Australian National University

Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Boom or gloom? Examining the Dutch disease in a two-speed economy

CAMA Working Paper 76/2013
December 2013

Hilde C. Bjørnland
BI Norwegian Business School
Norges Bank and
Centre for Applied Macroeconomics Analysis

Leif Anders Thorsrud
BI Norwegian Business School
Norges Bank

Abstract

Traditional studies of the Dutch disease do not typically account for productivity spillovers between the booming energy sector and non-oil sectors. This study identifies and quantifies these spillovers using a Bayesian Dynamic Factor Model (BDFM). The model allows for resource movements and spending effects through a large panel of variables at the sectoral level, while also identifying disturbances to the real oil price, global demand and non-oil activity. Using Norway as a representative case study, we find that a booming energy sector has substantial spillover effects on the non-oil sectors. Furthermore, windfall gains due to changes in the real oil price also stimulate the economy, but primarily if the oil price increase is caused by global demand. Oil price increases due to, say, supply disruptions, while stimulating activity in the technologically intense service sectors and boosting government spending, have small spillover effects on the rest of the economy, primarily because of reduced cost competitiveness. Yet, there is no evidence of Dutch disease. Instead, we find evidence of a two-speed economy, with non-tradables growing at a much faster pace than tradables. Our results suggest that traditional Dutch disease models with a fixed capital stock and exogenous labor supply do not provide a convincing explanation for how petroleum wealth a effects a resource rich economy when there are productivity spillovers between sectors.

Keywords

Resource boom, oil prices, Dutch disease, learning by doing, two-speed economy, Bayesian Dynamic Factor Model (BDFM)

JEL Classification

C32, E32, F41, Q33

Address for correspondence:

(E) cama.admin@anu.edu.au

Boom or gloom? Examining the Dutch disease in a two-speed economy*

Hilde C. Bjørnland† Leif Anders Thorsrud‡

December 6, 2013

Abstract

Traditional studies of the Dutch disease do not typically account for productivity spillovers between the booming energy sector and non-oil sectors. This study identifies and quantifies these spillovers using a Bayesian Dynamic Factor Model (BDFM). The model allows for resource movements and spending effects through a large panel of variables at the sectoral level, while also identifying disturbances to the real oil price, global demand and non-oil activity. Using Norway as a representative case study, we find that a booming energy sector has substantial spillover effects on the non-oil sectors. Furthermore, windfall gains due to changes in the real oil price also stimulates the economy, but primarily if the oil price increase is caused by global demand. Oil price increases due to, say, supply disruptions, while stimulating activity in the technologically intense service sectors and boosting government spending, have small spillover effects on the rest of the economy, primarily because of reduced cost competitiveness. Yet, there is no evidence of Dutch disease. Instead, we find evidence of a two-speed economy, with non-tradables growing at a much faster pace than tradables. Our results suggest that traditional Dutch disease models with a fixed capital stock and exogenous labor supply do not provide a convincing explanation for how petroleum wealth affects a resource rich economy when there are productivity spillovers between sectors.

JEL-codes: C32, E32, F41, Q33

Keywords: Resource boom, oil prices, Dutch disease, learning by doing, two-speed economy, Bayesian Dynamic Factor Model (BDFM)

*This paper is part of the research activities at the Centre for Applied Macro and Petroleum economics (CAMP) at the BI Norwegian Business School. The authors would like to thank Francesco Ravazzolo, Ragnar Torvik and Benjamin Wong for valuable comments. Leif Anders Thorsrud thanks Centre for Applied Macroeconomic Analysis (CAMA) at Australian National University for hospitality while he was visiting as PhD scholar. The usual disclaimer applies. The views expressed in this paper are those of the authors and do not necessarily reflect the views of Norges Bank.

†CAMP - BI Norwegian Business School, CAMA and Norges Bank. Email: hilde.c.bjornland@bi.no

‡CAMP - BI Norwegian Business School, CAMA and Norges Bank. Email: leif.a.thorsrud@bi.no

1 Introduction

Over the last decade, the value of the Norwegian oil and gas industry - including services - grew by approximately 90 percent, while employment in this industry grew by 70 per cent. No other industry exhibited such growth rates.

The oil and gas boom in the North Sea has been the principal, but by no means, only cause of this substantial growth. Strong rises in oil and gas prices have caused Norway's terms of trade to double since 2001. These price rises have profound effects on the economy, as they constitute both a large shift in relative prices, which induces resource movements between industries, and a large increase in real incomes, which boosts aggregate demand in the overall economy.

While the recent financial crisis has suggested that energy rich countries - such as Norway - have occupied a different and better position than many other indebted industrial countries,[1] it is not clear that the gains from the energy sector benefited domestic sectors equally. For instance, employment in the construction and business sectors in Norway has increased by 30-40 percent over the last decade, while employment in the manufacturing industry and the retail, hotel and service industry has either fallen or hardly grown.

The energy boom has prompted much discussion of Norway having become a two-speed economy. There are concerns that the gains from the boom largely accrue to the profitable sectors servicing the energy industry, such as the business services, financial and construction sectors, while the rest of the country is being negatively affected by increased wage costs, an appreciated exchange rate and a lack of competitiveness as a result of the boom. Such a phenomenon has commonly been referred to in the literature as the Dutch disease, based on similar experiences in the Netherlands in the 1960s.[2] Concerns are also raised in other resource rich countries recently, such as the petroleum producer Canada and the mineral rich Australia.[3]

Much theoretical work has analyzed the benefits and costs of energy discoveries (see, e.g., Corden (1984) for a survey), but there have been relatively few empirical studies. Those that have investigated the empirical relationship between a booming energy sector and the macro economy have typically employed a structural vector autoregression (SVAR), which only includes a single sector such as manufacturing in each model, see, e.g., Hutchison (1994) and Bjørnland (1998), or a panel data approach that studies common movements in manufacturing across numerous countries, see, e.g., Ismail (2010). However, neither of these approaches accounts for all of the cross-sectional co-movement of variables within a country. That is, spillovers between sectors of the economy can be substantial due to intermediate inputs between the sectors and induced effects through increased demand and income in the energy sector or the sectors that are indirectly af-

[1] Mehlum et al. (2006) argue that in countries with strong property rights protection and little corruption, natural resources may have contributed to growth.

[2] Following the discovery and development of natural gas industries in the 1960s, the Netherlands experienced a period of real exchange rate appreciation relative to other nations and a corresponding loss of competitiveness for traditional industries that eventually contracted.

[3] See e.g., Lama and Medina (2012) for a discussion of the usefulness of exchange rate stabilization in relation to Dutch disease in Canada, and Corden (2012) for a discussion of the fast growing Australian mining sector on the one side and the lagging manufacturing sectors on the other.

fected. In addition, there may be shared productivity dynamics. For Norway, where oil extraction is conducted offshore and with greater technical difficulties than for typical onshore extraction, productivity (knowhow) spillovers through high-tech industries might be substantial. Lastly, there are other sources of shocks that could be causing the economic boom that need to be controlled for, such as common global demand shocks.

We contribute to this area of the literature by explicitly identifying and quantifying the linkages between a booming energy sector and sectoral performance in the rest of the economy, while also allowing for independent disturbances to the real price of oil, world activity and domestic (non-oil) activity. Our main focus is to test the hypothesis of Dutch disease by separately examining the windfall gains associated with energy booms and real oil price changes for various sectors, while also controlling for changes in global and domestic activity. Having established the linkages, we analyze how the domestic economy responded to the energy boom and energy price changes in different periods.

To explore these questions, we estimate a Bayesian Dynamic Factor Model (BDFM), that includes separate activity factors for oil and non-oil sectors in addition to global activity and the real price of oil. The BDFM is particularly useful to answer the research questions we address.[4] First, the interdependence between the different branches of an economy - traditionally measured by the input-output tables from the National Accounts - do not account for the indirect spillover effects (productivity or demand) between different sectors. Thus, co-movement across sectors due to common factors, i.e., oil or non-oil, is not captured by observable variables alone. Conversely, in the BDFM, latent common factors can be identified and estimated simultaneously with the rest of the model's parameters. Thus, the size and sign of spillover effects can be derived and analyzed. Second, to quantify the spillover effects across a large cross section of sectors and variables, standard multivariate time series techniques are inappropriate due to the curse of dimensionality. The BDFM is designed for data rich environments such as ours. Third, macroeconomic data are often measured with noise and errors. In the factor model framework, we can separate these idiosyncratic noise components from the underlying economic signal.

We extend the literature in three ways. First, to the best of our knowledge, this is the first paper to explicitly analyze and quantify the linkages between a booming energy sector and sectoral performance in the domestic economy using a structural model, while also allowing for explicit disturbances in real oil prices, world activity and activity in the non-oil sector. Thus far, very little is known about the effect that energy booms have on the rest of the economy in a resource rich economy, and equally important, if it is the booms themselves or the windfall gains associated with real oil price changes that are the most important. Second, given the large number of variables and industries included in the analysis, this is also the most comprehensive analysis to date of the relationship between energy booms and macroeconomic activity at the industry level in a resource rich economy. We lastly show that standard multivariate methods do not adequately quantify resource booms in a resource rich country such as Norway. The BDFM does, and the use of this modeling framework to analyze the Dutch disease is novel in the literature.

[4] As discussed in, e.g., Boivin and Giannoni (2006), there is a close resemblance between theoretical DSGE models and Dynamic Factor Models. Moreover, Bai and Wang (2012) discuss how the DFM can be related to the Structural VAR literature.

Our main conclusion emphasizes that a booming energy sector has significant and large productivity spillovers on non-oil sectors, effects that have not been captured in previous analysis. In particular, we find that the energy sector stimulates investment, value added, employment and wages in most tradable and non-tradable sectors. The most positively affected sectors are construction, business services and real estate.

Furthermore, windfall gains due to changes in the real oil price also stimulate the economy, particularly if the oil price increase is associated with a boom in global demand. Oil price increases due to, say, supply disruptions, while stimulating activity in the technologically intense service sectors and boosting government spending, have small spillovers effects to the rest of the economy, in part because of substantial real exchange rate appreciation and reduced cost competitiveness. Yet, there is no evidence of Dutch disease as experienced in the Netherlands in the 1970s, where the manufacturing sector contracted. Instead, we find evidence of a two speed economy, with employment in the manufacturing sector lagging behind the booming service sectors.

Our results suggest that traditional Dutch disease models with a fixed capital stock and exogenous labor supply do not provide a convincing explanation for how petroleum wealth affects a resource rich economy when there are productivity spillovers between the various sectors.

The remainder of the paper is structured as follows. In Section 2, we briefly discuss the theoretical literature on Dutch disease and present some stylized facts. Section 3 and 4 describe the data and the model, the identification strategy and the estimation procedure in detail. Our main results are reported in Section 5, while in Section 6, we show that these results are robust to numerous specification tests. Section 7 concludes.

2 Macroeconomic impacts of an energy discovery

There is a substantial theoretical literature on the Dutch disease, see, for instance Bruno and Sachs (1982), Corden and Neary (1982), Eastwood and Venables (1982), Corden (1984), Van Wijnbergen (1984) and Neary and van Wijnbergen (1984). The general finding in most of these papers is that there is an inverse long run relationship between increased exploitation of natural resources and growth in the manufacturing sector, similar to what the Netherlands experienced in the 1960s.

Although the disease most often refers to the consequences of the discovery of natural resources, it can also refer to any development that results in a large inflow of foreign currency, such as a sharp increase in commodity prices. As such, the analysis of the effects of a commodity price shock on a resource rich economy is simply a special case of the Dutch disease.

The standard theory model that these papers build on assumes a non-traded goods and service sector and two traded goods sectors: the booming sector and the lagging sector, also called the non-booming tradable sector. The booming sector is usually the extraction of oil or natural gas, but can also be mining. The lagging sector generally refers to manufacturing, but can also be agriculture when traded. The non-traded goods and service sector includes the government sector and other non-traded sectors.

The direct impact of oil and gas resources (or any other sectoral boom) is experienced through an increased demand for resources and goods and services in the energy producing sector. This is usually referred to as a the *Resource Movement Effect*. The increased demand for goods and services by the energy sector will lead to an indirect (secondary) effect of increased demand for resources by the sectors that will produce goods and services for the energy sector. If income in the energy sector has increased, there will also be a further (induced) effect of increased demand for goods and services. These induced effects are usually described as the *Spending Effects*, and will cause a real appreciation that will hurt some sectors and benefit others.

More formally, Corden and Neary (1982) assume that the booming sector (B) and the tradeable sectors (T) produce tradeables given world prices, whereas the prices for non-tradables (N) are given by domestic factors. The energy boom is understood as an exogenous (unpredicted) technical improvement in B. The resource movement effect will increase demand for labor in B, as the marginal product of labor increases due to the boom, given constant wages in terms of the tradables. Thus, there will be a movement of labor out of T and N into B. The movement of labor from T to B will directly reduce output in T, whereas the movement of labor from N to B at constant prices will initially reduce the supply of N and create an excess demand for N. In response to this excess demand, the price for non-tradables in terms of tradables will rise, which will produce real appreciation and further movements of resources out of T into N.

The aggregate income of the factors initially employed in the booming sectors will also rise. This will lead to a spending effect, directly by the factor owners in B or indirectly by the government that collects (part of) the income through taxes. With positive income elasticity of demand for N, the price of N relative to the price of T must rise, yielding a further real appreciation. Given full employment of all resources, this real appreciation will induce additional movement of labor from T to N.

Although the simple model of Dutch disease predicts that manufacturing will eventually contract as the energy sector expands, there are several ways that the core model may be altered. By changing some of the underlying assumptions (for instance, by allowing the factors of production to be mobile), the predicted effects of energy booms on the manufacturing sector may be less severe, and in fact, in some cases there may not be Dutch disease at all. In particular, if one is initially in a situation where domestic resources are not fully employed prior to the energy boom, the boom may actually have a stimulative effect on industry.

Output in the manufacturing industry may also increase if one assumes that the energy sector has its own specific factor, labor is mobile between the three sectors but capital is only mobile between the non-tradable and the tradable sector. This constitutes a miniature Heckscher-Ohlin economy, where one sector will be labor intensive while the other will be capital intensive. In this case, the resource movement effect will cause the output of the capital intensive industry to expand (as labor is moving out of the labor intensive industry and into the booming energy sector during the boom). If the tradable sector is the capital intensive industry, and the (negative) spending effect on output in the tradable sector is smaller than the resource movement effects, output in the tradable sector may actually increase, see Corden (1984) for a further discussion.

