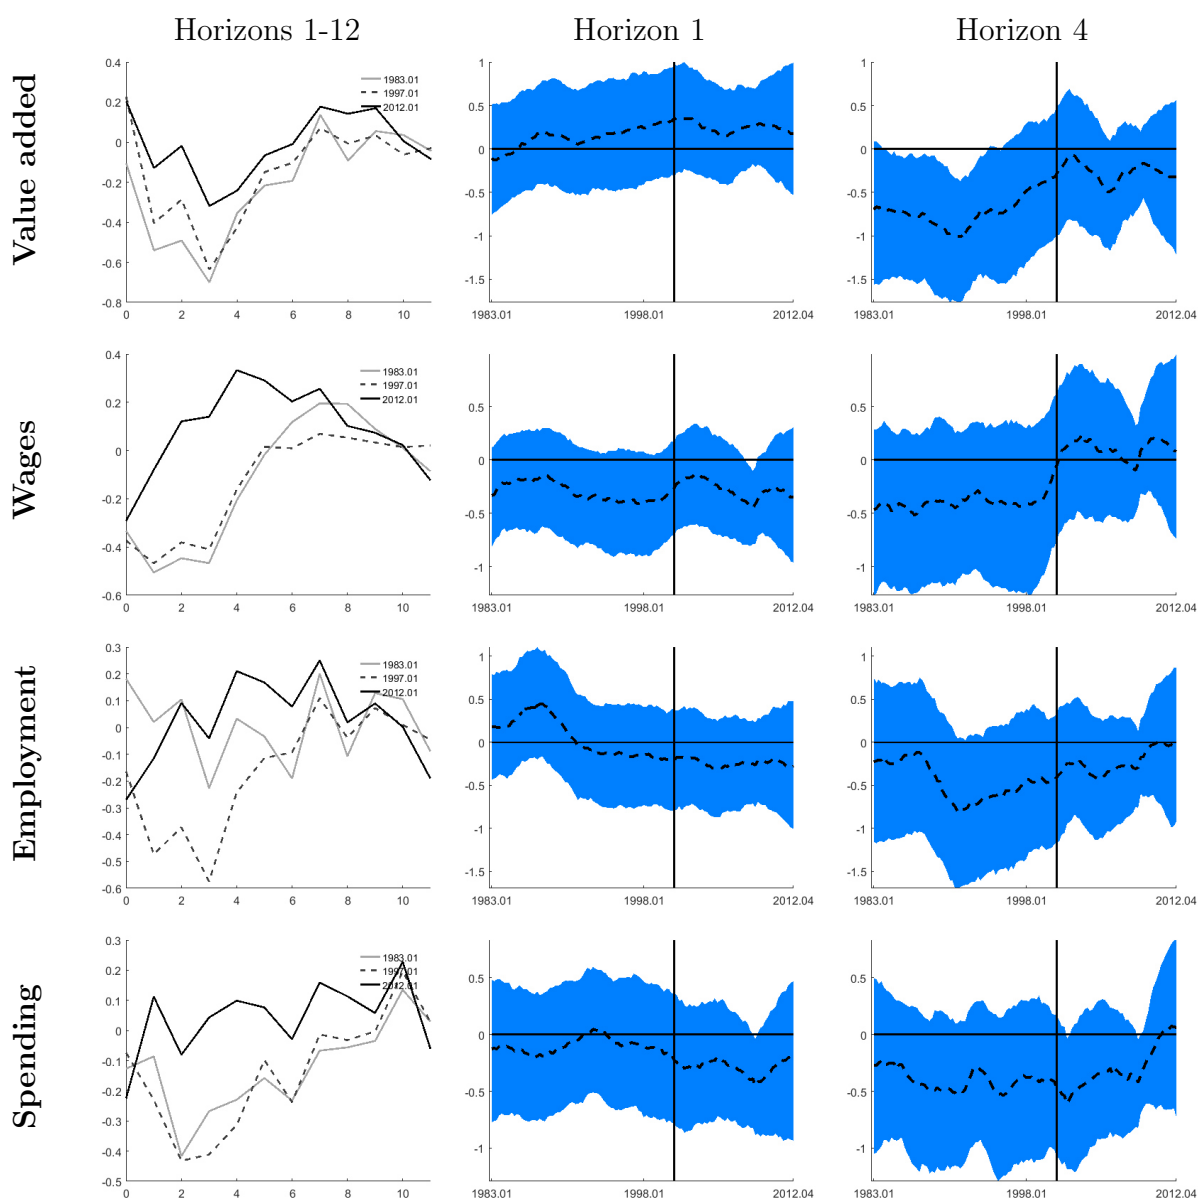


Figure 4. Global activity shock and the public sector: Time-varying responses. See Figure 3.

the last decade, see Figure 4. The takeaway is still that of a tendency for more (not less) fiscal procyclicality since the adoption of the spending rule. This suggests that following a global downturn accompanied by a contraction of the domestic economy, there is less room for fiscal policy maneuvering. Figures 9 and 10, in online Appendix B.3, provide more details on the procyclicality of fiscal policy by analysing the effect on additional public sector variables. The graphs show that the increase in fiscal procyclicality was particularly pronounced for spending (excluding pensions) and transfers to municipalities, while the increase in administrative expenses to oil price shocks has been more muted over the

sample.¹⁹

Summarizing, we find clear evidence of nonlinearities in the responses of fiscal policy to the oil price shocks in particular, and to some extent to the global activity shocks. In particular, public spending, employment, and wages have responded significantly more (not less) procyclically to these shocks since the adoption of the fiscal policy framework.²⁰

4.3 Transmission of shocks and the domestic economy

According to Pieschacon (2012), the fiscal framework adopted by Norway shields the economy from oil price fluctuations, implying only minor responses in domestic variables following an oil price shock. She further claims that had Norway been more like Mexico, without the fiscal policy framework in place, Norway would have had a larger share of variance in the domestic variables attributed to the oil price shock.²¹ Our focus here is not to compare Norway with Mexico, but to examine the extent to which the adoption of the fiscal rule lessened exposure to oil price shocks.²² This is done in Figure 5. The figure shows the contribution of the different shocks in the model to the variance in public value added and domestic GDP (measured as the average over all domestic industries except the public sector), over time. We focus on the role of oil price and global activity shocks for response horizons 1, 4, and 8.

Two features stand out. First, there is a marked difference in the role played by oil price shocks in explaining activity in the public sector since 2001. One year after the shock occurs (horizon 4), more than 40 percent of the variance in public value added is explained by oil price shocks by the end of the sample (2012), compared to 10 percent prior to the adoption of the fiscal rule (1998/1999). This pattern holds true for all public variables (results can be obtained on request) and clearly emphasizes the increased role of oil price shocks in fiscal policy since 2001. Global activity shocks, on the other hand, do

¹⁹While we have chosen to focus on the fiscal variables here, results for all variables can be obtained at request. However, as can be shown, the time variation observed in the fiscal variables is not due to time variation in variables such as the exchange rate.

²⁰In Figure 11 in online Appendix B.4, we show that the fiscal policy response to the two types of shocks are significantly different, at least for some horizons, and in some periods.

²¹The arguments are based on a counterfactual experiment, in which Norway and Mexico change parameters, but otherwise face their original shocks.

²²Clearly, there may be many reasons why Norway is less exposed to oil price shocks than Mexico, including, for instance, less corruption, more efficient bureaucracy, less mature oil sector etc.

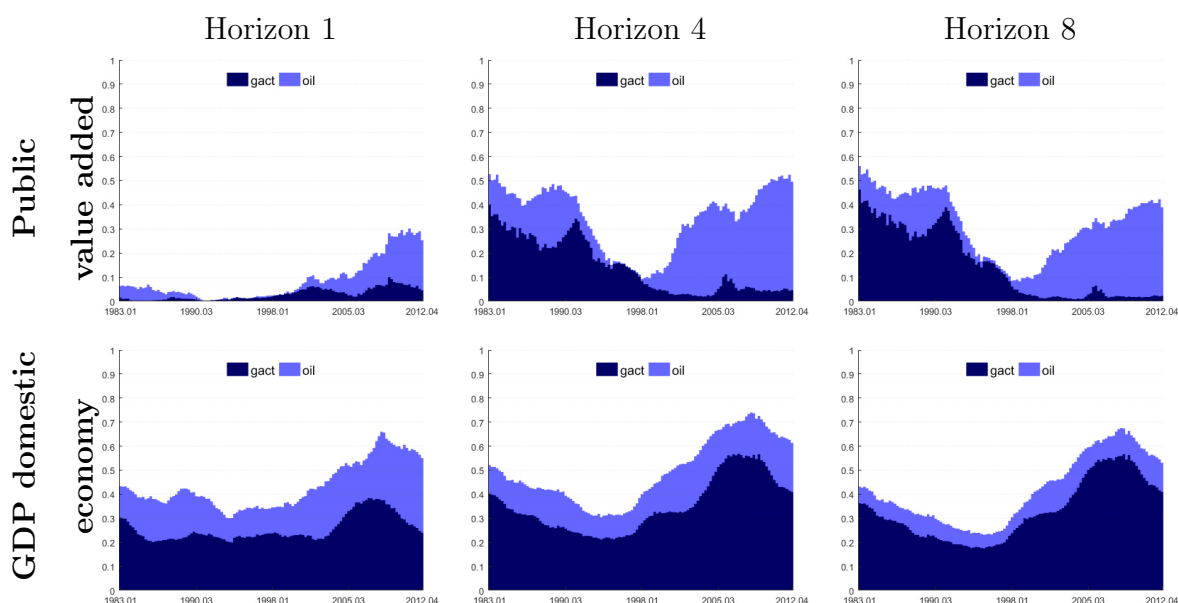


Figure 5. Time-varying variance decompositions. The plots report the median of the estimated variance decompositions associated with the levels response of the variables to the following shocks; gact = global activity shock; oil = oil price shock.

not explain much of the variation in the fiscal variables over the past decade, suggesting a more countercyclical or acyclical fiscal policy with regard to the global shocks.²³ It is noteworthy that our findings encompass those in Pieschacon (2012), although we reach opposite conclusions. That is, based on an estimation period from 1986 to 2005 she finds that approximately 10 percent of government purchases could be explained by oil price shocks after two years, which is not very different from what we also observe on average for the same period, see Figure 5. However, and as seen above, starting in 2001, the pattern changes markedly, suggesting increased exposure of the public sector to the oil price shocks.

Second, and turning to GDP in the domestic economy, we find that the oil price shocks explain more than 10 percent of the variance in the domestic variables at horizon 4. This is more than twice as much as is being explained in Pieschacon (2012). We also find an increase in the share explained by oil price shocks since 2001, albeit not to the extent of the public sector. Finally, the role played by the global activity shocks in explaining GDP increases throughout the sample, in particular during the last decade.

Impulse responses for non-oil GDP to an oil price shock and a global activity shocks are

²³For some public variables, i.e., public wages and spending, global activity shocks explain slightly more of the variance in the last decade, but always less than the variance explained by the oil price shock.

reported in Figures 12 and 13, in online Appendix B.5, respectively. In the same figures, we also report the responses in manufacturing production, emphasizing also the (non-oil) tradeable sector in the economy. As seen there, both GDP and manufacturing production respond slightly positively to an oil price shock. Furthermore, while the responses have increased slightly since the spending rule was adopted, suggesting that having a rule may have benefited the (non-oil) tradeable sector, by the end of the sample the positive effect is no longer significant. After a world activity shock, however, the response is clearly significant and positive, although the effect has declined by the end of the sample.

As the above discussion shows there is evidence of nonlinearity also in the response of the domestic economy to shocks that increase oil prices, suggesting a stronger pass-through of oil related shocks to the economy after the 2001 adoption of the fiscal rule. In short, then, the fiscal framework does not effectively shield the economy from oil price fluctuations. If anything, fiscal policy has exacerbated the effects of oil price shocks on the domestic economy, even more so in fact after the adoption of the fiscal rule. If Norway has a more muted response to oil price shocks than countries like Mexico, as argued in Pieschacon (2012), it must be for other reasons than the fiscal rule.

Having said that, one can easily argue that if the private sector is also stimulated by the oil price and global activity shocks, as indicated by the results in Figures 12 and 13, maybe the stimulus to the public sector is simply following the increase in the domestic economy. Some studies, i.e., Céspedes and Velasco (2014) and Husain et al. (2008), estimate the effect of a change in commodity prices on government expenditures relative to GDP, and find that measured in relative terms, fiscal policy has been countercyclical. Figure 6 addresses this issue, as well as highlights the importance of separating between the shocks driving the oil market and their domestic implications. In particular, the figure reports the response, across time and horizons, of value added, wages and employment in the public sector relative to the response in the domestic economy. A value above zero indicates the public sector responds more positively to the given shock than the private sector. The last row in the figure reports the effect on the primary balance.²⁴

We find that for a given oil price shock, the public sector has clearly grown at the expense of the private sector. This again suggests that fiscal policy exacerbates the effect

²⁴To enhance comparison across the graphs the global activity and real oil price shocks are normalized to 1 and 10 percent, respectively.

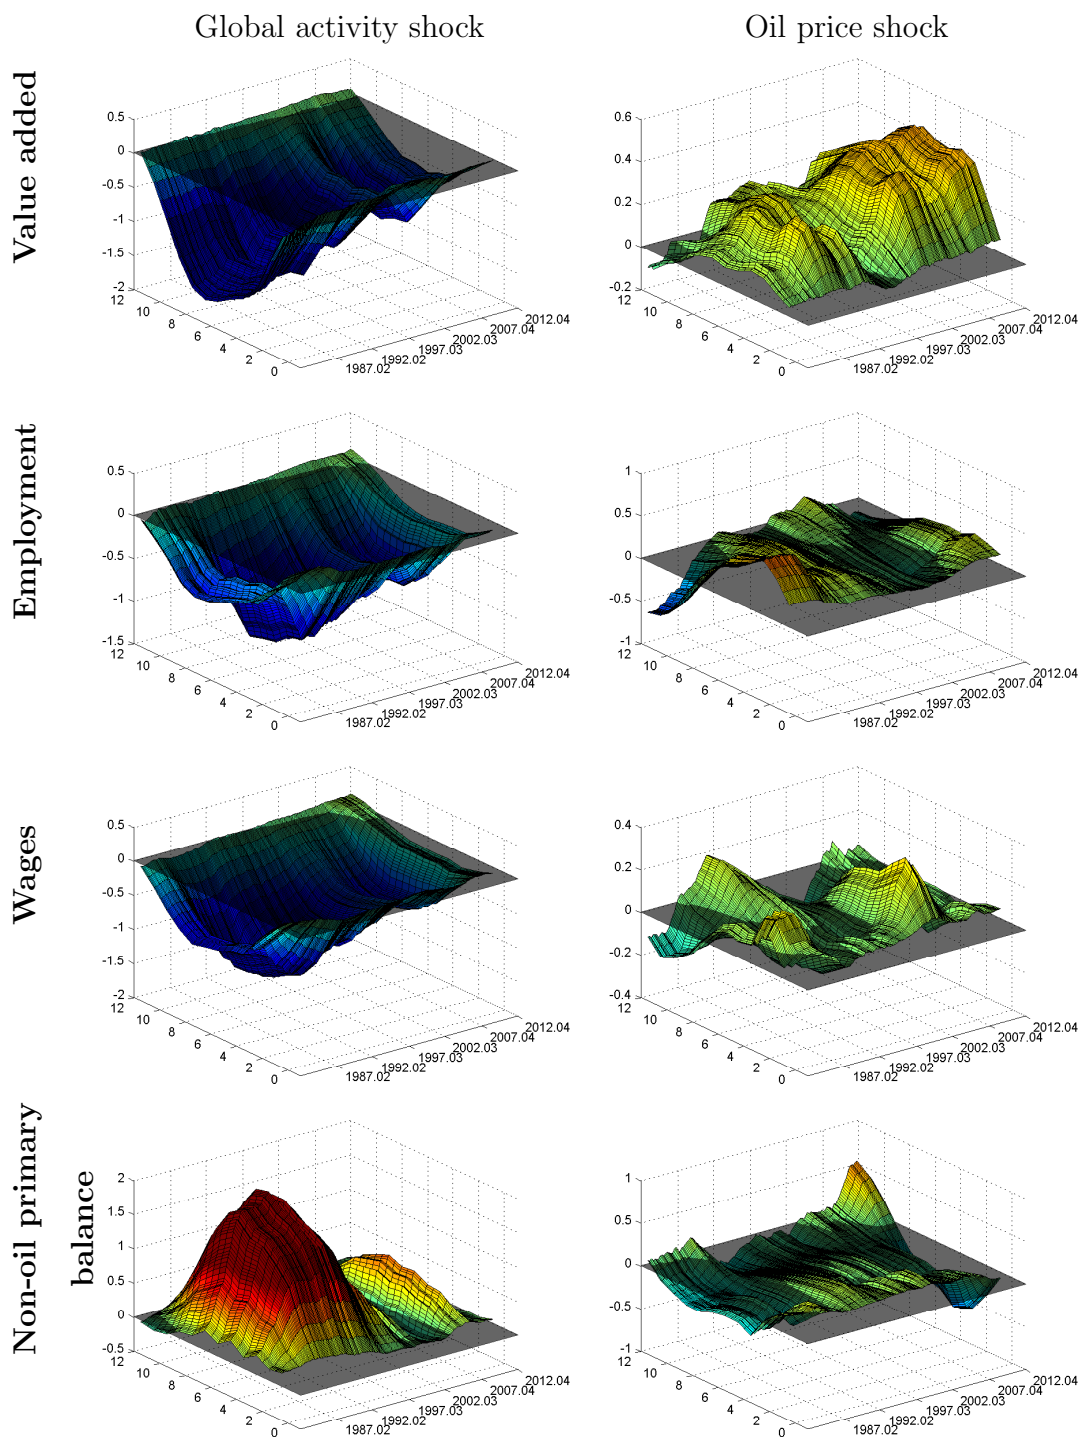


Figure 6. Public sector relative to the domestic economy and the primary balance. Each plot reports the response, across time (x-axis) and horizons (y-axis), of an outcome variable in the public sector relative to the response in the domestic economy. Here, the domestic economy is defined as the average response across all sectors, except the public sector. The initial shock is normalized to 1 percent. All responses are reported in levels. A value above zero indicates a more positive response by the public sector to the given shock than by the mainland economy as a whole. For readability the relative responses are also smoothed by applying a 3 quarter moving average transformation. See the text for the definition of the primary balance.

of the oil price shocks on the domestic economy. This is quite different to what [Pieschacon \(2012\)](#) and others have asserted. There it is argued that fiscal policy regulates the size of the pass-through.

