

CAMP Working Paper Series
No 6/2021

The Role of Precautionary and Speculative Demand in the Global Market for Crude Oil

Jamie L. Cross, Bao H. Nguyen and Trung Duc Tran



© Authors 2021 This paper can be downloaded without charge from the CAMP website bi.no/camp



The Role of Precautionary and Speculative Demand in the Global Market for Crude Oil*

Jamie L. Cross[†] Bao H. Nguyen[‡] Trung Duc Tran[§]

November 1, 2021

Abstract

Contemporary structural models of the global market for crude oil jointly specify precautionary and speculative demand as a composite shock, known as storage demand shocks, due to difficulties in identifying these distinct demand components. This difficulty arises because shifts in the underlying expectations are latent and operate through similar transmission mechanisms. In this paper, we resolve this identification problem and for the first time examine the relative effects of these two shocks, in addition to more conventional demand and supply shocks, on the global price of crude oil. Overall, we find that uncertainty driven precautionary demand for crude oil is, on average, the primary driver of fluctuations in the real price of oil that has previously been associated with the composite storage demand shock. Historically, we find that these shocks also shaped the real oil price dynamics since the 1970s. Precautionary demand for oil was the primary driver of the oil price spike in the 1979 oil crisis, the second most important driver of the price decline in the Great Recession of 2008 and provided significant contributions towards the price dynamics during the Iran-Iraq War of 1980, the Persian Gulf War of 1990, the collapse of OPEC in 1985. Speculative demand for oil largely shaped the oil price dynamics around the collapse of OPEC and contributed towards the oil price spike in the Persian Gulf War and the oil price decline of 2014.

JEL-codes: C32, C52, Q41, Q43

Keywords: Oil price uncertainty, Oil market, SVAR, Narrative sign restrictions

*Previously circulated as “The Role of Uncertainty in the Market for Crude Oil”. We thank Hilde Bjørnland, Knut Are Aastveit, Lutz Kilian, Xiaoqing Zhou, Ana Maria Herrera, Soojin Jo, James Morley, Benjamin Wong, Leif Anders Thorsrud, Thomas Størdal Gundersen, Even Comfort Hvinden, Felix Kapfhammer, Reinhard Ellwanger, Dimitris Korobilis, Gary Koop, Francesco Ravazzolo and Francesca Loria for their valuable discussions. We also benefited from the comments of members at the 2019 Workshop on Energy Economics at Sungkyunkwan University, and seminar participants at the Bank of Canada and University of Strathclyde. This paper is part of the research activities at the Centre for Applied Macroeconomics and Commodity Prices (CAMP) at the BI Norwegian Business School.

[†]BI Norwegian Business School.

[‡]University of Tasmania.

[§]University of Sydney.

1 Introduction

It is a common belief among academic economists and industry practitioners that uncertainty induced precautionary demand for oil is an important driver of the international price of crude oil. Despite this fact, contemporary structural models of the global market for crude oil jointly specify precautionary and speculative demand for oil as a composite shock due to difficulties in identifying these distinct demand components.¹ The primary difficulty is that the underlying expectation shifts are latent and operate through similar transmission mechanisms. On the one hand, an unanticipated increase in uncertainty about future market conditions causes agents to insure against possible shortfalls by increasing their holdings of above-ground oil inventories. This precautionary demand for oil results in an immediate increase in the real spot price of crude oil, followed by a gradual decline (Kilian, 2009; Alquist and Kilian, 2010). On the other hand, when speculators purchase a large quantity of oil inventories, they send a signal to oil producers that they expect higher prices in the future. This speculative demand results in producers increasing their holdings of inventories in order to sell it at the higher future price (Kilian and Murphy, 2014). This fact has historically lead scholars to either examine one of these two shocks in isolation (Elder and Serletis, 2010; Fattouh et al., 2013; Jo, 2014; Anzuini et al., 2015; Juvenal and Petrella, 2015), or to group these two shocks together into a single convolution (Kilian and Murphy, 2014; Zhou, 2019; Baumeister and Hamilton, 2019). The cost of either approach is that our understanding of the relative effects that these distinct demand shocks have on both the price of oil and the economy remains incomplete.