More recently, Torvik (2001) advanced a model in which there is learning by doing (LBD) in both the traded and non-traded sectors, as well as learning spillovers between the sectors. Under certain conditions, this will imply a real exchange rate depreciation in the long run, due to a shift in the steady state relative productivity between the traded and non-traded sectors. In contrast to the standard models of the Dutch disease, production and productivity in both sectors can then increase.[5]

Thus, while the traditional theory of Dutch disease implies that the tradable sector will eventually contract as the energy sector expands, there are several ways the dynamics of the core model may change such that the predicted effects of energy booms on the tradable sector may be less severe than in the basic case, and in some cases there may be no Dutch disease at all.

2.1 Dutch disease and stylized facts of Norway

Figure 1 depicts the evolution of the important variables involved in the debate on Dutch disease. Key to the discussion is the real oil price and the real exchange rate, depicted for the period 1983-2012 in Figures 1a and 1b, respectively. Two features stand out. The real exchange rate depreciated considerably between the beginning of the 1980s and 2000, after which it appreciated sharply.[6] Taking everything else as given, the prolonged period of real exchange rate depreciation in the first half of the sample fits nicely into the framework of a model that allows for productivity advances due to learning by doing within and between sectors, such as in Torvik (2001), discussed above. The timing of the strong appreciation in the latter half of the sample corresponds to the increase in the real oil price, and thus indicating a more classical Dutch disease pattern.

Figure 1c shows the evolution of employment by industry since 1996 (from which data are available). The figure suggests a two speed economy, with resources rapidly moving into both the booming oil and gas industry and the profitable service sectors, while employment in other sectors such as manufacturing is lagging behind.

Lastly, Figure 1d illustrates the importance of investments in the energy sector over the business cycle for GDP in Mainland Norway (value added of total GDP minus the oil and gas sector). The figure clearly shows a leading and pro-cyclical relationship between investment in the oil sector and GDP in Mainland Norway (the correlation coefficient is 0.6 when oil investment leads the business cycle by 4 quarters), except during the Norwegian banking crisis in the early 1990s, when other factors were at play. However, the figure also indicates that since 2003/2004, the dynamics of the economy are not all driven by oil. While oil investment is still pro-cyclical, the stimulus from the oil sector seems small compared to the stimulus during the booms in the early 1980s and mid 1990s. Other factors will have to explain the boom in the mainland economy in this period.

Thus, there is evidence that the energy sector has positive spillovers to the mainland economy, albeit possibly to a smaller extent in the most recent boom and bust. However,

[5]Traditional LBD models such as Van Wijnbergen (1984), which accounts for LBD by assuming that productivity in the tradable sector depends on production in the first period alone, or Sachs and Warner (1995), which employs an endogenous growth model, find unambiguously that productivity will decline.

[6]This is the effective exchange rate, where an increase implies appreciation.

Figure 1. Stylized facts

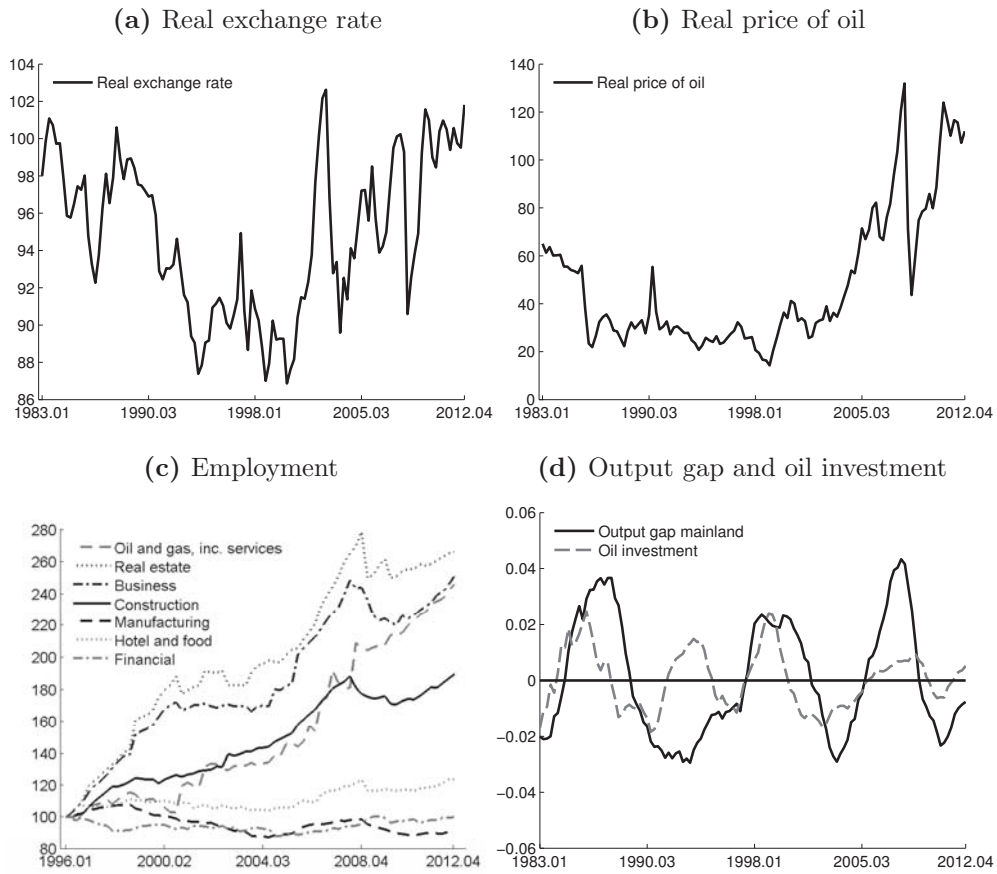

Note: All employment series are on a log scale, and normalized to 100 in 1996:Q1. Figure 1d displays the smoothed Hodrick-Prescott filtered output-gap in GDP Mainland Norway as well as the smoothed fraction between cyclically adjusted oil investments and the trend growth in GDP Mainland Norway.

three concurrent evolutions after 2001, the appreciation of the currency, the strong rise in commodity prices and strong growth in the oil sector relative to the manufacturing sector, suggest a typical case of Dutch disease, where some sectors are growing at the expense of others. We examine this subject below.

3 Theory meets data

How can one apply the theoretical model to the data? The approach we adopt relies on the standard model presented in Corden and Neary (1982), but augmented in some dimensions by allowing for productivity spillovers between sectors of the economy. In particular, we develop a framework where the energy sector uses its own factor of production and develops its own specific productivity dynamics, but there may be instantaneous spillovers to all the other domestic sectors. Thus, developments in the energy sector will be exogenous at time t, but after a period, it may respond to the other sectors of the economy. For instance, capacity constraint in the domestic economy could eventually also

affect the energy sector. Furthermore, we assume that the tradable and the non-tradable sectors of the economy have their own factors of production and develop their own productivity dynamics, but there may be instantaneous spillovers between the tradable and non-tradable sectors (in addition to the spillovers from the energy sector). Thus, we allow for learning by doing in both the traded and the nontraded sector and learning spillovers between these sectors, as suggested in Torvik (2001). Finally, we will allow for common shocks to the global oil market.

Given the framework described above, we can identify four factors with associated shocks that have the potential to affect all sectors: Two shocks will relate to the dynamics in domestic economy. The *energy boom (or oil activity shock)*[7] and the *non-oil activity shock*. We let energy booms represent an unexpected technical improvement or windfall discovery of new resources in the energy sector, while the non-oil activity shock controls for the remaining domestic impulses (tradable and non-tradable) contemporaneously unrelated to the oil sector. In addition, we allow for two shocks that relate to the dynamics specific to the global oil market, an *oil specific shock* and a *global demand shock*. The oil specific shock allows for a windfall gain due to higher real oil prices from, say, a supply disruption in oil production, while the global demand shock allows for higher oil prices due to increased global activity.

A central premise of the theory is that the energy sector supports many more jobs than it generates, directly owing to its long supply chains and spending by employees and suppliers. Thus, to accommodate resource movement and spending effects, we employ a broad range of sectoral employment, production, wage and investment series for the Norwegian economy. The intuition is as follows: First, energy extraction may stimulate value-added among downstream industries, such as refining, or industries that provide the energy sector with goods and services. This will generate additional jobs in excess of those directly produced in the energy sector. Furthermore, energy extraction can induce a reallocation of labor from the less profitable sectors into the booming sectors. We capture these effects by including data for value added and employment at the industry level.

Second, there will be induced spending effects through the wages paid to workers in the energy sector or the sectors that are indirectly affected. Moreover, as the booming sector also pays significant taxes on its increased income, these benefits will easily spread to the whole economy. However, as Norway has a centralized wage bargaining system, we do not include wage data for all sectors, which would be highly correlated. Instead, for wages, we separate between the booming sector (oil and gas), the mainland (non-oil) sector and the public sector. Note that the public sector is included to also account for the pass through of changes in oil income to the economy.

Third, specific sectors of the economy may benefit due to productivity spillovers when the patterns of domestic demand shifts in their favor. The loser are those producers that do not benefit from these spillovers, what Corden (2012) terms the lagging sector. To account for these productivity spillovers, we also include investment at the sectoral level. We separate investments in the same way as wages.

Naturally, we include the real price of oil and the real exchange rate, which are core

[7] We will use the terms energy booms and oil activity shocks interchangeably

factors in the Dutch Disease literature. The real price of oil is constructed based on Crude Oil-Brent prices, deflated using the US CPI. As such, it is meant to reflect the global real price of oil. The notion is that an increase in the real oil price will directly cause the exchange rate to appreciate via the terms of trade. This will have adverse effects on the tradable sector, leading to a period of de-industrialization. While this is only one part of the question we analyze, many papers have only focused on the effects of an oil price increase when analyzing the Dutch disease, see, e.g., Charnavoki and Dolado (2012) and the references therein.

The de-industrialization effect described above could be a feature of Dutch disease, but it could also be a common feature of many open economies. To control for the state of the international business cycle, we also include a measure of global activity. We measure global or world activity as the simple mean of four-quarter logarithmic changes in real GDP in: China, Denmark, Germany, Japan, the Netherlands, Sweden, the UK and the US. This set of countries includes Norway's most important trading partners and the largest economies in the world.

In sum, this gives a panel of 50 international and domestic data series, covering a sample period from 1996:Q1 to 2012:Q4.

Our focus is on quantifying economic fluctuations over the horizons relevant for medium term macroeconomic policy and over business cycle horizons. To capture the economic fluctuations of interest, we transform all variables to four-quarter logarithmic changes; $log(x_{i,t}) - log(x_{i,t-4}))$.[8] Lastly, all variables are demeaned before estimation. Further details on the data are provided in Appendix A.

3.1 Quantifying the resource boom - a simple attempt

The petroleum sector's share of total GDP in Norway has fluctuated around 20 percent the last decade. However, although the sector is capital intensive, it does not operate in isolation. According to Eika et al. (2010), the total use of (non-oil) resources in the petroleum sector was equivalent to 17 percent of the GDP of Mainland Norway (based on input-output tables from 2008).[9] However, this measure of petroleum dependency likely represents a lower bound on the Norwegian economy's oil dependence. Typically, it will underestimate the links across sectors, as it does not account for the effects induced over time from increased demand and income in the energy sector or the sectors that are indirectly affected (e.g., the government sector).

To obtain an initial impression of the oil dependence of the Norwegian economy, one can run a series of simple structural vector autoregressions (VARs) relating the oil sector to the mainland economy. The analysis below is an attempt in that direction, although as we will see, it is far from adequate in capturing the spillovers we seek.

Panels (a)-(c) of Figure 2 report the responses of GDP in Mainland Norway to three different shocks: Global activity, oil price (specific) and oil activity, respectively. Panels

[8] We experimented with specifying the model using data transformed to quarterly changes, i.e., $log(x_{i,t}) - log(x_{i,t-1}))$. However, for Norwegian data, such transformations yield a very weak factor structure, making the dynamic factor model, see Section 4, less appropriate.

[9] This number is calculated based on the intermediate inputs to the petroleum sector, adjusted for the indirect use of resources between the different sectors.

Figure 2. VAR (non) evidence

Impulse responses - GDP Mainland Norway

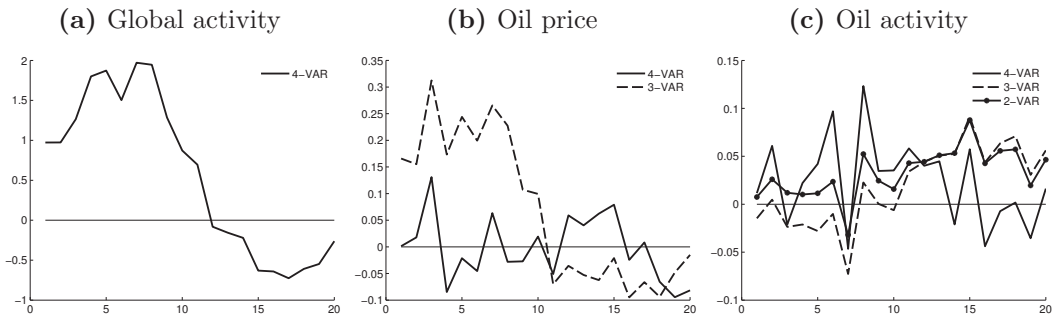

Variance decompositions - GDP Mainland Norway

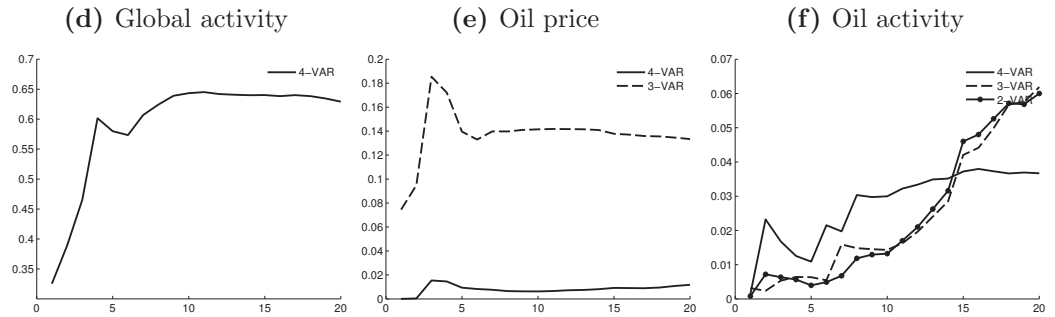

Note: The figures report impulse responses and variance decompositions of GDP in Mainland Norway to three structural shocks: An international activity shock, an oil price shock, and an activity shock to the petroleum sector. Three different VAR specifications are estimated: 4-VAR (world activity, real price of oil, oil activity, mainland activity), 3-VAR (real price of oil, oil activity, mainland activity), 2-VAR (oil activity, mainland activity). All variables are transformed to log year on year changes, and all VARs are specified with eight lags. The structural shocks are identified employing a recursive ordering.