Turning to global activity shocks, the results are reversed. The positive effect on GDP is stronger than that on the public sector; as seen by the negative effect in the figures. Hence, when the oil price rises due to global demand, the direct spillovers to the domestic economy are substantial, much stronger than the pass-through via government spending. Yet, we note that, since the start of the millennium, the spillovers to the private sector (relative to the public sector) have diminished, in line with the more procyclical fiscal policy responses reported above.

The transmission of shocks via the primary balance confirms our claim. Following an oil price shock that generates procyclical fiscal responses, we would expect to see negative numbers for the primary balance, all else being equal. We see evidence of this in the lower right frame of [Figure 6](#). However, all else is not equal, as emphasized above. In particular, the oil price shocks are also transmitted to the domestic economy, implying increased tax-receipts and an improvement in the primary balance: hence the more muted response in [Figure 6](#). For the global activity shock, however, we should observe positive numbers for the primary balance, since the stimulus to the domestic economy (and subsequent tax-receipts) is more substantial now than the effect via increased public spending. This is confirmed in the lower left frame in [Figure 6](#). However, note also the steep decline in the primary balance from 2000/2001, consistent with the findings of a more procyclical fiscal policy in recent times.

We therefore conclude that studies suggesting a countercyclical fiscal policy response, as a share of GDP or based on the primary balance, in the recent boom, should attribute it to global activity shocks and their domestic propagation, rather than the adopted fiscal framework.

4.4 Interpretations and extensions

After reading our conclusion, it would be reasonable to ask why the fiscal rule has induced, against the government's intentions (c.f. [Section 2](#)), a procyclical fiscal policy with regard to oil price fluctuations? The answer is not that the fiscal rule has been violated. In fact, the fiscal authorities have managed in large part to actually follow the rule by using only

roughly 4 percent of the Fund every year, see e.g. [Ministry of Finance \(2015\)](#). Instead, we believe the answer has to do with the design of the fiscal framework and bad luck in timing. When the rule was promulgated, in 2001, the Fund's market value amounted to roughly 20 percent of Norwegian GDP. Going forward, the fiscal authorities assumed the price of oil would remain more or less unchanged, at 200 NOK per barrel. This expectation turned out to be wrong. During the 10-year period from 2001 to 2011, the price of oil increased considerably, to over 600 NOK per barrel. Accordingly, the inflow of money to the Fund was much higher than expected, and by 2013 its market value exceeded 180 percent of Norwegian GDP. Thus, almost by construction, it has been difficult to restrain the close to automatic increase in spending that has followed from taking out a close to constant fraction (4 percent) of a Fund that, for long periods, has been highly correlated with the oil price.²⁵

Figure 7 illustrates this point, and shows the close connection between the growth in the price of oil and the revisions of the fiscal rule from 2001 to 2014. Over this period, public spending implied by the 4 percent rule has been revised upwards in line with growth in the oil price. According to our results, the discretionary deviations from the 4 percent rule that may have taken place during the recent commodity price boom have simply not been large enough to counteract the changes in structural policy parameters induced by the introduction of the rule and the associated increase in spending potential from the higher oil prices.

Our narrative about fiscal policy differs from earlier studies in the literature because we allow for time-varying parameters, and because we allow for two different oil market shocks. Additional exercises provide some further intuition, as well as a check for robustness. First, regarding the role of shocks, from the late 1990s the unconditional correlation between oil prices and public spending gradually increase, (c.f. Figure 1), suggesting a more procyclical fiscal policy. During the financial crisis, however, the positive unconditional correlation fell (again, c.f. Figure 1), suggesting a more countercyclical fiscal policy. However, as the financial crisis is a period of negative global demand shocks (reducing both oil prices and global activity), the countercyclical behaviour in this period is most likely due to the changing composition of oil market shocks, not to changing policy. Hence,

²⁵That developments in the Fund have been highly correlated with the oil price in the past ten years has also been emphasized by [Norges-Bank \(2012\)](#).




Figure 7. The price of oil and spending revisions. Spending revisions refer to the difference between the assumed path for the fiscal rule in the budget year 2014 (in percentage of GDP) minus the assumed path in the original budget in 2001 (in percentage of GDP).

even investigating simple correlations, we see that the composition of shocks matters.

Second, the results documented in this analysis are conditional on, but not driven by, the proposed time-varying DFM. In fact, when we estimate a constant parameter DFM model over two different samples, pre and post 2001, we find clear evidence of increased procyclicality post 2001 for both the structural oil price shocks and the global activity shocks, see Figure 14 in online Appendix B.6. Still, using a simple split sample framework is sensitive to the subjectively chosen break date, and may under/over estimate the true coefficients if there is variation within each sub-sample, as the analysis in Section 4.2 clearly suggests. To capture such behavior, one needs a flexible model allowing for time-varying parameters, such as the one used in our analysis.

Third, to further illustrate significance of results, we graph in Figure 15 in online Appendix B.6 impulse response distributions using the time-varying DFM at two different points in time: in 2012:Q1 (After fiscal rule) and in 1999:Q1 (Before fiscal rule). Thus, the figure is comparable to Figure 14. Again we find evidence of significant increased procyclicality, but now only following the oil price shocks. However, this is not very surprising, as posterior uncertainty is large due to the high dimensionality of the time-varying parameter specification.

Relatedly, the fact that we find a more procyclical fiscal policy after the introduction of the fiscal rule is not driven by changes in how global activity responds to oil price shocks or vice versa. In particular, when we normalize the oil price (world activity) shock to have the same effect on global activity (oil prices) throughout the sample, we confirm

the results we have already obtained, see Figure 16 in online Appendix B.7.

Fourth, as Figure 17 in online Appendix B.8 shows, changing the dynamic specification used in the baseline model does not change our qualitative conclusions. That is, estimating the time-varying DFM with $s = 0$ (no lags of the dynamic factors in the observation equation of the system), $h = 4$ (four lags in the transition equation of the system), or both $s = 0$ and $h = 4$, result in the same conclusion regarding procyclical fiscal policy after 2001.

Finally, the assumption that global economic activity does not react on impact to oil price shocks could be debated. To address this concern we identify the shocks with sign restrictions as a robustness check. Specifically, exogenous oil price shocks are identified as shocks that have the opposite (sign) contemporaneous impact on oil prices and global economic activity, whereas global activity shocks are shocks that have the same (sign) impact on both variables (see e.g. Peersman (2005)). Such an identification strategy does allow for an immediate effect of oil price shocks on global economic activity. Doing so, we find that our earlier results are robust, see Figures 18 and 19 in Appendix B.9.

5 Conclusion

This paper questions whether the adoption of fiscal spending rules insulates resource-rich economies from oil price fluctuations. In pursuing the question we develop a time-varying Dynamic Factor Model, in which both the volatility of structural shocks and the systematic fiscal policy responses are allowed to change over time. We focus on Norway. Unlike most oil exporters, Norway has established a sovereign wealth fund operated as a hybrid between a savings- and stabilization fund, and a fiscal rule designed specifically to shield the domestic economy from oil price fluctuations. We find that, contrary to common perceptions, fiscal policy has been more (not less) procyclical with oil price fluctuations since the adoption of the fiscal rule. That is to say, fiscal policy has not effectively sheltered the economy from oil price shocks. However, following a global activity shock that also increases oil prices, the picture is more nuanced, with some components of public spending being countercyclical relative to GDP. We suggest that studies that find a countercyclical fiscal policy response in the recent boom, should put it down to global activity shocks and their domestic propagation, rather than the adopted fiscal framework.

Still, although the fiscal rule has not managed to shelter the Norwegian economy from oil price fluctuations, the goal of saving resource revenue for future usage has been accomplished. The fiscal authorities have in large managed to actually follow the rule, and by only using roughly 4 percent of the Fund every year the Norwegian sovereign wealth fund is today the largest in the world. From a policy point of view, the implications of our findings are therefore of general interest since they highlight the strengths and weaknesses of the fiscal framework adopted in a resource rich economy whose handling of its resource wealth has been described as exemplary.

References

- Aastveit, K. A., H. C. Bjørnland, and L. A. Thorsrud (2015). What drives oil prices? Emerging versus developed economies. *Journal of Applied Econometrics* 30(7), 1013–28.
- Aguilar, O. and M. West (2000, July). Bayesian dynamic factor models and portfolio allocation. *Journal of Business & Economic Statistics* 18(3), 338–57.
- Alexeev, M. and R. Conrad (2009). The elusive curse of oil. *The Review of Economics and Statistics* 91(3), 586–598.
- Allcott, H. and D. Keniston (2018). Dutch disease or agglomeration? The local economic effects of natural resource booms in modern America. *The Review of Economic Studies* 85(2), 695–731.
- Bai, J. and S. Ng (2013). Principal components estimation and identification of static factors. *Journal of Econometrics* 176(1), 18 – 29.
- Bai, J. and P. Wang (2015). Identification and bayesian estimation of dynamic factor models. *Journal of Business & Economic Statistics* 33(2), 221–240.
- Barro, R. J. (1979). On the determination of the public debt. *The Journal of Political Economy*, 940–971.
- Barsky, R. B. and L. Kilian (2002). Do we really know that oil caused the great stagflation? A monetary alternative. In B. S. Bernanke and K. Rogoff (Eds.), *NBER Macroeconomics Annual 2001*, pp. 137–183. MIT Press: Cambridge, MA.

- Baumeister, C., P. Liu, and H. Mumtaz (2013). Changes in the effects of monetary policy on disaggregate price dynamics. *Journal of Economic Dynamics and Control* 37(3), 543–560.
- Baumeister, C. and G. Peersman (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.
- Bernanke, B., J. Boivin, and P. S. Elias (2005). Measuring the effects of monetary policy: A Factor-augmented Vector Autoregressive (FAVAR) approach. *The Quarterly Journal of Economics* 120(1), 387–422.
- Bjørnland, H. C. and L. A. Thorsrud (2016). Boom or gloom? Examining the Dutch disease in two-speed economies. *The Economic Journal* 126(598), 2219–2256.
- Blanchard, O. J. and J. Gali (2007). The macroeconomic effects of oil price shocks: Why are the 2000s so different from the 1970s? In *International Dimensions of Monetary Policy*, NBER Chapters, pp. 373–421. National Bureau of Economic Research, Inc.
- Carter, C. K. and R. Kohn (1994). On Gibbs Sampling for State Space Models. *Biometrika* 81(3), 541–553.
- Céspedes, L. F. and A. Velasco (2014). Was this time different?: Fiscal policy in commodity republics. *Journal of Development Economics* 106, 92–106.
- Charnavoki, V. and J. Dolado (2014). The Effects of Global Shocks on Small Commodity-Exporting Economies: Lessons from Canada. *American Economic Journal: Macroeconomics* 6(2), 207–237.
- Cogley, T. and T. J. Sargent (2005, April). Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S. *Review of Economic Dynamics* 8(2), 262–302.
- Dargay, J. M. and D. Gately (2010). World oil demand’s shift toward faster growing and less price-responsive products and regions. *Energy Policy* 38(10), 6261–6277.
- Del Negro, M. and C. Otrok (2008). Dynamic factor models with time-varying parameters: measuring changes in international business cycles. Staff Reports 326, Federal Reserve Bank of New York.

- Eickmeier, S., W. Lemke, and M. Marcellino (2011). The changing international transmission of financial shocks: Evidence from a classical time-varying FAVAR. CEPR Discussion Papers 8341.
- Ellis, C., H. Mumtaz, and P. Zabczyk (2014). What Lies Beneath? A Time-varying FAVAR Model for the UK Transmission Mechanism. *Economic Journal* 0(576), 668–699.
- Hamilton, J. D. (1983). Oil and the Macroeconomy since World War II. *Journal of Political Economy* 91(2), 228–48.
- Hamilton, J. D. (2003). What is an oil shock? *Journal of Econometrics* 113(2), 363–398.
- Hamilton, J. D. (2009). Causes and consequences of the oil shock of 2007–08. *Brookings Papers on Economic Activity* 40(1), 215–283.
- Hamilton, J. D. and A. M. Herrera (2004). Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy: Comment. *Journal of Money, Credit and Banking* 36(2), 265–86.
- Husain, A., K. Tazhibayeva, and A. Ter-Martirosyan (2008). Fiscal policy and economic cycles in oil-exporting countries. IMF Working Papers 253.
- Johnson-Calari, J. and M. Rietveld (2007). *Sovereign wealth management*. Central Banking Publications.
- Kaminsky, G. L., C. M. Reinhart, and C. A. Vegh (2004, September). When it Rains, it Pours: Procyclical Capital Flows and Macroeconomic Policies. NBER Working Papers 10780, National Bureau of Economic Research, Inc.
- 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 B. Hicks (2013). Did unexpectedly strong economic growth cause the oil price shock of 2003–2008? *Journal of Forecasting* 32(5), 385–394.
- Kilian, L. and D. Murphy (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics* 29(3), 454–478.

- Lie, E. (2013). Bukken og pengesekken: Om folketrygdfondet og oljefondets tilblivelse. *Nytt Norsk Tidsskrift* (4), 323–335.
- Lippi, F. and A. Nobili (2012). Oil and the macroeconomy: A quantitative structural analysis. *Journal of the European Economic Association* 10(5), 1059–1083.
- Liu, P., H. Mumtaz, and A. Theophilopoulou (2014). The transmission of international shocks to the UK. Estimates based on a time-varying factor augmented VAR. *Journal of International Money and Finance* 46(C), 1–15.
- Lopes, H. F. and C. M. Carvalho (2007). Factor Stochastic Volatility with Time Varying Loadings and Markov Switching Regimes. *Journal of Statistical Planning and Inference* (137), 3082–3091.
- Lopez-Murphy, P. and M. Villafuerte (2010). Fiscal policy in oil producing countries during the recent oil price cycle. IMF Working Papers 28.
- Ministry of Finance (2001). Report no. 29 to the Storting (2000-2001). White paper, Norway.
- Ministry of Finance (2015). Norwegian public report (NOU) 2015:9.
- Nakov, A. and A. Pescatori (2010, 03). Oil and the Great Moderation. *Economic Journal* 120(543), 131–156.
- Norges-Bank (2012). Økonomiske perspektiver. Sentralbanksjef Øystein Olsen foredrag til Norges Banks representantskap og inviterte gjester, torsdag 16. februar 2012.
- OECD (2005). OECD economic surveys: Norway. OECD Economic Surveys 2005/17.
- OECD (2007). OECD economic surveys: Norway. OECD Economic Surveys 2007/2.
- Peersman, G. (2005). What caused the early millennium slowdown? evidence based on vector autoregressions. *Journal of Applied Econometrics* 20(2), 185–207.
- Peersman, G. and I. Van Robays (2012). Cross-country differences in the effects of oil shocks. *Energy Economics* 34(5), 1532–1547.
- Pieschacon, A. (2012). The value of fiscal discipline for oil-exporting countries. *Journal of Monetary Economics* 59(3), 250–268.

- Portes, J. and S. Wren-Lewis (2014). Issues in the design of fiscal policy rules. CFM discussion paper series 12, Centre For Macroeconomics.
- Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies* 72(3), 821–852.
- Schaechter, A., T. Kinda, N. T. Budina, and A. Weber (2012). Fiscal rules in response to the crisis-toward the 'next-generation' rules: A new dataset. IMF working paper 187/12.
- Smith, B. (2014). Dutch Disease and the Oil and Boom and Bust. OxCarre Working Papers 133, Oxford Centre for the Analysis of Resource Rich Economies, University of Oxford.
- Takáts, E. (2012). Countercyclical policies in emerging markets. *BIS Quarterly Review*, 25–31.
- Velculescu, D. (2008). Norway's oil fund shows the way for wealth funds. *IMF Survey Magazine* 9.

Appendices

Appendix A Data

Table 2 summarizes the data entertained, their sources, and the transformations used. As described briefly in Section 3.2 of the main paper, we also remove the local mean (of the transformed data) and then standardize the resulting data before estimation. The local mean adjustment is done prior to the standardization to control for low frequent movements in the growth rates (changes in the mean) across time, see, e.g., [Stock and Watson \(2012\)](#). We have experimented with different methods of doing the local mean adjustment. In the benchmark case we simply subtract a deterministic linear time trend from the transformed data. In two alternatives we estimate the local mean as the average of the transformed data over a centered moving window of ± 30 quarters, and as the components of the time series with fluctuations between 150 and 200 quarters, obtained from a band-pass filter. Irrespective of which method we use the resulting time series are highly similar. The band-pass and the deterministic linear trend removal approaches result in time series with a correlation coefficients well above 0.95 for most series. Finally, some of the series also inhabit clear outliers or measurement errors (e.g., *Transfers*). We define outliers as observations being outside $3 \times$ *interquartile* range, and automatically remove them by using linear interpolation.

As in the factor model proposed in [Bjørnland and Thorsrud \(2016\)](#), all data are transformed to year-on-year growth. In Norway, and with the data used here, the factor structure is much stronger using this data transformation than for example using quarter-on-quarter growth. In particular, estimating two factors using the principal component method explains only roughly 20 percent of the variation in the data when quarter-on-quarter growth is used, but 40 percent when using year-on-year growth. Still, in Figure 20, in Appendix B.10, we show that our main results are qualitatively robust to estimating the DFM using quarter-on-quarter growth.