In this paper, we resolve this identification problem and for the first time investigate the relative effects of these two shocks, in addition to more conventional demand and supply shocks, on the global price of crude oil. To this end, we build on the workhorse structural VAR model of the global oil market developed in Kilian and Murphy (2014), as recently

¹See Kilian and Zhou (2020b) for a recent overview of the methodological developments in state-of-the-art oil market models and Herrera et al. (2019) for a recent survey of the empirical literature.

refined in [Zhou \(2019\)](#), in two ways.² First, we construct a newly observable monthly measure of real oil price uncertainty. Second, we extend the identification restrictions in [Zhou \(2019\)](#) to disentangle the relative effects of precautionary and speculative demand shocks. The key variation that we exploit is the fact that precautionary motives are necessarily associated with a high state of oil market uncertainty, while speculative motives are not.

To measure oil price uncertainty (OPU) we construct an observable monthly OPU index. In the spirit of [Diebold and Kilian \(2001\)](#) and [Jurado et al. \(2015\)](#), OPU is defined as the conditional volatility of the unpredictable component from a forecasting model of the real price of oil.³ Unlike commonly used volatility indicators, such as the Chicago Board Options Exchange’s (CBOEs) Oil Price Volatility Index (OVX), or model based measures, e.g. autoregressive heteroscedasticity (ARCH) and stochastic volatility (SV) models, this definition captures the fact what matters for economic decision making is not whether the real price of oil has become more or less variable, but rather, whether it has become more or less predictable, i.e. less or more uncertain. In this sense, the index also differs from alternative OPU indexes that are based on OPEC announcements ([Plante and Traum, 2012](#)) or media coverage ([Bonaparte, 2015](#)).

Our results provide new insights on the relative roles of precautionary and speculative demand in generating real price of oil dynamics since the 1970s. Overall, we find that uncertainty driven precautionary demand for crude oil is, on average, the primary driver of fluctuations in the real price of oil that has previously been associated with storage demand shocks. Historically, we find that shifts in precautionary demand account for much of the oil price variation during periods of adverse sociopolitical conditions in the Middle-East, such as the 1979 oil crisis, as well as the oil price collapse during the Great Recession of 2008. In addition to finding an important role for precautionary demand, we also find that

²We could alternatively build on the recently proposed oil market model in [Baumeister and Hamilton \(2019\)](#), but choose the model in [Kilian and Murphy \(2014\)](#) for two reasons. First, it has been shown that the conclusions from these models are very similar after conditioning on a specific prior distribution ([Herrera and Rangaraju, 2020](#)). Second, the [Kilian and Murphy \(2014\)](#) model has been the workhorse model in the oil market literature for the past decade and therefore serves as a useful benchmark.

³A modified version of the index has also been used in [Nguyen et al. \(2021\)](#) to examine the effects of oil uncertainty shocks in a non-linear SVAR model, and in [Tran \(2020\)](#) to investigate the macroeconomic effect of commodity price uncertainty in a small open economy setting.

speculative demand was an important driver of the real price of oil collapse associated with the disbandment of OPEC in 1985 as well as the oil price decline of 2014. In line with existing research, we also find no evidence of rising speculative demand during the early millennium price surge of mid-2003-08 (Kilian, 2009; Kilian and Murphy, 2014). Instead, flow demand shocks account for much of the oil price dynamics during this period, a result that is generally attributed to unexpectedly high demand from emerging Asia (Kilian, 2009; Kilian and Hicks, 2013; Aastveit et al., 2015).

The rest of the paper is organized as follows. In Section 2 we discuss the OPU index. We present the oil market model in Section 3, discuss results in Section 4 and conclude in Section 5.

2 Construction of the Oil Price Uncertainty Index

A key challenge in empirically examining the effects of uncertainty driven precautionary motives in the global market for crude oil is that they are not directly observable. For this reason, scholars interested in examining the effects of oil price uncertainty shocks have historically relied on model based proxies such as GARCH or SV models (Elder and Serletis, 2010; Jo, 2014). Despite the popularity of these approaches, an alternative view is that volatility based measures are not good proxies of uncertainty because they do not capture the fact that what matters for economic decision making is not whether particular economic variables have become more or less disperse, but whether the economy has become more or less predictable (Diebold and Kilian, 2001; Jurado et al., 2015). As a result, uncertainty should not be defined in terms of volatility, but instead, in terms of predictability. This is not to say that modeling volatility is irrelevant for measuring uncertainty, but rather, that it is important for the predictive model to be sufficiently informative, so that the measured forecast error is first “purged of predictive content” (Jurado et al., 2015, p.1184). Only then should a volatility model be applied to extract the underlying uncertainty component of the time series. With this idea in mind, the one-period ahead uncertainty, OPU_{t+1} , of an oil price series, y_t , is defined as forecast error variance of a forecasting model:

$$OPU_{t+1} = \sqrt{E [(y_{t+1} - E[y_{t+1}|I_t])^2 | I_t]}, \quad (1)$$

where the expectation $E(\cdot|I_t)$ is formed with respect to information available at time t . Note that the definition implies that uncertainty about oil prices will be higher when the expectation today of the squared error in forecasting y_{t+1} rises, and vice versa.

The forecast in (1) is obtained by a time series model of the form

$$y_{t+1} = \phi(L)y_t + \psi(L)X_t + \sigma_{t+1}\epsilon_{t+1}, \quad (2)$$

$$\log[(\sigma_{t+1})^2] = \alpha + \beta \log[(\sigma_t)^2] + \omega\eta_{t+1}, \quad (3)$$

$$\begin{bmatrix} \epsilon_{t+1} \\ \eta_{t+1} \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right), \quad (4)$$

where $\phi(L)$ and $\psi(L)$ are lag polynomials and X_t is a matrix of predictors which contain information that is considered robust in forecasting oil prices. The stochastic volatility parameters α, β, ω can be estimated using Bayesian methods (Kastner, 2016). Given these values, one-period-ahead uncertainty defined in (1) is then given by

$$\begin{aligned} OPU_{t+1} &= \sqrt{E[(\sigma_{t+1})^2|I_t]}, \\ &= \sqrt{\exp\left(\alpha + \beta \log(\sigma_t)^2 + \frac{\omega^2}{2}\right)}. \end{aligned} \quad (5)$$

It should be noted that with longer-horizon forecasts, uncertainty is not equal to stochastic volatility in residual σ_{t+1} . Instead, there are additional autoregressive terms, stochastic volatility in additional predictors and covariance terms (Jurado et al., 2015).

Two decisions must be made when constructing the index. First is the choice of an appropriate oil price series, and second is to select a matrix of variables that are useful predictors of the selected oil price series.

In the first stage, we use the U.S. refiners' acquisition cost for imported crude oil (IRAC) as a proxy for the price of crude oil in global markets. We use IRAC because that price is likely to be the most commonly used measure of the price of crude oil in academic studies that investigate the underlying drivers of the real price of crude oil. This includes both Kilian and Murphy (2014) and Zhou (2019), whom we build on in this paper.

In the second stage, we forecast the real price of crude oil using a set of additional variables from a state of the art oil price forecasting model in Alquist et al. (2013). Our set of variables includes the set of fundamental oil market variables suggested by Kilian and




Figure 1: Oil price uncertainty (OPU) index

Notes: The figure plots the oil price uncertainty index (OPU) constructed in Section 2 from 1975:2 to 2018:6.

Murphy (2014): oil production, real economic activity and above-ground oil inventories, and additional variables that have been shown to be important drivers of the price of oil: US CPI inflation and the M1 money supply, commodity currency exchange rates, and excess co-movement with other commodity prices. The number of lags for both the autoregressive and predictor polynomials is set to be 24. The choice of a long lag length is known to be essential when modeling oil prices as it allows for a richer dynamic relationship (Kilian and Lütkepohl, 2017). Finally, following Bai and Ng (2008), the predictors X_t that we ultimately use in the predictive equation (2) in each forecast is restricted to those that have significant predictive power, as defined by a $|t - stat| > 2.575$.

The resulting OPU index is plotted in Figure 1, along with major events associated with the oil market.⁴ Overall, the general trend is that the uncertainty index displays clear spikes around significant events. This includes sociopolitical events such as the Iranian revolution, the Iran-Iraq War, the disbandment of OPEC in 1986, and the Persian Gulf War. It also includes other well known episodes of significance, such as the Asian crisis of 1997/98, when the real price of oil fell to an all-time low, the large price decline during the Great Recession and the more recent 2014-15 price drop.

