

SVERIGES RIKSBANK
WORKING PAPER SERIES

332



Oil prices in a real-business-cycle model with precautionary demand for oil

Conny Olovsson

November 2016

WORKING PAPERS ARE OBTAINABLE FROM

www.riksbank.se/en/research

Sveriges Riksbank • SE-103 37 Stockholm

Fax international: +46 8 21 05 31

Telephone international: +46 8 787 00 00

The Working Paper series presents reports on matters in the sphere of activities of the Riksbank that are considered to be of interest to a wider public.

The papers are to be regarded as reports on ongoing studies and the authors will be pleased to receive comments.

The opinions expressed in this article are the sole responsibility of the author(s) and should not be interpreted as reflecting the views of Sveriges Riksbank.

Oil prices in a real-business-cycle model with precautionary demand for oil

Conny Olovsson*

Sveriges Riksbank Working Paper Series

No. 332

November 2016

Abstract

This paper analyzes the interaction between oil prices and macroeconomic outcomes by incorporating oil as an input in production alongside a precautionary motive for holding oil in a real-business-cycle model. The driving forces are factor-specific technology shocks and supply shocks that can be imprecisely forecasted by noisy news shock. These shocks explain most of the U.S. business cycle as well as the empirical distribution of oil prices. Oil shocks are mainly driven by increasing precautionary/smoothing demand, but supply shocks contribute substantially to both the oil-price volatility and the magnitude of oil shocks mainly through their effect on oil reserves.

Keywords: Oil price shocks, business cycles

JEL: E32; Q43

*Conny Olovsson: Sveriges Riksbank, SE-103 37 Stockholm. Phone: +46 8 787 04 33. Email: conny.olvsson@riksbank.se. I am grateful for comments from Per Krusell, participants at the Nordic Summer Symposium in Macroeconomics and seminar participants at Sveriges Riksbank. The opinions expressed in this article are the sole responsibility of the author and should not be interpreted as reflecting the views of Sveriges Riksbank.

The post-war period features highly volatile oil prices and several episodes of fast and dramatic swings in the oil price, so called oil shocks. The most famous examples are probably the oil shocks of 1973/74 and 1978/79, where the real oil price increased by almost 220 and 105 percent, respectively, from the first year to the next, but at least six more oil shocks have materialized since World War II.¹ In general, the volatility of the oil price is roughly an order of magnitude larger than that of output. These facts raise the joint questions of why the oil price is so volatile and if this high volatility has important consequences for our economies. Among other things, oil shocks have been suggested to reduce output, cause recessions, and increase job destruction.²

Early papers that studied the effects of oil-price changes on U.S. variables simply viewed oil-price changes as exogenous to the U.S. economy.³ That view, however, has been seriously challenged by Hamilton (1983, 2003) and Barsky and Kilian (2002). Specifically, Barsky and Kilian argue convincingly that there is reverse causality from macroeconomic aggregates to oil prices. Among other things, they point out that sustained oil-price increases only are possible under conditions of excess demand in the oil market and that the formation of an oil embargo is not an exogenous event. Hamilton (2003) also reports that oil prices have responded drastically to macroeconomic conditions since 1972. To the extent that oil prices are endogenous, cause and effect are not well-defined when relating changes in the oil price to macroeconomic outcomes. Without a well-defined cause, it is impossible to estimate its effect.

The insight of endogenous oil prices has been followed by a debate in the empirical literature about the extent to which oil shocks are then mainly caused by shocks to oil demand, oil supply, or increased risk of future short-falls in oil supply.⁴ While both sides admit that demand pressures matters, Hamilton argues that oil shocks are largely caused by direct supply shocks, whereas Kilian downplays these shocks and instead assigns the fear of future supply shocks as the most important factor.⁵

¹Examples include the Iran-Iraq war initiated in 1980, the first Persian Gulf War in 1990/91, and the oil-price spike of 2007/08. See also Hamilton (2011), which describes and discusses oil shocks that have occurred since the first one in 1862-1864.

²See Jones, Leiby and Paik (2004) for a review about oil price shocks and the economy. An additional need to to understand energy markets has emerged in recent years because of the need to deal with global warming.

³Examples include Pierce and Enzler (1974), Rasche and Tatom (1977), Bruno and Sachs (1982), Kim and Loungani (1992). Later contributions with a similar approach include Rotemberg and Woodford (1996) and Finn (2000).

⁴See, for example, Barsky and Kilian (2002), Hamilton (2003, 2009, 2011), Kilian (2008, 2009) and Kilian and Murphy (2014).

⁵Hamilton (2011) writes that “My overall conclusion thus supports the conventional interpretation: historical oil price shocks were primarily caused by significant disruptions in crude oil production that

This paper analyzes the interaction between oil prices and macroeconomic outcomes by setting up a general equilibrium model where output is produced with a capital-labor composite and oil and all factor prices are endogenous. A general equilibrium model is an ideal tool for controlling for reverse causality. It has the potential to simultaneously explain the movements in the oil price, identify how and when oil-shocks are realized and quantify the potential costs of volatile oil prices.

The basic framework is a modified real-business-cycle (RBC) model that takes oil as an input and allows for storage of oil above ground between periods. I introduce three types of shocks into this framework: factor-specific technology shocks, shocks to the supply of ground oil and noisy signals about future supply levels. Consistent with the discussion in the empirical literature, these shocks affect the demand and supply of oil as well as the risk of future supply levels.

The factor-specific technology shocks constitute demand shocks since they directly affect factor demands. These non-Hicks-neutral demand shocks are motivated by findings in Hassler, Krusell and Olovsson (2012) that show that technologies that save on capital/labor and energy typically grow at distinct rates and also have different volatilities.

Oil supply consists, in each period, of an exogenous amount that is extracted from the ground and is subject to stochastic shocks and an endogenous amount that has been kept in storage from the previous period. The forecasts of future supply levels, finally, are influenced by noisy signals - or news shocks - that agents receive in each period about the future supply of ground oil. Specifically, bad news reflects big and potentially important events, such as wars and political instability in the oil-supplying region, and they are allowed to affect the variance and/or the mean of the distribution of the future supply of ground oil.

I calibrate and estimate the shock processes, feed them into the model and analyze to what extent they can generate business-cycle properties similar to that found in the U.S. data, and to what extent the model economy goes through episodes with substantial oil shocks. Given that the model is successful along these dimensions, it can be used to analyze the feedback effects from oil shocks back to the economy.

The results provide contributions along three important dimensions. First, they show that the model closely matches a large number of observed business-cycle properties.

were brought about by largely exogenous geopolitical events.” In contrast, Kilian (2008a) writes “In general, exogenous oil production shortfalls are of limited importance in explaining oil price changes during crisis periods...Even if physical production does not move, expectations of future oil supply interruptions alone may have powerful effects, as is evident in my analysis of the 1990/91 Persian Gulf War episode.”

Specifically, it matches the volatilities of output, oil use and hours worked, thus implying that it also replicates the fact that oil use is more volatile than both output and hours worked. The model is able to explain between 80-95 percent of the volatility of the oil price, depending on how costly it is to store oil.⁶ Also, the correlation between output and hours worked and that between output and oil use are close to the data. The factor-specific shocks explain the bulk of the variation in inputs and output and about 45 percent of the overall volatility of the oil price. These findings are far from obvious, given that none of these properties are targeted in the calibration. In addition, even though standard RBC models have successfully shown that shocks to total factor productivity can explain a large fraction of observed properties for inputs and output, it has not yet been shown that shocks to factor productivities also can explain oil use and a large part of the oil price.⁷

Second, the distribution of oil prices is similar in the model and the data, implying that the model economy goes through episodes with realistic and endogenous oil-price shocks. In fact, the model can quantitatively match all observed oil shocks since 1900, with the exception of the episode in 1973/74.

Third, the analysis of the timing of oil shocks shows that the majority of oil shocks are driven by increasing demand to hold oil for precautionary/smoothing reasons. The median annual increase in the oil price after a bad news shock is 15 times larger than after a direct supply shortfall in the benchmark calibration. This difference is then reduced to a factor two when news shocks exclusively increase uncertainty and leave the expected value unchanged. These results are expected to increase with risk aversion, and they are broadly consistent with the findings in Kilian (2008a, 2009) that precisely show that fears of future shortfalls are important for understanding oil-price shocks.

The largest oil shocks generated by direct shortfalls in supply, however, are as large as the largest price increases generated by bad news. In addition, a variance decomposition reveals that direct supply shocks contribute substantially to both the overall volatility of the oil price and to the magnitude of the oil-price shocks.

The main influence of direct supply shocks on oil prices is somewhat more indirect in that it goes through the effect on oil reserves. These reserves vary over time due to shocks to demand and, in particular, to the supply. A series of relatively close supply shocks, in fact, will lead to a rundown of reserves and increasingly larger oil shocks each time a

⁶The marginal cost for storing oil is zero in the benchmark formulation of the model.

⁷In fact, Backus and Crucini (2000) consider Hicks-neutral technology shocks and find that these shocks have basically no impact on oil-related variables.

bad news shock hits the economy. The intuition is that, in the model, an oil shock is a signal of oil scarcity and exactly how scarce oil is when a shock hits depends on the size of oil reserves. The model's prediction of a negative relationship between annual oil-price changes and the stock of oil reserves is tested empirically and found to be statistically significant during the period 1972-2014.⁸

In the model as in the data, a contraction in economic activity is more likely than an expansion after an oil-price shock. This result is interesting in light of the long-standing question of whether oil-price hikes are contributing to U.S. recessions.⁹ Since the oil price is endogenous in the model, it never *causes* a recession. Instead, a sharp increase in the oil price that follows a bad news shock is a signal of potential future oil scarcity. The direct effect of this scarcity is a reduction in investments, which has a negative effect on future output. The direction of causality in the model, thus, runs from expectations of a contraction - due to oil scarcity - to the oil price and not from the oil shock to the contraction.

Finally, I use the model to compute the welfare costs associated with volatile oil prices. Following Lucas (1987), these costs are straightforwardly quantified by computing how much the representative agent would be willing to pay to get rid of the fluctuations in the oil price. The results from this exercise show that the welfare costs of volatile oil prices are negligible. In fact, the model does not provide any support for the notion that fluctuating oil prices are quantitatively harmful for the economy, at least not as long as it is possible to take precautionary actions against future short-falls in the oil supply. This result is maybe not too surprising; Lucas (1987) showed that the welfare costs associated with short-run fluctuations are negligible in a standard RBC model. Since the model in this paper is not too different from that in Lucas (1987), the results may be expected to be similar along this dimension.

The paper is structured as follows. Section 1 describes the data, the business cycle properties for the U.S. and some features of the oil shocks. Section 2 sets up the model, and the results are then presented in Section 3. Section 4 carries out a variance decomposition and empirically evaluates the relationship between price-changes and reserves, and Section 5 analyzes to what extent recessions are followed by oil-price hikes. The welfare costs associated with stochastic ground oil are computed in Section 6. The main assumptions are discussed in 7, and Section 8, finally, concludes.

⁸Over the longer period 1949-2014, the negative point estimate is only marginally significant if the state-owned oil reserves are excluded from the measure of reserves, which I argue is a reasonable assumption.

⁹See Hamilton (1983, 2003), and Kilian and Vigfusson (2014).

1 Data

This section describes the data to which the model is compared. All data is annual and the considered period is 1949-2014.¹⁰ Data on energy use is from the U.S. Energy Information Administration, and the data on oil prices is from BP.¹¹

The GDP data is denoted in chained (2009) dollars and is taken from the Bureau of Economic Analysis.¹² Hours worked per person between ages 15-64 is computed with data from the OECD database. It will, in the model, be possible to hold oil for precautionary and/or smoothing reasons. This requires a comparison with the actual oil reserves that are held in the data. Data on the ending stock of oil is taken from the Energy Information Administration. Specifically, the ending stock consists of the “primary stocks of crude oil and petroleum products held in storage as of 12 midnight on the last day of the month. Primary stocks include crude oil or petroleum products held in storage at (or in) leases, refineries, natural gas processing plants, pipelines, tank farms, and bulk terminals that can store at least 50,000 barrels of petroleum products or that can receive petroleum products by tanker, barge, or pipeline. Crude oil that is...in the Strategic Petroleum Reserve [SPR] is included.”¹³

The SPR is maintained by the United States Department of Energy and is the world’s largest supply of emergency crude oil. It was established in 1975 in the aftermath of the oil embargo in 1973/74 and was originally intended to hold reserves of up to one billion barrels of petroleum. The build-up of the SPR has mainly been politically motivated; it acts as a significant deterrent to oil import cutoffs and is an important tool of foreign policy. This fact is reflected in that the SPR only has been used three times. The first time occurred in 1991, at the beginning of the first Persian Gulf War. The second time was in September 2005 after Hurricane Katrina, and the third time was in 2011 to offset disruptions in global oil supplies caused by unrest in Libya and other countries.

¹⁰Papers on the business cycle typically concentrate on quarterly data. Since, I do not have access to quarterly oil-price data, however, this paper instead focuses on annual data.

¹¹BP reports that the oil price is Arabian Light posted at Ras Tanura for the period 1947-1983, and Brent-dated for the period 1984-2014.

¹²The data is found at http://www.bea.gov/histdata/Releases/GDP_and_PI/2009/Q4/Third_March-26-2010/Section1ALL_xls.xls.

¹³The data is available at http://www.eia.gov/dnav/pet/pet_sum_snd_a_epc0_mbbbl_m_cur.htm.

1.1 Business cycle properties

Most macroeconomic models with oil as an input completely abstract from other energy input and exclusively focus on oil. This simplification is typically done without any discussion.¹⁴ Since oil is not the only energy input in practice, however, a relevant question is to what extent such simplification is restrictive? Before the data is presented, a short discussion about different energy sources is then useful. Figure 1 plots consumption of the energy inputs that are used in the U.S. over the period 1949-2015. The top graph reveals that oil has consistently been the dominant energy source during this period. Oil is followed, in quantitative importance, by the fossil-energy sources of natural gas and coal. The use of nuclear power and renewable energy have been growing over time, but these sources have been of relatively marginal quantitative importance during the considered period.

Even though oil is the dominant energy input, a potential worry could be that alternative energy sources frequently substitute for oil in the data, whereas these alternatives are not available in the model. This is clearly not the case for the non-fossil energy sources. The amounts of energy generated by these sources have only been marginal, and the correlation between usage of renewable energy and oil is close to zero (at -0.14). It should therefore be without loss of generality to abstract from non-fossil energy inputs to understand the period of consideration, however important these inputs may become in the future.