Table 2. Data. The vintage of Norwegian data are collected from Statistics Norway (SSB). We use data from the quarterly national account and the central government fiscal account. In the official statistics, data for some of the series (e.g. for employment), do not cover the earlier part of our sample, but were generously made available to us by SSB. The international data, with the exception of the exchange rate, were sourced from the GVAR database constructed by Gang Zhang, Ambrogio Cesa Bianchi, and Alessandro Rebucci at the Inter-American Development Bank. In the column head “Empl.,” an “x” indicates that we use both value added and employment data for the variable at hand. All value added data are measured in real terms, as calculated by SSB. The following transformation codes applies: 9 = year-on-year percentage growth ($y_t = x_t/x_{t-4} \times 100 - 100$), 11 = year-on-year logarithmic difference ($y_t = \ln(x_t) - \ln(x_{t-4})$). See Section 3.2 and the text for additional details.

Source	Variable	Empl.	Trans.	Description
National Account	Res. extraction	x	11	Oil and natural gas extraction/mining
	Res. service	x	11	Service activities in oil and gas/mining
	Manufacturing	x	11	Manufacturing
	Construction	x	11	Construction
	Retail	x	11	Wholesale and retail trade
	Transp. ocean	x	11	Ocean transport
	Transportation	x	11	Transport activities excl. ocean transport
	Hotel and food	x	11	Accommodation and food service activities
	Financial	x	11	Financial and insurance activities
	Scientific	x	11	Professional, scientific and technical activities
	Business	x	11	Administrative and support service activities
	Non-resource	x	11	Total excl. oil and gas extraction/mining
	Public	x	11	General government
	Public consumption		11	General government
	Wages petroleum		11	Wages petroleum sector
	Wages public		11	Wages public sector
	Wages non-res.		11	Total excl. wages to petroleum sector
Fiscal account	Spending		11	Central government total expenditures
	Spending excl. pensions		11	Central government total expenditures excluding pensions
	Operating costs		11	Central government operating costs
	Transfers		11	Central government transfers to municipalities and county authorities
	Tax revenue		9	Tax revenue excl. petroleum
	Tax revenue petroleum		9	Tax revenue from petroleum
Int.	World activity		11	World economic activity indicator, see Section 3.2
	Oil price		11	Real price of oil, see Section 3.2
	Exchange rate		11	Bank of International Settlements (BIS) effective exchange rate index, broad basket

Appendix B Additional results

B.1 Business cycle factors

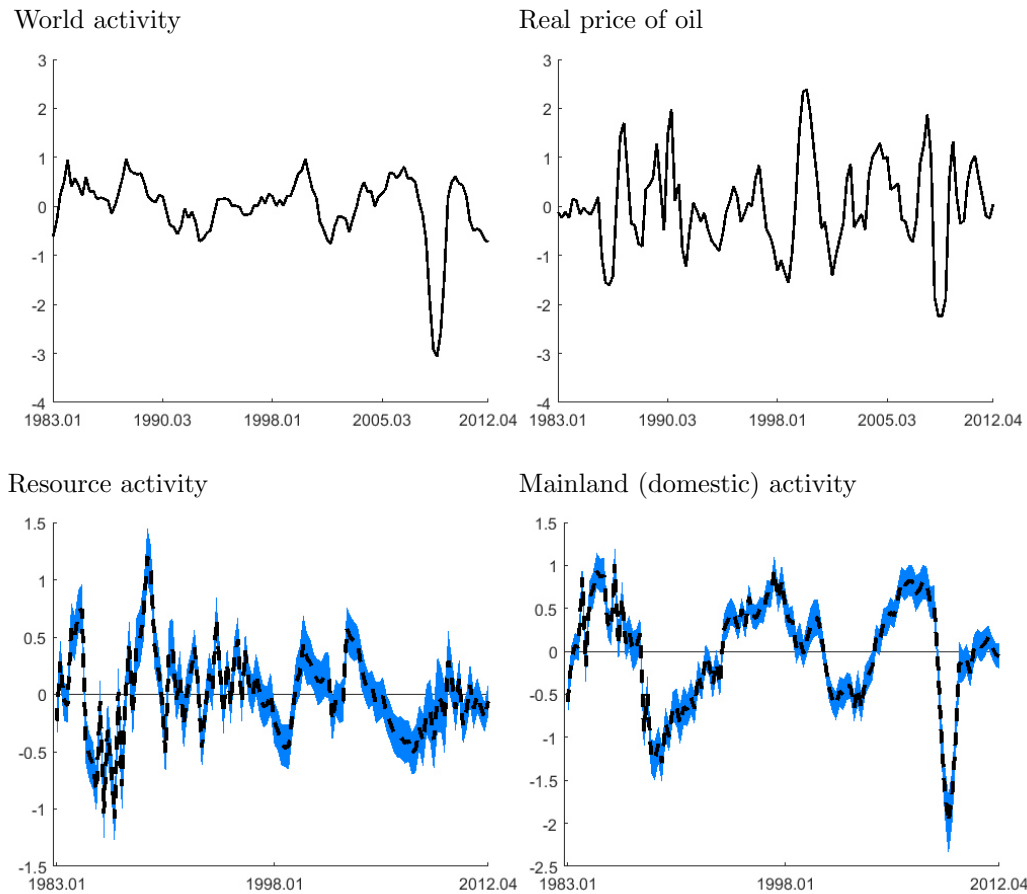


Figure 7. Observable and latent factors. The figures display the two observable factors together with the two estimated latent factors. The color shadings represent 68 percent posterior probability bands. The black line is the median estimate.

B.2 Volatility of shocks

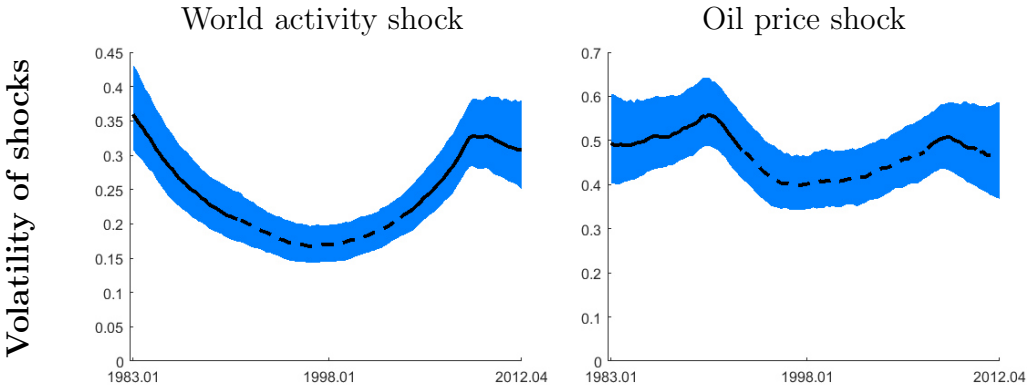


Figure 8. Time-varying volatility and oil price- and global activity shocks. The figures report the estimated standard deviation of the shocks across time. The color shadings represent 68 percent posterior probability bands. The black line is the median estimate. The line is solid (dotted) whenever the median estimate is outside (inside) the 68 percent area in 2001:Q1.

B.3 Components of public spending

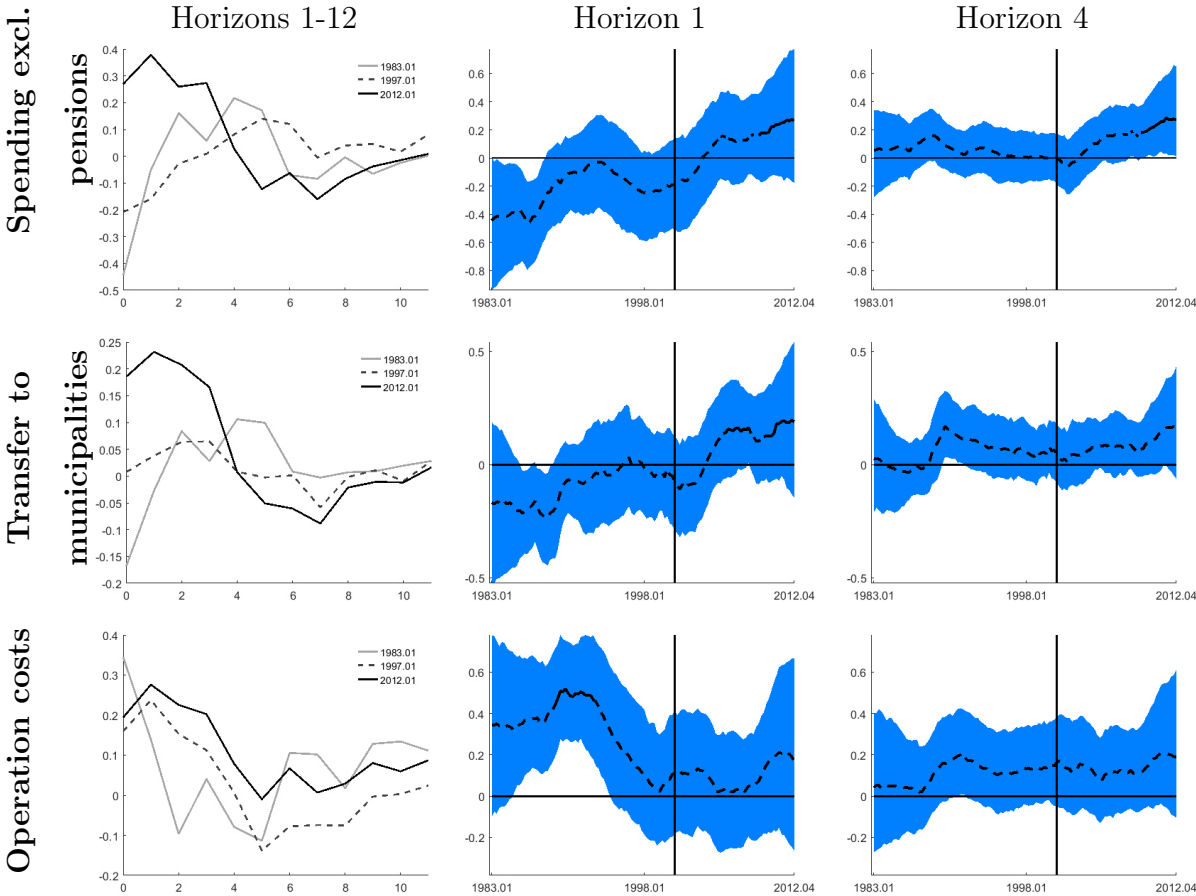


Figure 9. Oil price shock and the fiscal budget: Time-varying responses. The first column displays estimated impulse responses at three different periods of time. The initial shock is normalized to 1 percent (of the normalized data). All responses are reported in levels (of the normalized data). The subsequent two columns report a snapshot of the responses across the whole sample for two specific response horizons. The color shadings represent 68 percent posterior probability bands. The black line is the median estimate. The line is solid (dotted) whenever the median estimate is outside (inside) the 68 percent area in 2001:Q1.

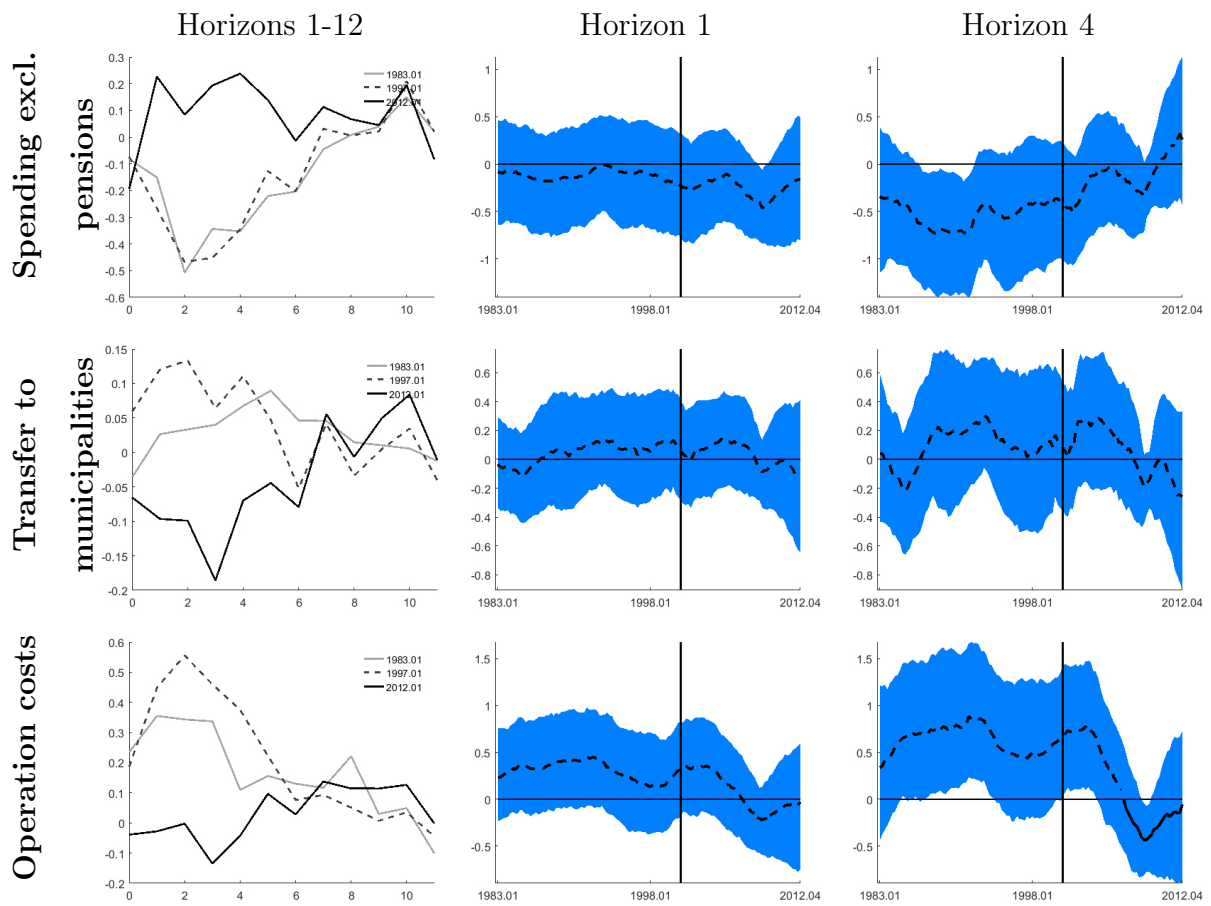


Figure 10. World activity shock and the fiscal budget: Time-varying responses. See Figure 9.

B.4 Response differences in public sector variables

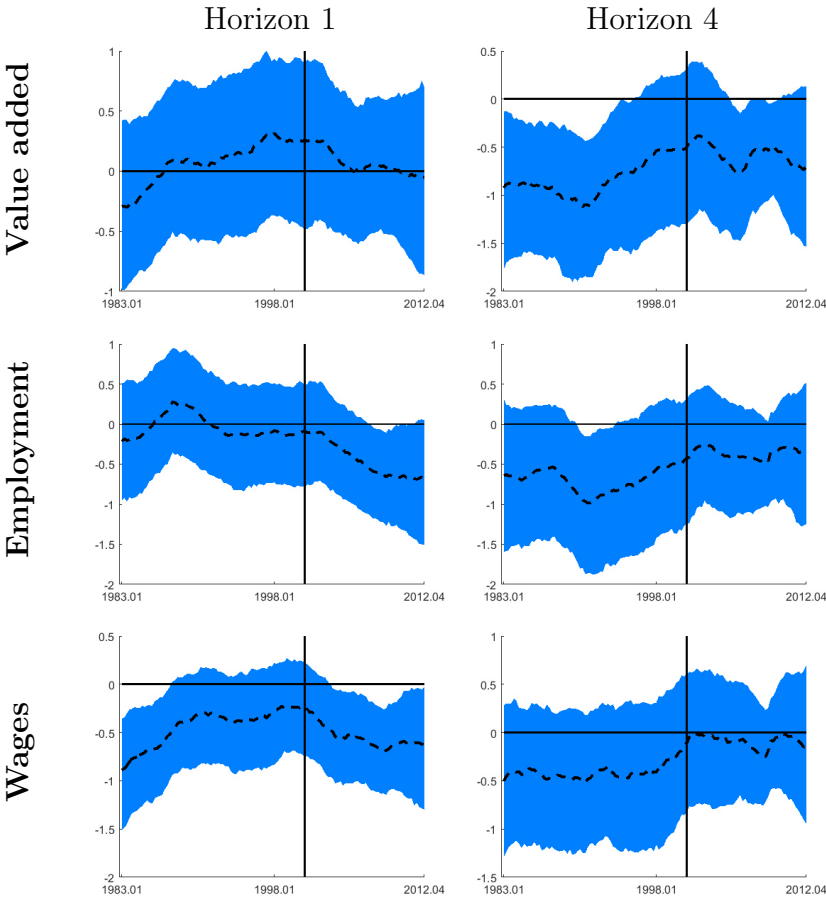


Figure 11. Response differences and public variables. The graph reports the response difference in value added, wages and employment in the public sector following a world activity relative to oil price shock (world activity - oil price) at horizon 1 (left column) or horizon 4 (right column). The color shadings represent 68 percent posterior probability bands. The dotted black line is the median difference.