2.1 Comparison with Other Uncertainty Measures

A natural question is whether our oil price uncertainty measure is distinct from alternative measures of uncertainty. To investigate this point, we compare our OPU measure with the CBOE Oil Price Volatility Index (OVX) and four widely used sources of alternative uncertainty measures: financial uncertainty, as measured by the CBOE (stock price) Volatility Index (VIX); the US Economic Policy Uncertainty (EPU) index proposed by Baker et al. (2016), the US macroeconomic uncertainty (JLN) index constructed by Jurado et al. (2015) and the Geopolitical Risk (GPR) index developed by Caldara and Iacoviello (2018).

The comparison in Figure 2 reveals that the dynamics of OPU are most consistent with the OVX index. Since option prices are driven by both precautionary and speculative

⁴We use full-sample estimation as in Jurado et al. (2015) to provide a full historical estimate of OPU. Instead of full sample estimation (and hence in-sample forecasts), we conduct out-of-sample forecasts as a robustness check. We find that estimate of OPU is not sensitive to whether we use out-of-sample or in-sample forecasts. The correlation between the two is 0.93.

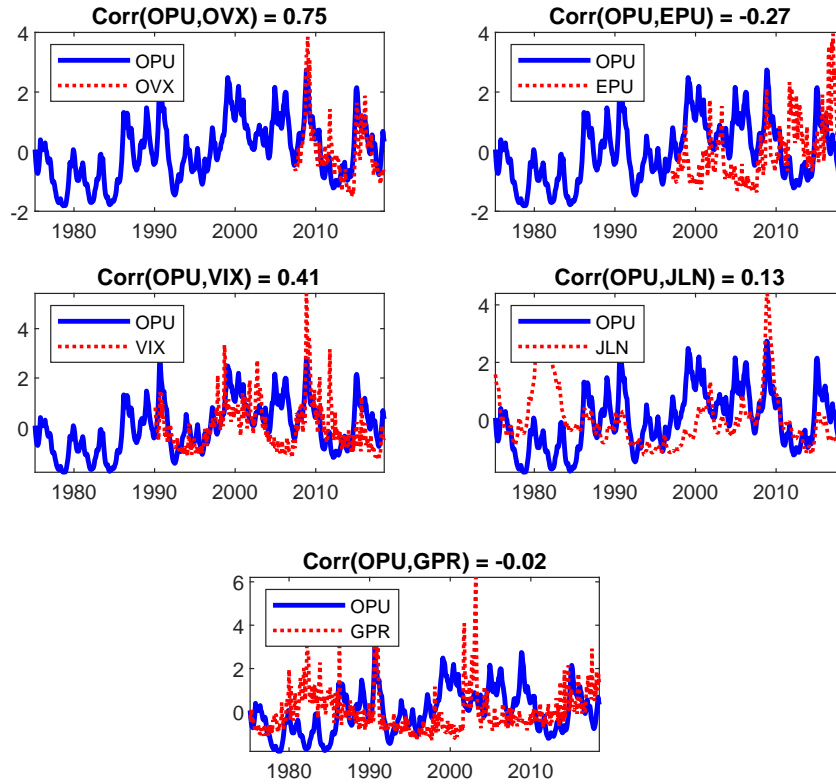


Figure 2: Oil price uncertainty (OPU) index: Comparison with other uncertainty indices

Notes: The figure compares the oil price uncertainty index (OPU) constructed in Section 2 from 1975:2 to 2018:6 to: (i) The CBOE Oil Price Volatility Index (OVX) from 2007:5 to 2017:6, (ii) The Global Economic Policy Uncertainty index (EPU) by Baker et al. (2016) from 1997:1 to 2017:6, (iii) The CBOE volatility index (VIX) from 1994:7 to 2017:6, (iv) The uncertainty index (JLN) for the U.S by Jurado et al. (2015) from 1975:2 to 2018:6 and (v) The geopolitical risk index by Caldara and Iacoviello (2018) from 1975:2 to 2018:6. All series are normalized to have means of zero and standard deviations of one.

motives, the moderately high correlation between the two series is expected. The major distinction between the OPU and the OVX is that the OPU does not report any heightened uncertainty around 2011. Next, the lack of correlation with the EPU index and the GPR suggests that oil price uncertainty is highly different from economic policy uncertainty and geopolitical risk. Last, although oil price uncertainty correlates moderately with both stock market (VIX) and macroeconomic uncertainty in the US (JLN), there are still some notable differences. While the OPU detects spikes following the collapse of the OPEC in 1986 and the 1990/91 Persian Gulf War, these high-uncertainty events are not reported by the JLN index. The OPU does not pick up high uncertainty about the Dotcom crisis or the European Debt Crisis that are otherwise detected by the VIX since those events are more relevant to the stock exchange. In addition, neither the VIX nor the JLN macro uncertainty index detects any surge in oil uncertainty during 2000/02 and during 2015/16. Taken together, this suggests that the OPU index is able to pick up uncertainty events that are highly specific to the oil market.