What about the fossil-energy sources of coal and natural gas? The two bottom graphs in Figure 1 plot the relationships between oil and total fossil-fuel use. Specifically, the middle graph shows annual changes in percent, whereas the lower graph focuses on annual changes in levels. Both graphs display highly correlated series; the correlation for annual changes in percent is 0.9 and the correlation for levels is 0.85. This implies that changes in oil use correspond with changes in total fossil-fuel use on an almost one-to-one basis, and that there is limited substitution between the fossil fuels.¹⁵ In fact, the lower graph illustrates that positive (negative) changes in oil use are generally accompanied by positive (negative) changes in gas and coal.

I therefore conclude that it is not restrictive to focus on oil as the sole energy input for the United States. An alternative would be to model energy as a fossil-fuel composite

¹⁴See, for example, Kim and Loungani (1992), Rotemberg and Woodford (1996), Backus and Crucini (2000), and Bodenstein, Erceg and Guerrieri (2011).

¹⁵One reason for this is that the prices of the fossil fuels are positively correlated. The correlation between the price of oil and the price of gas is as high as 0.82.

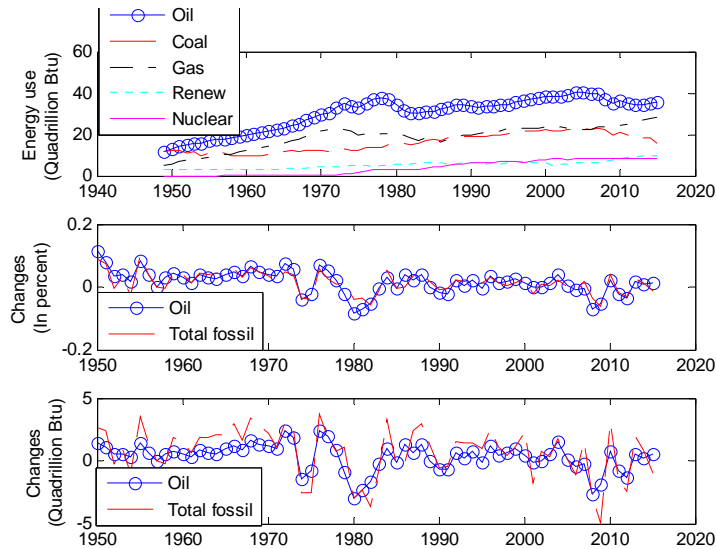


Figure 1: Oil and other energy inputs in the United States. Top graph: annual energy use of oil, coal, natural gas, renewable and nuclear energy. Middle graph: annual changes in percent of oil use and total fossil-energy use. Bottom graph: annual changes in levels of oil and total fossil-energy use. Source: the U.S. Energy Information Administration.

as in Hassler, Krusell and Olovsson (2012).¹⁶ Since it is more straightforward to calibrate the model to oil, however, I follow the previous literature and model the energy input as oil.¹⁷ Setting up a model that takes coal and gas in addition to oil, and that also can replicate the properties in Figure 1, is an interesting extension, but it is left for future research.

As stated in the Introduction, the model will be evaluated on its ability to match business-cycle properties as well as on its ability to produce oil shocks. We are now ready to present business cycle properties for the United States. The variables of interest are output (y), oil use (o), hours worked (l), the oil price (p_o) and oil kept in storage (R). I apply the HP filter with a smoothing parameter of 6.25 for annual data to the variables to produce cyclical components for these macroeconomic variables.¹⁸ Summary statistics for the selected series are provided in Table 1.

¹⁶In fact, the correlation between usage of oil and a fossil-fuel composite that is constructed as in Hassler, Krusell and Olovsson, is 0.99, and the correlation between the oil price and the fossil-fuel composite price is 0.94.

¹⁷Basically, the difference between oil and the composite is found in the calibration of the supply shortfall and in the resulting volatility of the price. Using oil has the advantage that there are empirical estimates that can be used in the calibration, whereas this is not the case for the composite.

¹⁸The value of 6.25 is based on the findings in Ravn and Uhlig (2002).

Table 1: Business cycle statistics for the U.S. economy

$std(y)$	$std(o)$	$std(l)$	$std(p_o)$	$std(R)$
1.63	2.10	1.22	14.53	2.41
$corr(y, l)$	$corr(y, o)$	$corr(y, \Delta R)$	$corr(y, p_o)$	$corr(p_o, \Delta R)$
0.88	0.78	-0.15	0.01	0.01

All variables are in logarithms and have been detrended with the HP filter. The time frequency is annual. y is output, o is oil, l is hours worked, p_o denotes the price of oil and R is oil kept in storage. If petroleum for passenger cars is deducted from the oil series as discussed in Section 2, the standard deviation for oil is 2.41. All other numbers in the table remain unaffected by this deduction.

The table shows that oil use is more volatile than both output and hours worked, with the latter being the least volatile of the three. The volatility of the oil price is almost an order of magnitude larger than for output. In addition, both energy and hours are highly correlated with output, but hours have a slightly higher correlation with output than energy. Interestingly, the statistics in Table 1 are close to those presented in Kim and Loungain (1992), even though their data only covers a period up to 1987.¹⁹

1.2 Oil price shocks in the data and in the literature

As is evident from the above section, the oil price is highly volatile. The post-war period has also demonstrated several periods with large price hikes, so called oil-price shocks. The first row in Table 6 presents the annual real price changes that occurred during five episodes with substantial annual oil price increases.²⁰

Table 2: Real oil-price changes

Year	1973/74	1978/79	1989/90	1999/00	2007/08
Δp_{oil} at t	217.18	102.57	24.12	53.58	38.23
Δp_{oil} at $t + 1$	-8.74	2.64	-19.18	-16.86	-38.12

Annual oil-price changes for six different episodes. The first row denotes the price change at the time of the “oil shock” and the second row shows the price changes realized in the year following the increase.

As is well known, the oil shocks of the 1970s are substantial. In fact, the real oil price increased by more than 200 percent in 1973/74 and by more than 100 percent in

¹⁹Their numbers are higher because they use a smoothing parameter of 400 in the Hodrick-Prescott filter, whereas I use a value of 6.25. However, the relative volatilities and the correlations are similar, except for the correlation between output and the oil price, which is lower in their paper.

²⁰It is somewhat standard in the literature to look at three-year net increases in the oil price. However, as shown in Figure 5 these two measures are highly correlated. In any case, the purpose of Table 6 is mainly to provide descriptive statistics and for this purpose it should not matter which of the two price series that is used.

1978/79. The other oil-price hikes are large, but relatively more modest: between 24 and 55 percent. The second row in the table displays the oil-price changes that were realized in the year following the initial increase. The 1989/90 episode and the 2007/08 episodes are interesting in that the oil price then reverts back by almost the same amount as the initial increase in the year following the increase. The table, thus, also reveals that large negative oil shocks occasionally occur.

2 The model

The ambition is now to set up a simple, realistically calibrated model that can replicate the business-cycle properties presented in Table 1, and also endogenously generate oil price shocks with similar properties to those presented in Table 6. The standard RBC framework is modified to incorporate three types of shocks: factor-specific technology shocks, shocks to the supply of ground oil and noisy signals about future supply levels. Consistent with the discussion in the empirical literature, these shocks affect the demand and supply of oil as well as the risk of future supply levels.

The inclusion of input-saving shocks is not a novel feature in itself in this setting. Backus and Crucini (2000) consider Hicks-neutral technology shocks, i.e., shocks that affect all inputs in the same way. They find that such shocks do not contribute much to the overall volatility of the oil price. There is, however, no reason to *a priori* assume that the efficiency of capital, labor and energy should evolve in exactly the same way and at exactly the same rate. In fact, Hassler, Krusell and Olovsson (2012) show that technologies that save on capital/labor and energy typically grow at distinct rates and also have different volatilities. The novel aspect of my formulation is precisely to consider factor-specific technology shocks. Variations in these shocks directly affect both factor demands and factor prices.

Oil-supply shocks are also not new in this setting. The new feature is that agents receive noisy signals, or news shocks, about the future supply of ground oil allowing for the possibility to take precautionary action against adverse outcomes.

Motivated by the findings in Section 1.1, there is one energy input in the model and it is thought of as oil. The model features a representative agent that derives utility from consumption and leisure. I omit time subscripts and use $'$ to denote next period's variables.

The agent has one unit of productive time per period, which must be divided between work and leisure. To specify the model as simply as possible, I abstract from households'

oil consumption. Oil is, thus, purely used as an input into production. This assumption is standard in the literature and made mainly for simplification.²¹ Assuming that oil is also used for consumption requires a more complicated utility function, preferable with durables that require oil.

The instantaneous utility function is given by

$$u(c, 1 - l) = (1 - \mu) \log(c) + \mu \log(1 - l), \quad (1)$$

where l denotes hours worked and $1 - l$ is leisure.

Output is produced with labor, capital (k), and oil (o), according to the following CES-function

$$y \equiv F(a, a_o, k, l, o) = \left[(1 - \gamma) (ak^\alpha l^{1-\alpha})^{\frac{\varepsilon-1}{\varepsilon}} + \gamma (a_o o)^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}}, \quad (2)$$

where $a \equiv \exp(z)$ and $a_o \equiv \exp(z_o)$, respectively, denote the separate shocks to the productivity of capital/labor and oil, and ε is the elasticity of substitution between capital/labor and oil. The production function in (2) ensures that the relative shares of capital and labor inherit their properties from the usual Cobb-Douglas form used in growth studies.²² Note that when $\varepsilon = 1$, the Cobb-Douglas composite and oil are perfect substitutes, when $\varepsilon = 1$, the production function collapses to being Cobb-Douglas in all input arguments; and when $\varepsilon = 0$ the Cobb-Douglas composite and oil are perfect complements, implying a Leontief function in the capital-labor composite and oil. The shocks a and a_o directly affect the demand for capital, labor and oil. They will, therefore, interchangeably be referred to as productivity and demand shocks.²³

²¹For the model to be directly comparable to the data, the direct usage of oil should then ideally be deducted from the data on total oil use. The main category of direct oil use is transportation services provided by the private use of cars. The other main category is housing services but they are, in fact, imputed into GDP. The oil used for heating of residential houses, therefore, should not be removed. Unfortunately, the only data on petroleum consumption that exists is an aggregate for all passenger cars, i.e., it also includes fuel for passenger cars that are used for professional services (and, thus, are measured in GDP). In addition, there are at least two breaks in the data series on petroleum consumption which invalidates comparisons over time. As a benchmark, I therefore follow Hassler, Krusell and Olovsson (2012) by not deducting the oil used by passenger cars. Because the correlation between overall oil use and petroleum for passenger cars is as high as 0.96, however, the only moment that is affected by its inclusion is the standard deviation of oil use. This standard deviation associated with a removal of the oil used for passenger cars is reported in Table 4.

²²A similar function is used in Hassler, Krusell and Olovsson (2012), Stern and Kander (2012) and Gars and Olovsson (2015).

²³The model is set up so that a and a_o can decrease from one period to the next. It is, however, straightforward to add positive trends to the productivity processes to make sure that the probability that productivity levels actually decrease is small. After a transformation to make the growth model

The processes for z and z_o are assumed to obey the following laws of motion:

$$z' = \rho_z z + \epsilon'_z, \quad (3)$$

$$z'_o = \rho_{z_o} z_o + \epsilon'_{z_o}, \quad (4)$$

with $0 \leq \rho_z \leq 1$, and $0 \leq \rho_{z_o} \leq 1$ and where ϵ'_z and ϵ'_{z_o} are both independent and normally distributed with mean zero and variances respectively given by σ_z and σ_{z_o} .

Aggregate capital evolves according to the following law of motion

$$k' = i + (1 - \delta) k, \quad (5)$$

where i_h denotes investment.

Finally, the aggregate resource constraint is given by

$$c + k' = y + (1 - \delta) k. \quad (6)$$

2.1 The supply of fossil fuel

The oil used in production can be sourced from two different suppliers: it can be pumped out from the ground and it can be sourced from inventories of stored oil. For simplicity, the supply of ground oil just follows an exogenous process. This assumption is standard in the literature and is, for instance, found in Backus and Crucini (2000), Bodenstein, Erceg and Guerrieri (2011) and Arezki, Ramey, and Sheng (2016). The assumption effectively implies a vertical short-run supply curve for ground oil (but not for the total supply of oil), and it is motivated by the fact that oil-producing countries seem to be slow to respond to demand shocks due to adjustment costs and uncertainty about the oil market.²⁴

Appendix G.4, however, shows that the results are not particularly sensitive to the assumption of inelastic ground oil by allowing for the possibility, at a cost, to expand this supply in some states of the world. The main effect of elastic ground oil is to reduce the amount of stored oil. Since inelastic ground oil is a simpler setup I abstract from elastic ground oil in the benchmark version of the model.

stationary over time, the transformed production function can be shown to be identical to 2. See also Hassler, Krusell and Olovsson (2016) for a more elaborate description of a version of the model with growth.

²⁴See Kilian (2009). He argues further that the unresponsive oil supply is consistent with evidence from interviews of Saudi officials in the early 1980s, and that the fact that state-owned Saudi oil company producers only produce forecasts for demand once a year.

2.1.1 Ground oil

In a given period, the supply of ground oil, o_g , is given by

$$o_g = \varphi \exp(\eta), \quad (7)$$

where φ is a parameter, and η is a random shock to the supply of ground oil. Next period's value of η' can be imprecisely forecasted with the noisy signal, ξ , that arrives in each period and contains imperfect information about the future supply of ground oil. Hence, ξ is a news shock about big and potentially important events that might affect the supply. Good news implies zero risk of a supply shortfall in the next period. Bad news, however, is thought of as wars and political instability in the oil-supplying region and they signal a positive probability of a supply shortfall in the subsequent period. The news shock is assumed to be state independent, thus implying that bad news is realized with the constant probability π_ξ :

$$\xi = \begin{cases} \xi_b & \text{w/prob } \pi_\xi \\ \xi_g & \text{w/prob } 1 - \pi_\xi. \end{cases} \quad (8)$$

The shock to the supply of ground oil then obeys the following process

$$\eta'(\xi) = \begin{cases} \eta_l \text{ w/prob } \pi_\eta, \text{ and } \eta_h \text{ w/prob } 1 - \pi_\eta & \text{if } \xi = \xi_b \\ \eta_m \text{ w/prob } 1 & \text{if } \xi = \xi_g. \end{cases} \quad (9)$$

The law of motion for ground oil implies that agents know with certainty that next period's supply equals η_m if news is good ($\xi = \xi_g$). If news is bad ($\xi = \xi_b$), however, next period's supply is low with probability π_η , and high with probability $1 - \pi_\eta$.