B.5 Responses in the non-oil economy

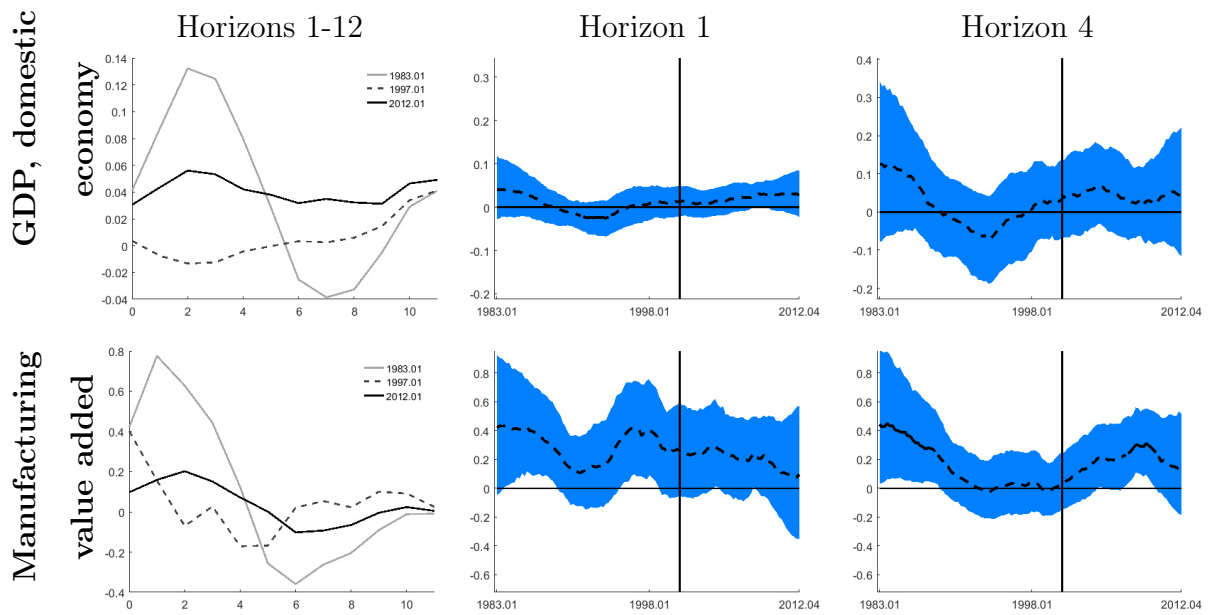


Figure 12. Oil price shock and the macroeconomy: Time-varying responses. See Figure 3.

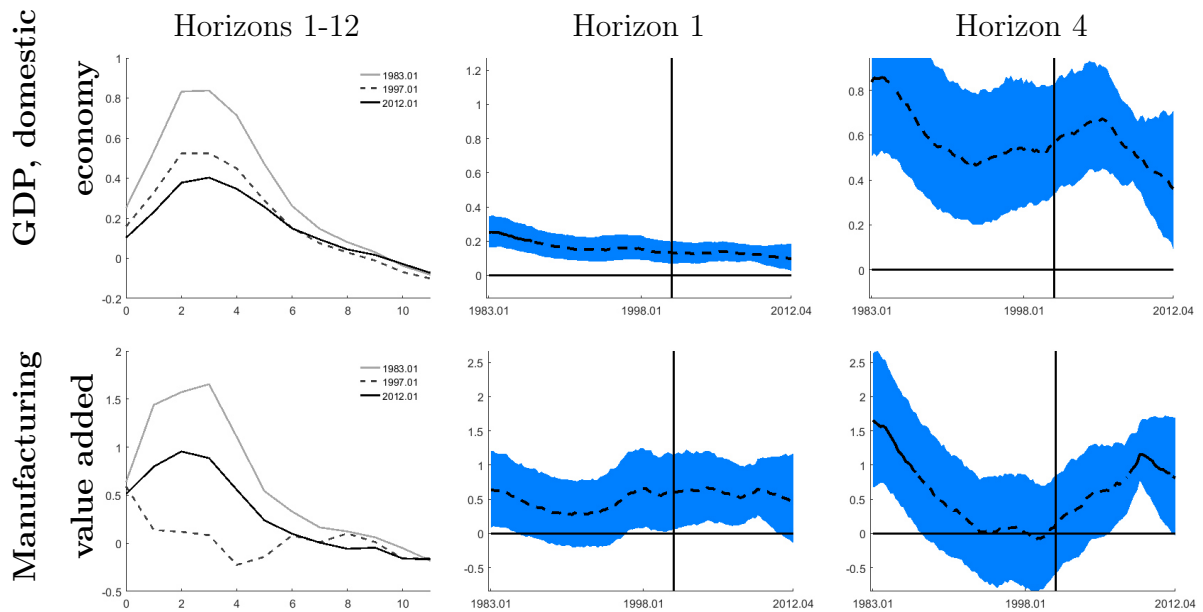


Figure 13. World activity shock and the macroeconomy: Time-varying responses. See Figure 3.

B.6 Sub-sample analysis and significance

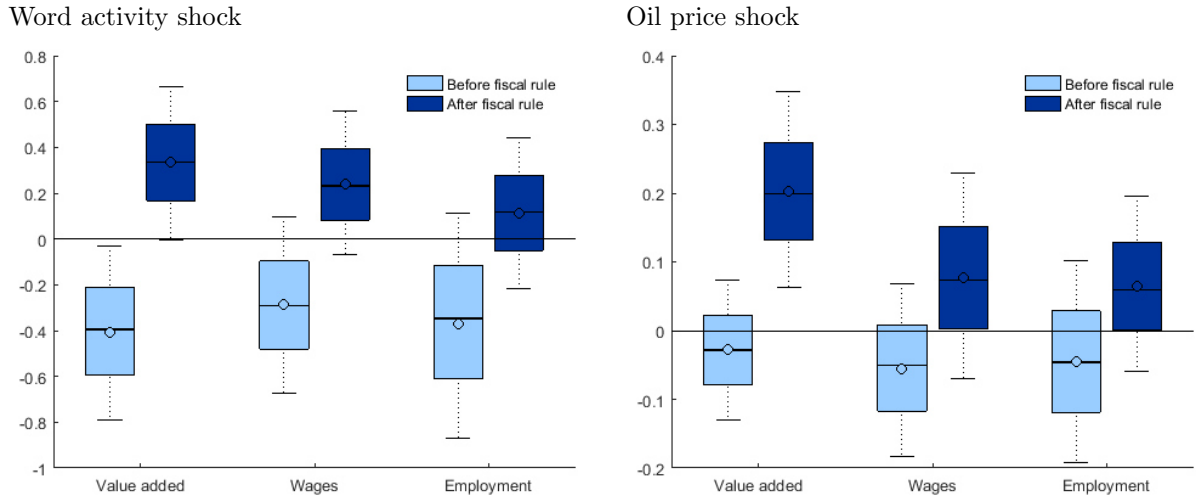


Figure 14. Box plots of impulse responses estimated using a constant parameter DFM on two sub-samples, pre and post 2001. The response horizon is 4 quarters. The width of each box reflects the interquartile range. The median (mean) estimate is displayed by a horizontal line (circle).

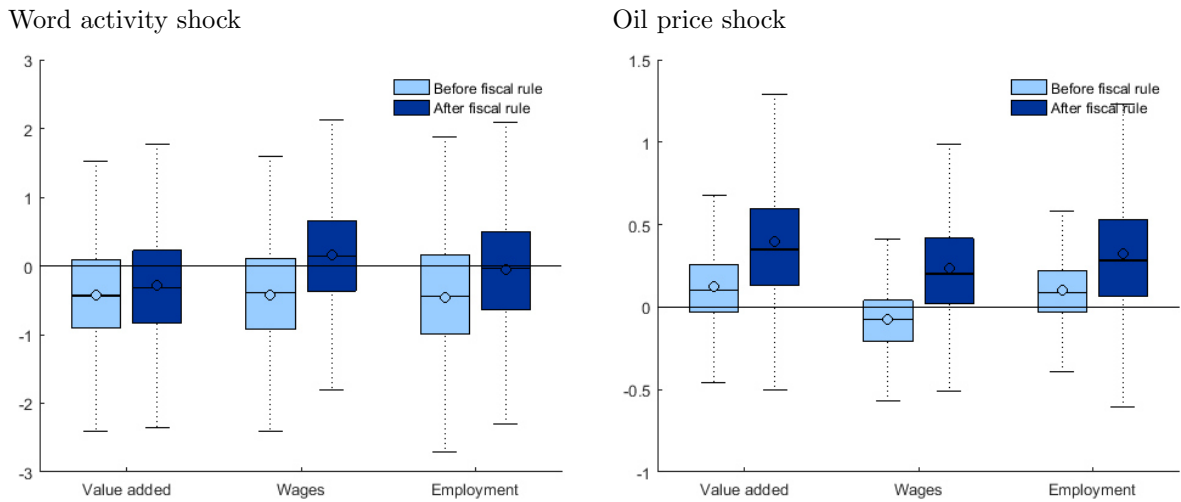


Figure 15. Box plots of impulse response distributions using the time-varying DFM at two different points in time: in 2012:Q1 (After fiscal rule) and in 1999:Q1 (Before fiscal rule). The initial shock is normalized to 1 percent (of the normalized data). The response horizon is 4 quarters. The width of each box reflects the interquartile range. The median (mean) estimate is displayed by a horizontal line (circle).

B.7 Alternative normalizations

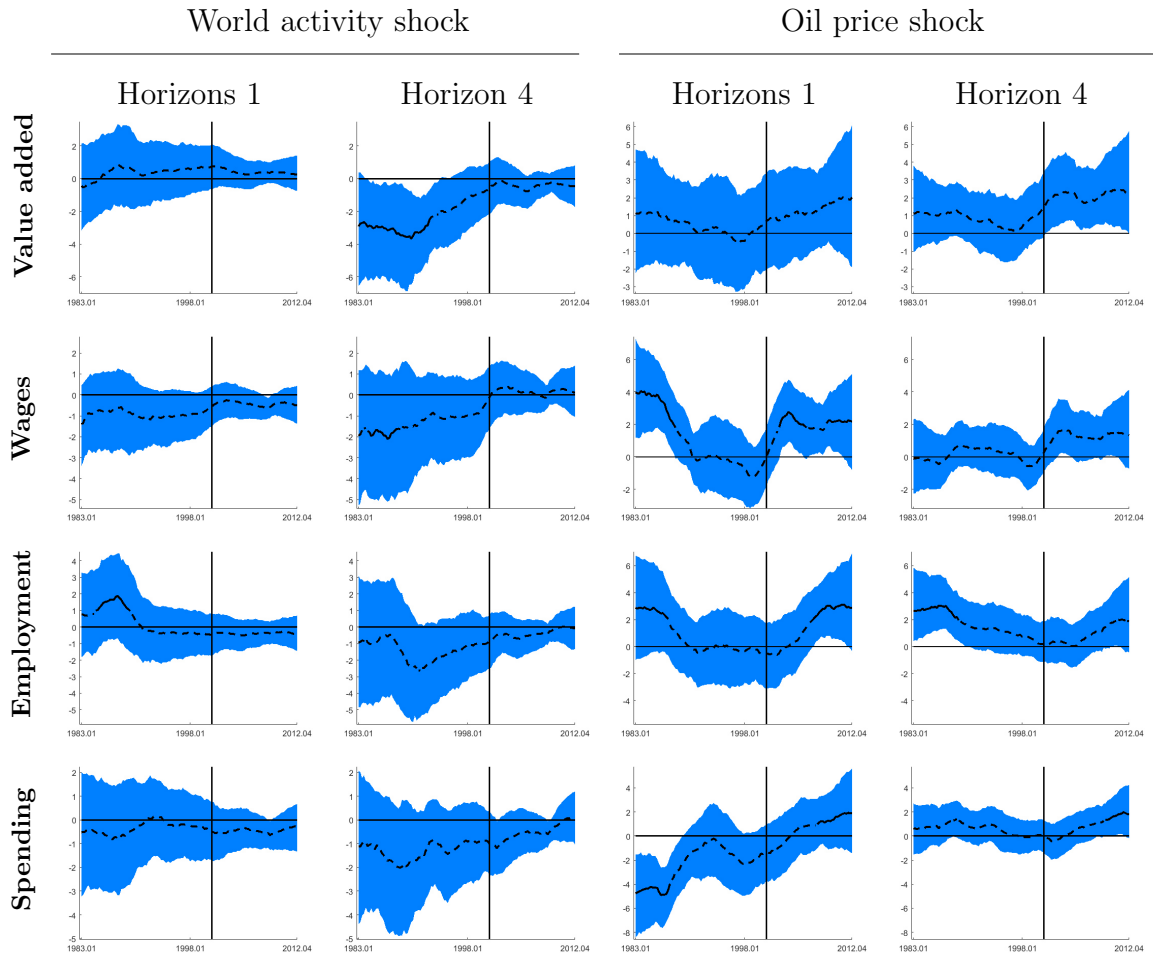


Figure 16. World activity and oil price shocks and the public sector: Time-varying responses. The first two columns display estimated impulse responses following a world activity shock. The initial shock is normalized to increase the price of oil with one percent (at each point in time). The subsequent two columns display estimated impulse responses following an oil price shock. The initial shock is normalized to decrease global activity by one percent after 8 quarters (at each point in time). All responses are reported in levels. The color shadings represent 68 percent posterior probability bands. The black line is the median estimate. The line is solid (dotted) whenever the median estimate is outside (inside) the 68 percent area in 2001:Q1. Finally, we plot a vertical line in 2001:Q1 to indicate the introduction of the fiscal rule.

B.8 Different model specifications

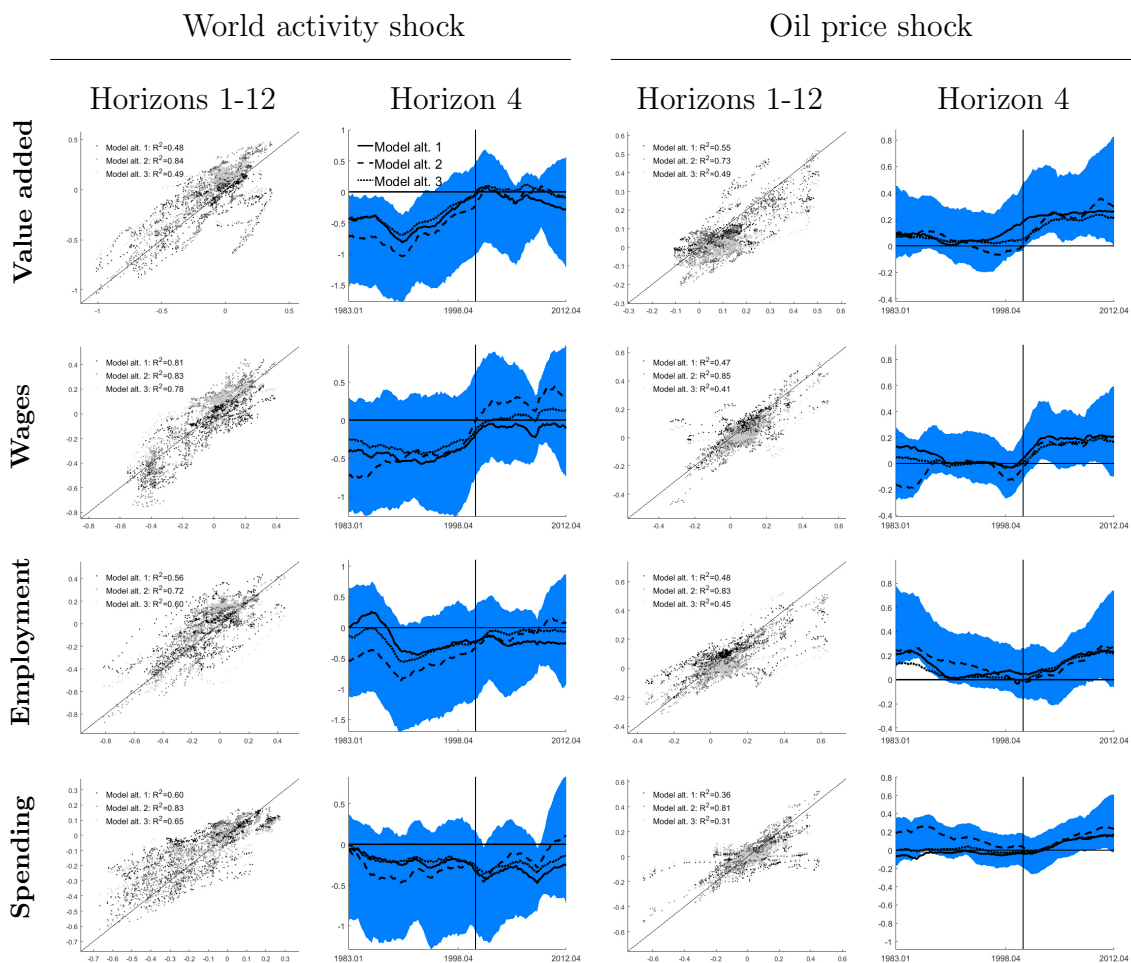


Figure 17. World activity and oil price shocks and the public sector: Scatter plots and time-varying responses for different model specifications. The initial shock is normalized to one percent (of the normalized data). All responses are reported in levels (of the normalized data). The scatter plots report time-varying responses at all horizons for different model specifications relative to the baseline model. The comparisons use the median estimate at each point in time and compares the baseline predictions (x-axis) against three alternative model specifications: Model alternative 1 is identical to the baseline model but with $s = 0$, i.e., no lags of the dynamic factors in the observation equation of the system; Model alternative 2 is identical to the baseline model but with $h = 4$, i.e., only four lags in the transition equation of the system; Model alternative 3 is identical to the baseline model but with $s = 0$ and $h = 4$. In the horizon specific impulse response plots, the color shadings represent 68 percent posterior probability bands from the baseline model. The black lines are the median estimates from three alternative model specifications. We plot a vertical line in 2001:Q1 to indicate the introduction of the fiscal rule.