2.2 Examining the Role of Additional Predictors

Another natural question is whether or not it is important to include additional predictive information when estimating the OPU. We consequently re-estimate the OPU index using potentially misspecified models in which we replace (2) with: (i) a constant conditional mean equation, i.e. $y_{t+1} = \mu + \hat{\sigma}_{t+1}\hat{\epsilon}_{t+1}$, and (ii) autoregressive terms only, i.e. $y_{t+1} = \phi(L)y_t + \hat{\sigma}_{t+1}\hat{\epsilon}_{t+1}$. A comparison of these models reveals how our measure of uncertainty are affected by the predictable variation in oil prices. The resulting indices are shown in Figure 3. We observe that there is a substantial predictable component in the selected oil price series. In particular, our OPU index is significantly lower than the misspecified measures of uncertainty in every peak, with the largest peak in the index coincides with the 1990/91 Persian Gulf War, as compared to the 2007/08 Great Recession. This suggests that an OPU with no predictors will not only overstate the degree of oil price uncertainty, but may also erroneously attribute oil price uncertainty to events that are not specific to the global market for crude oil.

To further examine the role of each individual predictor in driving the OPU index, we

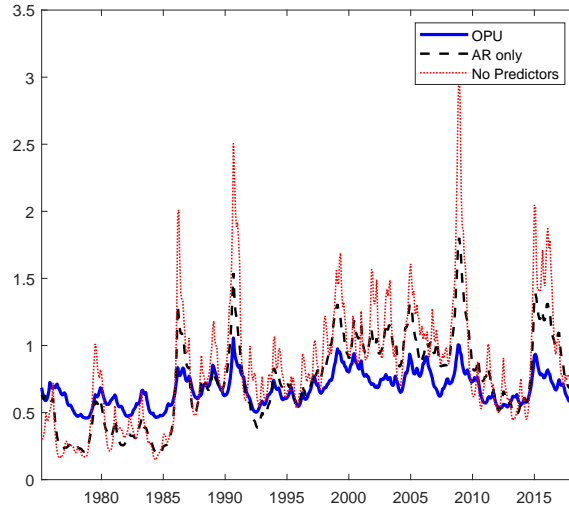


Figure 3: The role of incorporating predictive information

Notes: The figure contains (i) our oil price uncertainty index (OPU), (ii) a potentially misspecified model in which we only include an intercept (dotted line), and (iii) a model in which only autoregressive terms are included to forecast (dashed line) from 1975:2 to 2018:6.

now investigate how the index changes when we subsequently remove a variable in X_t .⁵ The results in Figure 4 show that the most frequently selected predictors are commodity exchange rates and excess co-movement terms as removing these variables affects the OPU most. Omitted information about above-ground oil inventories, US inflation and M1 money stock affects the OPU to a lesser extent, while we see little effect from removing the real economic activity index.

⁵For example, to see the contribution of commodity exchange rates, we start from the retained list of regressors that passes the hard threshold test. We then run a regression of y_{t+1} on a constant, the AR terms and the retained regressors without commodity exchange rates. Then we compute a measure of uncertainty that does not utilize information on commodity exchange rates.