Depending on the values η_l , η_m and η_h , a bad news shock can thus imply a higher variance and/or a different mean of the distribution of ground oil relative to a good news shock. The benchmark calibration sets $\eta_h = \eta_m$, so that bad news implies both more uncertainty as well as a lower expected value of ground oil. To isolate the effect of increased risk, however, I also report results for when news exclusively implies a mean-preserving spread in the distribution for the supply of ground oil.²⁵

The supply process for ground oil deserves some comments. First, it does not allow for serially correlated shocks. This assumption is supported by the empirical estimates of supply shortfalls that are provided in Kilian (2008a). Specifically, he finds exogenous oil

²⁵Naturally, a mean-preserving spread requires that the supply of ground oil is substantially above average levels in some states following bad news.

supply shocks to be serially uncorrelated. I also abstract from autocorrelated shocks to the supply of ground oil in good times. With elastic oil supply, however, deviations from trend are endogenously highly autocorrelated in good times. As shown in Appendix G.4, this does not change the results.

Second, the process does not allow for positive shocks to the supply. This is purely for simplicity, but it is straightforward to extend the (9) to also allow for positive shocks.²⁶

Third, there is a tight link between news and supply shocks in that all supply shocks are preceded by noisy news. This assumption is discussed and relaxed in the sensitivity analysis in Section 7.1.

Fourth and finally, Hamilton (2003) points out that exogenous drops in oil extraction often are followed by endogenous increases in oil supply from other sources. Such endogenous response can, in the model, potentially come from the supply of stored oil, which is the topic of the next section.

2.1.2 Stored oil

It is possible to store oil above ground between periods. The stock of stored fuel, R , must then obey the following law of motion:

$$R' = R - o_s, \quad (10)$$

where o_s is the net outtake of fossil fuel from storage. The amount of stored oil must, naturally, be positive, i.e., it is required that

$$R \geq 0. \quad (11)$$

Since the cost of storing a marginal unit of oil, κ , is likely negligible in the aggregate, this cost is set to zero in the benchmark calibration. As a comparison, a small positive marginal cost will also be considered. Note, however, that even though there is no physical cost to storing oil, there is an opportunity cost in the form of delayed production and consumption.

The total amount of oil available for production is then, in each period, equal to the

²⁶An extension of the model with positive shocks to the supply would allow for an analysis of effects of the shale oil revolution on the oil price. In fact, Baumeister and Kilian (2016) argues that the fall in the oil price since June 2014 has been driven by the same shocks that are already present in this paper: shocks to the aggregate demand for oil, the supply, and expectations about future supply levels.

sum of the oil from the ground and the net outtake from storage, i.e.,

$$o = o_g + \underbrace{R - R'}_{o_s}. \quad (12)$$

2.2 Equilibrium

Since there are no market failures in the model the problem can, without loss of generality, be formulated in recursive form as a social planning problem. The state variables are the factor specific shocks, the capital stock, the stock of oil reserves, the shock to the current supply of ground oil and the news shock. Denoting the Lagrange multiplier on (11) by λ , the planning problem is given by (13)

$$\begin{aligned} V(a, a_o, k, R, \eta, \xi) = \max_{k', R', l} & (1 - \mu) \log \left(\left[(1 - \gamma) (ak^\alpha l^{1-\alpha})^{\frac{\varepsilon-1}{\varepsilon}} + \gamma (a_o (o_g + R - R'))^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}} \right. \\ & + (1 - \delta) k - k' - \kappa R' + \mu \log(1 - l) \} + \beta \mathbb{E}_t V(a', a'_o, k', R', \eta', \xi') \\ & \left. - \lambda R' \right) \\ & \text{s.t. (3), (4), (8), and (9)}. \end{aligned} \quad (13)$$

The first order conditions to (13) deliver an Euler equation for capital, an equation that governs the return to storing oil and an intra-temporal equation for labor supply. These equations are laid out in Appendix C. The unknown variables are k' , R' , l , o_s , c_t and λ , and the equations that characterize the equilibrium are (6), (10), (11), (19), (20), (21), as well as laws of motions for factor productivity (3)-(4), and for news and the supply of ground oil (8)-(9).

The price of oil is, in each period, given by the marginal product of oil in that period, i.e.,

$$p_o(a, a_o, k, R, \eta, \xi) = \frac{y(a, a_o, k, R, \eta, \xi)}{o_g + R - R'(a, a_o, k, R, \eta, \xi)} \frac{\gamma (a_o (o_g + R - R'(a, a_o, k, R, \eta, \xi)))^{\frac{\varepsilon-1}{\varepsilon}}}{(1 - \gamma) (ak^\alpha l(a, a_o, k, R, \eta, \xi))^{1-\alpha} \frac{\varepsilon-1}{\varepsilon} + \gamma (a_o (o_g + R - R'(a, a_o, k, R, \eta, \xi)))^{\frac{\varepsilon-1}{\varepsilon}}}. \quad (14)$$

Note that even though the price is determined by the marginal product of oil in the current period, this marginal product depends crucially on the future through the news shock ξ . Specifically, a bad news shock in a specific period may induce more oil to be put into storage. That would reduce the amount available for current production and increase the marginal product of oil as well as the oil price.

2.3 Hotelling pricing with a storage technology?

Hotelling (1931) shows that the net price received after paying the cost of extraction of a resource in finite supply, such as oil, should grow at the rate of the interest rate. This is a prediction that seems to have weak empirical support for many natural resources.²⁷ A natural question is how the oil price evolves in the current setting, i.e., when oil supply has an exogenous part as well as an endogenous part that can be allocated across time? Section D in the Appendix shows that the existence of an exogenous amount of oil (or an exogenous substitute to oil, such as nuclear power) leads to a breakdown of the classical Hotelling relation where net oil prices simply grow at the rate of the interest rate.

2.4 Calibration

The model is calibrated as follows. The discount factor is set at $\beta = 0.98$ to generate an interest rate around two percent. The utility weight on leisure is set to $\mu = 0.65$, which implies that the representative agent dedicates roughly 1/3 of her time to work. There are three production parameters: α , γ and ε . The parameter α is set at $\alpha = 0.32$, to generate a capital share of income of 0.3. The weight on energy in production, γ , is set to 0.05. The elasticity of substitution between capital/labor and fossil energy, ε , has been analyzed in a large number of studies. Early empirical investigations are Hudson and Jorgenson (1974) and Berndt and Wood (1975), who both use time-series data to estimate the substitutability of energy with other inputs. Both studies find capital and energy to be gross complements. Griffin and Gregory (1976), instead, uses pooled international data to capture the long-run relationships between capital and energy. The finding is that capital and energy are substitutes in the long-run. Koetse et al. (2008) carries out a meta-regression analysis on empirical estimates of capital-energy substitution and conclude that the demand for energy-saving capital is affected by energy-price increases, but that it generally takes significant time before demand reacts.

Another approach to estimate the elasticity is applied in Hassler, Krusell and Olovsson (2016). Specifically, they consider a production function similar to (2), and allow the short-run elasticity to differ from the long-run elasticity. The short-run elasticity is governed by the parameter ε , whereas the possibility to substitute in the long-run is affected by endogenous investments into the *levels* of a and a_o . They then estimate the short-run, i.e., one-year elasticity of substitution between capital/labor and fossil energy and find it to be close to zero, in fact, not statistically significant from zero. In addition,

²⁷See, for example, Hart and Spiros (2011).

the results show that elasticities closer to one seem hard to reconcile with the data.²⁸ Since the setting in Hassler, Krusell and Olovsson is close to the one in this paper, in combination with the fact that this paper is concerned with the short run and all studies find capital and fossil energy to be short-run complements, ε is set to 0.03.²⁹ Results are also reported for higher elasticities. As will become apparent, however, for the model to successfully replicate the data it is necessary that the short-run elasticity of substitution is not too high. The elasticity of substitution is discussed further in Section 7.

The scarcity of energy is determined by the parameter φ . This value is set so that energy's share of income on average equals 2.5 percent.

I turn now to the calibration of the input-saving shocks (3)-(4). Hassler, Krusell and Olovsson (2012) shows that the level of the technologies that save on capital/labor and oil, respectively, can be computed from the following two equations:

$$\log A_t = \log y_t - \log (k_t^\alpha l_t^{1-\alpha}) + \frac{\varepsilon}{\varepsilon - 1} \log \left(\frac{w_t l_t}{y_t} \right) - \frac{\varepsilon}{\varepsilon - 1} \log ((1 - \alpha) (1 - \gamma)) \quad (15)$$

and

$$\log A_{o,t} = \log (y_t) - \log o_t + \frac{\varepsilon}{\varepsilon - 1} \log \left(\frac{p_{o,t} o_t}{y_t} \right) - \frac{\varepsilon}{\varepsilon - 1} \log \gamma. \quad (16)$$

The standard procedure when computing the shock processes for z and z_o , is to fit linear trends to (15) and (16) and then use the residuals from these regressions to estimate ρ_z , ρ_{z_o} , σ_z and σ_{z_o} . A potential problem here is Hassler, Krusell and Olovsson's finding of directed technical change. Specifically, they show that average growth rates for capital/labor and energy are changing over time in a negatively related fashion. The capital/labor-saving technology, for instance, grew relatively fast before the oil shock in 1973 and relatively slow for almost a decade after this shock with the opposite being true for the energy-saving technology. This negative tradeoff for the medium-term is present throughout the whole sample period, implying that the standard assumption of constant underlying growth rates for A and A_o is not applicable in a meaningful way the current setting. To remove the variation in the trend-components, I instead estimate the relevant parameters for four sub-periods with distinct average growth rates, 1949-1972, 1973-1985, 1986-1997, and 1998-2011, and set the parameters ρ_z , ρ_{z_o} , σ_z and σ_{z_o} to their

²⁸Specifically, higher elasticities imply nonsensical evolutions for the technology trends.

²⁹An elasticity of substitution close to zero is also employed in Rotemberg and Woodford (1996) and Backus and Crucini (2000).

weighted averages of the estimates over the periods.³⁰ Appendix E compares the main simulation results in the paper to those that instead results from averaging over three periods, before, during and after the oil shocks of the 1970s, and show that the results then are close to identical between these two specifications, even though the estimated parameters differ. Failure to account for changes in trend growth, however, produces very high and unrealistic volatilities for the two input-saving shocks.³¹

Table 3: Parameter estimates

	ρ_z	ρ_{z_o}	σ_z	σ_{z_o}	$corr(z, z_o)$
1949-1972	0.61	0.71	0.0260	0.0253	0.64
1973-1985	0.35	0.74	0.0209	0.0443	0.24
1986-1997	0.80	0.62	0.0114	0.0232	-0.13
1998-2011	0.47	0.52	0.0081	0.0195	-0.69
Weighted average	0.40	0.54	0.016	0.0250	0.11

Estimated parameters for different time periods, with the weighted averages.

The correlation between z and z_o is set to zero in the model, which is close to the correlation of 0.11 that is found in the data. Setting up a process with an intermediate correlation requires additional assumptions for the transition matrix and increases the already large state space.

The benchmark calibration sets $\eta_h = \eta_m$, so that bad news implies both more uncertainty and a lower expected value of ground oil. This seems realistic in that bad news signals a higher risk of a future shortfall in supply, but does not require the supply of ground oil to be substantially above average levels in some states following bad news. To explicitly add to the debate about the main causes of oil shocks, however, I also report results for when news exclusively produces a mean-preserving spread in the distribution for the supply of ground oil.

In the benchmark calibration, the supply of ground oil can then be high or low. Hamilton (2003) estimates the exogenous shortfalls in world production for five specific events and finds them to be between 7.2-10.1 percent. Specifically, the events are the Suez crisis in 1956, The Arab-Israeli war in 1973, the Iranian revolution in 1978/79, the

³⁰The growth rates for the oil-saving technology is for the four periods given by 0.24, 3.20, 1.26, and 1.96 percent per year.

³¹The four-period average is motivated by the fact that a is similar in concept to the standard TFP shock. This shock should therefore ideally have properties that are not too different from the standard shock, which is the case with four periods. Without taking changing trend growth into account the volatility of a and a_o are however, respectively, more than two and five times as large as output. See Section E for a more detailed discussion about the estimation process.

Iran-Iraq war in 1980, and the Persian Gulf War in 1990/91.^{32,33} I set $\eta_l = \eta_m - x$ and choose x and η_m so that the unconditional expected value of η equals zero and that conditional on bad news in period t , the expected short-fall in supply in the subsequent period is equal to 10.1 percent. This is the upper bound of Hamilton's estimates, but the interval is fairly small.³⁴

A mean-preserving spread leaves the calibration of η_l and η_m unchanged, but requires η_h to be set to the value that implies that, also conditional on bad news, the expected value of η equals zero.

Bad news is thought of as big events such as wars and political instability in the oil-supplying region and they signal a positive probability of a supply shortfall in the subsequent period. There have been roughly nine distinct episodes of net oil-price increases since 1949.³⁵ Even though most of these episodes were associated with, at least, some supply short-falls, there are also a few cases where bad news did not translate into a drop in oil supply.³⁶ Based on these facts, the two probabilities are set so that bad news occurs roughly every tenth year, whereas the probability that it actually materializes into a supply shortfall is ninety percent. This implies setting $\pi_\xi = 0.10$ and $\pi_l = 0.90$.

2.5 Numerical approximation of the equilibrium

The decision rules are approximated by cubic splines over a six-dimensional grid. Specifically, the grid is continuous in k and R but, given the curse of dimensionality, the variables η , ξ , a and a_o are all discretized and only allowed to take two values each.³⁷ The parameters that govern the processes for η and ξ are described in the previous section. The shocks a and a_o are discretized with the algorithm described in Tauchen and Hussey (1991). The model is simulated for 16,000 periods, but the first 1,000 periods are thrown

³²The exogenous shortfalls that are reported in Hamilton (2003) are related to world production in the respective years. The drop during the Iranian revolution in 1978, for instance, is said to represent 8.9 percent of the 1978 world production.

³³Kilian (2008a) considers an alternative measure of exogenous shortfalls in oil supply. His measure suggests smaller shocks on impact. The index, however, displays substantial volatility which could be modeled as smaller but more frequent shocks to the supply.

³⁴Note, however, that it cannot be ruled out that the *expected* short-falls at each of the five episodes analyzed by Hamilton were potentially substantially larger than the *realized* short-falls. In other words, just because a substantial supply short-fall has not yet occurred, this does not necessary lead to the conclusion that such short-fall should be ruled out.

³⁵Kilian and Vigfusson (2014) argues that there have been eight episodes since 1974.