(d)-(f) present the variance decomposition of the same three shocks. Three different VAR models are specified. In the 2-VAR, we jointly model oil activity and mainland activity, in the 3-VAR we add the real price of oil, while in the 4-VAR world activity is also included, see Figure 2 for more details. None of the VAR specifications yield results that provide an economic meaningful depiction for quantifying a resource boom in a two-speed economy. That is, an unexpected positive innovation in oil activity increases GDP in Mainland Norway in all VAR specifications (Panel c), but the shock explains a negligible share of the variance in the GDP of Mainland Norway (3-6 percent, see Panel f). This is at odds with conventional wisdom, earlier research (see, e.g., Bjørnland (1998) and Larsen (2006)), and most important, the National Account statistics described above.

However, the positive and large effects of a world activity shock (Panels a and d) is in accordance with new and existing evidence of international business cycle synchronization, see, e.g., Kose et al. (2003), Stock and Watson (2005b), and Thorsrud (2013). Furthermore, an unexpected increase in the real price of oil increases mainland activity, but primarily in the 3-VAR specification, see Panel (b). However, as shown in, e.g.,

Aastveit et al. (2012), a large fraction of the variation in the real price of oil can be attributed to global activity. Only the 4-VAR specification takes this into account by allowing the oil price to also respond to global activity. Thus, the oil price shock in the 3-VAR model is likely a combination of world activity innovations and pure unexpected oil price innovations. This renders the structural interpretation of this model dubious and suggests that the 4-VAR specification is more appropriate.[10]

Why do the structural VAR models fail to explain the resource boom in a two speed economy? The answer is simple. They do not take all the cross-sectional co-movement of main sectoral variables into account. That is, oil activity alone does not accurately measure the resource moving and spending effects induced by an oil boom, or any potentially shared productivity developments.

The Dynamic Factor Model (DFM) proposed in this study solves these issues. Within the DFM framework, the co-movement of a large cross section of variables is assumed to be driven by a few latent (or observable) factors. The factors and the unexpected innovations (shocks) to the factors can be identified, and structural analysis can be conducted. Geweke (1977) is an early example of the use of the DFM in the economic literature. Kose et al. (2003) and Mumtaz et al. (2011) are more recent examples, while Stock and Watson (2005a) provide a brief overview of the use of this type of models in economics. In the next section, we provide a more detailed description of the DFM, and identification and estimation within this framework, before turning to the results in section 5.

4 The Dynamic Factor Model

We specify a Dynamic Factor Model (DFM). As noted above, this model is particularly useful in a data rich environment such as ours, where common latent factors and shocks are assumed to drive the co-movements between economic variables in the Norwegian economy.

The DFM is given by equations 1 and 2:

$$y_t = \lambda_0 f_t + \cdots + \lambda_s f_{t-s} + \epsilon_t \tag{1}$$

where the $N \times 1$ vector y_t represents the observables at time t. λ_j is a $N \times q$ matrix with dynamic factor loadings for $j = 0, 1, \cdots, s$, and s denotes the number of lags used for the dynamic factors f_t. In our application the $q \times 1$ vector f_t contains both latent and observable factors. Lastly, ϵ_t is an $N \times 1$ vector of idiosyncratic errors.

The dynamic factors follow a VAR(h) process:

$$f_t = \phi_1 f_{t-1} + \cdots + \phi_h f_{t-h} + u_t \tag{2}$$

where u_t is a $q \times 1$ vector of VAR(h) residuals. The idiosyncratic and VAR(h) residuals are assumed to be independent:

$$\begin{bmatrix} \epsilon_t \\ u_t \end{bmatrix} \sim i.i.d.N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} R & 0 \\ 0 & Q \end{bmatrix} \right) \tag{3}$$

[10] The variables included in the VARs are noise measures of the underlying business cycles. However, the results reported in Figure 2 are robust to using HP-filtered data.

Further, in our application R is assumed to be diagonal.

The model described above can easily be extended to the case with serially correlated idiosyncratic errors. In particular, we consider the case where $\epsilon_{t,i}$, for $i = 1, \cdots, N$, follows independent AR(l) processes:

$$\epsilon_{t,i} = \rho_{1,i}\epsilon_{t-1,i} + \cdots + \rho_{l,i}\epsilon_{t-l,i} + \omega_{t,i} \tag{4}$$

where l denotes the number of lags, and $\omega_{t,i}$ is the AR(l) residuals with $\omega_{t,i} \sim i.i.d.N(0, \sigma_i^2)$. I.e.:

$$R = \begin{bmatrix} \sigma_1^2 & 0 & \cdots & 0 \\ 0 & \sigma_2^2 & \ddots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & \cdots & \sigma_N^2 \end{bmatrix}, \tag{5}$$

4.1 Identification

Equations 1 and 2 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 equation 1:

$$\lambda_0 = \begin{bmatrix} \lambda_{0,1} \\ \lambda_{0,2} \end{bmatrix} \tag{6}$$

where $\lambda_{0,1}$ is a $q \times q$ identity matrix, and $\lambda_{0,2}$ is left unrestricted. As shown in Bai and Ng (2010) and Bai and Wang (2012), 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, u_t, can be linked to structural shocks that are implied by economic theory.

In our application, we set $q = 4$ and identify four factors: global activity, the real price of oil, Norwegian oil specific activity, and Norwegian non-oil (Mainland) activity. The number of factors and names are motivated by the model as discussed in Section 3 above.[11] Of these four factors, the first two are observable and naturally load with one on the corresponding element in the y_t vector. The two latter factors must be inferred from the data. We require that the Norwegian oil specific activity factor loads with one on value added in the petroleum sector, and the Norwegian Mainland activity factor loads with one on value added in Mainland Norway. Note that while this identifies the factors, it does not mean that the factors and the observables are identical as we will use the full information set to extract the factors.

Based on a minimal set of identification restrictions, we identify four structural shocks: a global demand shock, an oil specific shock, a Norwegian oil activity shock (energy booms) and a Norwegian non-oil (domestic) activity shock. The shocks are identified by imposing a recursive ordering of the latent factors in the model, i.e. $f_t = [f_t^{gact}, f_t^{oilp}, f_t^{oact}, f_t^{noact}]'$,

[11] Moreover, as we show in Appendix C.1, four factors also explain a large fraction of the variance in the dataset.

such that $Q = A_0 A_0'$. Specially, the mapping between the reduced form residuals u_t and structural disturbances e_t, $u_t = A_0 e_t$, is given by:

$$\begin{bmatrix} u_t^{gact} \\ u_t^{oilp} \\ u_t^{oact} \\ u_t^{noact} \end{bmatrix} = \begin{bmatrix} a_{11} & 0 & 0 & 0 \\ a_{21} & a_{22} & 0 & 0 \\ a_{31} & a_{32} & a_{33} & 0 \\ a_{41} & a_{42} & a_{43} & a_{44} \end{bmatrix} \begin{bmatrix} e_t^{gdem} \\ e_t^{oils} \\ e_t^{oact} \\ e_t^{noact} \end{bmatrix} \tag{7}$$

where e_t^i are the structural disturbances for $i = [gdem, oils, oact, noact]$, with $e_t e_t' = I$, and $[gdem, oils, oact, noact]$ denote global demand, oil specific, Norwegain oil activity and non-oil activity, respectively.

For most energy importing countries, a higher price of oil causes production costs and inflation to gradually increase, thereby eventually affecting overall activity. We therefore follow the usual assumption from both theoretical and empirical models of the oil market, and restrict global activity to respond to oil specific disturbances with a lag. This restriction is consistent with the sluggish behavior of global economic activity after each of the major oil price increases in recent decades.

Furthermore, any unexpected news regarding global demand is assumed to affect the real price of oil contemporaneously. As such, and consistent with recent work, we do not treat the real price of oil as exogenous to the rest of the macro economy, see, e.g., Aastveit et al. (2012). In doing so, we confirm that both global demand and the oil specific shock can drive up oil prices significantly. However, whereas the global demand shock also stimulates global activity, the oil specific shock reduces global activity (with a lag) and can therefore be interpreted as an adverse supply shock to the oil market.

In the short run, disturbances originating in the Norwegian economy are exogenous to global activity and the real oil price. These are plausible assumptions, as Norway is a small open economy that only accounts for less than three percent of global oil production. However, both the oil and the non-oil domestic activity factors respond to unexpected disturbances in global activity and the real price of oil on impact. In a small open economy such as Norway, news regarding global activity will affect variables such as the exchange rate, the interest rate, asset prices and consumer sentiments contemporaneously, and thereby affect overall demand in the economy. Norway is also a net oil exporter. Thus, any disturbances to the real price of oil will most likely rapidly affect both the demand and supply side of the economy.

Lastly, in the short run, the oil activity factor is exogenous to the rest of the domestic economy but can affect the other sectors contemporaneously (for instance via productivity spillovers). However, and as discussed in Section 3, after a period we allow the energy sector to respond to the dynamics in the other sectors of the economy.

4.2 Estimation

Let $\tilde{y}_T = [y_1, \cdots, y_T]'$ and $\tilde{f}_T = [f_1, \cdots, f_T]'$, and define $H = [\lambda_0, \cdots, \lambda_s]$, $\beta = [\phi_1, \cdots, \phi_h]$, Q, R, and $p_i = [\rho_{1,i}, \cdots, \rho_{l,i}]$ for $i = 1, \cdots, N$, as the model's hyper-parameters.

Inference in our model can be performed using both classical and Bayesian techniques. In the classical setting, two approaches are available, two-step estimation, and maximum

likelihood estimation (ML). In the former, $\tilde{f}_T$, H and R are first typically estimated using the method of principal components analysis (PCA), then the dynamic components of the system, A and Q, are estimated conditional on $\tilde{f}_T$, H and R. Thus, the state variables are treated as observable variables. If estimation is performed using ML, the observation and state equations are estimated jointly. However, employing ML still involves some type of conditioning. That is, we first obtain ML estimates of the model's unknown hyper-parameters. Then, to estimate the state, we treat the ML estimates as if they were the true values for the model's nonrandom hyper-parameters. In a Bayesian setting, both the model's hyper-parameters and the state variables are treated as random variables.

We estimated the DFM using both the two-step procedure in the classical setting and Bayesian estimation. The results reported in section 5 are not qualitatively affected by the choice of estimation method. However, we prefer the Bayesian approach primarily due to: 1) In contrast to the classical approach, inferences regarding the state are based on the joint distribution of the state and the hyper-parameters, not a conditional distribution. 2) ML estimation would be computationally intractable given the number of states and hyper-parameters. 3) Our data are based on logarithmic year-on-year differences. This spurs autocorrelation in the idiosyncratic errors.

In a Bayesian setting, the model can readily be extended to accommodate these features of the error terms. In a classical two-step estimation framework, this is not the case. Furthermore, in the two-step estimation procedure, it is not straightforward to include lags of the dynamic factors in observation equation.

Thus, our preferred model is a Bayesian Dynamic Factor Model (BDFM). We set, $s = 2$, $h = 8$, and $l = 1$. That is, we include 2 lags for the dynamic factors in the observation equation (see equation 1), 8 lags in the transition equation (see equation 2), and let the idiosyncratic errors follow AR(1) processes (see equation 4).[12] In section C.1 we explain the choice of this particular specification and analyze its robustness.

4.2.1 The Gibbs sampling approach

Bayesian estimation of the state space model is based on Gibbs simulation, where the following three steps are iterated until convergence is achieved:
Step 1: Conditional on the data ($\tilde{y}_T$) and all the parameters of the model, generate $\tilde{f}_T$
Step 2: Conditional on $\tilde{f}_T$, generate β and Q
Step 3: Conditional on $\tilde{f}_T$, and data for the i-th variable ($\tilde{y}_{T,i}$), generate H_i, R_i and p_i for $i = 1, \cdots, N$

In Appendix D we describe each step in greater detail and document the employed prior specifications. We simulate the model using a total of 50000 iterations. A burn-in period of 40000 draws is employed, and only every 5th iteration is stored and used for inference.[13]

[12] Note that we let $s = 0$ and $l = 0$ when estimating the DFM using the two-step estimation procedure.
[13] Standard MCMC convergence tests confirm that the Gibbs sampler converges to the posterior distribution. Convergence statistics are available on request.

5 Results

Our results are presented in the following subsections. We first present the identified factors before investigating how GDP, investment, employment and wages in the Mainland economy and the real exchange rate respond to the various shocks. Then we examine the sectoral reallocation following the energy booms and oil price shocks, before investigating the implications for spending in the public sector in greater detail.

5.1 Factors and global shocks

The upper panel of Figure 3 displays, from the left, the global activity factor, the real price of oil, the oil activity factor and the non-oil (Mainland) activity factor. The two first factors are treated as observables in the estimation. Accordingly, they are measured without uncertainty.

Global activity declined during the Asian crisis in the latter part of the 1990s, following the dot com bubble that burst in 2000/2001, and during the recent financial crisis. The latter trough is by far the most severe. Turning to the real oil price, Figure 3 suggests that the most pronounced cycles in the real price of oil follow global activity cycles. There is significant growth in the real oil price during the economic booms in 1999/2000 and 2006/2007 and a decrease in the real price of oil during the Asian crisis and the recent financial crisis.