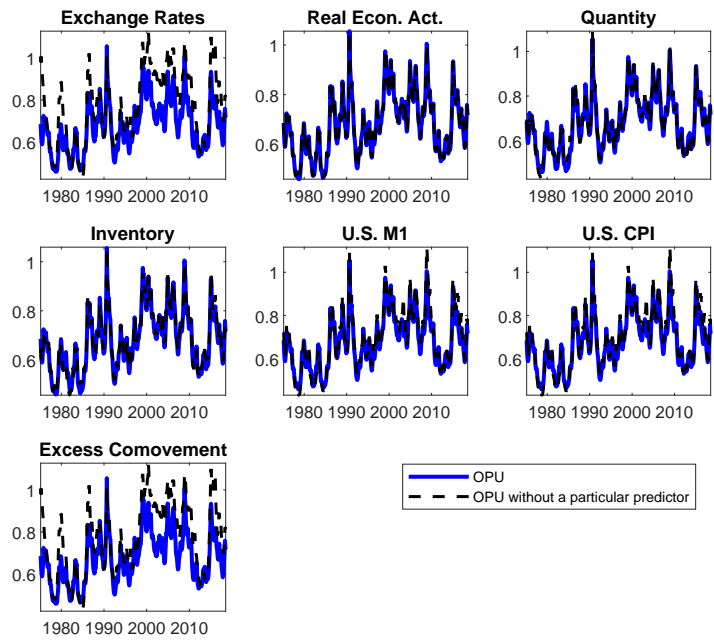


Figure 4: The role of individual predictor

Notes: The figure shows the oil price uncertainty index (OPU) (solid line) and an OPU in which only once particular predictor in X_t is removed at one time (dashed line) from 1975:2 to 2018:6.

3 Empirical Methodology

3.1 The Structural VAR Model

The structural VAR model of the global market for crude oil is given by

$$\mathbf{B}_0 \mathbf{y}_t = \mathbf{b} + \sum_{j=1}^{24} \mathbf{B}_j \mathbf{y}_{t-j} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (6)$$

where $\mathbf{y}_t = (\% \Delta prod_t, rea_t, rpo_t, \Delta inv_t, OPU_t)'$ in which: $\% \Delta prod_t$ is the percent change in global crude oil production, rea_t is a measure of global real economic activity, rpo_t is the natural logarithm of the global real price of oil, Δinv_t is the change in above-ground global crude oil inventories and OPU_t is the oil price uncertainty index discussed in Section 2. The lag order of 24 is in line with existing studies, e.g. [Kilian \(2009\)](#), [Kilian and Murphy \(2014\)](#) and [Zhou \(2019\)](#).

3.2 Data

The reduced form version of the above structural VAR model is estimated with a data set that contains monthly observations on the four fundamental oil market variables from 1973:1 to 2018:6, plus our oil price uncertainty index. The raw and transformed data for the oil market variables is in line with [Kilian and Murphy \(2014\)](#) and [Zhou \(2019\)](#). First, *crude oil production* is taken from the U.S. Energy Information Administration (EIA) and converted to percent changes. Second, *real economic activity* is taken to be the index of global real activity constructed from bulk dry cargo ocean shipping freight rates proposed in [Kilian \(2009\)](#) and subsequently revised in [Kilian \(2019\)](#). This index is stationary by construction. Third, the *real price of crude oil* is defined as the US refiners' acquisition cost for imported crude oil (IRAC), as reported by the EIA, extrapolated from 1974:1 back to 1973:1 as in [Barsky and Kilian \(2004\)](#) and deflated by the US consumer price index (all items), which are obtained from the FRED database. Fourth, *above-ground crude oil inventories* are measured using the first difference of the total US crude oil inventories scaled by the ratio of OECD petroleum stocks over US petroleum stocks, all of which are obtained from the EIA. Finally, we deseasonalized each of these variables before estimating the model.

3.3 Identification

To identify the global oil market shocks, we extend the four sets of identifying assumptions used in [Zhou \(2019\)](#). This consists of the (1) static sign restrictions, (2) elasticity restrictions (3) dynamic sign restrictions, and (4) narrative restrictions. In particular, we make modifications to steps (1), (3) and (4), and leave step (2) unchanged. We discuss each of these changes in the proceeding subsections.⁶

3.3.1 Static Sign Restrictions

The first stage of identification utilizes a set of static sign restrictions to obtain a set of admissible models in which the variables contemporaneously respond to the four structural shocks are in line with economic theory.⁷ Any remaining variation after identifying these four shocks is then treated as an unexplained residual. As noted in [Zhou \(2019\)](#), this residual component contains all other shocks to the demand for oil such as preference shocks, shocks to the storage of technology, or politically motivated changes in the Strategic Petroleum Reserve.