³⁶This episode is the oil-price hike in the mid-1990s. See Kilian (2008b) for a discussion.

³⁷The model has also been solved with four states for the technology shocks and the results are then very similar to the results with two states.

away. Statistics are then computed from the remaining periods.

3 Results

The results from the model simulations are now presented and compared to the data. Section 3.1 presents business cycle properties, whereas Section 3.2 illustrates how oil-price shocks are realized in the model.

3.1 Business cycle statistics

The output from the simulations of the model is presented in Table 4 along with the data. Results are shown for two specifications of the model: “Model 1” features zero marginal costs for storing oil, whereas “Model 2” feature a small positive marginal cost for storing oil.^{38,39}

Table 4: Business cycle statistics in the data and the model

	$std(y)$	$std(o)$	$std(l)$	$std(p_o)$	$std(R)$
Data	1.63	2.10	1.22	14.53	2.41
Model 1	1.68	2.08	1.05	11.75	1.46
Model 2	1.68	2.06	1.10	13.76	1.40
	$corr(y, l)$	$corr(y, o)$	$corr(y, \Delta R)$	$corr(y, p_o)$	$corr(p_o, \Delta R)$
Data	0.88	0.78	-0.15	0.01	0.01
Model 1	0.91	0.72	-0.09	0.17	0.06
Model 2	0.89	0.72	-0.10	0.16	0.07

All variables are in logarithms and have been detrended with the HP filter. The upper half of the table provides standard deviations and the lower half presents correlations. Model 1 has zero marginal cost for storing oil, whereas Model 2 features a small positive cost for storing oil. ΔR is the net change in the reserve stock of oil.

It is striking how well the model explains the data. This is far from obvious since none of the ten moments that are presented in the table were targeted in the calibration. The model matches the observed volatilities of output and oil use. Specifically, it replicates the fact that oil use is more volatile than both output and hours worked, with the latter being the least volatile of the three. As in standard RBC models, this model also suffers from the fact that the standard deviation for hours worked relative to the standard deviation

³⁸The marginal cost is not calibrated to any specific number in the data, but the total resources spent on storing oil as a share of output in the model is around one-hundredth of a promille.

³⁹The results for a mean-preserving spread are reported in Appendix F. As can be seen, the business cycle properties are similar between these two specifications.

for output is too small in the model.⁴⁰ An extension of the model with variable capacity utilization will potentially improve the model in this dimension.

The specification without costs for storing oil explains roughly 80 percent of the volatility of the oil price. However, the model with storing costs increases this factor up to 95 percent. Intuitively, a marginal cost for storing oil leads to lower quantities of stored oil, which implies less available resources for smoothing oil between periods. It is interesting to note that the oil price is so volatile even though it is possible to store oil between periods. As we will see in Section 6, this fact is explained by the finding that the high volatility is not costly in terms of consumption units.

The standard deviation that the model is least successful in matching is that for reserves. In fact, the model only accounts for about 70 percent of the volatility in the data. As discussed in Section 1, the buildup of the SPR has to a large extent been politically motivated in that it partly acts as an important tool of foreign policy. Since oil is not traded for political and strategic reasons in the model, the volatility of the reserve stock is expected to be different than in the data.⁴¹

The model predicts a correlation between output and hours worked of around 0.9, which is close to the true correlation of 0.88. Standard RBC models typically produce substantially higher values for this correlation: around 0.97. Complementarity between energy and capital/labor and separate input-saving technology shocks brings the correlation down to that in the data. Also, the correlations between oil use and output, between changes in reserves and output, and between changes in reserves and the oil price are close to the data. Only the correlation between output and the oil price is slightly off.

The model in this paper is basically a real-business-cycle model and even though standard RBC models have successfully shown that shocks to total factor productivity can explain a large fraction of observed properties for inputs and output, it has not until now been shown that shocks to factor productivities also can explain important features of oil use and part of the oil price.

Figure 2 displays the evolution of energy's share of income in the model and the data. The share in the data is located around two percent for most of the time, but increases substantially in response to oil shocks like those in the 1970s. The bottom graph plots the model-share in a simulation of the model and illustrates that it resembles the data: it is also mostly concentrated around two percent and then increases dramatically when

⁴⁰See, for instance, King and Rebelo (1999).

⁴¹Not shown in the table is the standard deviation for investments, which is 6.24 percent in the model and 5.30 percent in the FRED database. Hence, the model predicts a volatility for this moment that is just a bit high.

adverse shocks hit the economy at a bad time.




Figure 2: Top graph: the oil share of income in the data. Bottom graph: the oil share of income in a simulation of the model.

It is important to note here that the model is not calibrated to match any specific feature or time period in the data. Instead, the bottom graph just shows a simulation of the model and the peak is placed in the middle of the interval where it also happens to be in the data. Even though Figure 2 just shows an example, the properties for energy's share of income in the model are similar to the data due to the fact that p_o , o , and y all match the data.

3.2 Oil-price shocks in the model

To what extent does the model economy go through episodes with oil-price shocks? Figure 3 answers that question by plotting the distributions of annual oil-price changes in the data as well as in the model.

The two distributions are similar, thus implying that the model economy occasionally experiences substantial oil-price shocks with realistic properties. The main differences are that the model distribution has slightly more mass for large negative oil shocks, but also slightly more mass for shocks that are somewhat larger than 100 percent relative to the data. In contrast to the data, however, the model has zero mass for price changes of more than 200 percent. Consequently, the model can quantitatively match all the observed oil shocks since 1900 with the exception of the episode in 1973/74. It can also account for




Figure 3: The distribution of annual real oil-price changes in the data and the model.

the large negative oil-shocks that are a noticeable feature of the data.

4 Variance decomposition

This section decomposes the individual contributions of the different shocks to the considered volatilities and the oil-price shocks. The analysis starts once more with the business cycle statistics and then moves on to oil shocks.

4.1 The business cycle

To put the contribution of factor-specific shocks into contexts, the results are also presented here for Hicks-neutral technology shocks, i.e., shocks that affect all inputs in the same way. For this specification, the production function is given by

$$y \equiv F(z, k, l, o) = \exp(\hat{z}) \left[(1 - \gamma) (k^\alpha l^{1-\alpha})^{\frac{\varepsilon-1}{\varepsilon}} + \gamma o^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}}, \quad (17)$$

where $\exp(\hat{z})$ is the Hicks-neutral technology shock. As in the benchmark specification, the process for \hat{z} follows an AR(1)-process

$$\hat{z}' = \rho_{\hat{z}} \hat{z} + \epsilon'_{\hat{z}}, \quad (18)$$

where $0 \leq \rho_{\hat{z}} \leq 1$, and $\epsilon'_{\hat{z}}$ is normally distributed with mean zero and variance given by $\sigma_{\hat{z}}$. The parameters $\rho_{\hat{z}}$ and $\sigma_{\hat{z}}$ are set to 0.95 and 0.1, respectively.

The variance decomposition is presented in Table 5.

Table 5: Decomposing the contributions of each shock

	$std(y)$	$std(o)$	$std(l)$	$std(p_o)$	$std(R)$
Only shocks to a and a_o	1.69	2.12	0.94	5.85	1.09
Only shocks to \hat{z}	0.79	0.02	0.22	4.36	0.02
Only shocks to ξ and η	0.26	0.49	0.39	7.80	1.04
Only shocks to ξ	0.001	0.002	0.001	4.31	0.001

All variables are in logarithms and have been detrended with the HP filter. a and a_o are shocks to the productivity of capital/labor and oil, respectively; \hat{z} is a Hicks-neutral productivity shock; ξ is a news shock and η is a supply shock.

The first row in the table shows that the demand shocks a and a_o explain the bulk of the variation in inputs and output. The volatilities for output and oil consumption are basically unchanged relative to the full model. The standard deviations for hours worked and oil reserves are, however, slightly lower with only demand shocks.

The second row provides the effects associated with Hicks-neutral technology shocks. This specification does a particularly poor job matching the data. Output is only 50 percent as volatile as in the data, and the corresponding numbers for hours worked, oil use and reserves are respectively 18, 1, and 1 percent. This finding is consistent with Backus and Crucini (2000), who consider Hicks-neutral technology shocks and also find them to not contribute much to the volatility of the oil price.

Results for news and supply shocks are stated in the third row. In line with the finding that demand shocks alone explain the bulk of the variation in inputs and output, news and supply shocks add little to the variation in these variables. However, these shocks do contribute substantially to the oil-price volatility. This volatility is 33 percent higher with news and supply shocks than with demand shocks. Consequently, two-thirds of the oil-price volatility originate from news and supply shocks, whereas one-third comes from demand shocks.⁴² The fourth row features only news shocks, implying that supply shocks never actually materialize. This specification reveals that 55 percent of the two-thirds come purely from changing expectations due to news shocks. The contribution of supply shocks to the overall volatility of the oil price is thus 45 percent of the two-thirds, i.e., 30 percent in total. Section 4.3 further analyzes the respective contributions of news and

⁴²Due to the fact that the shocks interact in a non-linear fashion, it is not possible to sum the volatilities and get exactly the same volatility as in the benchmark economy.

supply shocks. Before that it is done, however, it is necessary to first understand how and when oil shocks are generated in the model

4.2 The timing of oil shocks

The results are reported here for the benchmark calibration as well as for when news exclusively produces a mean-preserving spread in the distribution for the supply of ground oil. The first two columns under the respective headings in Table 6 show the changes in oil prices as a percent that are respectively realized after a bad news shock and a supply short-fall.

The results show that even though maximum increases are similar for the two shocks, the median price increase associated with bad news is substantially higher than that of a direct supply short-fall. This is particularly true in the benchmark calibration where the effect of news shocks is almost 15 times larger than that of direct supply shocks. Also, when news only signals a mean-preserving spread they are more important than direct supply shocks, even though the difference then is reduced to factor of two. In this case, fewer reserves are accumulated when a news shock hits, which results in relatively less impact from the news shock and relatively more if/when a supply shock hits. These effects, however, are expected to be even larger with higher risk aversion. Hence, even though direct supply short-falls occasionally can and do generate substantial oil-price increases, the most common reason for oil shocks is fears of future short-falls in the oil supply.

Table 6: Changes in oil prices and reserves between two consecutive years (in percent)

	Higher uncertainty and a lower mean			A mean-preserving spread		
	News	Supply	$\eta_t = \eta_h \xi_{t-1} = \xi_b$	News	Supply	$\eta_t = \eta_h \xi_{t-1} = \xi_b$
$Md(p_o)$	18.80	1.28	-17.28	12.49	6.10	-53.03
$Max(p_o)$	132.99	128.51	-71.40	140.64	138.11	-88.40

The first column under each heading presents the oil-price changes that are realized when bad news hit the economy. The second column presents the changes that are realized when a supply shock hits the economy. The third column, finally, shows the changes that are realized when the supply of ground oil turns out to be high even though news was bad in the previous period. All changes are relative to the previous period.

The third column under the respective headings in Table 6 focuses on the mechanisms behind negative oil shocks by showing oil-price changes for events where bad news was never followed by a drop in the supply of ground oil. The median fall in the oil price in the benchmark calibration is then 17 percent. A real-world example of such an event is

the sharp increase in the oil price in the mid-1990s. This increase is explained by the “news” that Iraq might invade Saudi Arabia and occupy its oil fields. The occupation never happened, however, and the price fell sharply after the U.S. had moved enough troops to Saudi Arabia to prevent the occupation.⁴³ In both the model and the data the oil price can, thus, fall sharply after uncertainty has been resolved.

The results in Table 6 are broadly consistent with the findings in Kilian (2008a, 2009) that show that fears of future short-falls are important for understanding oil-price shocks. At the same time, the previous section reveals that direct supply shocks contribute substantially to the overall oil-price volatility. The next section explains exactly how supply short-falls influence the volatility of the oil price and the magnitude of oil shocks.

4.3 Oil-price shocks and the level of oil reserves

The top graph in Figure 4 plots the evolution of reserves and oil-price changes in a simulation of the model without factor-specific shocks.




Figure 4: Top graph: reserves and the oil price in a simulation of the economy without demand shocks. Bottom graph: same as in the top graph except that reserves are artificially increased to the mean level after each time they have been run down.

It shows three distinct spikes in the oil price at periods 10, 14, and 16. All three spikes coincide with bad news. The first oil shock (in period 10) is relatively modest, due to the fact that the spare capacity of oil is high when the news shock hits the economy. The consecutive drop in supply, however, leads to a run-down of reserves. The second

⁴³See Kilian (2008b).

news shock (in period 14) then hits the economy at a more vulnerable state, resulting in a somewhat larger spike in the oil price. Reserves are then reduced even further so that the third news shock (in period 16) generates an oil shock that is more than twice as large as the initial shock in period 10.

As a comparison, the bottom graph shows the evolution of the oil price in a simulation where reserves are artificially increased back to the mean level each time after that they have been run down. All oil shocks then remain modest and below 20 percent.

The magnitude of an oil-price shock in a given period, thus, depends crucially on the level of reserves in the same period and this level depends directly on past supply shocks. The intuition is straightforward: a sharp increase in the oil price that follows bad news is a signal of potential oil scarcity. Exactly how scarce it is depends on the level of current reserves. It is also when oil reserves are low that supply shocks can directly generate substantial oil shocks. Hence, even though most oil shocks are associated with increasing precautionary/smoothing demand for oil, the magnitudes of these shocks depend in a quantitatively important way on the recent history of supply shocks. The influence of supply shocks on the volatility of the oil price and on oil-price shocks is, however, more indirect; they (together with demand shocks) affect reserves which influence the magnitude of the oil shocks. The above analysis verifies the empirical observation that the oil price can react in substantially different ways to the same shock. The reason is that the state of the economy might be different when the shock hits.

4.3.1 Oil-price shocks and the level of oil reserves in the data

To what extent is the negative correlation between reserves and oil-price changes found above also present in the data? To answer this question I regress annual oil-price changes on the stock of oil reserves.⁴⁴ Specifically, I run the following regression

$$\Delta p_{o,t} = a + b * Effective_reserves_{t-1} + u_t,$$

where $\Delta p_{o,t} \equiv p_{o,t}/p_{o,t-1}$ is the measure of annual oil-price changes, *Effective_reserves* is a measure of the stock of oil in inventories and u_t is an error term. Important to point out is that I am here only looking for a correlation between two variables. Hence, I do not take a stand on whether one variable could be considered exogenous. Since annual oil consumption has been growing over time during the considered period, the measure of

⁴⁴See also Kilian and Murphy (2014), who take inventories of oil into account in a structural VAR specification.

reserves must be scaled to control for the fact that a reserve stock of 100 million barrels in 2014 was a relatively smaller buffer than in 1949.⁴⁵ For this reason, *Effective _ reserves* are defined as follows

$$Effective_reserves_t = \frac{Actual_reserves_t}{\frac{1}{x+1} \sum_{j=t-x}^t O_j}.$$

The definition of effective reserves relates actual reserves in year t to average consumption between year $t - x$ and t . I set $x = 2$, which implies averaging over three years, but results are similar for alternative values for x . Actual reserves are measured in January each year.