B.9 Sign restrictions

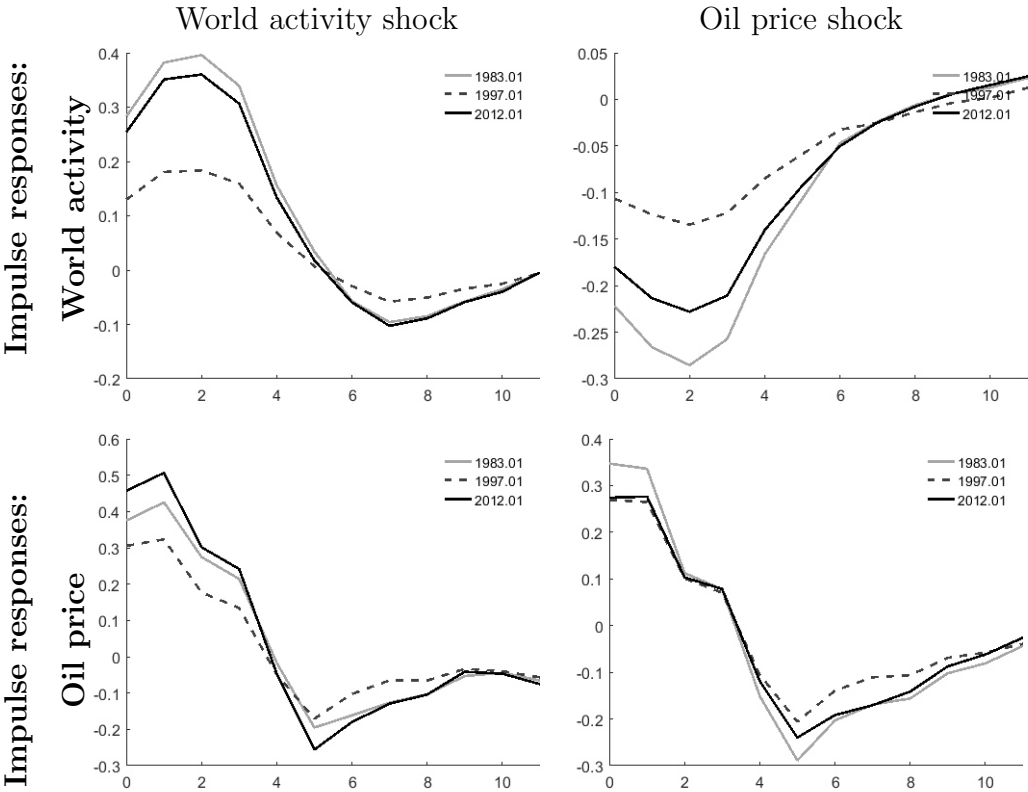


Figure 18. The figure reports the estimated impulse responses at three different periods of time. The shocks are identified with sign restrictions. That is, the world activity shock is restricted to increase world activity and the oil price on impact. The oil price shock is restricted to decrease world activity and increase the oil price on impact. The lines reflect the median estimate from the sign identified admissible sets. The initial shock corresponds to a one standard deviation innovation (of the normalized data).

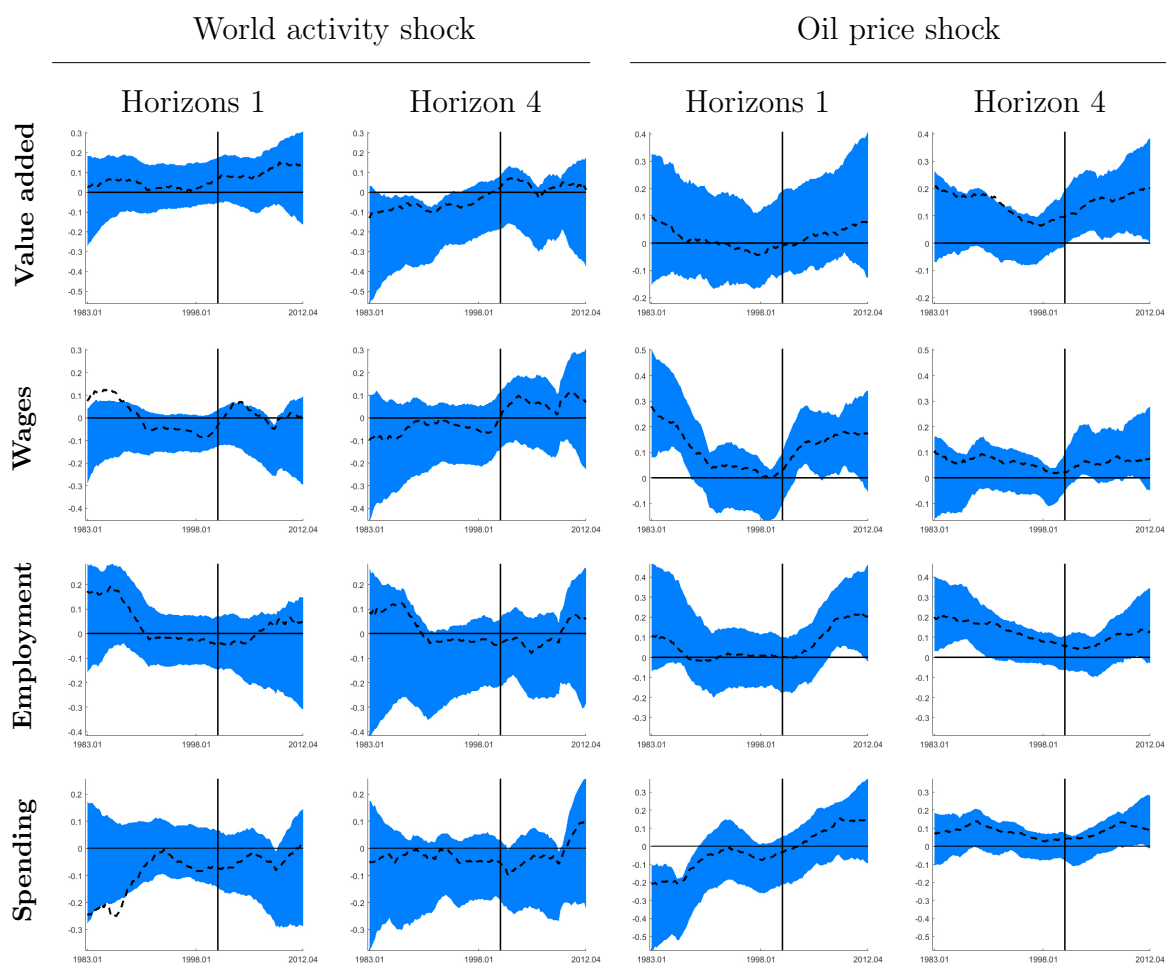


Figure 19. World activity and oil price shocks and the public sector: Time-varying responses using sign restrictions. The first two columns display estimated impulse responses following a one standard deviation world activity shock. The shock is restricted to increase world activity and the oil price on impact. The subsequent two columns display estimated impulse responses following a one standard deviation oil price shock. The shock is restricted to decrease world activity and increase the oil price on impact. All responses are reported in levels. The color shadings represent 68 percent posterior probability bands from the baseline model (identified using recursive ordering). The dotted black line is the median estimate from the sign identified admissible set. Finally, we plot a vertical line in 2001:Q1 to indicate the introduction of the fiscal rule.

B.10 QoQ model specification

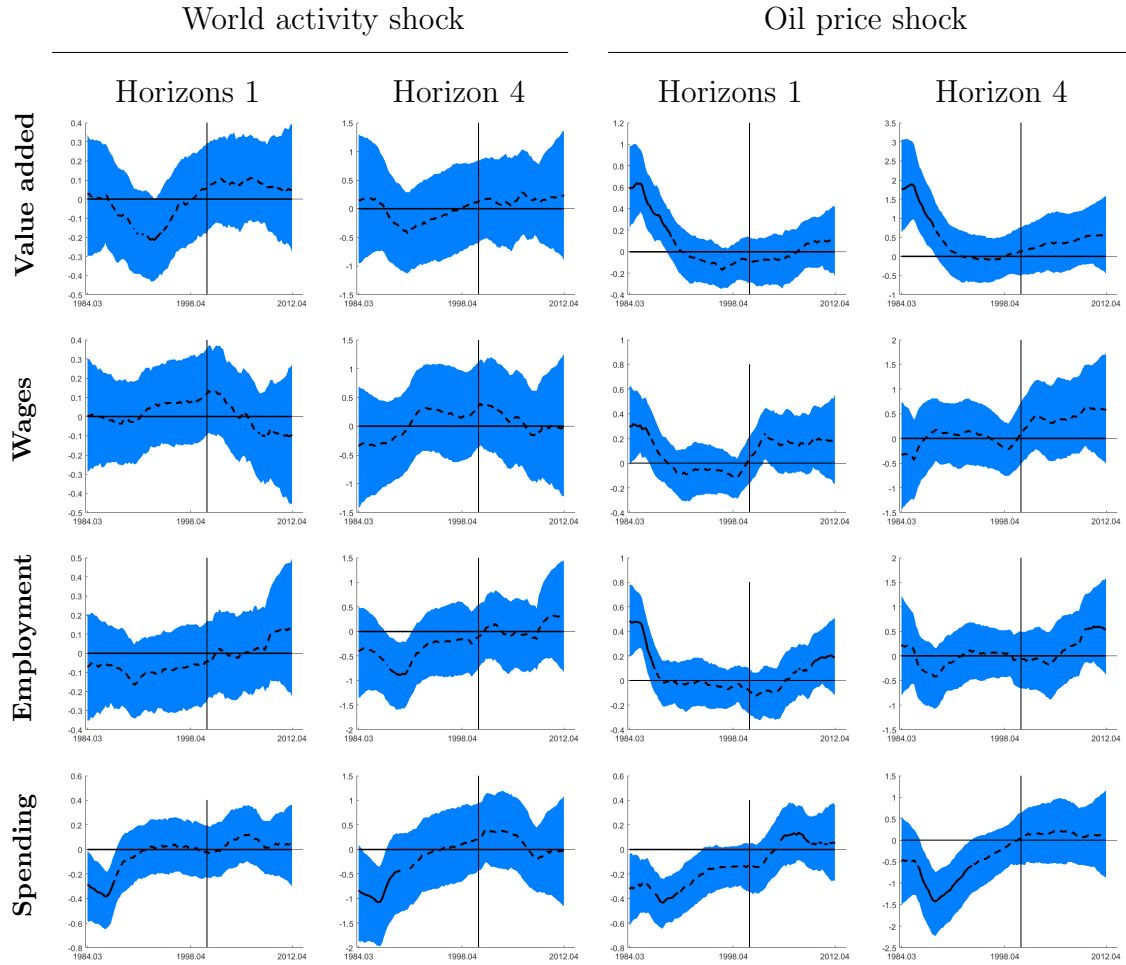


Figure 20. World activity and oil price shock and the public sector: Time-varying responses. The model is specified using quarterly growth rates (as opposed to year-on-year growth as in the baseline model). The initial shock is normalized to one percent (of the normalized data). All responses are reported in levels (of the normalized data). The color shadings represent 68 percent posterior probability bands. The black line is the median estimate. The line is solid (dotted) whenever the median estimate is outside (inside) the 68 percent area in 2001:Q1. Finally, we plot a vertical line in 2001:Q1 to indicate the introduction of the fiscal rule.

Appendix C Convergence of the Markov Chain Monte Carlo Algorithm

Table 3 summarizes the main convergence statistics used to check that the Gibbs sampler mixes well. In the table we first report the mean, as well as the minimum and maximum, of the 10th-order sample autocorrelation of the posterior draws across all parameters. A

Table 3. Convergence statistics. The *AutoCorr* row reports the 10th-order sample autocorrelation of the draws, the *RNE* row reports the relative numerical efficiency measure, proposed by Geweke (1992), while the *IRL* row reports the i-statistic, proposed by Raftery and Lewis (1992). For each entry we report the mean value together with the minimum and maximum value obtained across all parameters in parentheses. Finally, for computational convenience, the scores for Q_t^* , H_t^* , and Z_t^* are only computed for observations 10, 40, 70, and 100.

Statistic	Parameters							
	W	S	B	V	Q_t^*	H_t^*	Z_t^*	a_t
AutoCorr	-0.0 (-0.1,0.1)	-0.0 (-0.1,0.0)	-0.0 (-0.0,0.1)	0.0 (-0.1,0.1)	0.0 (-0.1,0.1)	0.0 (-0.1,0.1)	0.0 (-0.1,0.1)	0.0 (-0.1,0.1)
RNE	0.9 (0.4,1.7)	0.3 (0.3,0.5)	0.7 (0.4,1.0)	0.2 (0.1,0.3)	0.5 (0.2,1.5)	1.0 (0.6,1.8)	0.8 (0.3,2.1)	0.4 (0.1,1.2)
IRL	1.0 (1.0,1.1)	1.3 (1.3,1.3)	1.1 (1.1,1.1)	1.5 (1.0,2.1)	1.2 (1.1,1.3)	1.0 (1.0,1.0)	1.0 (0.9,1.6)	1.0 (1.0,1.0)

low value indicates that the draws are close to independent. The second row of the table reports the relative numerical efficiency measure (RNE), proposed by Geweke (1992). Here we use an RNE version controlling for autocorrelation in the draws by employing a 4 percent tapering of the spectral window used in the computation of the RNE. The RNE measure provides an indication of the number of draws that would be required to produce the same numerical accuracy if the draws represented had been made from an i.i.d. sample drawn directly from the posterior distribution. An RNE value close to or below unity is regarded as satisfactory. The last row, labeled IRL, reports the mean of the i-statistic. This statistic was proposed by Raftery and Lewis (1992). In essence it measures the ratio of two other statistics: the total number of draws needed to achieve the desired accuracy for each parameter, and the number of draws that would be needed if the draws represented an i.i.d. chain, see Raftery and Lewis (1992) for details.¹ Values of IRL exceeding 5 indicate convergence problems with the sampler.

As can be seen from the results reported in Table 3, the sampler seems to have converged. That is, the mean autocorrelations are all very close to zero, and the minimum or maximum values obtained seldom exceed 0.1 in absolute value. Moreover, the mean RNE statistic does not exceed unity by a large margin for any of the parameters. However, for W , Q_t^* , H_t^* , and Z_t^* there are signs that some of the parameters have higher scores.

¹The parameters used for computing these diagnostics are as follows: quantile = 0.025; desired accuracy = 0.025; required probability of attaining the required accuracy = 0.95.

For example, for Z_t^* , the maximum obtained score is 2.1, indicating that only roughly 20 percent of the numbers of draws would be required to achieve the same accuracy from an i.i.d. set of draws. Finally, the IRL statistics are always well below 5. Additional convergence results can be obtained on request.

Appendix D Prior specification and sensitivity

As noted in Section 3, to estimate the model, we need priors for the initial state variables a_0 , z_0 , ao_0 , h_0^σ , and h_0^η , and for the hyper-parameters Φ , S , B , W , Φ and V . Below, we first describe the prior specification used in the benchmark model, i.e., the model on which the results in Section 4 build. We then discuss in greater detail the reasons for why this specific prior specification was adopted, and subsequently also discuss the sensitivity of our main results to other prior specifications.

The priors for the initial states take the following form:

$$\begin{aligned} a_0 &\sim N(\underline{y}, I_{qh}) & z_0 &\sim N(\hat{z}_{OLS}, I_{\tilde{m}}) \\ ao_0 &\sim N(0, I_{qq}) & h_0^\sigma &\sim N(0, I_q) \\ h_0^\eta &\sim N(0, I_N) \end{aligned}$$

where $\tilde{m} = q(s + 1)N$, $qq = \frac{q(q-1)}{2}$, and \underline{y} is a stacked column vector of the observed values for the first q variables in y_t for $t = 0, \dots, -h$. \hat{z}_{OLS} are constant parameter OLS estimates of the matrix Z (stacked by rows) in equation (1a), covering the sample 1981:Q3-1990:Q1. In these initial estimates the unknown elements of a_t are approximated by principal components estimates of the panel of observables in \tilde{y}_T .²

The priors for the hyper-parameters Φ and Φ are set to:

$$\begin{aligned} \Phi &\sim N(\hat{\Phi}_{OLS}, V(\hat{\Phi}_{OLS})) \\ \Phi_i &\sim N(0, I_p \cdot 0.5) && \text{for } i = 1, \dots, N \end{aligned}$$

where $\hat{\Phi}_{OLS}$ are OLS estimates of equation (1b), covering the sample 1982:Q1-2012:Q4. As above, when estimating the OLS quantities, the unknown elements of a_t are approximated by principal components estimates of the panel of observables in \tilde{y}_T . $V(\hat{\Phi}_{OLS})$ is

²These estimates do not take into account the potential autocorrelation and stochastic volatility associated with the idiosyncratic errors.

a diagonal matrix where the non-zero entries are the variance terms associated with the $\hat{\Phi}_{OLS}$ elements.

The priors for the remaining hyper-parameters are all from the Inverse-Wishart distribution:

$$\begin{aligned}
W_i &\sim IW(\underline{T}^W, \underline{T}^W \cdot I_m \cdot \kappa_W^2) & \underline{T}^W = 125, \kappa_W = 0.1 & \text{ for } i = 1, \dots, N \\
S_l &\sim IW(\underline{T}^S, \underline{T}^S \cdot I_l \cdot \kappa_S^2) & \underline{T}^S = 25, \kappa_S = 0.05 & \text{ for } l = 1, \dots, q - 1 \\
B &\sim IW(\underline{T}^B, \underline{T}^B \cdot I_q \cdot \kappa_B^2) & \underline{T}^B = 100, \kappa_B = 0.1 & \\
V_i &\sim IW(\underline{T}^V, \underline{T}^V \cdot I_p \cdot \kappa_V^2) & \underline{T}^V = 10, \kappa_V = 0.1 & \text{ for } i = 1, \dots, N
\end{aligned}$$

where the first element in each prior distribution is the degrees of freedom parameter, and the second the scale parameter. We note that for the Inverse-Wishart distribution the prior scale matrix has the interpretation of the prior sum of squared residuals.³

D.1 Prior discussion

In the following we elaborate our reasons for choosing the prior specification described above. We focus our discussion of alternative specifications of \underline{T}_ν and κ_ν for $\nu = \{W, S, B, V\}$, since the prior specifications for the other priors, \hat{z}_{OLS} , $\hat{\Phi}_{OLS}$, and $V(\hat{\Phi}_{OLS})$, seem to be of minor empirical importance.⁴

Before going into the details it is worth considering a simplified example. Assume a parameter α_t follows a random walk like $\alpha_t = \alpha_{t-1} + e_t \sim N(0, Q)$, where $Q \sim IW(\underline{T}, \underline{T} \cdot \kappa^2)$. Then, for a given \underline{T} , varying the size of κ will result in very different prior beliefs about the amount of time variation in α_t . For example, letting $\kappa = \{0.05, 0.1, 0.15\}$, will result in a roughly 95 percent prior probability of a 100 percent, 200 percent and 300 percent cumulative change in α_t , respectively, over a period of 100 observations. Accordingly, the priors on V , S , B and W , defines our prior belief on the amount of time variation in the parameters. We will discuss the latter three first, since our results do not seem to be sensitive to the prior settings for \underline{T}_V and κ_V .