It is more interesting to investigate the cyclical patterns of the estimated latent factors, i.e., oil activity and non-oil activity. Statistically, both factors are identified. As seen in the figure, they are also economically meaningful. The latent oil activity factor shows booms and busts that relate to the petroleum sector, such as the investment boom in the North Sea in the middle of the 1990s, the decline in activity from 2000 (when oil production peaked) and the decline in new investments in the period after the financial crisis. The non-oil factor shows cyclical patterns that are well in line with the conventional view of the Norwegian business cycle over the last two decades. The bust in 2002/2003, the subsequent boom, and the recent bust during the financial crisis stand out. As expected, the volatility of the oil activity factor is larger than that of the non-oil activity factor.

The estimation procedure we employ, see section 4, is inherently a smoothing algorithm. Thus, it is unsurprising that the oil and non-oil activity factors resemble the cyclical patterns of oil investment (cyclical contribution) and the GDP of Mainland Norway, respectively, both displayed in Figure 1d. Importantly, however, the factors and the observables are not identical. As stressed in section 2, the oil sector's contribution to the domestic economy comes through many more channels than investments alone. The information set used to extract the two latent factors reflects this, as do the estimated factors.

As discussed in section 3.1 above, we do not wish to treat the oil price as exogenous and allow for reverse causality from global activity to the oil price. This implies that both supply and demand shocks can affect oil prices. Figure 3, lower panel, illustrates this. It displays the effect of a global demand shock to global activity and the real oil price and subsequently the effect of an oil specific shock to the same two variables. While the global

Figure 3. Factors and global impulse responses

Factors

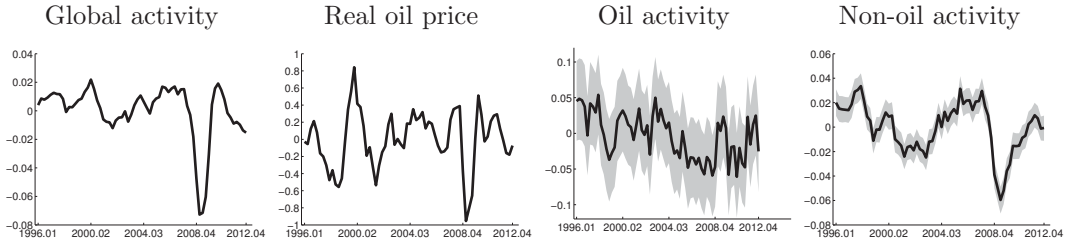

Global impulse responses

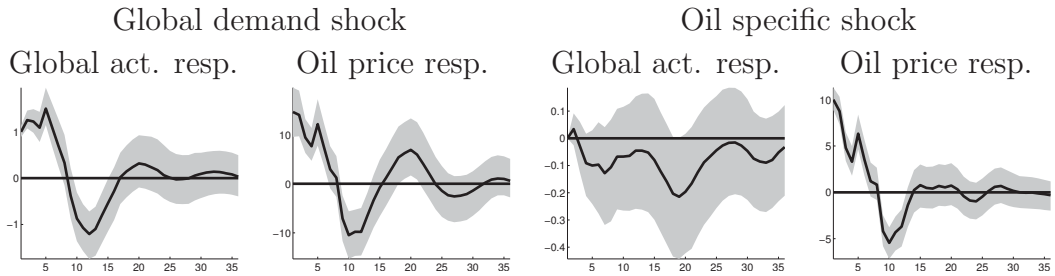

Note: The first row in the figure displays the observed variables and the estimated latent factors. The second row displays impulse responses. The responses are displayed in levels of the variables. The Global demand shock is normalized to a 1 percent increase, while the oil specific shock is normalized to increase the real price of oil with 10 percent. The black solid lines are median estimates. The gray shaded areas are 68 percent probability bands.

demand shock increases both activity and the real oil price, the oil specific shock generates a temporary inverse relationship between the oil price and global activity, equivalent to a supply type disturbance. Again, this is consistent with recent studies that have found that a large fraction of the variation in the real price of oil can be attributed to global demand, see e.g. Lippi and Nobili (2012) and Aastveit et al. (2012) among many others.

5.2 A resource rich economy

Table 1 displays the variance decomposition to the four identified shocks: oil activity (energy booms), oil specific, global demand and non-oil activity, for GDP, employment, investment and wages in the oil sector, the non-oil sector (Mainland Norway) and the public sector, as well as for the real exchange rate. Figure 4 then displays the impulse responses to the four identified shocks for the mainland economy and the real exchange rate.

As expected, the oil activity and oil specific shocks together explain 60-70 percent of the variation in production, employment, wages and investment in the petroleum sector. However, while the investment dynamics in the petroleum sector are strongly associated with oil specific shocks (that drive up oil prices), oil activity shocks are most important for value added and employment. Lastly, global demand shocks (that drive up oil prices) also affect the oil sector, and in particular petroleum investment. More than 20 percent of the variation in petroleum investment refers back to global demand and its effect via

Table 1. Variance decompositions

Variable	Sector & Horizon	Shock: Oil activity 4, 8	Oil specific 4, 8	Global demand 4, 8	Non-oil activity 4, 8
GDP	Oil	0.82, 0.69	0.13, 0.12	0.04, 0.13	0.02, 0.06
	Mainland	0.25, 0.32	0.06, 0.04	0.49, 0.44	0.20, 0.20
	Public	0.06, 0.05	0.48, 0.40	0.01, 0.05	0.45, 0.50
Employment	Oil	0.66, 0.54	0.24, 0.20	0.06, 0.12	0.04, 0.14
	Mainland	0.08, 0.04	0.12, 0.16	0.20, 0.28	0.59, 0.52
	Public	0.21, 0.15	0.18, 0.23	0.05, 0.08	0.56, 0.54
Wages	Oil	0.46, 0.41	0.36, 0.29	0.15, 0.17	0.03, 0.13
	Mainland	0.19, 0,08	0.05, 0.08	0.26, 0.38	0.49, 0.47
	Public	0.66, 0.37	0.08, 0.15	0.05, 0.15	0.21, 0.32
Other	Investment Oil	0.01, 0.03	0.74, 0.61	0.21, 0.20	0.04, 0.15
	Investment Mainland	0.17, 0.28	0.28, 0.16	0.49, 0.49	0.06, 0.06
	Real Exchange Rate	0.11, 0.22	0.67, 0.58	0.23, 0.20	0.00, 0.00

Note: Each row-column intersection reports median variance decompositions for horizons 4 (left) and 8 (right)

higher oil prices.

What are the implications for the rest of the economy? Clearly, the oil boom stimulates the mainland economy. In particular, Figure 4 shows that a boom in the energy sector that increases oil activity by one percent increases GDP and investment in the mainland sector by 0.4 and 0.7 percent, respectively, after 1-2 years. The effect is substantial; approximately 30 percent of the variation in each of these variables is explained by energy booms (see Table 1).

The spillovers from the energy sector to the labor market are more gradual. Employment and wages eventually increase after a year, peaking after 2-3 years. Ultimately, energy booms are more important for wage dynamics than for employment, explaining more than 20 percent of the changes in wages versus less than 10 percent of the employment variation in the mainland economy. The evidence is consistent with the view that productivity increases in the energy sector worked to raise labor income in all sectors via the centralized system of pay determination.

Lastly, the response in the real exchange rate is small and mostly insignificant, if anything, showing evidence of real depreciation. This helps to explain why energy booms can have such stimulative effects on the mainland economy.

There are two structural shocks that increase oil prices, an oil specific shock and a global demand shock. Figure 4 shows that an oil specific shock is strongly associated with real exchange rate appreciation. In fact, 60-70 percent of the variation in the real exchange rate is explained by oil specific shocks, see Table 1. However, after 2-3 years, the currency appreciation effect no longer operates.

Figure 4. Domestic impulse responses

GDP Mainland Norway

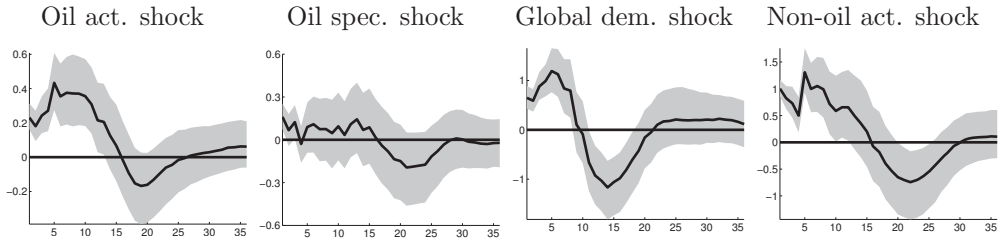

Investment Mainland Norway

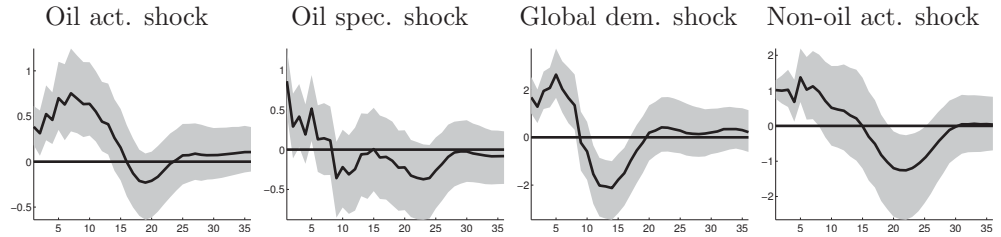

Employment Mainland Norway

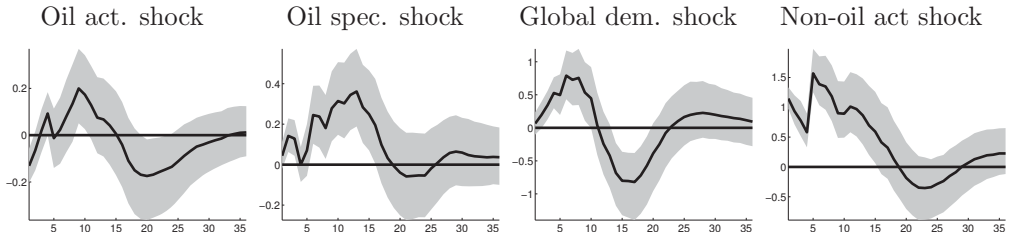

Wages Mainland Norway

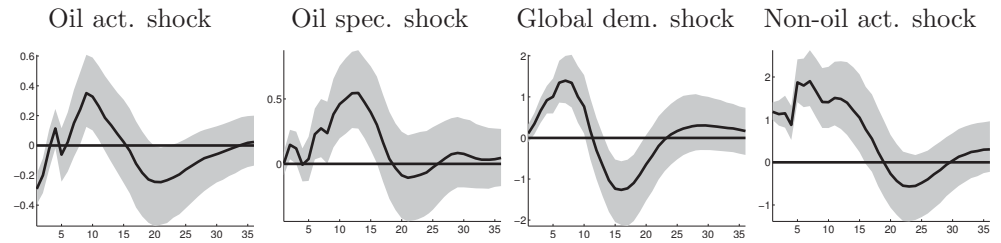

Real exchange rate

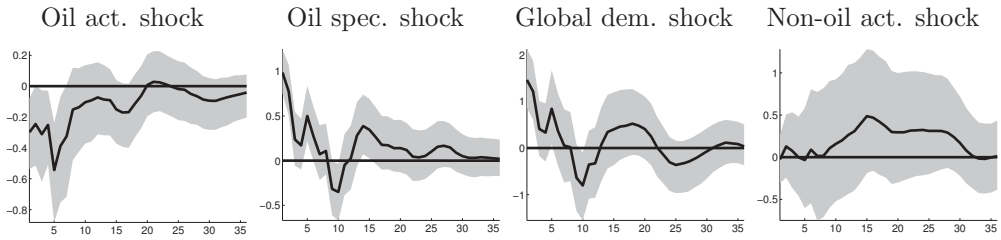

Note: The responses are displayed in levels of the variables. All shocks are normalized to a 1 percent increase, except for the oil specific shock, which is normalized to increase the real price of oil with 10 percent. The gray shaded area represent 68 percent probability bands, while the black solid lines are median estimates.

Table 2. Productivity

Shock	Horizon		
	4	**8**	**16**
Oil activity	0.36	0.25	0.22
Oil specific	0.01	-0.01	-0.00

Note: The numbers show the difference between the response in value added and employment for Mainland, interpreted as labour productivity.

The oil specific shock also has spillovers to the rest of the economy, although to a lesser extent than the oil activity shock. In particular, following an oil specific shock that increases oil prices by 10 percent, GDP and investment in Mainland Norway increase temporarily by 0.25 and 1 percent, respectively, most likely as petroleum investment also increases, see Table 1. Furthermore, employment and wages gradually increase, suggesting that there are spending effects owing to the windfall gains associated with increased oil prices.

The second shock that can potentially increase oil prices, a global demand shock, also causes the Norwegian currency to appreciate. However, the response in the exchange rate is less pronounced than for the oil specific shock, explaining approximately 20 percent of the real exchange rate variation. As a consequence, the effect on GDP and investment, as well as the spillovers to employment and wages, are more substantial. Between 40 and 50 percent of the variation in mainland GDP and investment activities can be explained by global demand.[14] The finding that foreign factors are important for the Norwegian business cycles is consistent with Aastveit et al. (2011) and Furlanetto et al. (2013).

Lastly, a non-oil (domestic) activity shock increases GDP, employment and wages in the mainland economy. The effect on investment is also positive, but the variation explained by the domestic shock is modest (less than 10 percent). The effect on the real exchange rate is negligible.

It is too early to make any conclusions regarding any evidence (or lack thereof) of Dutch disease. To do so, we need to examine sectoral reallocation, which we do below. However, it is obvious that the Norwegian economy has benefitted from having a highly profitable oil and gas sector: Both windfall gains due to energy booms and higher oil prices had positive spillover effects on the mainland economy. What are the mechanisms behind these spillovers? While we have seen that labor input clearly increased following this shock, Table 2, which measures productivity gains after 4, 8 and 16 quarters, suggests that productivity spillovers are also of first order importance for energy booms. As productivity measures the efficiency of production, this also explains why investment in the mainland economy increased substantially following this shock. This is interesting, as it highlights the empirical relevance of alternative theoretical Dutch disease models, see, e.g., Torvik

[14] An one-percent increase in global demand, increases real oil prices by approximately 10-12 percent, see Figure 3. Compared to a similar sized oil price increase due to an oil specific shock, the effects on GDP and investment in Mainland Norway are more than twice as large; GDP increases by 0.7-1 percent after a year, while investment increases by 2 percent.