The regressions are carried out in the full sample, as well as for the more recent period 1972-2014. The motivation for looking specifically at this period is that Hamilton (1985) and Barsky and Kilian (2002) both argue that the period 1949-1972 differs substantially from that after 1972 in two important dimensions. First, in the former period oil delivery contracts were long-term agreements between OPEC producers and oil companies. During this regime, the producers agreed to supply oil at prices that remained fixed in nominal terms for several years in advance. Second, between 1948-1972, the U.S. had the capacity to act as a supplier of last resort to both Europe and Japan. Not only did this excess oil-capacity limit the need for holding oil reserves, but it also provided a hedge against supply short-falls in other oil-producing countries. Both these features changed in the early 1970s.⁴⁶

Results are also reported for reserve measures that exclude the SPR. The rationale for excluding these reserves is that the SPR mainly is a political tool that only has been used three times over a 40-year period. It therefore seems reasonable to assume that the market does not actually rely on the SPR when deciding how much precautionary oil to hold, since it is unlikely that these reserves will be used should there be supply short-fall.

The regression results are presented in Table 17 in Appendix I. All regressions deliver a negative point estimate for the effect of the level of effective reserves on oil-price changes. The point estimate, however, is not statistically significant from zero in the full sample when the SPR is included in the measure of effective reserves. This estimate then becomes marginally significant when the SPR is excluded from effective reserves.

For the period 1972-2014, however, the negative point estimates are statistically sig-

⁴⁵Oil consumption was three times higher in 2014 than in 1949. A reserve stock of 100 million barrels thus provided a completely different spare capacity in 1949 compared to in 2014.

⁴⁶See Hamilton (1983) and Barsky and Kilian (2002) for details.

nificant irrespective of whether SPR is included or not. This negative effect is strongest when the SPR is excluded from effective reserves. A reduction of the reserve stock by 50 percent then leads to five times higher oil-price changes (compared to 2.8 times higher when the SPR is included).

I conclude that the negative correlation between oil-price changes and reserves identified in Section 4.3 is present in the data since 1972. This is particularly true if the market does not rely on SPR when deciding how much precautionary oil to hold. Given that these reserves are rarely used, this seems like a natural assumption.

Note, finally, that even though the level of reserves matters for the size of the oil shock, changes in reserves are basically uncorrelated with the oil price as is shown in the last bottom column in Table 4. The reason for the zero correlation in the model is that there are several interacting shocks that affect reserves in different directions.

5 Are oil price increases followed by economic contractions?

A long-standing question in the literature concerned with oil prices is to what extent oil-price hikes are contributing to U.S. recessions. In a seminal paper, Hamilton (1983) shows that all but one of the U.S. recessions between World War II and 1983 was preceded by a dramatic increase in the price of crude petroleum. Hamilton (2011) extends the period of examination and argues that all but one of the postwar recessions up to 2011 were preceded by substantial oil-price increases.

Kilian and Vigfusson (2014), instead, look at the period 1974-2014 and arrive at a slightly different conclusion. Specifically, they argue that there have been eight distinct episodes of net oil-price increases during the observed period and that only five of them were followed by recessions.

The model in this paper features both business cycles and oil shocks, which makes it ideal for analyzing the relationship between oil prices and economic activity. I, here, focus on unfiltered data since this is always the focus in the empirical literature. Specifically, I analyze to what extent substantial oil price increases are followed by contractions in economic activity. NBER defines a recession as “a significant decline in economic activity spread across the economy, lasting more than a few months, normally visible in real GDP, real income, employment, industrial production, and wholesale-retail sales.” Following the NBER definition, the measure of economic activity considered is one where both GDP

and hours worked decrease relative to the previous period.⁴⁷ Table 7 presents the share of substantial oil-price increases that are succeeded by a contraction in economic activity.

Table 7: Are oil price increases followed by economic contractions

$\frac{\text{Oil price}_t}{\text{Oil price}_{t-1}}$	1.35	1.50	1.75
Share of contractions in $t + 1$	0.61	0.66	0.68
$\frac{i_t}{\mathbb{E}[i]}$	0.98	0.97	0.95

The first row in the table gives oil-price increases of different sizes, the second row presents the share of these increases that are followed by a contraction in economic activity, and the third row gives investments (i) at the time of the price increase relative to average investments.

The table shows that a contraction in economic activity is indeed more likely after large oil-price hikes. It also shows that even with large shocks, far from all price increases are followed by a contraction. Kilian and Vigfusson (2014) argues that five out of eight price increases are followed by a recession, which implies a share of 62.5%. The numbers presented in Table 7 are closely in line with that number. However, the numbers are substantially lower than the estimates in Hamilton (1983, 2011).

It is a striking result, though, that oil shocks are followed by economic contractions also in the model, even though these price hikes are clearly not the cause of the subsequent contraction.

The intuition for why a contraction is more likely after an oil shock in the model is straightforward. A sharp increase in the oil price is a signal of future (and current) oil scarcity. Larger oil shocks, thus, signal higher scarcity. The direct effect of future oil scarcity is to reduce investments as is shown in the third row of Table 7. Investments are between 2-5 percent lower relative to mean investments during episodes of sharp oil-price increases, and they decrease with the size of the oil shock.⁴⁸ Lower investments have a direct negative effect on future labor supply and output. With a high probability, the news shock then also translates into an actual drop in ground oil with adverse effects on output and hours worked as a result.

Hamilton (2011) argues that, “[T]he correlation between oil shocks and economic recessions appears to be too strong to be just a coincidence”.⁴⁹ This is true in the model

⁴⁷One could easily make the measure of contraction more strict by requiring that output and hours worked decrease by a certain amount. For instance, if we require the reduction in output to be at least one percent, then this reduces the numbers in the second row of Table 7 somewhat for the two “smaller” shocks. The effect associated with the largest shock, however, is not affected.

⁴⁸The decrease in investments due to oil scarcity follows from the assumption that capital/labor and oil are gross complements in production.

⁴⁹The quote is found on page 26.

as well, but here the direction of causality runs from expectations of a contraction - due to oil scarcity - to the oil price, and not from the oil shock to the contraction.

6 How costly are the oil shocks?

Because the model is successful in matching the data and also can explain the volatility of the oil price, it can be used to address the question raised after the oil shocks in the 1970s, i.e., to what extent the volatility of the oil price hurt the U.S. economy. Following Lucas (1987), the welfare costs associated with volatile oil prices are straightforwardly quantified by computing how much the representative agent would be willing to pay to get rid of the fluctuations in the oil price. There are, however, two distinct sources for the volatility of the oil price: the input-saving technology shocks and variations in ground oil.

I therefore, first, quantify the welfare costs associated with fluctuations in ground oil, i.e., by computing the amount that the representative agent would be willing to pay to get rid of the fluctuations in ground oil as well as the news shocks. This number is found to amount to less than one-hundredth of a percent of consumption in the model. Second, I compute how much the agent would consider paying to get rid of all fluctuations in the model and find a number that is only marginally higher.

These results are maybe not that surprising. Lucas (1987) showed that the welfare costs associated with short-run fluctuations are negligible in a standard RBC model. The model in this paper is not too different from that in Lucas (1987), thus implying that the result may be expected to be similar. The welfare effects could, potentially, be larger if households' direct oil consumption were incorporated into the model.⁵⁰ This exercise is left for future research.

7 Sensitivity analysis

As all economic models, the model here is a stylized version of the real world and it incorporates several simplifying assumptions. This discussion focuses on some of the key assumptions and the potential implications of relaxing them.

⁵⁰Hitzemann and Yaron consider oil consumption by households and find somewhat larger, but still modest, costs of oil shocks.

7.1 The link between news and supply shocks

In the model, supply shocks are preceded by news shocks that signal potential short-falls. Naturally, this assumption matters for the relative importance of supply versus news shocks. If supply shocks are not predictable, there is no room for oil shocks to originate from increasing precautionary demand for oil. The assumption that supply shocks are, at least, somewhat predictable seems reasonable as a first approximation. Kilian (2008) computes the exogenous short-falls associated with the Iranian Revolution of 1978/79, the Iran-Iraq war of 1980-1988, the Gulf War of 1990/91, the Iraq War of 2002/03 and the civil unrest in Venezuela in 2002/03. In all these cases, noisy information indicating that what actually happened could happen were available prior to the specific events. The drop in oil supply associated with the Iranian Revolution was, for instance, initiated in October 1978 (when striking workers paralyzed oil installations across the country). Reports about demonstrations and civil unrest, however, started to show up already one year before. This conflict then intensified further in January 1978, i.e., eight months prior to the initial drop in oil supply. Similar conditions apply to the other events referred to in Kilian (2008a). The Iran-Iraq War, for example, followed a long history of border disputes. Iraq's invasion of Kuwait were preceded by disputes over the financial debt, as well as by accusations of Kuwait stealing oil from Iraq through slant drilling.

One way to evaluate whether the benchmark specification where supply shocks are preceded by news shocks imposes too little risk in the model is to compare future prices for oil in the model and in the data. One period is one year in the model, so I compare future prices at a 12-month horizon.⁵¹ Annual data on oil futures at the 12-month horizon that were traded on the New York Merchantile Exchange (NYMEX) over the period 1983-2015 is available from the U.S. Energy Information Administration. The standard deviation of these HP-filtered prices is 13.81, which is just slightly below the standard deviation for the spot price over the same period (14.80). In the model, the standard deviation of the 12-month future price is 11.70, i.e., roughly the same as for the spot price. The fact that the volatility of future prices is not significantly lower than spot prices in the model suggests that supply shocks do not impose too little risk in the model relative to the data.

There might still be a worry that news shocks reduce the potential contribution of supply shocks too much. I therefore allow for the possibility of supply shocks that are

⁵¹Ahlquist and Kilian (2010) shows that trading volumes fall sharply at maturities longer than one year, which makes longer maturities less reliable.

not preceded by any signal in Appendix G.3. This alternative specification features a five-percent possibility of a supply short-fall when news is good.

As shown in Table 13, unexpected supply shocks do not reduce the model's ability to match the observed business-cycle statistics. They do, however, change the relative importance of news and supply shocks as shown in Table 14. The median increase is still larger for news than supply shocks, but the difference is now substantially smaller. In addition, maximum price increases are larger for both types of shocks, but the maximum increase associated with a drop in supply is now larger than that of a news shock. In fact, the maximum oil-price hike from a short-fall in supply is now slightly larger than the oil-price shock in 1973/74.

The conclusion from this exercise is that it matters whether supply shocks are predictable or not. If shortfalls in supply are preceded by any information (that can be noisy), oil shocks tend to coincide with increasing precautionary/smoothing demand for oil but completely unexpected supply shocks can generate substantial oil-price shocks.

7.2 The parameters

The explanation of oil-shocks that relies on potential oil scarcity relies on limited possibilities to substitute oil for other inputs. If oil easily could be substituted, then there would be no need to hold precautionary quantities of oil. Instead, firms would just use the alternatives when oil becomes expensive and/or scarce. The value assigned to the elasticity of substitution between capital/labor and energy is based on empirical estimates that find the short-run elasticity to be low. A relevant question, however, is how the results depend on the value of ε . This is answered in Appendix G.1, where results are presented for different values of ε . These results show that the model gets closer to the data as ε is reduced. The results also show that an elasticity above 0.50 is at odds with the data, but that an elasticity equal to 0.10 generates results that are relatively similar to those in the benchmark economy. A value of 0.10 is also able to generate substantial oil-price shocks. Hence, the results go through as long as the short-run elasticity of substitution is not too high.

Section G.2 presents data for the lower discount factor $\beta = 0.96$. As can be seen, the results from that exercise are similar to the benchmark calibration.

8 Conclusions

This paper analyzes the interaction between oil prices and macroeconomic outcomes by setting up a general equilibrium model where output is produced with a capital-labor composite, and oil and all factor prices are endogenous. The basic framework is a modified real-business-cycle model that takes oil as an input and allows for storage of oil above ground between periods. The driving forces are factor-specific technology shocks and supply shocks that can be imprecisely forecasted by noisy news shock.

The results, first, show that the model closely matches a large number of observed business-cycle properties. These findings are far from obvious, given that not any of these properties are targeted in the calibration. The model economy also goes through episodes with substantial oil-price shocks. The distribution of oil prices is similar in the model and the data, implying that the model can quantitatively match all the observed oil shocks since 1900 with the exception of the episode in 1973/74.

The median annual increase in the oil price after a bad news shock is 15 times larger than after a direct supply short-fall in the benchmark calibration. This difference is then reduced to a factor of two when news shocks only increase uncertainty. These results are consistent with the findings in Kilian (2008a, 2009) that precisely shows that fears of future short-falls are important for understanding oil-price shocks.

Direct supply shocks, however, are also quantitatively important for understanding changes in the oil price, but their main influence is more indirect and goes through the effect on oil reserves. These reserves vary over time due to shocks to demand and, in particular, to shocks to the supply. Specifically, a series of relatively close supply shocks will lead to a run-down of reserves and increasingly larger oil shocks each time that bad news hits the economy. The model's prediction of a negative relationship between annual oil-price changes and the stock of oil reserves is tested empirically and found to be statistically significant during the period 1972-2014.

In the model as in the data, a contraction in economic activity is more likely than an expansion after an oil-price shock. Finally, the welfare costs of volatile oil prices, however, are found to be negligible.

References

- [1] Ahlquist, Ron, and Lutz Kilian (2010), "What Do we Learn from the Price of Crude Oil Futures", *Journal of Applied Econometrics* 25, 539-573.