³Therefore, each scale matrix is multiplied by the degrees of freedom parameter. Also, for the Inverse-Wishart prior to be proper, the degrees of freedom parameter must be larger than the dimension of the scale matrix. This is the case in all our prior specifications.

⁴This finding is common in the literature entertaining time-varying parameter models, and is also found in, e.g., [Primiceri \(2005\)](#) and [Del Negro and Otrok \(2008\)](#), all of whom estimate models that are related to ours.

The setting of S and B defines our prior beliefs about the amount of time variation in equations (7b) and (7c), and ultimately the stochastic volatility part of the transition equation in (1b). The setting of W defines our prior belief about the amount of time variation in (7a), i.e., the time-varying factor loadings (z_t). Unfortunately, there is a trade-off between these two. For example, for a similarly sized prior belief on \underline{T}_B , \underline{T}_S , and \underline{T}_W , setting κ_B and κ_S very low, but κ_W very high, will force most of the model fit to be explained by the time-varying factor loadings. In the reverse case, setting κ_B and κ_S very high and κ_W very low, will almost remove the time variation in z_t .

The trade-off faced in setting the priors for κ_B , κ_S and κ_W make their role important. However, our research question and earlier literature can guide us in setting these priors. The time-varying parameters and stochastic volatilities are introduced in the Dynamic Factor Model to capture important “stylized facts” associated with global business cycles in general and the Norwegian domestic business cycle in particular. That is, we want to allow for: 1) A changing dependence structure, where the panel of domestic variables have a time-varying exposure to the aggregate business cycle factors, due to, e.g., changes in systematic fiscal policy, and 2) Great Moderation and Recession effects, where the volatility in aggregate business cycle variables seems to have fallen and then increased again over the last decades.

To allow for 1) we set $\kappa_W = 0.1$ and $\underline{T}_W = 125$. This prior belief permits the factor loadings to vary considerably across time. For our purpose, which is to uncover any potential changes in the parameters due to structural changes in the conduct of fiscal policy, this seems reasonable. Importantly, as described in Appendix E, the sampling algorithm used to estimate the time path for the structural parameters is essentially a smoothing algorithm. However, across the time period evaluated in this analysis many fiscal regimes have been present, see Section 2. Each new regime will plausibly be associated with a new set of policy parameters, that should not be smoothed out. Thus, by allowing for a high degree of variability in the process driving the time evolution of the structural factor loadings, we ensure these parameters are free to jump in response to new policy regimes. The downside of imposing this prior belief is, of course, that the factor loadings might change considerably over time, but just in order to explain outliers and push the in-sample errors to zero. As noted in Primiceri (2005), this type of behavior by the time-varying parameters is typical of very narrow likelihood peaks in possibly uninteresting regions of

the parameter space, where the level of the likelihood is not informative of the model’s fit. However, our focus is not on, e.g., forecasting, where the above problem might be a bigger concern, but on uncovering jumps in the systematic policy parameters across time and the associated implications for the Norwegian macro economy.

To allow for 2), we set $\kappa_S = 0.05$ and $\kappa_B = 0.1$. This belief is in accordance with a large literature that have already established that the volatility of international business cycle shocks have indeed changed significantly in recent decades. For example, both [Stock and Watson \(2005\)](#) and [Del Negro and Otrok \(2008\)](#) document drops in volatility among G7 countries of over 50 percent since the late 1970s. The findings in [Del Negro and Otrok \(2008\)](#) suggest moreover that the fall in the volatility in the Norwegian business cycle is even bigger, close to 150 percent over the period from the early 1980s to the mid 2000s. Likewise, according to findings in [Baumeister and Peersman \(2013\)](#), the conditional standard deviation in the change in the real price of oil has moved from around 20 in the mid 1980s, to 10 in the mid 1990s, and back again to above 20 at the late 2000s, reflecting changes of over 100 percent within a period of 10 years. Our setting of $\underline{T}_B = 100$ reflects our confidence in this evidence. Conversely, our setting of $\underline{T}_S = 25$ reflects our lack of strong prior beliefs about time-variation in aO_t , at least for the Norwegian economy. Moreover, as described above, both κ_S , κ_B , and their associated degrees of freedom parameters must be set in relation to κ_W and \underline{T}_W . Since we allow for a large degree of variation in the factor loadings, we must also allow for a large degree of variation in the volatilities. If not, our experience is that most of the model fit is tilted toward variation in one of them, which is not a desirable property.⁵

We have neither good evidence nor prior beliefs as to the amount of time variation to expect in the stochastic volatilities associated with the idiosyncratic errors. Therefore, κ_V is set equal to κ_B for consistency, and $\underline{T}_V = 10$, reflecting that we are reasonably uninformative about this parameter.

⁵As a check of whether or not these priors are sensible we have also estimated a Dynamic Factor Model with constant parameters over two different sub-samples, pre and post 2001. By computing the absolute change in the factor loadings and the standard deviation of the errors in equation (1b) across those two sub-samples we find that the absolute change in z is well above 100 percent, for many variables, and that the absolute change in Σ is in fact close to 150 percent for the Norwegian business cycle factors. Thus, although somewhat on the high end, our prior belief on the amount of time variation in z_t (reflected by the W prior) and Σ_t (reflected by the B prior) seem reasonable also according to this criteria.

D.2 Prior sensitivity

To gauge the extent to which our results are sensitive to the specific prior specification discussed above, we have estimated the time-varying DFM with a set of alternative prior specifications. In the light of our discussion in Section D.1, we focus on the priors for B and W , and the setting of κ . Especially, we estimate the model letting κ be in the set $\kappa = \{0.05, 0.1, 0.15\}$ for all combinations of $B(\kappa)$ and $W(\kappa)$. In total this amounts to 9 different model estimates, encompassing our benchmark model but also allowing for models which a-priori allow for somewhat lower and higher parameter variability. We evaluate the appropriateness of these models both informally and formally.

In a Bayesian setting, the natural formal scoring metric is the marginal likelihood. However, for high dimensional and complex time-varying factor models such as ours, computing this statistic is difficult, and we are not aware of any good agreed upon method for how to do so. For this reason we developed a Reversible Jump Markov Chain Monte Carlo (RJCMCMC) algorithm to assess the marginal likelihood implied by the different model and prior specifications. A full description of how our implementation of the RJCMCMC algorithm is provided in Appendix F. Here we note that we in a simulation experiment have validated that the algorithm seems to be able to select the correct model among a set of competing specifications, but that the convergence properties of the algorithm are poor and that the estimates have a large degree of uncertainty.⁶ Still, conditional on these shortcomings, the marginal likelihood assessment seems to favor models with prior specifications where $\kappa = 0.05$ for the W prior and $\kappa \geq 0.05$ for the B prior, or in other words, a somewhat lower variability in the factor loadings than what we believe to be true in the benchmark model, but more or less the same variability for the time-varying volatilities. Given the attendant caveats explained above, however, we do not put too much confidence in these results.

More important, then, is the informal evaluation in which we assess the extent to which our main results change depending on the prior specification. As explained in Section D.1, we want to allow for substantial time variation in the factor loadings and the volatilities, but not enforce it such that the results are solely driven by our prior beliefs.

⁶This might be because the likelihood surface is highly complex, or because our implementation of the algorithm is inefficient. Another reason might be that we have to be parsimonious regarding the number of simulations due to computational issues, see the discussion in Appendix F.1.

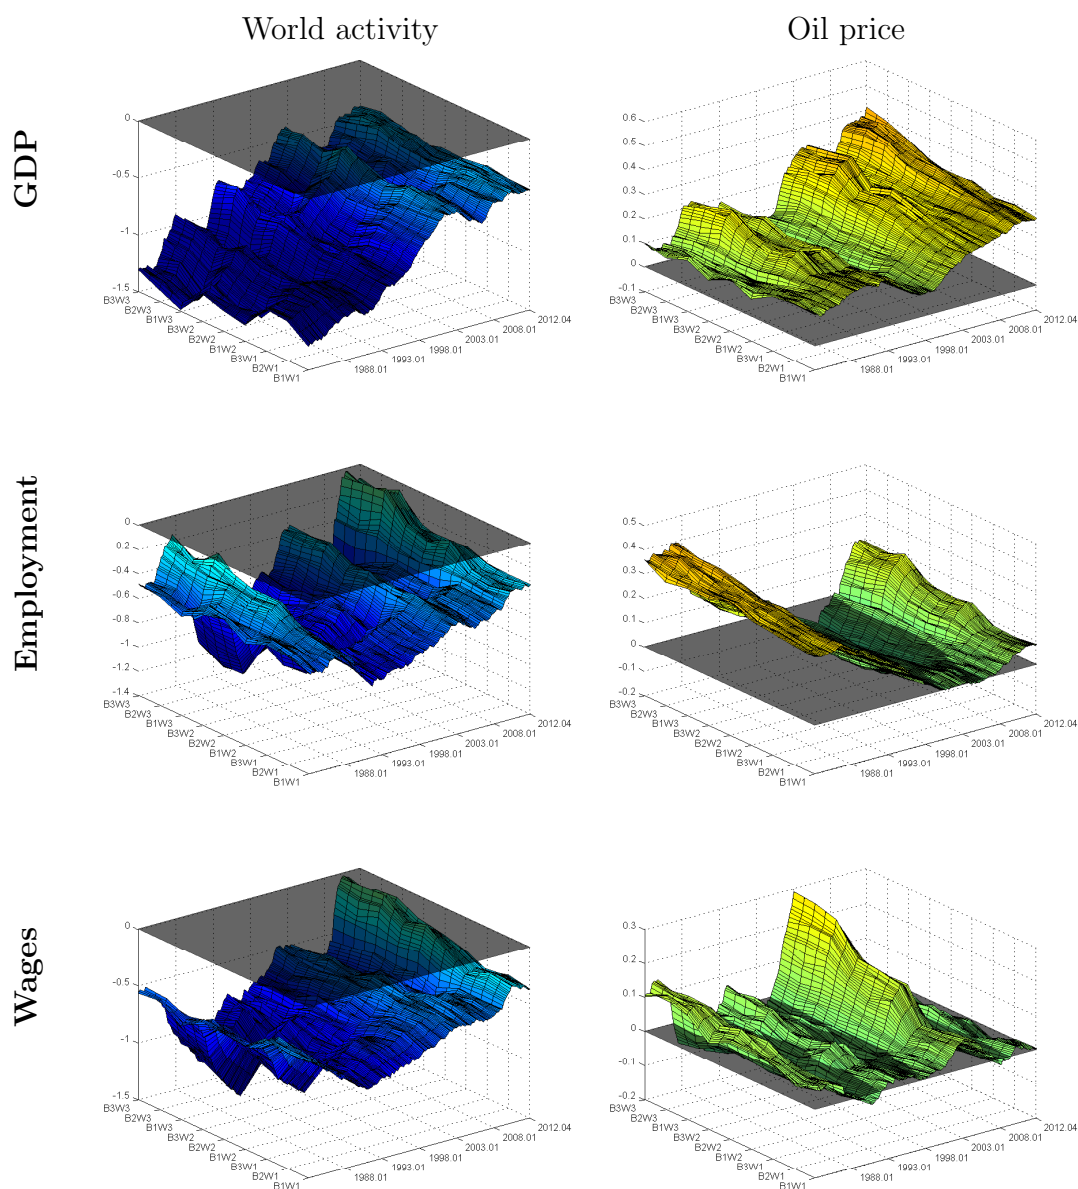


Figure 21. Public sector relative to the mainland economy using different priors. Each plot reports the response, across time (x-axis) and prior specification (y-axis), of an outcome variable in the public sector relative to the response in the mainland economy. Here, the mainland economy is defined as the average response across all sectors, except the public sector. The initial shock is normalized to 1 percent (of the normalized data). All the relative responses are reported as averages across impulse response horizons 1 to 8. A value above zero indicates that the public sector responds more positively to the given shock than the mainland economy as a whole. The different prior specifications listed on the y-axis correspond to different combinations of κ for the B and W priors. In particular, we compare models letting κ be in the set $\kappa = \{0.05, 0.1, 0.15\}$ for all combinations of $B(\kappa)$ and $W(\kappa)$. Thus, $B1W1$ corresponds to setting $\kappa = 0.05$ for both B and W , $B1W2$ corresponds to setting $\kappa = 0.05$ for B and $\kappa = 0.1$ for W , etc. The main results reported in Sections 4.1 and 4 correspond to using a model with the $B2W2$ prior specification.

As documented in Section 4, when allowing for a large degree of time variation in both the volatilities and factor loadings, i.e., setting κ_W and κ_B high, the results point to a large degree of time-varying impulse responses. However, the main conclusion regarding increased procyclical fiscal policy after the implementation of the fiscal rule also holds for models with priors that allow for much less time variation. Indeed, even for combinations of $B(\kappa)$ and $W(\kappa)$ where $\kappa = 0.05$, we observe time variation in the result implying a more procyclical fiscal policy after the adoption of the fiscal rule. These additional results are summarized in Figure 21. The figure reports the same type of results as reported in Figure 6, but for different prior specifications.⁷ As is clearly seen in the figure, after a positive world activity shock, fiscal policy has become less countercyclical over time. After a positive oil price shock, the public sector grows relative to the mainland economy, and particularly so after the adoption of the fiscal rule. Both findings confirm what have already documented in Section 4.

In sum, the sensitivity analysis shows that our main results are not driven by the prior specification. We leave it to future research to devise better ways of formally computing posterior model probabilities, or marginal likelihoods, for high dimensional and complex models as the one entertained here.

Appendix E The Gibbs sampling approach

Section 3.3 of the main paper gives a short overview of how the DFM is estimated. Here we provide a more detailed overview. For convenience, we repeat the main system equations:

$$y_t = z_{0,t}a_t + \dots + z_{s,t}a_{t-s} + e_t \quad (16a)$$

$$a_t = \Phi_1 a_{t-1} + \dots + \Phi_h a_{t-h} + A0_t^{-1} \Sigma_t \epsilon_t \quad (16b)$$

$$e_t = \Phi_1 e_{t-1} + \dots + \Phi_p e_{t-p} + \Upsilon_t u_t \quad (16c)$$

where (16a) is the observation equation, (16b) the transition equation, and finally, (16c) the equation describing the law of motion for the idiosyncratic errors. Moreover, the

⁷To make the results across different prior specifications presentable in one figure, we report average impulse responses across horizons 1-8 for each time period. Additional results for each prior specification and all impulse response horizons can be obtained on request.

time-varying parameters and covariances of the model follow random walk processes:

$$z_t = z_{t-1} + w_t \sim N(0, W) \quad (17a)$$

$$ao_t = ao_{t-1} + s_t \sim N(0, S) \quad (17b)$$

$$h_t^\sigma = h_{t-1}^\sigma + b_t \sim N(0, B) \quad (17c)$$

$$h_t^\eta = h_{t-1}^\eta + v_t \sim N(0, V) \quad (17d)$$

where $Z_t = [z_{0,t}, \dots, z_{s,t}]$ and $z_t = \text{vec}(Z_t')$ (the matrix Z_t stacked by rows). ao_t is the vector on non-zero and non-one elements of the matrix $A0_t$ (stacked by rows), and $h_t^\sigma = \log(\sigma_t)$ and $h_t^\eta = \log(\eta_t)$, see equations (2), (3), (5) and (6). Thus, the model's hyper-parameters are defined by Φ , Φ , W , S , B , and V , while the model's unknown state variables are defined by a_t , e_t , z_t , ao_t , h_t^σ , and h_t^η .

This system is then estimated using Gibbs simulations, which draw the conditional posterior utilizing 7 blocks. Blocks 1 to 4 draws the states and hyper-parameters associated with equations (16b), (17b) and (17c). Block 5 draws the state and hyper-parameter associated with equation (17a), and Blocks 6 and 7 draw the state and hyper-parameters associated with equations (16c) and (17d). Below we describe each block in greater detail. For future reference and notational simplicity it will prove useful to define the following: $\tilde{y}_T = [y_1, \dots, y_T]'$, $\tilde{a}_T = [a_1, \dots, a_T]'$, $\tilde{z}_T = [z_1, \dots, z_T]'$, $\tilde{e}_T = [e_1, \dots, e_T]'$, $\tilde{ao}_T = [ao_1, \dots, ao_T]'$, $\tilde{h}_T^\sigma = [h_1^\sigma, \dots, h_T^\sigma]'$, $\tilde{h}_T^\eta = [h_1^\eta, \dots, h_T^\eta]'$, $\Phi = [\Phi_1, \dots, \Phi_h]$, and $\Phi = [\Phi_1, \dots, \Phi_p]$.