Table 3. Residual regressions

Shock	Variable	Lag 1	2	3	4	R^2
	CPI	-0.00 (0.55)	-0.00 (0.92)	-0.00 (0.52)	0.00 (0.25)	0.02
	PPI	-0.03 (0.01)	0.01 (0.69)	0.00 (0.98)	-0.01 (0.20)	0.06
Oil activity	OSEBX	0.08 (0.27)	0.11 (0.03)	0.10 (0.19)	-0.02 (0.67)	0.16
	⇀ Energy	0.13 (0.10)	0.14 (0.00)	0.06 (0.39)	-0.06 (0.49)	0.17
	ToT	-0.01 (0.56)	0.01 (0.22)	0.00 (0.95)	0.00 (0.76)	0.02
	CPI	0.00 (0.60)	0.00 (0.13)	0.00 (0.01)	0.00 (0.24)	0.15
	PPI	0.04 (0.00)	0.02 (0.00)	0.03 (0.00)	0.00 (0.81)	0.34
Oil specific	OSEBX	0.09 (0.02)	0.03 (0.50)	0.03 (0.48)	-0.05 (0.31)	0.13
	⇀ Energy	0.12 (0.01)	0.07 (0.23)	0.07 (0.14)	0.00 (0.95)	0.19
	ToT	0.01 (0.02)	0.01 (0.17)	0.01 (0.22)	-0.01 (0.21)	0.16

Note: For each variable the rows show coefficient estimates and Newey-West estimated p-values (in parenthesis) from simple OLS regressions:

$$y_{t,i} = \alpha_{i,j} + \sum_{p=1}^{P} \beta_{p,i,j} e_{t-p,i,j} + u_{t,i,j}$$

where i denotes variable $i = 1, \cdots, 4$, j denotes structural shocks $j = [$Oil activity and Oil specific$]$, and p are the number of lags with $P = 4$. All dependent variables, $y_{t,i}$, are transformed to four quarter logarithmic differences. $e_{t-p,i,j}$ are the median estimates of the structural shocks. The sample is $1997Q1-2012Q4$.

(2001), which emphasize learning by doing mechanisms and productivity spillovers.

Conversely, the oil specific shocks (that increase oil prices) have virtually no effect on productivity, see Table 2. As such, our results show that is important to distinguish between windfall gains due to volume and price changes when analyzing the Dutch disease hypothesis. To the best of our knowledge, this is the first paper to explicitly separate and quantify these two channels, while also allowing for explicit disturbances to world activity and the non-oil sector.

Table 3 adds further evidence to the structural interpretation. In the table, we separately regressed the lags of the median structural shocks on consumer price inflation (CPI), producer price inflation (PPI), total stock returns (OSEBX), stock returns for the energy firms (Energy) and the terms of trade (ToT). Although simple, these regressions not only confirm that the structural identification of our benchmark model is sound but also shed light on the additional channels through which the energy sector affects the economy.

First, as asset prices are the present discounted values of the future net earnings of the firms in the economy, unexpected energy booms that enhance the production possibilities for the whole economy should be positively related to stock returns. This is confirmed in our regressions, where the oil activity shock explains a considerable share of the variation in stock returns (both OSEBX and Energy). We find no evidence that the shock increases costs, as energy booms do not explain a substantial amount of the variation in CPI and

PPI. Furthermore, the effect on terms of trade is insignificant, confirming that the windfall gains associated with energy booms are not related to energy prices. Instead, energy booms change the distribution of wealth due to productivity spillovers, the subsequent movement of resources, higher income and increased spending in the overall economy.

Moreover, we find that the oil specific shock leads to a general rise in production costs (PPI). This erodes the real effect of spending and may explain why this shock has less stimulating effects on the economy. However, we confirm that the terms of trade are positively affected by oil price increases in an oil exporting economy, which explains the pronounced effect on the real exchange rate we observed above. Furthermore, oil specific shocks also explain a substantial share of the variation in energy specific stock returns.[15]

The results presented thus far reflect average responses over the sample analyzed. In Figure 5, we show that the structural shocks are also well identified in terms of timing. In particular, the figure displays the model's historical decomposition of the domestic factor representing the non-oil economy. As seen in the figure, oil activity shocks stimulated the Norwegian economy, particularly from the middle of the 1990s and until 2000 (after which there was a temporary cyclical decline in oil activity, see also Figure 1d), and again during the economic upswing beginning around 2004. However, while the period of high economic growth in the middle and late 1990s can in large part be explained by increased oil activity, the high growth period predating the financial crisis was primarily driven by increased global demand and oil specific shocks, which both drove up oil prices. The windfall gain from higher oil prices stimulated investment in the petroleum sector and thereby also the mainland economy through spillover effects. However, by the end of 2008, Norway was affected by the financial crisis. The subsequent downturn was primarily caused by negative global demand as well as by oil specific shocks (that lowered oil prices) and oil activity shocks. The return to trend growth was gradual, with positive contributions from oil specific shocks. From 2011, global demand again contributed positively to the mainland economy (again via higher oil prices).

For the reader with detailed knowledge of the Norwegian economy, Figure 5 presents a reasonable story of a country that has benefited from increased activities in the North Sea, albeit with cyclical up and downturns. However, the negative or only mildly positive contribution from the oil activity shocks since 2006/2007 provides some cause for concern. To the extent that an oil boom is associated with productivity dynamics (that positively affect value added in the overall economy), the muted role of these shocks suggests that that productivity spillovers have declined recently. This is consistent with the view portrayed in Olsen (2013) of a slow down in productivity since 2005. Furthermore, labor input per hour worked has also declined in recent years relative to Norway's trading partners. Thus, while the enhanced linkages from both the oil sector and energy prices have been positive for growth and employment in the Norwegian economy for nearly two decades, the declining productivity spillovers coupled with increased costs could be a

[15] Without reading too much into these simple regressions, we also observe that the oil specific shocks explain more of the energy specific returns than overall returns in the economy (measured by OSEBX), which is consistent with the view that the increased costs eroded the value added from the oil specific shock. I.e., while firms in the Energy sector benefit from increased oil prices, the spillover to value added in the overall economy was small.

Figure 5. Historical shock decomposition: Non-oil activity Norway

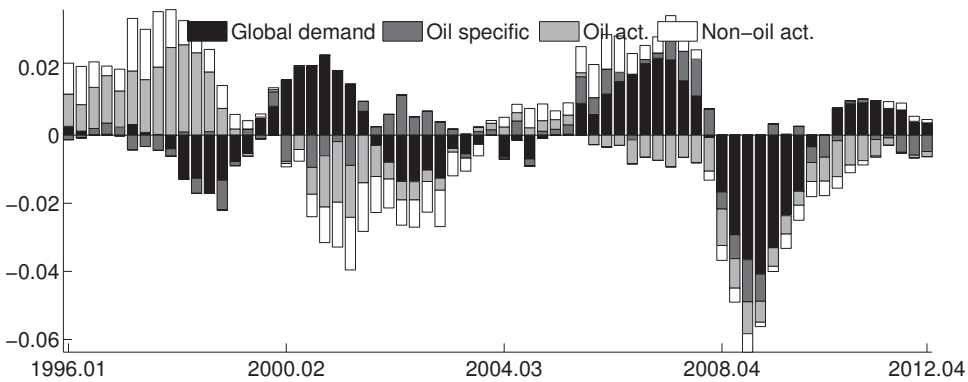

Note: The figures report the accumulated contribution of each structural shock to the growth in the non-oil activity factor.

major concern in the long run.

5.3 Sectoral performance - Two speed boom?

Figure 6 displays the responses in value added and employment to energy booms (left column) and oil specific shocks (right column). The figure displays the quarterly average of each sector's response (in levels) to the different shocks. The oil activity shock is normalized to increase oil activity by 1 percent, while the oil specific shock increases oil prices by 10 percent (which is customary in the literature). Note that the white bars indicate that the shock explains less than 10 percent of the variation in a sector.

The figure emphasizes that energy booms stimulate value added in all industries in the private sector, but to a varying degree. The construction and business sectors are among the most positively affected. Between 30 and 40 percent of the variance in these sectors is explained by energy booms, see Table 4 in Appendix B. These are industries with moderate direct input into the oil sector, but the indirect effects are large. Value added in manufacturing is also positively affected, but less so than in the non-tradable sectors. Yet, there is no evidence of Dutch disease wherein the sector eventually contracts.

Turning to the labor market, our model confirms the stylized facts presented above in Figure 1c. Norway has become a two speed economy, with employment in non-tradable sectors such as construction, the business service sector and real estate growing at a much faster pace than tradables such as manufacturing. However, and as above, there is no evidence of Dutch disease; manufacturing does not contract. Interestingly, the effect on the public sector (value added and employment) is negligible, suggesting only a minor government spending effect following this shock.

Are these numbers reasonable? Compared to Eika et al. (2010), who calculate the direct and indirect effects based on input-output tables, our numbers are more substantial. Yet, Eika et al. (2010) also found the service sector (e.g., as business industries) to be the most affected, once accounting for indirect effects such as inputs between the sectors. Where we diverge is in the size of the spillovers and the number of sectors involved.

Figure 6. Relative responses

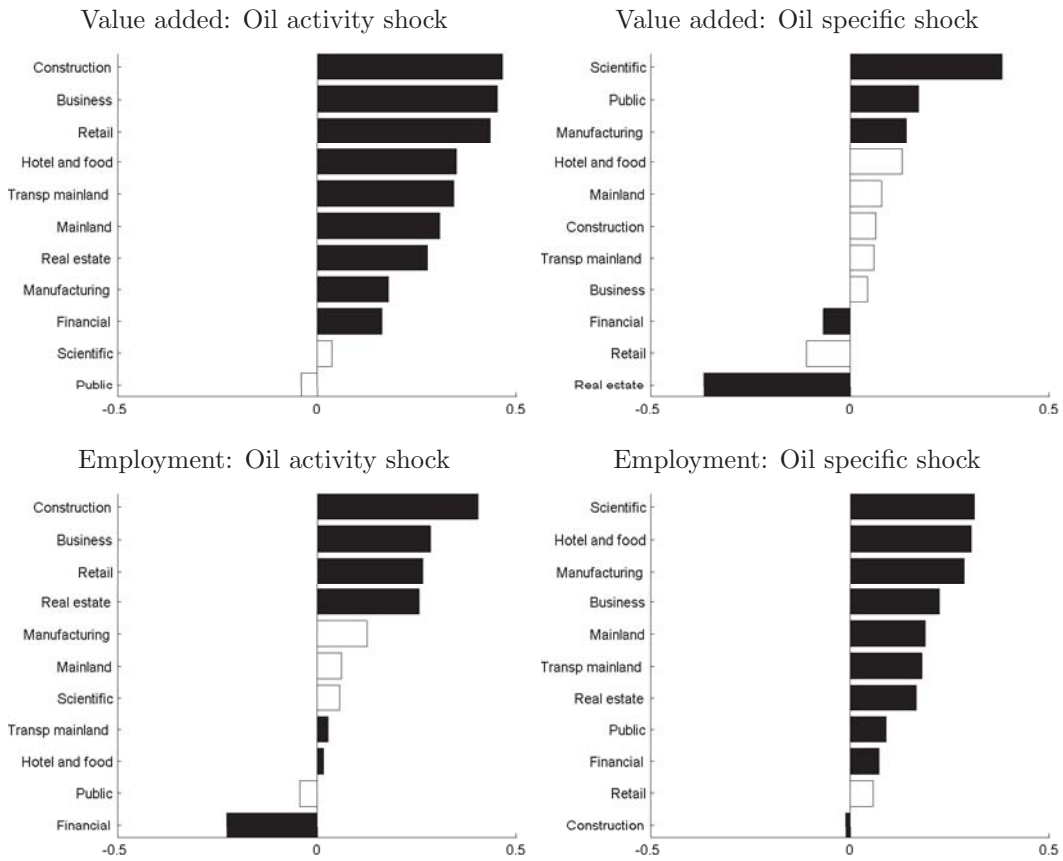

Note: Each plot displays the quarterly average of each sector i's response (in levels) to the different shocks. The averages are computed over horizons 1 to 12. The oil activity shock is normalizes to increase oil activity by 1 percent, while the oil specific shock is normalized to increase the real price of oil with 10 percent. White bars indicate that the shock explains less than 10 percent of the variation in the sector.

However, this should come as no surprise, as we also allow for induced spending effects via income and wage growth, see Table 1. Moreover, in our framework the input-output table becomes endogenous, as we allow for shared productivity dynamics across sectors.

As seen in Table 1, and indicated by the white bars in Figure 6, the oil specific shock generally explains a substantially smaller share of the variance in the sectoral variables than the oil activity shock. The responses to the oil specific shock also present a more diverse picture. Now sectors such as scientific services and manufacturing are among the most positively affected. This is interesting, as these sectors are also technology intensive and enjoy spillovers from the significant boost in petroleum investment that follows the oil specific shock. As offshore oil often demands complicated technical solutions, the oil specific shock generated positive knowledge externalities that benefited employment in these sectors in particular. Thus, the theory of Dutch disease is turned on its head following this shock. However, compared to the responses reported for the oil activity shock, the public sector is now also positively affected, suggesting the presence of a spending effect. We examine this in greater detail in the next section.

Lastly, the global demand shock is important for all industries in the private sector, but

most so for manufacturing (relative plots are available on request).[16] Thus, the stylized fact that manufacturing is lagging behind the other sectors, in particular in the financial crisis (see Figure 1c), also refers to manufacturing's substantial exposure to foreign shocks (which were all negative in the financial crisis).

In summary, we find no evidence of Dutch disease as experienced in the Netherlands in the 1970s. Instead, we find positive spillovers between the energy sector and both the tradable and non-tradable sectors. As discussed, an important channel for these spillovers could be productivity and learning by doing. As such, our results highlight the empirical relevance of alternative theoretical Dutch disease models, such as that proposed by, e.g., Torvik (2001). Moreover, our model successfully replicates the stylized facts portrayed in Figure 1 indicating a two speed economy. Importantly, however, the observed two speed pattern is not a function of resource wealth in isolation; global factors need to be taken into account.