- [2] Arezki, Rabah, Valery Ramey, and Liugang Sheng (2016), “News Shocks in Open Economies: Evidence from Giant Oil Discoveries”, mimeo.
- [3] Backus, David, and Mario Crucini (1998). “Oil Prices and the Terms of Trade”, *Journal of International Economics* 50, 185-213.
- [4] Barsky, Robert, and Lutz Kilian (2002), “Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative” In *NBER Macroeconomics Annual 2001*. Volume 16, ed. Ben S. Bernanke and Kenneth Rogoff, 137–83. Cambridge, MA: MIT Press.
- [5] Baumeister, Christiane, and Lutz Kilian (2016) “Understanding the Decline in the Price of Oil since June 2014”, *Journal of the Association of Environmental and Resource Economists*, 3, 131-158.
- [6] Berndt, Ernst and David Wood (1975), “Technology, Prices, and the Derived Demand for Energy”, *The Review of Economic Studies*, 57, 259-268.
- [7] Bruno, Michael, and Jeffrey Sachs (1982), “Input Price Shocks and the Slowdown in Economic Growth: The Case of U.K. Manufacturing”, *Review of Economic Studies*, 49, 679–705.
- [8] Bodenstein, Martin, Christopher J. Erceg, and Luca Guerrieri (2011), “Oil Shocks and Eexternal Adjustment”, *Journal of International Economics*, 83(2), 168–184.
- [9] Data on the ending stock of oil is from the Energy Information Administration, available at https://www.eia.gov/dnav/pet/pet_sum_snd_d_nus_mbbldpd_m_cur.htm.
- [10] Finn, Mary (2000). “Perfect Competition and the Effects of Energy Price Increases on Economic Activity”, *Journal of Money, Credit, and Banking*, 32 (3), 400-416.
- [11] Gars, Johan, and Conny Olovsson (2015) “Fuel for Economic Growth?”, Sveriges Riksbank, WP No. 299.
- [12] Griffin, James, and Paul Gregory (1976), “An Intercountry Translog Model of EnergySubstitution Responses”, *American Economic Review* 66(5), 845-857.
- [13] Hamilton, James D. (1983), “Oil and the Macroeconomy since World War II.”, *Journal of Political Economy*, 91 (2), 228-48.

- [14] Hamilton, James D. (2003), “What Is an Oil Shock”, *Journal of Econometrics*, 113, 363-98.
- [15] Hamilton, James D. (2009), “Causes and Consequences of the Oil Shock of 2007-2008”, NBER working paper w15002.
- [16] Hamilton, James D. (2011), “Historical Oil Shocks”, NBER working paper w16790.
- [17] Hart, Rob, and Daniel Spiro (2011), “The Elephant in Hotelling’s Room”, *Energy Policy*, Vol. 39 (12), 7834–7838.
- [18] Hassler, John, Per Krusell, and Conny Olovsson (2012) “Energy-Saving Technichal Change”, NBER working paper w18456.
- [19] Hitzemann, Stefen, and Amir Yaron (2016) “Welfare Costs of Oil Shocks”, mimeo.
- [20] Hudson, Edward, and Dale Jorgenson (1974), “U.S. Energy Policy and Economic Growth, 1975-2000”, *Bell Journal of Economics* 5, 461-514.
- [21] Jones, Donald W., Paul N. Leiby, and Inja K. Paik (2004), “Oil Price Shocks and the Macroeconomy: What has been Learned Since 1996”, *Energy Journal*, 25(2), 1–32.
- [22] King, Robert, and Sergio Rebelo (1999), “Resuscitating Real Business Cycles”, *Handbook of Macroeconomics*, 927-1007
- [23] Kilian, Lutz (2008a), “Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?”, *Review of Economics and Statistics*, 90, 216-240.
- [24] Kilian, Lutz (2008b), “The Economic Effects of Energy Price Shocks”, *Journal of Economic Literature*, 46, 871-909.
- [25] Kilian, Lutz. (2009), “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market”, *American Economic Review*, 99(3), 1053-69.
- [26] Kilian, Lutz, and Daniel Murphy (2014), “The Role of Inventories and Speculative Trading in the Global Market for Crude Oil?”, *Journal of Applied Econometrics*, 29, 454-478.

- [27] Kilian, Lutz and Robert J. Vigfusson (2014), “The Role of Oil Price Shocks in Causing U.S. Recessions”, Board of Governors of the Federal Reserve System, International Finance Discussion Papers, Nr. 1114.
- [28] Kim, In-Moo & Loungani, Prakash (1992), “The Role of Energy in Real Business Cycle Models”, *Journal of Monetary Economics*, 29(2), 173-189.
- [29] Koetsea, Mark J., Henri L.F. de Groota, and Raymond J.G.M. Florax (2008), “Capital-Energy Substitution and Shifts in Factor Demand: A Meta-Analysis”, *Energy Economics*, 30(5), 2236–2251.
- [30] Lucas, Robert E. (1987) *Models of business cycles*, Oxford: Basil Blackwell.
- [31] Pierce, James L., and Jared J. Enzler (1974), “The Effects of External Inflationary Shocks”, *Brookings Papers on Economic Activity*, 1, 13–61.
- [32] Rasche, Robert H., and John Tatom (1977), “The Effects of the New Energy Regime on Economic Capacity, Production, and Prices”, *Federal Reserve Bank of St. Louis Review*, 2–12.
- [33] Ravn Morten, and Harald Uhlig (2002), “On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations”, *Review of Economics and Statistics*, 84(2), 371–375.
- [34] Rotemberg, Julio, and Michael Woodford (1996), “Imperfect Competition and the Effects of Energy Price Increases on Economic Activity”, *Journal of Money, Credit, and Banking*, 28(4), 549-577.
- [35] Short-term energy outlook (2016), U.S. Energy Information Administration.
- [36] Tauchen George, and Robert Hussey (1991), “Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models”, *Econometrica*, 59(2), 371-396.

A Appendix

B Annual price changes versus three-year net changes

It is common in the literature on oil-price shocks to look at three-year net increases in the oil price. As shown in Figure 5 annual and three-year net increases are highly correlated.




Figure 5: Annual and three-period oil-price changes over the period 1949-2014.

C First order conditions to the social planning problem

The first order conditions to (13) deliver an Euler equation for capital, an equation that governs the return to storing oil and an intra-temporal equation for labor supply. For notational simplicity, I here omit the fact that all variables are functions of the state variables. The equations are then, respectively, given by

$$\frac{1}{c} = \beta \mathbb{E}_t \frac{1}{c'} (F_{k'} + 1 - \delta), \quad (19)$$

$$\frac{1}{c} = \beta \mathbb{E}_t \frac{1}{c'} \frac{F_{o'} - \kappa \frac{\partial R''}{\partial R'}}{F_o - \kappa} + \lambda, \quad (20)$$

and

$$\frac{F_l}{c} = \frac{1}{1-l} \frac{\mu}{1-\mu}, \quad (21)$$

where F_k , F_l and F_e respectively refers to the derivative of the production function with respect to k , l , and o , and we used the fact that $R'' = R' - o'_s$.

In steady-state, the values of all shocks equal zero. There is then no need for storing oil. In fact, the amount of stored oil in steady state has to be exactly zero so R' and o_s both equals zero in steady state. The total supply of oil is then equal to φ . The additional equations that characterize the steady state are given by

$$\begin{aligned} 1 &= \beta (F_k + 1 - \delta); \\ 1 &= \beta \frac{F_o + \lambda}{F_o}; \\ c &= F_l (1-l) \frac{1-\mu}{\mu}; \\ c &= y - \delta k. \end{aligned}$$

These four equations solve uniquely for the steady state levels of c , k , l and λ .

D Hotelling pricing with a storage technology?

To derive some interpretable expressions, I here abstract from uncertainty. In addition, the expressions in this section are derived for a two-period model to enhance transparency, but the results generalize straightforward to an infinite horizon.

Combining (19) and (20) in Appendix C delivers the following modified Hotelling formula

$$\frac{p'_o}{p_o} = r' - \lambda \frac{c'}{\beta F_o}. \quad (22)$$

The above equation reveals that the oil price should grow at the rate of the interest minus a correction term that depends on the multiplier on (11). The standard Hotelling formula, thus, only holds when λ is exactly zero. If the multiplier instead is positive, the

oil price should grow at a rate that is lower the interest rate. Price growth can, in fact, be negative in this setting.

When will λ be binding, however? It is not possible to give a general answer to that question, but if we are willing to assume that capital depreciates fully between periods, and that the elasticity of substitution between capital/labor and oil equals one, then the optimal depletion rate for stored oil can be shown to be governed by the following expression

$$\frac{o_{g,1} + R_0 - o_{s,0}}{o_{g,0} + o_{s,0}} - \frac{\lambda(\eta_1 + R_0 - o_{s,0})}{(1 - \alpha)(1 + \alpha\beta)} = \frac{\beta}{1 + \alpha\beta}. \quad (23)$$

Equation (23) shows that an interior solution (with $\lambda = 0$) requires total oil use to fall at the rate $\frac{\beta}{1 + \alpha\beta} < 1$ between the two periods.⁵² In this case, the standard Hotelling formula applies. Note, however, that an interior solution only is possible when $o_{g,0} + R_0$ is large relative to $o_{g,1}$. If, instead, $o_{g,1}$ is large relative to $o_{g,0} + R_0$ the planner would like to borrow from the future endowment. This is not allowed, so the solution is to set $o_{s,0} = R_0$, which implies $\lambda > 0$ and that the Hotelling formula does not hold.

Also in the special case when ground oil is constant between periods (which is close to the assumption in this paper) is it possible for total oil consumption to fall between periods if R_0 is large relative to the endowments. Otherwise, the multiplier on the constraint that stored oil cannot be negative will again be binding and the growth rate for the oil price will be strictly lower than the interest rate.

In a setting with an infinite horizon, the Hotelling formula can hold for a finite number of periods if R_0 is large relative to the endowments. It is, however, straightforward to verify that without uncertainty, the stock of stored oil will be monotonically depleted over time. As R becomes smaller and smaller, the stock of stored oil will at some point reach a threshold level after which total oil use will fall at rate larger than β . From this point on, the oil price will grow at a rate that is strictly below the interest rate.

Note, finally, that it is not necessary that o_g is interpreted as oil for the Hotelling formula to break down. This will also be true in a setting where oil is given by o_s , whereas o_g is a substitute to oil such as nuclear power. The existence of an exogenous amount of oil, or an exogenous substitute, thus, leads to a breakdown of the classical Hotelling relation.

⁵²This rate equals β with an infinite horizon.

E Estimating the factor-specific technology shocks

Figure 6 shows the evolution of the technology trends A and A_o alongside with the fitted regression lines for the four periods. Consistent with the findings in Hassler, Krusell and Olovsson (2012), the technology trends seems to be subject to directed technical change: when A is growing relatively fast, A_o grows relatively slowly and vice versa. Completely ignoring to take changing trend growth into account results in volatilities for a and a_o that respectively are more than two and five times as large as output. With that specification, a_o has not been below trend since 1982, which also indicates that a straight line is a bad approximation for the whole period.

As explained in Section 2.4, I estimate the relevant parameters for the four sub-periods plotted in Figure 6 to remove the variation in the trend components. The results from this estimation was presented in Table 3. Four periods is motivated by the fact that this reduces the volatility of the a -shock down to typical estimates of total-factor-productivity shocks.

An alternatively strategy could have been to instead average over three periods: before, during and after the oil shocks of the 1970s (which simply implies merging period three and four in Figure 6). The estimated parameter are then given by $\rho_z = 0.56$, $\rho_{z_o} = 0.60$, $\epsilon_z = 0.0184$ and $\epsilon_{z_o} = 0.0256$, i.e., slightly different values than in the benchmark calibration. The results from averaging over three periods are presented in Table 8, along with the result from the benchmark calibration in the paper. As can be seen, the results are basically unchanged, which is interesting due to the fact that the parameters ρ_z , ρ_{z_o} , ϵ_z and ϵ_{z_o} differ somewhat across the specifications.

Table 8: Business cycle statistics in the data and the model

	$std(y)$	$std(o)$	$std(l)$	$std(p_o)$	$std(R)$
Model based on 4-period averages	1.68	2.10	1.05	11.75	1.46
Model based on 3-period averages	1.68	2.04	1.02	12.13	1.49
	$corr(y, l)$	$corr(y, o)$	$corr(y, \Delta R)$	$corr(y, p_o)$	$corr(p_o, \Delta R)$
Model based on 4-period averages	0.91	0.72	-0.09	0.17	0.06
Model based on 3-period averages	0.88	0.73	-0.05	0.28	0.05

All variables are in logarithms and have been detrended with the HP filter. ΔR is the net change in the reserve stock of oil.




Figure 6: The evolution of the technology trends A and A_o alongside with the fitted regression lines for the four periods

Table 9: Business cycle statistics in the data and the mode when news signals a mean preserving spread

	$std(y)$	$std(o)$	$std(l)$	$std(p_o)$	$std(R)$
Data	1.63	2.10	1.22	14.53	2.41
Model 1	1.70	2.20	1.07	12.75	3.03
	$corr(y, l)$	$corr(y, o)$	$corr(y, \Delta R)$	$corr(y, p_o)$	$corr(p_o, \Delta R)$
Data	0.88	0.78	-0.15	0.01	0.01
Model 1	0.91	0.72	-0.09	0.17	0.06

All variables are in logarithms and have been detrended with the HP filter. The upper half of the table provides standard deviations and the lower half presents correlations.

F Results for a mean-preserving spread

G Sensitivity analysis

G.1 A higher elasticity of substitution

Table 10 present results for different values of the elasticity of substitution between capital/labor and energy ε .

Basically all variables are brought closer to the data as ε is decreased. Table 11 display how median and maximal price increases from bad news depend on ε .

Table 10: Business cycle properties for different elasticities

	$std(y)$	$std(o)$	$std(l)$	$std(p_o)$	$std(R)$
Data	1.63	2.10	1.22	14.53	2.13
$\varepsilon = 0.75$	1.71	1.79	0.99	2.98	5.83
$\varepsilon = 0.5$	1.71	1.76	0.98	3.95	1.57
$\varepsilon = 0.1$	1.69	1.92	0.97	9.15	1.35
$\varepsilon = 0.05$	1.69	2.04	1.02	10.94	1.42
	$corr(y, l)$	$corr(y, o)$	$corr(y, \Delta R)$	$corr(y, p_o)$	$corr(p_o, \Delta R)$
Data	0.88	0.78	-0.15	0.01	0.01
$\varepsilon = 0.75$	0.99	0.22	-0.2	0.60	0.17
$\varepsilon = 0.5$	0.99	0.41	-0.17	0.51	0.15
$\varepsilon = 0.1$	0.96	0.72	-0.11	0.28	0.08
$\varepsilon = 0.05$	0.93	0.73	-0.10	0.21	0.07

All variables are in logarithms and have been detrended with the HP filter. The upper half of the table provides standard deviations and the lower half present correlations. ΔR is the net change in the reserve stock of oil.