E.1 Block 1: $\tilde{a}_T | \tilde{y}_T, \tilde{z}_T, \tilde{e}_T, \tilde{h}_T^\eta, \tilde{ao}_T, \tilde{h}_T^\sigma, \Phi, \Phi$

Equations (16a) and (16b) constitute a state space system we can use to draw the unobserved state a_t using the Carter and Kohn's multimove Gibbs sampling approach, see Section E.8. However, to do so we need to make the errors in the observation equation conditionally i.i.d. This is easy, given knowledge of equation (16c) and Φ , we can define $\Phi(L) = (I - \sum_{k=1}^p \Phi_k L^k)$ and pre-multiply equation (16a) by $\Phi(L)$ to obtain the system:

$$y_t^* = z_{0,t}^* a_t + \dots + z_{s,t}^* a_{t-s} + \Upsilon_t u_t \quad \Upsilon_t u_t \sim N(0, H_t) \quad (18a)$$

$$a_t = \Phi_1 a_{t-1} + \dots + \Phi_h a_{t-h} + A0_t^{-1} \Sigma_t \epsilon_t \quad A0_t^{-1} \Sigma_t \epsilon_t \sim N(0, \Omega_t) \quad (18b)$$

where $y_t^* = (I - \sum_{k=1}^p \Phi_k L^k) y_t$ and $z_{j,t}^* = (I - \sum_{k=1}^p \Phi_k L^k) z_{j,t}$ for $j = 0, \dots, s$.

Since all hyper-parameters and state variables, less \tilde{a}_T , are known (or conditionally known), it follows from equations (2), (3), (5) and (6) that Ω_t and H_t are also known for all t . Accordingly, we can use the equations in (18) together with Carter and Kohn's multimove Gibbs sampling approach, to sample a_t from:

$$a_T | \dots \sim N(a_{T|T}, P_T^a), \quad t = T \quad (19a)$$

$$a_t | \dots \sim N(a_{t|t, a_{t+1}}, P_{t|t, a_{t+1}}^a), \quad t = T-1, T-2, \dots, 1 \quad (19b)$$

to get \tilde{a}_T .

E.2 Block 2: $\Phi | \tilde{a}_T, \tilde{a}\tilde{o}_T, \tilde{h}_T^\sigma$

Conditional on \tilde{a}_T , the transition equation in (16b) is independent of the rest of the model. As above, conditional on knowing $\tilde{a}\tilde{o}_T$ and \tilde{h}_T^σ , also makes Ω_t known. Accordingly, we can draw Φ based on a conditional posterior that accounts for the heteroscedasticity in the error terms in (16b). This can be achieved by putting the transition equation on SUR form.⁸ To do so, we define:

$$Y_t = \begin{bmatrix} a_{1,t} \\ a_{2,t} \\ \vdots \\ a_{q,t} \end{bmatrix} \quad X_t = \begin{bmatrix} x_{t,1} & 0 & \cdots & 0 \\ 0 & x_{t,2} & \ddots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & x_{t,q} \end{bmatrix} \quad \varepsilon_t = \begin{bmatrix} \omega_{1,t}^A \\ \omega_{2,t}^A \\ \vdots \\ \omega_{q,t}^A \end{bmatrix} \quad \beta^\Phi = \begin{bmatrix} \beta_1^\Phi \\ \beta_2^\Phi \\ \vdots \\ \beta_q^\Phi \end{bmatrix} \quad (20)$$

where $\beta_l^\Phi = [\Phi_{l,1}, \dots, \Phi_{l,h}]'$ and $x_{t,l} = [Y'_{t-1}, \dots, Y'_{t-h}]$ for $l = 1, \dots, q$, i.e., the autoregressive coefficients from the l th equation in the transition equation and the lagged dependant variables, and $\omega_t^A = A0_t^{-1}\Sigma_t\epsilon_t$.

Stacking Y_t , X_t and ε_t together across time lets us write the transition equation as:

$$Y = X\beta^\Phi + \varepsilon \sim N(0, \Psi) \quad (21)$$

where $Y = [Y'_1, \dots, Y'_T]'$, $X = [X_1, \dots, X_T]'$, $\varepsilon = [\varepsilon'_1, \dots, \varepsilon'_T]'$, and Ψ is a $(T \times q) \times (T \times q)$ block diagonal matrix given by:

$$\Psi = \begin{bmatrix} \Omega_1 & 0 & \cdots & 0 \\ 0 & \Omega_2 & \ddots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & \Omega_T \end{bmatrix} \quad (22)$$

⁸With the transition equation specified in SUR form it becomes easy to adjust the VAR(h) model such that different regressors enter the q equations of the VAR(h).

The conditional posterior draws of β^Φ , and thus Φ , are:

$$\beta^\Phi | \dots \sim N(\bar{\beta}^\Phi, \bar{V}_{\beta^\Phi})_{I[s(\beta^\Phi)]} \quad (23)$$

where $I[s(\beta^{varphi})]$ is an indicator function used to denote that the roots of β lie outside the unit circle, and:

$$\bar{V}_{\beta^\Phi} = (V_{\beta^\Phi}^{-1} + X'\Psi^{-1}X)^{-1} \quad (24a)$$

$$\bar{\beta}^\Phi = \bar{V}_{\beta^\Phi}(V_{\beta^\Phi}^{-1}\underline{\beta}^\Phi + X'\Psi^{-1}Y) \quad (24b)$$

E.3 Block 3: $\tilde{a}_T | \tilde{a}_T, \tilde{h}_T^\sigma, \Phi, S$ and $S | \tilde{a}_T$

Conditional on \tilde{a}_T and Φ we can define $\hat{a}_t = a_t - (\Phi_1 a_{t-1} + \dots + \Phi_h a_{t-h})$, and write equation (16b) as:

$$A0_t \hat{a}_t = \Sigma_t \epsilon_t \quad (25)$$

Since $A0_t$ is a lower triangular matrix with ones on the diagonal, equation (25) together with equation (17b) can be written as the state space system:

$$\hat{a}_t = \tilde{Z}_t a o_t + \Sigma_t \epsilon_t \quad \Sigma_t \epsilon_t \sim N(0, \Sigma_t' \Sigma_t) \quad (26a)$$

$$a o_t = a o_{t-1} + s_t \quad s_t \sim N(0, S) \quad (26b)$$

where \tilde{Z}_t is the following $q \times \frac{q(q-1)}{2}$ matrix:

$$\tilde{Z}_t = \begin{bmatrix} 0 & \dots & \dots & 0 \\ -\hat{a}_{1,t} & 0 & \dots & 0 \\ 0 & -\hat{a}_{[1,2],t} & \ddots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & -\hat{a}_{[1,\dots,q-1],t} \end{bmatrix} \quad (27)$$

where $\hat{a}_{[1,\dots,j],t}$ denotes the row vector $[\hat{a}_{1,t}, \hat{a}_{2,t}, \dots, \hat{a}_{j,t}]$.

Now, although the equations in (26) have a (conditional) Gaussian state space representation, the system is nonlinear since \hat{a}_t essentially shows up on both sides of the equality sign in equation (26). Still, under the assumption that the S matrix is block diagonal, see equation (9), we can apply Carter and Kohn's multimove Gibbs sampling approach, see Section E.8, and draw \tilde{a}_T in a recursive manner that is consistent with the assumptions on $A0_t$. That is:

$$ao_T | \dots \sim N(ao_T | T, P_T^{ao}), \quad t = T \quad (28a)$$

$$ao_t | \dots \sim N(ao_t | t, ao_{t+1}, P_{t|t, ao_{t+1}}^{ao}), \quad t = T-1, T-2, \dots, 1 \quad (28b)$$

Once $\tilde{a}\tilde{o}_T$ has been drawn in this manner, the innovations in (17b) are observable, and we can compute the residual sums of squares. Thus, the conditional posterior of S_l for $l = 1, \dots, q-1$ can be sampled from the Inverse-Wishart distribution:

$$S_l | \dots \sim IW(\bar{v}^S, \bar{S}_l) \quad (29)$$

where $\bar{v}^S = T + \underline{T}^S$, $\bar{S}_l = [\underline{S}_l + \sum_{t=1}^T \xi_{l,t}^{S'} \xi_{l,t}^S]$, and $\xi_{l,t}^S = ao_{l,t} - ao_{l,t-1}$ are the errors associated with the l th block.⁹

E.4 Block 4: $\tilde{h}_T^\sigma | \tilde{a}_T, \tilde{a}\tilde{o}_T, \Phi, B$ and $B | \tilde{h}_T^\sigma$

Conditional on $\tilde{a}\tilde{o}_T$, Φ and \tilde{a}_T , the L.H.S of equation (25) is known, and can be written as:

$$\hat{a}_t^* = \Sigma_t \epsilon_t \quad (30)$$

where $\hat{a}_t^* = A0_t \hat{a}_t$ is an observable.

Together with the transition equation in (17c), the observation equation in (30) constitutes a nonlinear state space system. The nonlinearity can be converted into a linear one by squaring and taking logarithms of every element of (30), yielding:

$$\hat{a}_t^{**} = 2h_t^\sigma + d_t^\sigma \quad (31a)$$

$$h_t^\sigma = h_{t-1}^\sigma + b_t \quad (31b)$$

where $d_{l,t}^\sigma = \log(\epsilon_{l,t}^2)$, $h_{l,t}^\sigma = \log(\sigma_{l,t})$, $\hat{a}_{l,t}^{**} = \log[(\hat{a}_{l,t}^*)^2 + \bar{c}]$ for $l = 1, \dots, q$. $\bar{c} = 0.001$ is an offsetting constant added to the latter expression to avoid potentially taking the log of zero.

Now the system in (31) is linear, but it has a non-Gaussian state space form, because the innovations in the observation equation are distributed as $\log \chi^2(1)$. In order to further transform the system into a Gaussian one, a mixture of normals approximation of the $\log \chi^2(1)$ distribution is used. Following Kim et al. (1998), we select a mixture of seven

⁹Remember that the $l = 1$ elements of $\xi_{l,t}^S$ and S_l are associated with the $l + 1$ row of $A0_t$. Accordingly, for $l = 1$, S_1 will be a 1×1 matrix, for $l = 2$, S_2 will be a 2×2 matrix, etc. See also Section 3.1.

normal densities with component probabilities q_γ , mean $m_\gamma - 1.2704$, and variances v_γ^2 , for $\gamma = 1, \dots, 7$. The constants $q_\gamma, m_\gamma, v_\gamma^2$ are chosen to match a number of moments of the $\log \chi^2(1)$ distribution. Since the covariance matrix of ϵ is an identity matrix, this implies that the covariance matrix of d^σ is also a diagonal, and we can use the same (independent) mixture of normals approximation for any element of d^σ . Accordingly, conditionally on $\hat{a}_{j,t}^{**}$ and h_t , we can sample a selection matrix $\tilde{s}_T = [s_1, \dots, s_T]'$ as:

$$Pr(s_{l,t} = \gamma | \hat{a}_{l,t}^{**}, h_{l,t}^\sigma) \propto q_\gamma f_N(\hat{a}_{l,t}^{**} | 2h_{l,t}^\sigma + m_\gamma - 1.2704, v_\gamma^2) \quad \gamma = 1, \dots, 7 \quad l = 1, \dots, q \quad (32)$$

and use this to select which member of the mixture of the normal approximations that should be used to construct the covariance matrix of d^σ and adjust the mean of \hat{a}_t^{**} at every point in time.

Denoting the adjusted observations and covariances as $\hat{a}_t^{***} = \hat{a}_t^{**} - m_{\gamma=s_t} + 1.2704$ and D_t^σ , respectively, the system in (31) finally has an approximate linear and Gaussian state space form. Again, like above, this allows one to recursively recover h_t^σ for $t = 1, \dots, T$ using the Carter and Kohn algorithm:

$$h_T^\sigma | \dots \sim N(h_{T|T}^\sigma, P_{T|T}^{h^\sigma}), \quad t = T \quad (33a)$$

$$h_t^\sigma | \dots \sim N(h_{t|t, h_{t+1}^\sigma}^\sigma, P_{t|t, h_{t+1}^\sigma}^{h^\sigma}), \quad t = T-1, T-2, \dots, 1 \quad (33b)$$

Likewise, conditional on \tilde{h}_T^σ , the posterior of B is drawn from the Inverse-Wishart distribution:

$$B | \dots \sim IW(\bar{v}^B, \bar{B}) \quad (34)$$

where $\bar{v}^B = T + \underline{\mathbf{T}}^B$, $\bar{B} = [\underline{\mathbf{B}} + \sum_{t=1}^T \xi_t^{B'} \xi_t^B]$, and $\xi_t^B = h_t^\sigma - h_{t-1}^\sigma$.

E.5 Block 5: $\tilde{z}_T | \tilde{y}_T, \tilde{a}_T, \tilde{e}_T, \tilde{h}_T^\eta, \Phi, W$ and $W | \tilde{z}_T$

Conditionally on \tilde{a}_T the errors in (16a) are independent across i . Moreover, we have assumed that the covariance matrix of z_t in equation (17a) is block diagonal. Consequently, we can draw \tilde{z}_T one equation at a time. As in Appendix E.1 we deal with the fact that the errors in the observation equation are not conditionally i.i.d. by applying the quasi differencing operator, $\Phi(L) = (I - \sum_{k=1}^p \Phi_k L^k)$, to each equation. Thus, for each $i = q+1, \dots, N$, we define $\hat{z}_{j,t}$ as the i th row of $z_{j,t}$ and $\hat{w}_{j,t}$ as the errors in (17a) associated with $\hat{z}_{j,t}$ (for $j = 0, \dots, s$), and obtain the following Gaussian state space system:

$$y_{i,t}^* = a_t^* \hat{z}_{0,t} + \dots + a_{t-s}^* \hat{z}_{s,t} + \eta_{i,t} u_{i,t} \quad \eta_{i,t} u_{i,t} \sim N(0, \eta'_{i,t} \eta_{i,t}) \quad (35a)$$

$$\hat{z}_t = \Xi \hat{z}_{t-1} + \hat{w}_t \quad \hat{w}_t \sim N(0, W_i) \quad (35b)$$

where $a_t^* = [(I - \sum_{k=1}^p \Phi_k L^k) a_t]'$, $\hat{z}_t = [\hat{z}_{0,t}, \dots, \hat{z}_{s,t}]'$, $\hat{w}_t = [\hat{w}_{0,t}, \dots, \hat{w}_{s,t}]'$ and Ξ is a $(s+1) \times (s+1)$ identity matrix. Since W_i and $\eta'_{i,t} \eta_{i,t}$ are conditionally known for each i , the Carter and Kohn algorithm is implemented on (35), in the same manner as before, to sample:

$$\hat{z}_T | \dots \sim N(\hat{z}_{T|T}, P_{T|T}^{\hat{z}}), \quad t = T \quad (36a)$$

$$\hat{z}_t | \dots \sim N(\hat{z}_{t|h_{t+1}}, P_{t|h_{t+1}}^{\hat{z}}), \quad t = T-1, T-2, \dots, 1 \quad (36b)$$

Conditionally on \tilde{z}_T , we sample W_i from the Inverse-Wishart distribution:

$$W_i | \dots \sim IW(\bar{v}^W, \bar{W}_i) \quad (37)$$

where $\bar{v}^W = T + \mathbb{T}^W$, $\bar{W}_i = [\underline{W}_i + \sum_{t=1}^T \xi_{i,t}^{W'} \xi_{i,t}^W]$, and $\xi_{i,t}^W = \hat{z}_t - \hat{z}_{t-1}$.

Repeating this algorithm for $i = q+1 \dots, N$, gives us \tilde{z}_T and W .¹⁰

E.6 Block 6: $\tilde{e}_T | \tilde{y}_T, \tilde{a}_T, \tilde{z}_T$ and $\Phi | \tilde{e}_T, \tilde{h}_T^\eta$

For each observation we have that:

$$e_t = y_t - z_{0,t} a_t + \dots + z_{s,t} a_{t-s} \quad (38)$$

Thus, conditional on \tilde{y}_T, \tilde{a}_T and \tilde{z}_T, \tilde{e}_T is observable.

As above, since \tilde{e}_T is independent across i , we can sample Φ in (16c) one equation at the time. Conditional on \tilde{h}_T^η , this is done in the same manner as in Appendix E.2, with the difference that now the definitions in (20) are replaced by: $Y_t = [e_{1,t}, \dots, e_{N,t}]'$, and $\varepsilon_t = [\omega_{1,t}^E, \dots, \omega_{N,t}^E]'$, with $\omega_t^E = \mathcal{Y}_t u_t$. Further, $\beta^\Phi = [\beta_1^\Phi, \dots, \beta_N^\Phi]'$ with $\beta_i^\Phi = [\Phi_{i,1}, \dots, \Phi_{i,p}]'$, and

$$\Psi = \begin{bmatrix} H_1 & 0 & \dots & 0 \\ 0 & H_2 & \ddots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & H_T \end{bmatrix} \quad (39)$$

¹⁰For the first q elements in y_t , the identification restriction given by (10) implies that the factor loading is equal to one for all t .

The conditional posterior draws of β^Φ , and thus Φ , are therefore:

$$\beta^\Phi | \dots \sim N(\bar{\beta}^\Phi, \bar{V}_{\beta^\Phi})_{I[s(\beta^\Phi)]} \quad (40)$$

where $I[s(\beta^\Phi)]$ is an indicator function used to denote that the roots of β lie outside the unit circle, and:

$$\bar{V}_{\beta^\Phi} = (\underline{V}_{\beta^\Phi}^{-1} + X' \Psi^{-1} X)^{-1} \quad (41a)$$

$$\bar{\beta}^\Phi = \bar{V}_{\beta^\Phi} (\underline{V}_{\beta^\Phi}^{-1} \beta^\Phi + X' \Psi^{-1} Y) \quad (41b)$$

E.7 Block 7: $\tilde{h}_T^\eta | \tilde{e}_T, \Phi, V$ and $V | \tilde{h}_T^\eta$

Conditionally on having sampled \tilde{e}_T and Φ , we can continue to sample \tilde{h}_T^η for each equation independently. This is done as in Appendix E.4.