5.4 Public sector

One aspect of the results presented above that we have not discussed in detail thus far is the role of the public sector. Norway has a large public sector, and much of the petroleum income is directly managed through the Norwegian Petroleum Fund, which was specially designed with the express purpose of shielding the domestic economy from potential spending effects caused by the resource endowment. Through a fiscal rule, which permits the government to spend approximately 4 percent of the fund (expected return) every year, the income from the oil and gas sector should only gradually be phased into the economy, and thus ensure fiscal discipline.

Very few studies have analyzed the effects of oil price changes on government spending in Norway. Those that do find very small effects, see, e.g., Pieschacon (2012). However, Pieschacon (2012) does not control for the different sources that may affect the oil price. As we have shown here, oil price increases can be due to either global demand or oil specific shocks, and the mechanisms by which they affect the economy will not be identical.

Although we do not explicitly examine fiscal policy in this study, the results presented above reveal two interesting points regarding government spending in a resource rich economy. First, energy booms do not explain a large share of the variance in value added or employment in the public sector. As such, governmental arrangements to ensure fiscal discipline seem to work.[17] However, the results presented in Figure 6 suggested that the public sector is positively affected by the oil specific shock. Furthermore, 40 percent of the variation in government spending is explained by oil specific shocks (see Table 1). This suggests evidence of a spending effect from increased oil prices via the public sector, even though the fiscal rule is in place. To explore this further, we augment the dataset with the value added in the central and local governments, and re-estimate the model.[18]

[16] As seen in tables 1 and 4, the variance explained by the global demand shock is substantial for all sectors except the public one. However, the manufacturing sector is by far the most affected; 60 percent of the variation in value added in the manufacturing sector can be explained by foreign shocks.

[17] However, if the oil activity shocks are pure productivity spillovers, the public sector does not seem to benefit from these in the same manner as the other sectors of the economy.

[18] The baseline results are not quantitatively affected.

Figure 7. Impulse responses: Oil specific shocks

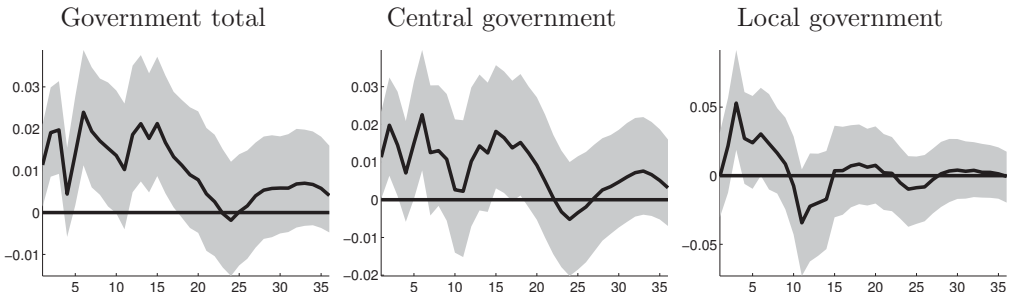

Note: The responses are displayed in levels of the variables. The oil specific shock is normalized to increase the real price of oil with 10 percent. The gray shaded area represent 68 percent probability bands, while the black solid lines are median estimates.

Norway has had an active population maintenance policy for rural districts. We therefore expect the increased income from the North Sea to have benefited local governments in particular. The impulse responses of the newly added government variables are displayed in Figure 7.

The results emphasize that there is a positive link between increased oil prices and government spending, in particular at the local government. While 30-40 percent of the variation in value added at the state level is determined by oil specific shocks, the corresponding number for local government is close to 60 percent. The results are consistent with Norway having an active government policy of investing in rural development. These could be regions that may not directly benefit from oil-related developments.

On a final note, in our model, government spending will respond to the various shocks affecting the economy. As oil specific shocks are generally beneficial for an oil exporter such as Norway, but less so for other oil importing countries (see Figure 3), increased spending by the government could also be a way to shelter the economy from a decline in foreign demand due to higher oil prices. However, our analysis shows that the consequences of increased spending will be manifested in an appreciated exchange rate and eventually increased costs, as can be seen in Figure 6 and Table 3, respectively. Both mechanisms deteriorate competitiveness, which could be a concern for the Norwegian economy in the long run.

6 Additional results and robustness

As mentioned in the main text, our results are robust to estimating the model using classical two-step estimation techniques. Furthermore, as descried in Section D.0.4, our results seem robust to different prior specifications. We have also conducted a series of other robustness checks. These are described fully in Appendix C. Below, we provide a brief summary.

First, global activity is not observed. We have approximated global activity by taking the simple mean across eight countries thought to be important to the global business cycle and Norway in particular. Qualitatively, the results reported in Section 5 are not

affected by excluding countries from this set or changing its composition. Details are provided in Appendix C.2.

Second, running the analysis on a different sample does not change the main conclusions reported in Table 1. That is, on average across sectors, a booming oil sector explains approximately 20-30 percent of the variation in the disaggregated series, irrespective of whether we estimate the model on a sample period from 1986:Q1 to 2012:Q4 or 1996:Q4 to 2012:Q4. However, as described in Appendix C.3, the subsample analysis should be interpreted with caution due to differences in data availability.

Third, the model specification is uncertain. The number of factors and lags employed in the model should be tested. We do this primarily by running a quasi-real-time forecasting experiment. The results reported in Appendix C.1 show that our benchmark model, outlined in Section 4, performs superior to simple univariate autoregressive processes. The Benchmark specification is also among the best performing specifications and is the best model specification over shorter forecasting horizons.

7 Conclusion

This study examines the empirical validity of the classical Dutch disease theory in a small and open oil and gas producing economy. Using Norway as a case study, we provide a novel contribution on the subject by explicitly identifying and quantifying windfall gains from a booming energy sector or higher oil prices and the associated sectoral performance in the rest of the economy.

We estimate a Bayesian Dynamic Factor Model that includes separate activity factors for oil and non-oil sectors in addition to global activity and the real price of oil. The model is particularly useful in a data rich environment such as ours, where common latent factors and shocks are assumed to drive the co-movements between economic variables in the economy.

We have two main results: First, booms in the energy sector have substantial productivity spillovers on the non-oil sectors, effects that have not been captured in previous analysis. In particular, we find that the energy sector stimulates investment, production, employment and wages in nearly all non-oil industries. Construction, business services and real estate are the most stimulated sectors. Second, windfall gains due to changes in the real oil price also stimulate the economy, but primarily if the oil price increase relates to a boom in global demand. Oil price increases due to, say, supply disruptions, while stimulating activity in the technologically intense service sectors and boosting government spending, have small spillover effects to the rest of the economy, in part because of a substantial real exchange rate appreciation and reduced cost competitiveness. Yet, there is no evidence of Dutch disease as experienced in the Netherlands in the 1970s, where natural gas discoveries had adverse effects on the Dutch manufacturing sector. Instead, there is evidence of a two speed economy, with the manufacturing sector lagging behind the booming service sectors. Importantly, however, the observed two speed pattern is not a function of resource wealth in isolation; global factors need to be taken into account.

Our results suggest that traditional Dutch disease models with a fixed capital stock

and exogenous labor supply do not provide a convincing explanation for how petroleum wealth affects a resource rich economy when there are productivity spillovers between the various sectors.

References

Aastveit, K. A., H. C. Bjørnland, and L. A. Thorsrud (2011). The world is not enough! Small open economies and regional dependence. Working Paper 2011/16, Norges Bank.

Aastveit, K. A., H. C. Bjørnland, and L. A. Thorsrud (2012). What drives oil prices? Emerging versus developed economies. Working Paper 2012/11, Norges Bank.

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

Bai, J. and S. Ng (2010). Principal components estimation and identification of the factors. Unpublished manuscript.

Bai, J. and P. Wang (2012). Identification and estimation of dynamic factor models. MPRA Paper 38434, University Library of Munich, Germany.

Bjørnland, H. C. (1998). The economic effects of North Sea oil on the manufacturing sector. *Scottish Journal of Political Economy 45*(5), 553–85.

Boivin, J. and M. Giannoni (2006, December). DSGE Models in a Data-Rich Environment. Working Paper 12772, National Bureau of Economic Research.

Bruno, M. and J. Sachs (1982). Energy and resource allocation: a dynamic model of the "Dutch Disease". *The Review of Economic Studies 49*(5), 845–859.

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

Charnavoki, V. and J. Dolado (2012). The effects of global shocks on small commodity-exporting economies: New evidence from Canada. Discussion Paper DP8825, CEPR.

Corden, W. M. (1984). Booming sector and dutch disease economics: survey and consolidation. *Oxford Economic Papers 36*(3), 359–380.

Corden, W. M. (2012). Dutch disease in australia: Policy options for a three-speed economy. *Australian Economic Review 45*(3), 290–304.

Corden, W. M. and J. P. Neary (1982). Booming sector and de-industrialisation in a small open economy. *The Economic Journal 92*(368), 825–848.

Eastwood, R. and A. J. Venables (1982). The macroeconomic implications of a resource discovery in an open economy. *The Economic Journal 92*(366), 285–299.

Eika, T., J. Prestmo, and E. Tveter (2010). Etterspørselen fra petroleumsvirksomheten. Betydningen for produksjon og sysselsetting i Norge. Økonomiske analyser 3/2010, Statisics Norway.

Furlanetto, F., S. Sarferaz, and F. Ravazzolo (2013). Business cycles in commodity-exporting countries. Mimeo, Norges Bank.

Geweke, J. (1977). The dynamic factor analysis of economic time series. In *latent Variables in Socio-economic Models*. D.J. Aigner and A.S. Goldberger (Eds.).

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.

Hutchison, M. M. (1994). Manufacturing sector resiliency to energy booms: empirical evidence from Norway, the Netherlands, and the United Kingdom. *Oxford Economic Papers 46*, 311–329.

Ismail, K. (2010). The structural manifestation of the "Dutch Disease": The case of oil exporting countries. *IMF Working Papers*, 1–36.

Kim, C.-J. and C. R. Nelson (1999). *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications* (1 ed.), Volume 1. The MIT Press.

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

Kose, M. A., E. S. Prasad, and M. E. Terrones (2003). How does globalization affect the synchronization of business cycles? *American Economic Review 93*(2), 57–62.

Lama, R. and J. P. Medina (2012). Is Exchange Rate Stabilization an Appropriate Cure for the Dutch Disease? *International Journal of Central Banking 8*(1), 5–46.

Larsen, E. R. (2006). Escaping the Resource Curse and the Dutch Disease? *American Journal of Economics and Sociology 65*(3), 605–640.

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

Mehlum, H., K. Moene, and R. Torvik (2006). Institutions and the Resource Curse. *The Economic Journal 116*(508), 1–20.

Mumtaz, H., S. Simonelli, and P. Surico (2011). International Comovements, Business Cycle and Inflation: A Historical Perspective. *Review of Economic Dynamics 14*(1), 176–198.

Neary, J. P. and S. van Wijnbergen (1984). Can an oil discovery lead to a recession? a comment on eastwood and venables. *The Economic Journal 94*(374), 390–395.

Olsen, Ø. (2013). Address by Governor Øystein Olsen to the Supervisory Council of Norges Bank and invited guests on Thursday 14 February 2013. Economic Perspectives, Norges Bank.

Pieschacon, A. (2012). The value of fiscal discipline for oil-exporting countries. *Journal of Monetary Economics 59*(3), 250–268.

Sachs, J. D. and A. M. Warner (1995). Natural resource abundance and economic growth. Working papers 5398, National Bureau of Economic Research.

Stock, J. H. and M. W. Watson (2002). Macroeconomic forecasting using diffusion indexes. *Journal of Business & Economic Statistics 20*(2), 147–162.

Stock, J. H. and M. W. Watson (2005a). Implications of dynamic factor models for var analysis. NBER Working Papers 11467, National Bureau of Economic Research, Inc.

Stock, J. H. and M. W. Watson (2005b). Understanding changes in international business cycle dynamics. *Journal of the European Economic Association 3*(5), 968–1006.

Thorsrud, L. A. (2013). Global and regional business cycles. Shocks and propagations. Working Papers 0012, Centre for Applied Macro- and Petroleum economics (CAMP), BI Norwegian Business School.

Torvik, R. (2001). Learning by doing and the Dutch disease. *European Economic Review 45*(2), 285–306.

Van Wijnbergen, S. (1984). The "Dutch Disease": a disease after all? *The Economic Journal 94*(373), 41–55.

Appendices

Appendix A Data and Sources

Sector	Abbreviation	Moments		Variable in National Accounts
		Mean	Std.	
GDP	Oil extraction	-0.63	6.28	Oil and natural gas extraction
	Oil service	6.89	25.65	Service activities incidental to oil and gas
	Manufacturing	1.56	3.53	Manufacturing
	Construction	3.04	5.02	Construction
	Retail	4.39	3.46	Wholesale and retail trade, repair of motor vehicles
	Transp ocean	-5.32	16.22	Ocean transport
	Transp mainland	1.03	5.26	Transport activities excl. ocean transport
	Hotel and food	1.76	4.84	Accommodation and food service activities
	Financial	3.90	7.88	Financial and insurance activities
	Real estate	9.20	10.81	Real estate activities
	Scientific	4.12	4.75	Professional, scientific and and technical activities
	Business	6.79	6.91	Administrative and support service activities
	Mainland	2.79	1.95	Mainland Norway
	Public	1.64	1.46	General government
Empl	Oil extraction	2.43	5.32	See above
	Oil service	11.14	14.30	
	Manufacturing	-0.59	3.35	
	Construction	3.95	4.05	
	Retail	1.28	1.99	
	Transp ocean	0.91	2.79	
	Transp mainland	0.70	2.22	
	Hotel and food	1.10	2.70	
	Financial	-0.13	2.92	
	Real estate	5.98	6.61	
	Scientific	3.61	3.75	
	Business	5.66	6.32	
	Mainland	1.33	1.54	
	Public	1.33	0.92	
Other	Wages oil	9.90	6.87	Wages petroleum
	Wages public	6.04	1.77	Wages public
	Wages mainland	6.06	2.38	Wages mainland
	Investment oil	4.52	22.62	Investment petroleum
	Investment mainland	4.06	8.60	Investment mainland
	Exchange rate	0.57	4.79	BIS effective exchange rate index, broad basket
Int.	World activity	2.78	1.90	See text, section 3
	Oil Price	9.01	33.11	Crude Oil-Brent, deflated using US CPI

Note: The table lists all the variables used in the Benchmark model. All activity, investment, wages and employment series for Norway are collected from the Quarterly National Accounts database of Statistics Norway. The international series were downloaded from Datastream. The real exchange rate is collected form BIS. All series are seasonally adjusted by their source. Std. denotes standard deviation. Int. denotes international. The moments are computed based on the transformed variables, i.e. $log(x_{i,t}) - log(x_{i,t-4})$).