Table 11: Median and maximum oil-price changes after a news shock for different elasticities of substitution (in percent)

	Median	Max
$\varepsilon = 0.75$	5.58	12.18
$\varepsilon = 0.5$	8.07	18.41
$\varepsilon = 0.1$	15.83	74.47
$\varepsilon = 0.05$	17.89	112.08

ε is the elasticity of substitution between capital/labor and energy.

G.2 A lower discount factor

Table 12 present results for a lower discount factor.

G.3 Unexpected supply shocks

This section allows for unexpected supply shocks, i.e., shocks that may hit without and warning from a news shock. The benchmark setting is changed so that the probability of a supply shock is five percent instead of zero. The results are presented in Table 13-14.

The results with unexpected supply shocks are close to identical to those without such shocks. The main difference is that reserves are more volatile.

Table 12: Business cycle properties with a discount factor equal to 0.96

	$std(y)$	$std(o)$	$std(l)$	$std(p_o)$	$std(R)$
Data	1.63	2.10	1.22	14.53	2.13
Model	1.70	2.05	1.18	15.91	1.28
	$corr(y, l)$	$corr(y, o)$	$corr(y, \Delta R)$	$corr(y, p_o)$	$corr(p_o, \Delta R)$
Data	0.88	0.78	-0.15	0.01	0.01
Model	0.88	0.73	-0.10	0.13	0.06

All variables are in logarithms and have been detrended with the HP filter. The upper half of the table provides standard deviations and the lower half present correlations. ΔR is the net change in the reserve stock of oil.

Table 13: Business cycle properties with unexpected supply shocks

	$std(y)$	$std(o)$	$std(l)$	$std(p_o)$	$std(R)$
Data	1.63	2.10	1.22	14.53	2.13
Model	1.71	2.12	1.11	11.82	1.58
	$corr(y, l)$	$corr(y, o)$	$corr(y, \Delta R)$	$corr(y, p_o)$	$corr(p_o, \Delta R)$
Data	0.88	0.78	-0.15	0.01	0.01
Model	0.90	0.73	-0.03	0.13	-0.12

All variables are in logarithms and have been detrended with the HP filter. The upper half of the table provides standard deviations and the lower half present correlations. ΔR is the net change in the reserve stock of oil.

G.4 Elastic ground oil

In this section, I analyze to what extent the assumption that ground oil is inelastic matters for the results. It is now, instead, assumed that ground oil in the good state can be increased at a cost. Supply in the bad state is still inelastic however. Specifically, the amount of ground oil can now be increased by setting the variable \varkappa to a value larger than zero:

$$o_g = \begin{cases} \varphi e^{\eta_l} (1 + \varkappa) & \text{if } \eta = \eta_l \\ \varphi e^{\eta_h} & \text{if } \eta = \eta_h. \end{cases} \quad (24)$$

The cost for increasing ground oil is increasing and convex:

$$\Pi(\varkappa) = \frac{\omega_1}{\bar{\varkappa} - \varkappa} - \omega_2,$$

where $\bar{\varkappa}$ is the capacity constraint.

The first order condition with respect to \varkappa becomes

$$\varkappa = \bar{\varkappa} - \sqrt{\omega_1} (y F_o \varphi e^{\eta_l})^{-1/2}.$$

Table 14: Changes in oil prices between two consecutive years with unexpected supply shocks (in percent)

	Bad news	Low supply
Median change in p_o	14.37	8.97
Maximum change in p_o	185.04	229.42

The first column presents the changes that are realized at the time when bad news hit the economy. The second column presents the changes that are realized at the time when a low supply of ground oil hits the economy.

The parameters ω_1 and ω_2 are calibrated so that $\varkappa_{ss} = 0$, and that $\Pi(\varkappa_{ss}) = 0$. This implies that if it was possible to increase the production of ground oil at the steady state, the social planner would not choose to do that. The cost of extraction is then also zero at the steady state. The parameter $\bar{\varkappa}$, finally, denotes the capacity constraint. It is set so ground oil, at most, can be increased by ten percent. This number is based on estimates in Kilian (2008a) that argues that the capacity utilization rate for world production of oil typically is running at 90-95 percent.

$$\omega_1 = \bar{\varkappa}^2 y_{ss} \frac{\gamma \varphi^{\frac{\sigma-1}{\sigma}}}{(1-\gamma)(k_{ss}^\alpha h_{ss}^{1-\alpha})^{\frac{\sigma-1}{\sigma}} + \gamma \varphi^{\frac{\sigma-1}{\sigma}}}, \quad \text{and} \quad \omega_2 = \frac{\omega_1}{\bar{\varkappa}}.$$

The results are presented in Table 15.

Table 15: Business cycle properties with elastic oil supply

	$std(y)$	$std(o)$	$std(l)$	$std(p_o)$	$std(R)$
Inelastic oil	1.68	2.08	1.05	11.75	1.46
Elastic oil	1.69	2.23	1.00	11.63	1.50
	$corr(y, l)$	$corr(y, o)$	$corr(y, \Delta R)$	$corr(y, p_o)$	$corr(p_o, \Delta R)$
Inelastic oil	0.91	0.72	-0.09	0.17	0.01
Elastic oil	0.94	0.67	-0.08	0.23	0.12

All variables are in logarithms and have been detrended with the HP filter. The upper half of the table provides standard deviations and the lower half present correlations. ΔR is the net change in the reserve stock of oil.

H Variance decomposition: the demand shocks

Table 16 presents the individual contributions to volatilities and correlations from the shocks a and a_o .

Table 16: Business cycle statistics with individual demand shocks

	$std(y)$	$std(o)$	$std(l)$	$std(p_o)$	$std(R)$
Only shocks to a	0.0177	0.0172	0.0101	0.0209	0.0086
Only shocks to a_o	0.0033	0.0109	0.0049	0.0858	0.0057
	$corr(y, l)$	$corr(y, o)$	$corr(y, \Delta R)$	$corr(y, p_o)$	$corr(p_o, \Delta R)$
Only shocks to a	0.99	1.00	-0.06	0.87	-0.36
Only shocks to a_o	0.99	-0.73	0.32	-0.99	-0.31

All variables are in logarithms and have been detrended with the HP filter. ΔR is the net change in the reserve stock of oil.

I Regression results

To make the coefficients comparable across definitions, I also normalize effective reserves by the mean over the sample period in each case, which only changes the interpretation of the coefficients. The results are displayed in table 17.

Table 17: Regression results

	1949-2014		1972-2014	
	SPR included	SPR excluded	SPR included	SPR excluded
<i>Effective_reserves</i>	-14.50 (10.98)	-27.59* (16.02)	-37.81** (18.39)	-88.99** (42.87)
<i>Constant</i>	21.27* (11.81)	34.36** (16.59)	48.11** (19.42)	99.3** (43.32)
<i>N</i>	63	63	42	42
<i>R</i> ²	0.0278	0.0464	0.0956	0.0973

The standard errors are provided in parentheses; * significant at 10%; ** significant at 5%; *** significant at 1%.

The results for the period 1972-2014 are similar if effective reserves are not normalized by oil consumption. The reason is that oil-consumption growth was limited during this period.

Earlier Working Papers:

For a complete list of Working Papers published by Sveriges Riksbank, see www.riksbank.se

Estimation of an Adaptive Stock Market Model with Heterogeneous Agents <i>by Henrik Amilon</i>	2005:177
Some Further Evidence on Interest-Rate Smoothing: The Role of Measurement Errors in the Output Gap <i>by Mikael Apel and Per Jansson</i>	2005:178
Bayesian Estimation of an Open Economy DSGE Model with Incomplete Pass-Through <i>by Malin Adolphson, Stefan Laséen, Jesper Lindé and Mattias Villani</i>	2005:179
Are Constant Interest Rate Forecasts Modest Interventions? Evidence from an Estimated Open Economy DSGE Model of the Euro Area <i>by Malin Adolphson, Stefan Laséen, Jesper Lindé and Mattias Villani</i>	2005:180
Inference in Vector Autoregressive Models with an Informative Prior on the Steady State <i>by Mattias Villani</i>	2005:181
Bank Mergers, Competition and Liquidity <i>by Elena Carletti, Philipp Hartmann and Giancarlo Spagnolo</i>	2005:182
Testing Near-Rationality using Detailed Survey Data <i>by Michael F. Bryan and Stefan Palmqvist</i>	2005:183
Exploring Interactions between Real Activity and the Financial Stance <i>by Tor Jacobson, Jesper Lindé and Kasper Roszbach</i>	2005:184
Two-Sided Network Effects, Bank Interchange Fees, and the Allocation of Fixed Costs <i>by Mats A. Bergman</i>	2005:185
Trade Deficits in the Baltic States: How Long Will the Party Last? <i>by Rudolfs Bems and Kristian Jönsson</i>	2005:186
Real Exchange Rate and Consumption Fluctuations following Trade Liberalization <i>by Kristian Jönsson</i>	2005:187
Modern Forecasting Models in Action: Improving Macroeconomic Analyses at Central Banks <i>by Malin Adolphson, Michael K. Andersson, Jesper Lindé, Mattias Villani and Anders Vredin</i>	2005:188
Bayesian Inference of General Linear Restrictions on the Cointegration Space <i>by Mattias Villani</i>	2005:189
Forecasting Performance of an Open Economy Dynamic Stochastic General Equilibrium Model <i>by Malin Adolphson, Stefan Laséen, Jesper Lindé and Mattias Villani</i>	2005:190
Forecast Combination and Model Averaging using Predictive Measures <i>by Jana Eklund and Sune Karlsson</i>	2005:191
Swedish Intervention and the Krona Float, 1993-2002 <i>by Owen F. Humpage and Javiera Ragnartz</i>	2006:192
A Simultaneous Model of the Swedish Krona, the US Dollar and the Euro <i>by Hans Lindblad and Peter Sellin</i>	2006:193
Testing Theories of Job Creation: Does Supply Create Its Own Demand? <i>by Mikael Carlsson, Stefan Eriksson and Nils Gottfries</i>	2006:194
Down or Out: Assessing The Welfare Costs of Household Investment Mistakes <i>by Laurent E. Calvet, John Y. Campbell and Paolo Sodini</i>	2006:195
Efficient Bayesian Inference for Multiple Change-Point and Mixture Innovation Models <i>by Paolo Giordani and Robert Kohn</i>	2006:196
Derivation and Estimation of a New Keynesian Phillips Curve in a Small Open Economy <i>by Karolina Holmberg</i>	2006:197
Technology Shocks and the Labour-Input Response: Evidence from Firm-Level Data <i>by Mikael Carlsson and Jon Smedsaas</i>	2006:198
Monetary Policy and Staggered Wage Bargaining when Prices are Sticky <i>by Mikael Carlsson and Andreas Westermark</i>	2006:199
The Swedish External Position and the Krona <i>by Philip R. Lane</i>	2006:200

Price Setting Transactions and the Role of Denominating Currency in FX Markets <i>by Richard Friberg and Fredrik Wilander</i>	2007:201
The geography of asset holdings: Evidence from Sweden <i>by Nicolas Coeurdacier and Philippe Martin</i>	2007:202
Evaluating An Estimated New Keynesian Small Open Economy Model <i>by Malin Adolfson, Stefan Laséen, Jesper Lindé and Mattias Villani</i>	2007:203
The Use of Cash and the Size of the Shadow Economy in Sweden <i>by Gabriela Guibourg and Björn Segendorf</i>	2007:204
Bank supervision Russian style: Evidence of conflicts between micro- and macro-prudential concerns <i>by Sophie Claeys and Koen Schoors</i>	2007:205
Optimal Monetary Policy under Downward Nominal Wage Rigidity <i>by Mikael Carlsson and Andreas Westermarck</i>	2007:206
Financial Structure, Managerial Compensation and Monitoring <i>by Vittoria Cerasi and Sonja Daltung</i>	2007:207
Financial Frictions, Investment and Tobin's q <i>by Guido Lorenzoni and Karl Walentin</i>	2007:208
Sticky Information vs Sticky Prices: A Horse Race in a DSGE Framework <i>by Mathias Trabandt</i>	2007:209
Acquisition versus greenfield: The impact of the mode of foreign bank entry on information and bank lending rates <i>by Sophie Claeys and Christa Hainz</i>	2007:210
Nonparametric Regression Density Estimation Using Smoothly Varying Normal Mixtures <i>by Mattias Villani, Robert Kohn and Paolo Giordani</i>	2007:211
The Costs of Paying – Private and Social Costs of Cash and Card <i>by Mats Bergman, Gabriella Guibourg and Björn Segendorf</i>	2007:212
Using a New Open Economy Macroeconomics model to make real nominal exchange rate forecasts <i>by Peter Sellin</i>	2007:213
Introducing Financial Frictions and Unemployment into a Small Open Economy Model <i>by Lawrence J. Christiano, Mathias Trabandt and Karl Walentin</i>	2007:214
Earnings Inequality and the Equity Premium <i>by Karl Walentin</i>	2007:215
Bayesian forecast combination for VAR models <i>by Michael K. Andersson and Sune Karlsson</i>	2007:216
Do Central Banks React to House Prices? <i>by Daria Finocchiaro and Virginia Queijo von Heideken</i>	2007:217
The Riksbank's Forecasting Performance <i>by Michael K. Andersson, Gustav Karlsson and Josef Svensson</i>	2007:218
Macroeconomic Impact on Expected Default Frequency <i>by Per Åsberg and Hovick Shahnazarian</i>	2008:219
Monetary Policy Regimes and the Volatility of Long-Term Interest Rates <i>by Virginia Queijo von Heideken</i>	2008:220
Governing the Governors: A Clinical Study of Central Banks <i>by Lars Frisell, Kasper Roszbach and Giancarlo Spagnolo</i>	2008:221
The Monetary Policy Decision-Making Process and the Term Structure of Interest Rates <i>by Hans Dillén</i>	2008:222
How Important are Financial Frictions in the U S and the Euro Area <i>by Virginia Queijo von Heideken</i>	2008:223
Block Kalman filtering for large-scale DSGE models <i>by Ingvar Strid and Karl Walentin</i>	2008:224
Optimal Monetary Policy in an Operational Medium-Sized DSGE Model <i>by Malin Adolfson, Stefan Laséen, Jesper Lindé and Lars E. O. Svensson</i>	2008:225
Firm Default and Aggregate Fluctuations <i>by Tor Jacobson, Rikard Kindell, Jesper Lindé and Kasper Roszbach</i>	2008:226