For concreteness, we define, for $i = 1, \dots, N$:

$$\hat{e}_{i,t}^* = \eta_{i,t} u_{i,t} \quad (42)$$

where $\hat{e}_{i,t}^* = [e_{i,t} - \Phi_{i,1} e_{i,t-1} + \dots + \Phi_{i,p} e_{i,t-p}]$ is now an observable. Squaring and taking logarithms on each element in (42) and using the law of motion for the stochastic volatilities in (17d), we get the following non-Gaussian state space system:

$$\hat{e}_{i,t}^{**} = 2h_{i,t}^\eta + d_{i,t}^\eta \quad (43a)$$

$$h_{i,t}^\eta = h_{i,t-1}^\eta + v_{i,t} \quad (43b)$$

where $d_{i,t}^\eta = \log(u_{i,t}^2)$, $h_{i,t}^\eta = \log(\eta_{i,t})$, $\hat{e}_{i,t}^{**} = \log[(\hat{e}_{i,t}^*)^2 + \bar{c}]$, and $\bar{c} = 0.001$.

The state space system in (43) is linear but non-Gaussian, and simulation of:

$$h_T^\eta | \dots \sim N(h_{T|T}^\eta, P_{T|T}^{h^\eta}), \quad t = T \quad (44a)$$

$$h_t^\eta | \dots \sim N(h_{t|h_{t+1}}^\eta, P_{t|h_{t+1}}^{h^\eta}), \quad t = T-1, T-2, \dots, 1 \quad (44b)$$

is conducted as described in Appendix E.4.

Finally, conditionally on \tilde{h}_T^η , we sample V_i , for $i = 1, \dots, N$, from the Inverse-Wishart distribution:

$$V_i | \dots \sim IW(\bar{v}^V, \bar{V}_i) \quad (45)$$

where $\bar{v}^V = T + \mathbf{T}^V$, $\bar{V}_i = [\underline{V}_i + \sum_{t=1}^T \xi_{i,t}^{V'} \xi_{i,t}^V]$, and $\xi_{i,t}^V = h_{i,t}^\eta - h_{i,t-1}^\eta$ are the errors associated with the i th equation.

E.8 The Carter and Kohn algorithm

Consider a generic state space system, written in companion form, and described by:

$$y_t = Z_t a_t + e_t \sim N(0, H_t) \quad (46a)$$

$$a_t = \Gamma a_{t-1} + G u_t \sim N(0, \Omega_t) \quad (46b)$$

where we assume Z_t , Γ , G , H_t and Ω_t are known, and we wish to estimate the latent state a_t for all $t = 1, \dots, T$. To do so, we can apply Carter and Kohn's multimove Gibbs sampling approach (see [Carter and Kohn \(1994\)](#)).

First, because the state space model given in equation (46) is linear and (conditionally) Gaussian, the distribution of a_t given \tilde{y}_T and that of a_t given a_{t+1} and \tilde{y}_t for $t = T-1, \dots, 1$ are also Gaussian:

$$a_T | \tilde{y}_T \sim N(a_{T|T}, P_{T|T}), \quad t = T \quad (47a)$$

$$a_t | \tilde{y}_t, a_{t+1} \sim N(a_{t|t, a_{t+1}}, P_{t|t, a_{t+1}}), \quad t = T-1, T-2, \dots, 1 \quad (47b)$$

where

$$a_{T|T} = E(a_T | \tilde{y}_T) \quad (48a)$$

$$P_{T|T} = Cov(a_T | \tilde{y}_T) \quad (48b)$$

$$a_{t|t, a_{t+1}} = E(a_t | \tilde{y}_t, a_{t+1}) = E(a_t | a_{t+1}, a_{t|t+1}) \quad (48c)$$

$$P_{t|t, a_{t+1}} = Cov(a_t | \tilde{y}_t, a_{t+1}) = Cov(a_t | a_{t+1}, a_{t|t+1}) \quad (48d)$$

Given $a_{0|0}$ and $P_{0|0}$, the unknown states $a_{T|T}$ and $P_{T|T}$ needed to draw from (47a) can be estimated from the (conditionally) Gaussian Kalman Filter as:

$$a_{t|t-1} = \Gamma a_{t-1|t-1} \quad (49a)$$

$$P_{t|t-1} = \Gamma P_{t-1|t-1} \Gamma' + G \Omega_t G' \quad (49b)$$

$$K_t = P_{t|t-1} Z_t' (Z_t P_{t|t-1} Z_t' + H_t)^{-1} \quad (49c)$$

$$a_{t|t} = a_{t|t-1} + K_t (y_t - Z_t a_{t|t-1}) \quad (49d)$$

$$P_{t|t} = P_{t|t-1} - K_t Z_t P_{t|t-1} \quad (49e)$$

That is, at $t = T$, equation (49d) and (49e) above, together with equation (47a), can be used to draw $a_{T|T}$. Moreover, $a_{t|t, a_{t+1}}$ for $t = T-1, T-2, \dots, 1$ can also be simulated

based on (47b), where $a_{t|t,a_{t+1}}$ and $P_{t|t,a_{t+1}}$ are generated from the following updating equations:

$$a_{t|t,a_{t+1}} = a_{t|t} + P_{t|t}\Gamma'(\Gamma P_{t|t}\Gamma' + G\Omega_t G')^{-1}(a_{t+1} - \Gamma a_{t|t}) \quad (50a)$$

$$P_{t|t,a_{t+1}} = P_{t|t} + P_{t|t}\Gamma'(\Gamma P_{t|t}\Gamma' + G\Omega_t G')^{-1}\Gamma P_{t|t} \quad (50b)$$

Appendix F Marginal Likelihood computation and the Reversible Jump Markov Chain Monte Carlo (RJMCMC) algorithm

The RJMCMC was first proposed by Green (1995), and has since been applied, and modified, in a number of different settings, including model selection. Dellaportas et al. (2002) and Lopes and West (2004) provide two, of many extant, examples. The algorithm derived here extends that presented in Primiceri (2005) to a dynamic factor model setting, accounting for time-varying parameters.

Following the notation in Primiceri (2005) we consider a set of M competing models. In our setting these models differ in the prior assumptions, see Section D.2. Essentially, the RJMCMC algorithm is nothing more than a Metropolis-Hastings (MH) sampler, where the goal is to sample the joint posterior distribution of model $m \in M$ and the associated model parameters, here denoted θ^m . This is done by generating a proposal value of $(m', \theta^{m'})$ from a proposal distribution $q_p(m', \theta^{m'}) = q(\theta^{m'} | m') \cdot J(m')$. The new proposal is accepted using a MH acceptance probability. If the proposal is rejected, it is replaced by the previous element of the chain. After a sufficient number of draws, an approximation to the posterior of $(m', \theta^{m'})$ can be computed, as can posterior probabilities over the models' space.

More specifically, the sampler used is an independence chain Metropolis Hastings algorithm, and we proceed as follows:

1. For every $m \in M$ we approximate the posterior based on the Gibbs sampler algorithm explained in Section E. These approximate posteriors are then used as

proposal distributions for the elements in θ^m :

$$\begin{aligned} q(\Phi^m|m) &= N(\overline{\Phi^m}, \overline{var(\Phi^m)} \cdot 1) \\ q(\Phi^m|m) &= N(\overline{\Phi^m}, \overline{var(\Phi^m)} \cdot 1) \\ q(W^m|m) &= IW(125, \overline{W^m} \cdot 125) \\ q(S^m|m) &= IW(25, \overline{S^m} \cdot 25) \\ q(B^m|m) &= IW(100, \overline{B^m} \cdot 100) \\ q(V^m|m) &= IW(10, \overline{V^m} \cdot 10) \end{aligned}$$

The variables denoted with an upper bar are the posterior means and variances from the initial Gibbs sampler. The variances are made more diffuse than the exact ones to facilitate the convergence of the MH algorithm.¹¹

2. Initialize m and draw θ^m from the proposal distribution $q(\theta^m|m)$.
3. Draw m' from an unconditional proposal distribution $J(m')$ over the models (here we use the uniform distribution), and draw $\theta^{m'}$ from the conditional proposal distribution $q(\theta^{m'}|m')$, such that $q_p(m', \theta^{m'}) = q(\theta^{m'}|m') \cdot J(m')$.
4. Accept the new couple $(m', \theta^{m'})$ with probability:

$$\alpha(m, m') = \min \left\{ 1, \frac{p(y_T|m', \theta^{m'})p(\theta^{m'}|m')p(m')q_p(m, \theta^m)}{p(y_T|m, \theta^m)p(\theta^m|m)p(m)q_p(m', \theta^{m'})} \right\}$$

where $p(y_T|m, \theta^m)$ is the likelihood of the model, computed using a particle filter, see Appendix F.1. $p(\theta^m|m)$ being the prior of θ^m within model m , with $p(m)$ being the prior model probability. We employ equal prior probabilities, so these terms cancel out. If the new draw is not accepted, keep the previous couple (m, θ^m) .

5. Go to 3.

The MH sampler is run using a handful of independent chains. The approximate posteriors, constructed in step 1, are based on 8000 iterations of the Gibbs sampler for each $m \in M$. The first 4000 iterations are discarded and only every fourth of the remaining iterations is used for inference.

¹¹However, as noted in Section D.2, the convergence properties of the proposed algorithm are not good. Although we have not been able to do this, it could probably be improved upon by a better specification of the proposal distributions. See also the discussion at the end of Appendix F.1.

F.1 The particle filter and likelihood computation

In step 4 of the reversible jump algorithm described above we need to calculate the likelihood of the proposed model; $p(y_T|m, \theta^m)$. For traditional (conditional) Gaussian state space models the likelihood can easily be computed through the Kalman filter, which integrates out the dependence of the stochastic latent factors. However, computing the likelihood of the state space model described by the equations in (16) and (17) is more complicated since the parameters are stochastic and the expressions involve products of stochastic variables. To evaluate the likelihood we therefore employ a Rao-Blackwellized particle filter. This particle filter is particularly suitable for state space systems, like ours, where part of the problem can be solved analytically, see Creal (2009) for a short overview and further references. Below we provide a brief description of the algorithm.

Consider a non-linear state space system with hyper-parameters θ , state variable x (in our model the state variable x contains both the time-varying factor loadings, the stochastic volatilities, and the factors), and observable data given by y_t . Then, the goal is to estimate the joint smoothing distribution of the latent state, given by:

$$p(x_{0:t}|y_{1:t}; \theta) = \frac{p(y_{1:t}, x_{0:t}; \theta)}{p(y_{1:t}; \theta)} = \frac{p(y_t|x_t; \theta)p(x_t|x_{t-1}; \theta)}{p(y_t|y_{1:t-1}; \theta)}p(x_{0:t-1}|y_{1:t-1}; \theta) \quad (52)$$

However, solving (52) analytically is difficult due to the assumed non-linearity of the system. This motivates the use of Sequential Monte Carlo methods, such as the particle filter. Instead of solving (52) directly, these methods utilize the recursive structure of the joint smoothing distribution, as highlighted by the last equality sign in (52), and two of its marginal distributions, namely the predictive distribution $p(x_t|y_{1:t-1}; \theta)$ and the filtering distribution given by:

$$p(x_t|y_{1:t}; \theta) = \frac{p(y_t, x_t|y_{1:t-1}; \theta)}{p(y_t|y_{1:t-1}; \theta)} = \frac{p(y_t|x_t; \theta)p(x_t|y_{1:t-1}; \theta)}{p(y_t|y_{1:t-1}; \theta)} \quad (53)$$

Here, $p(y_t, x_t|y_{1:t-1}; \theta) = p(y_t|x_t; \theta)p(x_t|y_{1:t-1}; \theta)$ is the joint distribution of the data and the state variables, and $p(y_t|y_{1:t-1}; \theta)$ is the contribution to the likelihood function (or the normalizing constant). Accordingly, to sequentially solve 52, starting from an initial draw of the state, the last period's filtering distribution is projected forward using the predictive distribution and then updated using the filtering distribution. These iterations continue until the end of the sample.

Difficulty arises because the predictive distribution ($p(x_t|y_{1:t-1}; \theta)$) and the contribution to the likelihood ($p(y_t|y_{1:t-1}; \theta)$) involves integrals that typically cannot be calculated

analytically. To see this, we can re-write the two terms as:

$$p(y_t|y_{1:t-1}; \theta) = \int p(y_t|x_t; \theta)p(x_t|y_{1:t-1}; \theta)dx_t \quad (54)$$

and

$$p(x_t|y_{1:t-1}; \theta) = \int p(x_t|x_{t-1}; \theta)p(x_{t-1}|y_{1:t-1}; \theta)dx_{t-1} \quad (55)$$

Still, these integrals can be approximated using Monte Carlo integration. Draw N particles from $p(x_0|y_0; \theta)$ and use (55) to calculate the predicted value of the state. Then update the value of the state variables based on the information in the data using (53). The latter step is conducted as an importance sampling step, where the particle draws are re-weighted. By conducting these prediction and updating steps for $t = 1, \dots, T$, the joint smoothing distribution in (52) can be obtained. Importantly for our purpose, the contribution to the likelihood at each time period, equation (54), can in most cases be obtained directly from the estimated importance weights.

Partly due to the importance sampling step, which creates options regarding importance distribution, many different particle filters have been proposed. In our model, part of the joint smoothing distribution can be solved analytically, and we take advantage of this fact when designing the filter by decomposing the state x_t into two blocks; $x_t = (x'_{1,t}, a'_t)'$. That is, we group the time-varying factor loadings and the stochastic volatilities into $x_{1,t}$ while the factors are grouped into a_t . In short, we have the state space system:

$$y_t = Z_t(x_t)a_t + e_t \sim N(0, H_t(x_t)) \quad (56a)$$

$$a_t = \Gamma a_{t-1} + G u_t \sim N(0, \Omega_t(x_t)) \quad (56b)$$

The marginal filtering distribution can then be decomposed as:

$$p(x_{1,t}, a_t|y_{1:t}; \theta) = p(a_t|x_{1,t}, y_{1:t}; \theta)p(x_{1,t}|y_{1:t}; \theta)$$

Particles are only simulated randomly from $p(x_{1,t}|y_{1:t}; \theta)$ while conditional on each draw of $x_{1,t}^i$, the distribution of $p(a_t|x_{1,t}^i, y_{1:t}; \theta)$ can be evaluated analytically. In sum, we proceed as follows:

1. At $t = 0$ and for $i = 1, \dots, N$, draw $x_{1,0}$ and a_0 from some unconditional distributions and set $w_0^i = \frac{1}{N}$
2. Set $t = t+1$. For $i = 1, \dots, N$, run the prediction step of the Kalman filter to obtain the conditional likelihood (using the prediction error decomposition), and calculate

the importance weights as $\hat{w}_t^i = \frac{w_t^i}{\sum_{j=1}^N w_t^j}$, where w_t^i is the conditional likelihood associated with particle i .

3. By the law of large numbers, the contribution to the likelihood can be approximated as $\log(p(y_t|y_{t-1}; \theta)) = \log(\frac{\sum_{i=1}^N w_t^i}{N})$
4. Re-sample the N particles $\{x_{1,t-1|t-1}^i, a_{t-1|t-1}^i\}_{i=1}^N$ with probabilities $\{\hat{w}_t^i\}_{i=1}^N$, and set $w_t^i = \frac{1}{N}$
5. For $i = 1, \dots, N$, draw $x_{1,t}$ conditional on $x_{1,t-1}$ and run the Kalman filter on each particle to obtain $a_{t|t}$.
6. Return to 2.

We confirmed in a simulation experiment the ability of the particle filter approach described above to estimate the latent state variables (the time-varying factor loadings, the stochastic volatilities, and the factors) with a high degree of precision. That said, in systems with a high number of states, as is the case here, a substantial number of particles needs to be entertained to obtain reliable estimates of the joint smoothing distribution. This, however, makes the use of the particle filter within the RJMCMC algorithm described above infeasible. The computation time is simply too large. However, it is our experience that a substantially lower number of particles is needed to obtain reasonably stable estimates of the contribution to the likelihood function. We take advantage of this when we employ the particle filter within the RJMCMC sampler, but emphasize that this downscaling of the number of particles likely contributes to increased sampling variation and thus worse convergence properties of the algorithm as a whole, cf. the discussion in Section [D.2](#).

References

- Baumeister, C. and G. Peersman (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.
- Bjørnland, H. C. and L. A. Thorsrud (2016). Boom or gloom? Examining the Dutch disease in two-speed economies. *The Economic Journal* 126(598), 2219–2256.

- Carter, C. K. and R. Kohn (1994). On Gibbs Sampling for State Space Models. *Biometrika* 81(3), 541–553.
- Creal, D. (2009). A survey of sequential Monte Carlo methods for economics and finance. Serie Research Memoranda 0018, VU University Amsterdam, Faculty of Economics, Business Administration and Econometrics.
- Del Negro, M. and C. Otrok (2008). Dynamic factor models with time-varying parameters: measuring changes in international business cycles. Staff Reports 326, Federal Reserve Bank of New York.
- Dellaportas, P., J. Forster, and I. Ntzoufras (2002). On Bayesian model and variable selection using MCMC. *Statistics and Computing* 12(1), 27–36.
- Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In *In BAYESIAN STATISTICS*.
- Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* 82(4), 711–732.
- Kim, S., N. Shephard, and S. Chib (1998). Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models. *Review of Economic Studies* 65(3), 361–93.
- Lopes, H. and M. West (2004, JAN). Bayesian model assessment in factor analysis. *STATISTICA SINICA* 14(1), 41–67.
- Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies* 72(3), 821–852.
- Raftery, A. E. and S. Lewis (1992). How many iterations in the Gibbs sampler. In *In BAYESIAN STATISTICS*.
- Stock, J. and M. Watson (2012). Disentangling the channels of the 2007-2009 recession. *Brookings Papers on Economic Activity Spring 2012*, 81–135.
- Stock, J. H. and M. W. Watson (2005). Understanding changes in international business cycle dynamics. *Journal of the European Economic Association* 3(5), 968–1006.

Centre for Applied Macro - and Petroleum economics (CAMP)
will bring together economists working on applied macroeconomic issues, with special emphasis on petroleum economics.

BI Norwegian Business School
Centre for Applied Macro - Petroleum economics (CAMP)
N-0442 Oslo

<http://www.bi.no/camp>