Appendix B Figures and tables

Table 4. Variance decompositions

Variable	Sector & Horizon	Shock: Oil activity	Oil specific	Global demand	Non-oil activity
		4, 8	4, 8	4, 8	4, 8
GDP	Manufacturing	0.04, 0.08	0.27, 0.18	0.65, 0.69	0.04, 0.04
	Construction	0.26, 0.35	0.03, 0.01	0.51, 0.46	0.20, 0.18
	Retail	0.74, 0.81	0.01, 0.02	0.21, 0.15	0.03, 0.02
	Transport ocean	0.27, 0.37	0.62, 0.48	0.10, 0.11	0.01, 0.04
	Transport mainland	0.21, 0.32	0.18, 0.10	0.55, 0.51	0.06, 0.06
	Hotel and food	0.10, 0.19	0.13, 0.06	0.57, 0.54	0.20, 0.20
	Financial	0.19, 0.31	0.58, 0.43	0.22, 0.24	0.01, 0.02
	Real estate	0.08, 0.14	0.77, 0.70	0.15, 0.15	0.00, 0.00
	Scientific	0.07, 0.03	0.43, 0.34	0.41, 0.50	0.08, 0.12
	Business	0.32, 0.42	0.19, 0.09	0.37, 0.36	0.12, 0.13
Employment	Manufacturing	0.07, 0.04	0.14, 0.17	0.23, 0.31	0.56, 0.48
	Construction	0.28, 0.39	0.14, 0.06	0.40, 0.37	0.18, 0.17
	Retail	0.32, 0.35	0.04, 0.02	0.32, 0.35	0.32, 0.28
	Transport ocean	0.10, 0.11	0.34, 0.18	0.22, 0.13	0.34, 0.59
	Transport mainland	0.31, 0.14	0.11, 0.11	0.03, 0.16	0.55, 0.59
	Hotel and food	0.29, 0.11	0.14, 0.17	0.30, 0.39	0.28, 0.33
	Financial	0.52, 0.63	0.05, 0.03	0.42, 0.30	0.01, 0.04
	Real estate	0.10, 0.12	0.22, 0.08	0.52, 0.61	0.16, 0.18
	Scientific	0.28, 0.09	0.03, 0.11	0.20, 0.35	0.49, 0.45
	Business	0.04, 0.08	0.16, 0.10	0.61, 0.62	0.19, 0.20

Note: Each row-column intersection reports median variance decompositions for horizons 4 (left) and 8 (right)

Appendix C Robustness

C.1 Model specification

The correct model specification is uncertain. Different test statistics, see Bai and Ng (2002), suggest between 3 and 8 static factors. Using 4 factors explains approximately 60 percent of the variation in the dataset. Including an additional 4 static factors increases the variance explained by a modest 17 percent. Although informative, the tests for the number of static factors are far from conclusive.

To fully test our preferred model relative to alternative specifications, we run a quasi-real-time forecasting experiment. The experiment is conducted as follows: For the sample period from 1996.01 to 2012.04, we estimate the BDFM with different lag specifications. in particular, we allow for up to 2 lags of the vector of factors in the observation equation ($s = 0, \ldots, 2$). For each lag specification we also estimate the model with and without autocorrelated idiosyncratic errors ($l = 0, 1$). Ultimately, this yields 6 different specifications. Lastly, for each of these combinations we estimate the model with 4 and 8 lags in the transition equation ($h = 4, 8$).

We compute the model's out of sample forecasting performance over the period from 1996.02 to 2012.04. The performance is scored by root mean forecasting errors (RMSE) and log scores (logScore).[19] The forecasting experiment is quasi-real-time, as we do not re-estimate the models for each new vintage of data we forecast, and we also do not use real-time vintage data when estimating the models or in the evaluation of forecasting performance. Thus, the distribution of the model parameters used to forecast is assumed to be constant throughout the evaluation period. For our purpose, which is to make comparison among nested structural models, this is an innocuous assumption. Furthermore, an advantage of a quasi-real-time forecasting experiment, as opposed to a real-time forecasting experiment, is that we can evaluate the forecasting performance over a much longer sample.[20]

Table 5 reports the results.[21] Panels (a) and (b) reports the results for $h = 4$ and $h = 8$, respectively. At the two step ahead horizon, and evaluated across all variables, our preferred model specification, BDFM s(2)a(1) (denoted Benchmark in the table), performs substantially better than any other model specification. In Panel (a) (Panel (b)), for 20 (19) and 19 (21) out of 39 variables, the Benchmark model performs best in terms of respectively RMSE and average logScore, respectively. At the four step ahead horizon, the ranking of the different model specifications changes, and the BDFM s(2)a(0) model receives a better score than the other models in approximately 40 to 50 percent of the cases.

[19] The RMSE is a quadratic loss function that is often used to evaluate point forecasts. If the focus is on the whole forecast distribution, the RMSE is not appropriate and log scoring is a better metric. The logScore is the logarithm of the probability density function evaluated at the outturn of the forecast. As such it provides an intuitive measure of density fit.

[20] I.e., in a real-time experiment, we would have to re-estimate the models for each new vintage and use a substantial part of the sample to estimate the initial parameter distributions.

[21] To save space, we only report the results for forecasting horizons 2 and 4. The conclusions do not change for horizons 1 and 3. These results are available on request.

Table 5. Forecast performance

Horizon	Model	Variable				All
		Y	E	I	W	variables
Panel a:						
2	BDFM s(0)a(0)	0.90, 0.97	1.00, 1.03	0.96, 0.96	1.00, 1.00	0, 1
	BDFM s(1)a(0)	1.01, 1.00	1.02, 1.01	0.98, 0.99	1.06, 1.00	1, 3
	BDFM s(2)a(0)	1.05, 1.02	1.05, 1.01	1.00, 1.01	1.08, 1.01	11, 7
	BDFM s(0)a(1)	0.99, 0.99	0.97, 1.02	0.97, 0.98	1.00, 1.01	0, 5
	BDFM s(1)a(1)	0.98, 1.00	0.99, 1.01	0.98, 0.99	0.99, 1.00	5, 2
	AR(1)	0.77, 0.90	0.62, 0.84	0.86, 0.87	0.51, 0.73	2, 2
	Benchmark	0.01, 2.92	0.01, 3.58	0.07, 1.30	0.01, 3.16	20, 19
4	BDFM s(0)a(0)	1.01, 1.01	1.01, 1.01	1.00, 0.97	0.98, 1.01	5, 7
	BDFM s(1)a(0)	1.07, 1.02	1.02, 1.02	1.02, 1.01	1.04, 1.02	9, 11
	BDFM s(2)a(0)	1.06, 1.03	1.03, 1.02	1.04, 1.03	1.07, 1.02	19, 16
	BDFM s(0)a(1)	1.00, 1.00	1.00, 1.00	0.97, 0.96	0.98, 1.01	0, 0
	BDFM s(1)a(1)	0.99, 1.00	1.02, 1.01	0.99, 0.98	1.00, 1.01	2, 3
	AR(1)	0.75, 0.75	0.65, 0.53	0.85, 0.78	0.47, 0.40	0, 0
	Benchmark	0.02, 2.63	0.01, 3.17	0.07, 1.19	0.01, 2.84	4, 2
Panel b:						
2	BDFM s(0)a(0)	0.90, 0.96	0.91, 1.02	0.97, 0.98	0.92, 0.98	1, 1
	BDFM s(1)a(0)	1.04, 1.02	0.96, 1.00	0.98, 0.99	1.02, 0.99	0, 2
	BDFM s(2)a(0)	1.08, 1.03	0.97, 0.98	1.00, 1.01	1.13, 1.00	13, 5
	BDFM s(0)a(1)	0.98, 0.98	0.96, 1.03	0.97, 0.98	0.99, 1.01	4, 5
	BDFM s(1)a(1)	0.99, 1.00	0.99, 1.01	0.97, 0.99	1.00, 1.01	0, 3
	AR(1)	0.66, 0.87	0.49, 0.82	0.85, 0.87	0.49, 0.71	2, 2
	Benchmark	0.01, 3.05	0.00, 3.69	0.07, 1.30	0.01, 3.21	19, 21
4	BDFM s(0)a(0)	1.01, 0.99	0.95, 1.02	1.00, 1.00	0.94, 0.99	6, 5
	BDFM s(1)a(0)	1.07, 1.03	0.98, 1.01	1.02, 1.01	1.03, 1.00	2, 8
	BDFM s(2)a(0)	1.09, 1.03	0.99, 0.99	1.04, 1.03	1.13, 1.01	22, 17
	BDFM s(0)a(1)	1.00, 0.99	0.99, 1.02	0.96, 0.96	1.00, 1.01	1, 2
	BDFM s(1)a(1)	1.00, 1.00	0.99, 1.01	0.97, 0.98	1.00, 1.01	0, 0
	AR(1)	0.59, 0.69	0.45, 0.48	0.82, 0.78	0.37, 0.36	0, 1
	Benchmark	0.01, 2.84	0.01, 3.45	0.07, 1.22	0.01, 3.04	8, 6

Note: Panel a) reports the results for $h = 4$, and Panel b) reports the results for $h = 8$, where h refers to the number of lags used in equation 2. Benchmark is BDFM s(2)a(1). s() denotes the number of lags used for the factors in the observation equation, a() denotes the number of lags used for the idiosyncratic AR process. The abbreviations Y, E I and W are respectively GDP, employment, investment and wages in mainland Norway. AR(1) is a univariate AR(1) models for each variable. For each model, variable, and horizon the reported numbers are relative RMSE (left) and - average logScore (right) scores, i.e. $\frac{M_{i,H,v}}{BDFMs(2)a(1)_{H,v}}$ for $i = 1, \cdots 6$ and $v = (Y, E, I, W)$. For the BDFM s(2)a(1) model the numbers reported are the actual scores. The numbers in the last column show how many times model i, at horizon H, is ranked as the best model when the performance across all variables $v = 1, \cdots, N$ is evaluated.

Generally, forecasting performance increases with the number of lagged factors, while the inclusion of autocorrelated idiosyncratic errors seems to be less important for forecasts four quarters into the future. Viewed from a bias-variance trade-off perspective, this is intuitive. The richer specified Benchmark model has a better in sample fit, thus a lower bias, but may have a higher degree of variance. At longer forecasting horizons, this reduces forecast accuracy.

Bai and Wang (2012) show in a simulation study that specifying a BDFM without autocorrelated idiosyncratic errors, although the underlying data generating process has this feature, generally produces estimates of the latent factors that are less reliable than specifying a BDFM with autocorrelated idiosyncratic errors, despite the underlying data generating process lacking this feature. Thus, although the BDFM s(2)a(0) specification also performs well in terms of forecasting, we prefer the Benchmark model.

Evaluating the Benchmark model across Panel (a) and (b), i.e., with $h = 4$ and $h = 8$, we see that the results are somewhat better in Panel (b). That is, the logScore is generally higher, indicating a better density fit (while the RMSE is essentially unchanged). The findings reported in section 5 are qualitatively similar, irrespective of whether we use $h = 4$ or $h = 8$. However, as documented in Hamilton and Herrera (2004), when modeling the oil market, an overly restrictive lag structure might lead to misleading results. Accordingly, we report the results for the $h = 8$ specification.

For many variables, e.g., GDP, simple time series models such as AR processes are often difficult to outperform with respect to forecasting performance. We therefore also compare the performance of the Benchmark model with that of a simple univariate AR(1) model.[22] As can be seen from Panel (a) and (b) in Table 5, the forecasting performance of the dynamic factor model is substantially better than the AR(1). For example, at horizon 2, and for GDP in Mainland Norway (Y), the performance of the Benchmark model is over 20 and 10 percent better than the AR(1) model when evaluated using RMSE and average logScores, respectively. For wages in Mainland Norway (W), the Benchmark model is even more superior, with an improvement of over 50 percent relative to the AR(1) model at horizon 4.

In summary, the results reported in table 5 support our Benchmark model specification. The highly parameterized, and structural, factor model is also superior to simple AR(1) models for most variables and at most horizons. As such, our findings confirm a voluminous literature documenting the usefulness of factor models for forecasting, see, e.g., Stock and Watson (2002).

C.2 What is global activity?

As described in Section 3, we construct the observable world activity series based on the mean across 8 different countries. These countries are not chosen ad-hoc: they represent Norway's most important trading partners and the largest economies in the world. That being said, world activity is not an observable variable. Thus, we have attempted to estimate the world activity factor as a latent factor in the same manner as we estimate

[22] We estimate one AR(1) model for each observable variable, $v = 1, \cdots, N$, and conduct the same quasi-real-time forecasting experiment as described above.

the latent oil and domestic activity factors. This did not work well. Employing reasonable uninformative priors, and without restrictions on the hyper-parameters, the model is not able to distinguish the different factors from each other in any meaningful manner. Our approach of approximating world activity as the mean across 8 different countries could be regarded as employing more informative priors and placing restrictions on the hyper-parameters. Ideally, this should have been performed within the modeling framework. However, as the extraction of the world activity factor is not the main research question of this study, we have not pursued the issue further.

Importantly, our main results are robust to different world activity approximations, with one exception. China should not be excluded from the set. As shown in Aastveit et al. (2012), growth in emerging economies (here represented by China), has been fundamental in explaining the surge in oil prices over the last two decades. To capture this important driver of the oil market, China should not be excluded from the construction of the global activity factor. Including or excluding countries other than the US and China from the international set, does not alter our main conclusions.