Re-Evaluating Swedish Membership in EMU: Evidence from an Estimated Model <i>by Ulf Söderström</i>	2008:227
The Effect of Cash Flow on Investment: An Empirical Test of the Balance Sheet Channel <i>by Ola Melander</i>	2009:228
Expectation Driven Business Cycles with Limited Enforcement <i>by Karl Walentin</i>	2009:229
Effects of Organizational Change on Firm Productivity <i>by Christina Håkanson</i>	2009:230
Evaluating Microfoundations for Aggregate Price Rigidities: Evidence from Matched Firm-Level Data on Product Prices and Unit Labor Cost <i>by Mikael Carlsson and Oskar Nordström Skans</i>	2009:231
Monetary Policy Trade-Offs in an Estimated Open-Economy DSGE Model <i>by Malin Adolfson, Stefan Laséen, Jesper Lindé and Lars E. O. Svensson</i>	2009:232
Flexible Modeling of Conditional Distributions Using Smooth Mixtures of Asymmetric Student T Densities <i>by Feng Li, Mattias Villani and Robert Kohn</i>	2009:233
Forecasting Macroeconomic Time Series with Locally Adaptive Signal Extraction <i>by Paolo Giordani and Mattias Villani</i>	2009:234
Evaluating Monetary Policy <i>by Lars E. O. Svensson</i>	2009:235
Risk Premiums and Macroeconomic Dynamics in a Heterogeneous Agent Model <i>by Ferre De Graeve, Maarten Dossche, Marina Emiris, Henri Sneessens and Raf Wouters</i>	2010:236
Picking the Brains of MPC Members <i>by Mikael Apel, Carl Andreas Claussen and Petra Lennartsdotter</i>	2010:237
Involuntary Unemployment and the Business Cycle <i>by Lawrence J. Christiano, Mathias Trabandt and Karl Walentin</i>	2010:238
Housing collateral and the monetary transmission mechanism <i>by Karl Walentin and Peter Sellin</i>	2010:239
The Discursive Dilemma in Monetary Policy <i>by Carl Andreas Claussen and Øistein Røisland</i>	2010:240
Monetary Regime Change and Business Cycles <i>by Vasco Cúrdia and Daria Finocchiaro</i>	2010:241
Bayesian Inference in Structural Second-Price common Value Auctions <i>by Bertil Wegmann and Mattias Villani</i>	2010:242
Equilibrium asset prices and the wealth distribution with inattentive consumers <i>by Daria Finocchiaro</i>	2010:243
Identifying VARs through Heterogeneity: An Application to Bank Runs <i>by Ferre De Graeve and Alexei Karas</i>	2010:244
Modeling Conditional Densities Using Finite Smooth Mixtures <i>by Feng Li, Mattias Villani and Robert Kohn</i>	2010:245
The Output Gap, the Labor Wedge, and the Dynamic Behavior of Hours <i>by Luca Sala, Ulf Söderström and Antonella Trigari</i>	2010:246
Density-Conditional Forecasts in Dynamic Multivariate Models <i>by Michael K. Andersson, Stefan Palmqvist and Daniel F. Waggoner</i>	2010:247
Anticipated Alternative Policy-Rate Paths in Policy Simulations <i>by Stefan Laséen and Lars E. O. Svensson</i>	2010:248
MOSES: Model of Swedish Economic Studies <i>by Gunnar Bårdsen, Ard den Reijer, Patrik Jonasson and Ragnar Nymoén</i>	2011:249
The Effects of Endogenous Firm Exit on Business Cycle Dynamics and Optimal Fiscal Policy <i>by Lauri Vilmi</i>	2011:250
Parameter Identification in a Estimated New Keynesian Open Economy Model <i>by Malin Adolfson and Jesper Lindé</i>	2011:251
Up for count? Central bank words and financial stress <i>by Marianna Blix Grimaldi</i>	2011:252

Wage Adjustment and Productivity Shocks <i>by Mikael Carlsson, Julián Messina and Oskar Nordström Skans</i>	2011:253
Stylized (Arte) Facts on Sectoral Inflation <i>by Ferre De Graeve and Karl Walentin</i>	2011:254
Hedging Labor Income Risk <i>by Sebastien Betermier, Thomas Jansson, Christine A. Parlour and Johan Walden</i>	2011:255
Taking the Twists into Account: Predicting Firm Bankruptcy Risk with Splines of Financial Ratios <i>by Paolo Giordani, Tor Jacobson, Erik von Schedvin and Mattias Villani</i>	2011:256
Collateralization, Bank Loan Rates and Monitoring: Evidence from a Natural Experiment <i>by Geraldo Cerqueiro, Steven Ongena and Kasper Roszbach</i>	2012:257
On the Non-Exclusivity of Loan Contracts: An Empirical Investigation <i>by Hans Degryse, Vasso Ioannidou and Erik von Schedvin</i>	2012:258
Labor-Market Frictions and Optimal Inflation <i>by Mikael Carlsson and Andreas Westermarck</i>	2012:259
Output Gaps and Robust Monetary Policy Rules <i>by Roberto M. Billi</i>	2012:260
The Information Content of Central Bank Minutes <i>by Mikael Apel and Marianna Blix Grimaldi</i>	2012:261
The Cost of Consumer Payments in Sweden <i>by Björn Segendorf and Thomas Jansson</i>	2012:262
Trade Credit and the Propagation of Corporate Failure: An Empirical Analysis <i>by Tor Jacobson and Erik von Schedvin</i>	2012:263
Structural and Cyclical Forces in the Labor Market During the Great Recession: Cross-Country Evidence <i>by Luca Sala, Ulf Söderström and Antonella Trigari</i>	2012:264
Pension Wealth and Household Savings in Europe: Evidence from SHARELIFE <i>by Rob Alessie, Viola Angelini and Peter van Santen</i>	2013:265
Long-Term Relationship Bargaining <i>by Andreas Westermarck</i>	2013:266
Using Financial Markets To Estimate the Macro Effects of Monetary Policy: An Impact-Identified FAVAR* <i>by Stefan Pitschner</i>	2013:267
DYNAMIC MIXTURE-OF-EXPERTS MODELS FOR LONGITUDINAL AND DISCRETE-TIME SURVIVAL DATA <i>by Matias Quiroz and Mattias Villani</i>	2013:268
Conditional euro area sovereign default risk <i>by André Lucas, Bernd Schwaab and Xin Zhang</i>	2013:269
Nominal GDP Targeting and the Zero Lower Bound: Should We Abandon Inflation Targeting?*	2013:270
<i>by Roberto M. Billi</i>	
Un-truncating VARs* <i>by Ferre De Graeve and Andreas Westermarck</i>	2013:271
Housing Choices and Labor Income Risk <i>by Thomas Jansson</i>	2013:272
Identifying Fiscal Inflation* <i>by Ferre De Graeve and Virginia Queijo von Heideken</i>	2013:273
On the Redistributive Effects of Inflation: an International Perspective* <i>by Paola Boel</i>	2013:274
Business Cycle Implications of Mortgage Spreads* <i>by Karl Walentin</i>	2013:275
Approximate dynamic programming with post-decision states as a solution method for dynamic economic models <i>by Isaiah Hull</i>	2013:276
A detrimental feedback loop: deleveraging and adverse selection <i>by Christoph Bertsch</i>	2013:277
Distortionary Fiscal Policy and Monetary Policy Goals <i>by Klaus Adam and Roberto M. Billi</i>	2013:278
Predicting the Spread of Financial Innovations: An Epidemiological Approach <i>by Isaiah Hull</i>	2013:279

Firm-Level Evidence of Shifts in the Supply of Credit <i>by Karolina Holmberg</i>	2013:280
Lines of Credit and Investment: Firm-Level Evidence of Real Effects of the Financial Crisis <i>by Karolina Holmberg</i>	2013:281
A wake-up call: information contagion and strategic uncertainty <i>by Toni Ahnert and Christoph Bertsch</i>	2013:282
Debt Dynamics and Monetary Policy: A Note <i>by Stefan Laséen and Ingvar Strid</i>	2013:283
Optimal taxation with home production <i>by Conny Olovsson</i>	2014:284
Incompatible European Partners? Cultural Predispositions and Household Financial Behavior <i>by Michael Haliassos, Thomas Jansson and Yigitcan Karabulut</i>	2014:285
How Subprime Borrowers and Mortgage Brokers Shared the Piecial Behavior <i>by Antje Berndt, Burton Hollifield and Patrik Sandås</i>	2014:286
The Macro-Financial Implications of House Price-Indexed Mortgage Contracts <i>by Isaiah Hull</i>	2014:287
Does Trading Anonymously Enhance Liquidity? <i>by Patrick J. Dennis and Patrik Sandås</i>	2014:288
Systematic bailout guarantees and tacit coordination <i>by Christoph Bertsch, Claudio Calcagno and Mark Le Quement</i>	2014:289
Selection Effects in Producer-Price Setting <i>by Mikael Carlsson</i>	2014:290
Dynamic Demand Adjustment and Exchange Rate Volatility <i>by Vesna Corbo</i>	2014:291
Forward Guidance and Long Term Interest Rates: Inspecting the Mechanism <i>by Ferre De Graeve, Pelin Ilbas & Raf Wouters</i>	2014:292
Firm-Level Shocks and Labor Adjustments <i>by Mikael Carlsson, Julián Messina and Oskar Nordström Skans</i>	2014:293
A wake-up call theory of contagion <i>by Toni Ahnert and Christoph Bertsch</i>	2015:294
Risks in macroeconomic fundamentals and excess bond returns predictability <i>by Rafael B. De Rezende</i>	2015:295
The Importance of Reallocation for Productivity Growth: Evidence from European and US Banking <i>by Jaap W.B. Bos and Peter C. van Santen</i>	2015:296
SPEEDING UP MCMC BY EFFICIENT DATA SUBSAMPLING <i>by Matias Quiroz, Mattias Villani and Robert Kohn</i>	2015:297
Amortization Requirements and Household Indebtedness: An Application to Swedish-Style Mortgages <i>by Isaiah Hull</i>	2015:298
Fuel for Economic Growth? <i>by Johan Gars and Conny Olovsson</i>	2015:299
Searching for Information <i>by Jungsuk Han and Francesco Sangiorgi</i>	2015:300
What Broke First? Characterizing Sources of Structural Change Prior to the Great Recession <i>by Isaiah Hull</i>	2015:301
Price Level Targeting and Risk Management <i>by Roberto Billi</i>	2015:302
Central bank policy paths and market forward rates: A simple model <i>by Ferre De Graeve and Jens Iversen</i>	2015:303
Jump-Starting the Euro Area Recovery: Would a Rise in Core Fiscal Spending Help the Periphery? <i>by Olivier Blanchard, Christopher J. Erceg and Jesper Lindé</i>	2015:304
Bringing Financial Stability into Monetary Policy* <i>by Eric M. Leeper and James M. Nason</i>	2015:305

SCALABLE MCMC FOR LARGE DATA PROBLEMS USING DATA SUBSAMPLING AND THE DIFFERENCE ESTIMATOR <i>by MATIAS QUIROZ, MATTIAS VILLANI AND ROBERT KOHN</i>	2015:306
SPEEDING UP MCMC BY DELAYED ACCEPTANCE AND DATA SUBSAMPLING <i>by MATIAS QUIROZ</i>	2015:307
Modeling financial sector joint tail risk in the euro area <i>by André Lucas, Bernd Schwaab and Xin Zhang</i>	2015:308
Score Driven Exponentially Weighted Moving Averages and Value-at-Risk Forecasting <i>by André Lucas and Xin Zhang</i>	2015:309
On the Theoretical Efficacy of Quantitative Easing at the Zero Lower Bound <i>by Paola Boel and Christopher J. Waller</i>	2015:310
Optimal Inflation with Corporate Taxation and Financial Constraints <i>by Daria Finocchiaro, Giovanni Lombardo, Caterina Mendicino and Philippe Weil</i>	2015:311
Fire Sale Bank Recapitalizations <i>by Christoph Bertsch and Mike Mariathasan</i>	2015:312
Since you're so rich, you must be really smart: Talent and the Finance Wage Premium <i>by Michael Böhm, Daniel Metzger and Per Strömberg</i>	2015:313
Debt, equity and the equity price puzzle <i>by Daria Finocchiaro and Caterina Mendicino</i>	2015:314
Trade Credit: Contract-Level Evidence Contradicts Current Theories <i>by Tore Ellingsen, Tor Jacobson and Erik von Schedvin</i>	2016:315
Double Liability in a Branch Banking System: Historical Evidence from Canada <i>by Anna Grodecka and Antonis Kotidis</i>	2016:316
Subprime Borrowers, Securitization and the Transmission of Business Cycles <i>by Anna Grodecka</i>	2016:317
Real-Time Forecasting for Monetary Policy Analysis: The Case of Sveriges Riksbank <i>by Jens Iversen, Stefan Laséen, Henrik Lundvall and Ulf Söderström</i>	2016:318
Fed Liftoff and Subprime Loan Interest Rates: Evidence from the Peer-to-Peer Lending <i>by Christoph Bertsch, Isaiah Hull and Xin Zhang</i>	2016:319
Curbing Shocks to Corporate Liquidity: The Role of Trade Credit <i>by Niklas Amberg, Tor Jacobson, Erik von Schedvin and Robert Townsend</i>	2016:320
Firms' Strategic Choice of Loan Delinquencies <i>by Paola Morales-Acevedo</i>	2016:321
Fiscal Consolidation Under Imperfect Credibility <i>by Matthieu Lemoine and Jesper Lindé</i>	2016:322
Challenges for Central Banks' Macro Models <i>by Jesper Lindé, Frank Smets and Rafael Wouters</i>	2016:323
The interest rate effects of government bond purchases away from the lower bound <i>by Rafael B. De Rezende</i>	2016:324
COVENANT-LIGHT CONTRACTS AND CREDITOR COORDINATION <i>by Bo Becker and Victoria Ivashina</i>	2016:325
Endogenous Separations, Wage Rigidities and Employment Volatility <i>by Mikael Carlsson and Andreas Westermark</i>	2016:326
Renovatio Monetæ: Gesell Taxes in Practice <i>by Roger Svensson and Andreas Westermark</i>	2016:327
Adjusting for Information Content when Comparing Forecast Performance <i>by Michael K. Andersson, Ted Aranki and André Reslow</i>	2016:328
Economic Scarcity and Consumers' Credit Choice <i>by Marieke Bos, Chloé Le Coq and Peter van Santen</i>	2016:329
Uncertain pension income and household saving <i>by Peter van Santen</i>	2016:330
Money, Credit and Banking and the Cost of Financial Activity <i>by Paola Boel and Gabriele Camera</i>	2016:331



Sveriges Riksbank
Visiting address: Brunkebergs torg 11
Mail address: se-103 37 Stockholm

Website: www.riksbank.se
Telephone: +46 8 787 00 00, Fax: +46 8 21 05 31
E-mail: registratorn@riksbank.se