C.3 Subsample analysis

For production variables we have data going back until the beginning of the 1980's. Thus, for comparison we estimate the BDFM with production data only, i.e. excluding employment, wage and investment series, on the two samples 1986:Q1 to 2012:Q4 and 1996:Q1 to 2012:Q4. We stress that extending the sample all the way back to the 1980s is not uncontroversial. In the 1980s the Norwegian exchange rate was more or less fixed, and the central bank was not targeting inflation. Further, as numerous papers have documented, both the volatility of foreign shocks and the degree of business cycle synchronization was different in the 1980s compared to today. Thus, the comparison between the two samples is only conducted as part of our robustness analysis. Further, the information set used to extract the latent factors effectively becomes much smaller when employment, wage and investment series are excluded from the analysis.[23]

Nevertheless, Table 6 compares the average non-mainland variance decompositions for the two periods. One finding stand out. The domestic economy's dependence on oil specific and international shocks have increased over the sample. Domestic activity shocks explain roughly 45 percent of the variation of key domestic sectors when we estimate the model over longest sample, and only around 25 percent when we estimate the model over the shorter sample. The results also show that world activity shocks are more important today than in previous periods, while the variance explained by the oil activity shock is more or less unchanged.

[23]Due to the smaller information set we also use the $h = 4$ model specification in this exercise.

Table 6. Variance decompositions: Short versus long sample comparison

Sample	Sector	Shock			
		Oil activity	**Oil specific**	**Global demand**	**Non-oil activity**
1986:Q1-2012:Q4	Average	0.27, 0.26	0.20, 0.16	0.14, 0.12	0.40, 0.46
1996:Q1-2012:Q4	Average	0.22, 0.23	0.23, 0.16	0.32, 0.35	0.23, 0.25

Note: Each row-column intersection reports median variance decompositions for horizons 4 (left) and 8 (right)

As alluded to above, the results reported in Table 6 are not directly comparable to the once reported in Table 1. Still, the results for the shorter sample are consistent with our earlier findings, although the fraction of variance explained attributed to the different oil specific and international shocks differ.

Appendix D The Gibbs sampling approach

The three steps of the Gibbs sampler, described in Section 4.2.1, are iterated until convergence. Below we describe the three steps in more detail. The exposition follows Kim and Nelson (1999) closely, and we refer to their book for details.

For convenience, we repeat some notation: $\tilde{y}_T = [y_1, \cdots, y_T]'$, $\tilde{f}_T = [f_1, \cdots, f_T]'$, $H = [\lambda_0, \cdots, \lambda_s]$, and $p_i = [\rho_{1,i}, \cdots, \rho_{l,i}]$ for $i = 1, \cdots, N$, and rewrite the state space model defined in equation 1 and 2 as:

$$y_t = \Lambda F_t + \epsilon_t \tag{8}$$

and

$$F_t = AF_{t-1} + e_t \tag{9}$$

where $F_t = [f'_t, \cdots, f'_{t-h}]'$, $e_t = Gu_t$, with $u_t \sim i.i.d.N(0, Q)$ and:

$$A = \begin{pmatrix} \phi_1 & \phi_2 & \cdots & \phi_h \\ I_q & 0 & \cdots & 0 \\ 0 & I_q & \ddots & \vdots \\ 0 & 0 & I_q & 0 \end{pmatrix}, \quad G = \begin{pmatrix} I_q \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \quad \Lambda = \begin{pmatrix} H & 0_{N,h-s} \end{pmatrix} \tag{10}$$

Note that $h > s$ in our application.

We also allow for serially correlated idiosyncratic errors. In particular, we consider the case where $\epsilon_{t,i}$, for $i = 1, \cdots, N$, follows independent AR(l) processes:

$$\epsilon_{t,i} = p_i E_{t,i} + \omega_{t,i} \tag{11}$$

where $\omega_{t,i}$ is the AR(l) residuals with $\omega_{t,i} \sim i.i.d.N(0, \sigma_i^2)$. I.e.:

$$R = \begin{bmatrix} \sigma_1^2 & 0 & \cdots & 0 \\ 0 & \sigma_2^2 & \ddots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & \cdots & \sigma_N^2 \end{bmatrix}, \tag{12}$$

and $E_{t,i} = [\epsilon_{t-1,i}, \cdots, \epsilon_{t-l,i}]'$.

D.0.1 Step 1: $\tilde{f}_T | \tilde{y}_T, \Lambda, A, R, Q, p$

We employ Carter and Kohn's multimove Gibbs sampling approach (see Carter and Kohn (1994)). Because the state space model given in equations 8 and 9 is linear and Gaussian, the distribution of F_T given $\tilde{y}_T$ and that of F_t given F_{t+1} and $\tilde{y}_t$ for $t = T-1, \cdots, 1$ are also Gaussian:

$$F_T | \tilde{y}_T \sim N(F_{T|T}, P_{T|T}) \tag{13}$$

$$F_t | \tilde{y}_t, F_{t+1} \sim N(F_{t|t,F_{t+1}}, P_{t|t,F_{t+1}}), \quad t = T-1, T-2, \cdots, 1 \tag{14}$$

where

$$F_{T|T} = E(F_T|\tilde{y}_T) \tag{15}$$
$$P_{T|T} = Cov(F_T|\tilde{y}_T) \tag{16}$$
$$F_{t|t,F_{t+1}} = E(F_t|\tilde{y}_t, F_{t+1}) = E(F_t|F_{t|t}, F_{t|t+1}) \tag{17}$$
$$P_{t|t,F_{t+1}} = Cov(F_t|\tilde{y}_t, F_{t+1}) = Cov(F_t|F_{t|t}, F_{t|t+1}) \tag{18}$$

Given $F_{0|0}$ and $P_{0|0}$, we obtain $F_{T|T}$ and $P_{T|T}$ from the last iteration of the Gaussian Kalman filter:

$$F_{t|t-1} = AF_{t-1|t-1} \tag{19}$$
$$P_{t|t-1} = AP_{t-1|t-1}A' + GQG' \tag{20}$$
$$K_t = P_{t|t-1}\Lambda'(\Lambda P_{t|t-1}\Lambda' + R)^{-1} \tag{21}$$
$$F_{t|t} = F_{t|t-1} + K_t(y_t - \Lambda F_{t|t-1}) \tag{22}$$
$$P_{t|t} = P_{t|t-1} - K_t\Lambda P_{t|t-1} \tag{23}$$

I.e., at $t = T$, equation 22 and 23 above, together with equation 13, is used to draw $F_{T|T}$.

We draw $F_{t|t,F_{t+1}}$ for $t = T-1, T-2, \cdots, 1$ based on 14, where $F_{t|t,F_{t+1}}$ and $P_{t|t,F_{t+1}}$ are generated from the following updating equations:

$$\begin{aligned} F_{t|t,F_{t+1}} &= E(F_t|F_{t|t}, F_{t|t+1}) \\ &= F_{t|t} + P'_{t|t}A(AP_{t|t}A' + GQG')^{-1}(F_{t+1} - AF_{t|t}) \end{aligned} \tag{24}$$

$$\begin{aligned} P_{t|t,F_{t+1}} &= Cov(F_t|F_{t|t}, F_{t|t+1}) \\ &= P_{t|t} + P_{t|t}A'(AP_{t|t}A' + GQG')AP_{t|t} \end{aligned} \tag{25}$$

D.0.2 Step 2: $A, Q|\tilde{y}_T, \tilde{f}_T, \Lambda, R, p$

Conditional on $\tilde{f}_T$, equation 9 is independent of the rest of the model, and the distribution of A and Q are independent of the rest of the parameters of the model, as well as the data.

By abusing notation, we put the transition equation in SUR form and define:

$$y = X\beta + \epsilon \tag{26}$$

where $y = [f_1, \cdots, f_T]'$, $X = [X_1, \cdots, X_T]'$, $\epsilon = [\epsilon_1, \cdots, \epsilon_T]'$ and $\beta = [\beta_1, \cdots, \beta_q]'$, with $\beta_k = [\phi_{1,k}, \cdots, \phi_{h,k}]$ for $k = 1, \cdots, q$. Further,

$$X_t = \begin{pmatrix} x_{t,1} & 0 & \cdots & 0 \\ 0 & x_{t,2} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & \cdots & x_{t,q} \end{pmatrix}$$

with $x_{t,k} = [f'_{t-1}, \cdots, f'_{t-h}]$. Finally, $\epsilon \sim i.i.d.N(0, I_q \otimes Q)$.[24]

[24] With 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).

To simulate β and Q, we employ the independent Normal-Whishart prior:

$$p(\beta, Q) = p(\beta)p(Q^{-1}) \tag{27}$$

where

$$p(\beta) = f_N(\beta|\underline{\beta}, \underline{V}_\beta) \tag{28}$$

$$p(Q^{-1}) = f_W(Q^{-1}|\underline{v}_Q, \underline{Q}^{-1}) \tag{29}$$

The conditional posterior of β is:

$$\beta|y, Q^{-1} \sim N(\overline{\beta}, \overline{V}_\beta)_{I[s(\beta)]} \tag{30}$$

with

$$\overline{V}_\beta = (\underline{V}_\beta^{-1} + \sum_{t=1}^{T} X_t' Q^{-1} X_t)^{-1} \tag{31}$$

and

$$\overline{\beta} = \overline{V}_\beta(\underline{V}_\beta^{-1}\underline{\beta} + \sum_{t=1}^{T} X_t' Q^{-1} y_t) \tag{32}$$

$I[s(\beta)]$ is an indicator function used to denote that the roots of β lie outside the unit circle.

The conditional posterior of Q^{-1} is:

$$Q^{-1}|y, \beta \sim W(\overline{v}_Q, \overline{Q}^{-1}) \tag{33}$$

with

$$\overline{v}_Q = \underline{v}_Q + T \tag{34}$$

and

$$\overline{Q} = \underline{Q} + \sum_{t=1}^{T}(y_t - X_t\beta)(y_t - X_t\beta)' \tag{35}$$

D.0.3 Step 3: $\Lambda, R, p|\tilde{y}_T, \tilde{f}_T, A, Q$

Conditional on $\tilde{f}_T$, and given our assumption of R being diagonal, equation 8 result in N independent regression models.

However, to take into account serially correlated idiosyncratic errors, and still employ standard Bayesian techniques, we need to transform equation 8 slightly.

Thus, for $i = 1, \cdots, N$, conditional on p, and with $l = 1$, we can rewrite equation 8 as:

$$y_{t,i}^* = \Lambda_i F_t^* + \omega_{t,i} \tag{36}$$

with $y_{t,i}^* = y_{t,i} - p_{1,i}y_{t-1,i}$, and $F_t^* = F_t - p_{1,i}F_{t-1}$, and Λ_i being the i-th row of Λ.

From 36 we can then simulate the parameters Λ_i and $R_{i,i} = \sigma_i^2 = \frac{1}{h_i}$ using standard independent Normal-Gamma priors (for notational convenience we drop the subscript i from the expressions below):[25]

$$p(\Lambda, h) = p(\Lambda)p(h) \tag{37}$$

[25] Note that with $l = 0$, we could have simulated the parameters Λ_i and σ_i^2 without doing the transformation of variables described above.

where

$$p(\Lambda) = f_N(\Lambda|\underline{\Lambda}, \underline{V}_\Lambda) \tag{38}$$

$$p(h) = f_G(h|\underline{s}^{-2}, \underline{v}_h) \tag{39}$$

The conditional posterior of Λ is:

$$\Lambda|\tilde{y}, h, p \sim N(\overline{\Lambda}, \overline{V}_\Lambda) \tag{40}$$

with;

$$\overline{V}_\Lambda = (\underline{V}_\Lambda^{-1} + h\sum_{t=1}^{T} F_t^{*\prime} F_t^*)^{-1} \tag{41}$$

and

$$\overline{\Lambda} = \overline{V}_\Lambda(\underline{V}_\Lambda^{-1}\underline{\Lambda} + h\sum_{t=1}^{T} F_t^{*\prime} y_t^*) \tag{42}$$

The conditional posterior for h is:

$$h|\tilde{y}, \Lambda, p \sim G(\overline{v}_h, \overline{s}^{-2}) \tag{43}$$

with

$$\overline{v}_h = \underline{v}_h + T \tag{44}$$

and

$$\overline{s} = \frac{\sum_{t=1}^{T}(y_t^* - \Lambda F_t^*)'(y_t^* - \Lambda F_t^*) + \underline{v}_h \underline{s}^2}{\overline{v}_h} \tag{45}$$

Finally, conditional on Λ and h, the posterior of p depends upon its prior, which we assume is a multivariate Normal, i.e.:

$$p(p) = f_N(p|\underline{p}, \underline{V}_p) \tag{46}$$

Accordingly, the conditional posterior for p is:

$$p|\tilde{y}, \Lambda, h \sim N(\overline{p}, \overline{V}_p)_{I[s(p)]} \tag{47}$$

with

$$\overline{V}_p = (\underline{V}_p^{-1} + h\sum_{t=1}^{T} E_t' E_t)^{-1} \tag{48}$$

and

$$\overline{p} = \overline{V}_p(\underline{V}_p^{-1}\underline{p} + h\sum_{t=1}^{T} E_t'\epsilon_t) \tag{49}$$

D.0.4 Prior specifications and initial values

The Benchmark model is estimated using two-step parameter estimates (see Section 4.2) as priors. We label these estimates OLS. In particular, for equations 28 and 29 we set $\underline{\beta} = \beta^{OLS}$, $\underline{V}^{\beta} = V_{\beta}^{OLS} \times 3$, $\underline{Q} = Q^{OLS}$ and $\underline{v}_Q = 10$.

For equations 38, 39 and 46 we set $\underline{v}_h = 10$, $\underline{s}^2 = s^{2,OLS}$, $\underline{\Lambda} = [\lambda_0^{OLS} : 0_{N,h-s-1}]$ and $\underline{V}^{\Lambda} = [(I_s \times 3) \otimes V_{\lambda_0^{OLS}}]$. $\underline{p} = 0$, and $\underline{V}^p = 0.5$.

In sum, these priors are reasonable uninformative, but still proper. We have also experimented with other prior specifications, e.g. using Minnesota style prior for the transition equation parameters, and setting $\underline{\Lambda} = 0$. This yields similar results as the once reported in the main text. However, the variables in our sample display very different unconditional volatilities. The prior specification should accommodate this feature.

The Gibbs sampler is initialized using parameter values derived from the two-step estimation procedure. Parameters not derived in the two-step estimation (i.e. p and $\lambda_1, \cdots, \lambda_s$) are set to 0.

In this model, a subtle issue arises for the $t = 0$ observations (i.e. lags of the dynamic factors and the idiosyncratic errors at time $t = 1$). However, since we assume stationary errors in this model, the treatment of initial conditions is of less importance. Accordingly, we follow common practice and work with the likelihood based on data from $t = h + 1, \cdots, T$.