The set of static sign restrictions is reported in [Table 1](#). A flow supply shock (column 1), reduces real economic activity and oil inventories, while increasing the real price of oil.⁸ Such shocks represent unanticipated supply disruptions associated with exogenous political events in oil-producing countries and politically motivated supply decisions by OPEC members ([Hamilton, 2003](#); [Kilian, 2008, 2009](#)). In contrast, a flow demand shock (column 2) increases each of oil production, real economic activity and the real price of oil. In both cases, the inventory response is unspecified, thereby allowing the data to determine the reaction.

⁶Step (2) amounts to imposing a lower bound on the short-run demand elasticity of -0.8 and an upper bound on the short-run supply elasticity of 0.04.

⁷The sign restrictions are implemented with the algorithm in [Antolín-Díaz and Rubio-Ramírez \(2018\)](#) which builds on the widely used procedure in [Rubio-Ramírez et al. \(2010\)](#) narrative restrictions.

⁸The restriction that inventories responds negatively to a flow supply shock was not used in [Kilian and Murphy \(2014\)](#), but is in line with more recent research, e.g. [Kilian and Zhou \(2020a\)](#) and [Känzig \(2021\)](#).

Table 1: Sign restrictions

	flow supply shock	flow demand shock	speculative demand shock	precautionary demand shock
Oil production	–	+	+	×
Real Economic Activity	–	+	–	–
Real oil price	+	+	+	+
Inventories	–	–	+	+
Uncertainty	×	×	×	+

Notes: + and – respectively indicate positive and negative responses, while × leaves the effect unrestricted. In the event that the signs in column four (the precautionary demand shock) are the same as columns three (speculative demand shock), we assume that the precautionary demand shock induces a larger response in uncertainty (i.e. element (5,4) is larger than elements (5,3)).

Since precautionary and speculative motives share similar transmission mechanisms, we use similar sign restrictions on each of these two shocks (columns 3 and 4). The one exception is that we remain agnostic about the contemporaneous response of oil production to a precautionary demand shock. This is motivated by the fact that higher uncertainty may increase the quantity of oil produced, however it may instead elicit a *real options* effect on oil producers—i.e., they delay production as they wait and see what happens to oil prices in the future (Bernanke, 1983). This point is also in line with the general equilibrium model of Alquist and Kilian (2010) who argue that oil producers may sell oil futures to protect against endowment uncertainty, however the strength of this mechanism remains an empirical question. In light of this theoretical mechanism, we think it is prudent to not impose or prevent such an *a priori* response, and instead allow the data to inform us about the empirical validity of such behavior.⁹

Finally, it is important to note that our decision to remain agnostic does not come without costs. In particular, the precautionary and speculative demand shock may elicit the same sign pattern in the contemporaneous response. To achieve identification of these

⁹Technically, real options theory relates to long-run uncertainty, however it is common to use short-run uncertainty in empirical studies (see, e.g. Castelnovo (2019) and references therein).

shocks, we exploit the fact that precautionary motives are, by definition, associated with high uncertainty while speculative demand shocks are not. This fact translates to a restriction that oil price uncertainty is necessarily more responsive to a precautionary demand shock relative to a speculative demand shock. In other words, element (5,4) of Table 1 is restricted to be greater than element (5,3).¹⁰

3.3.2 Dynamic Sign Restrictions

It is well known that static sign restrictions, such as those presented in Table 1, are by themselves insufficient to identify the underlying shocks in the global market for crude oil (Kilian and Murphy, 2012, 2014). Following Kilian and Murphy (2014, p.462), we therefore impose the additional dynamic sign restrictions that the responses of oil production and global real activity to an unanticipated flow supply disruption to be negative for the first 12 months, while the real price of oil response is restricted to be positive over the same period. As they note, these restrictions are necessary to rule out structural models in which unanticipated flow supply disruptions cause a decline in the real price of oil below its starting level. Such a decline would be at odds with conventional views of the effects of unanticipated oil supply disruptions.

In a similar spirit, we impose the additional dynamic sign restrictions that any increase in the real price of oil caused by either a precautionary and speculative demand shock must have an associated increase in oil inventories, and that precautionary demand shocks do not increase global economic activity for the first 12 months. The intuition for the first restrictions is that any expectation of a shortfall of future oil supply relative to future oil demand not already captured by flow demand and flow supply shocks necessarily causes an increase in the demand for above-ground oil inventories and hence in the real price of oil (Hamilton, 2009; Alquist and Kilian, 2010).¹¹ Next, ruling out responses in which precau-

¹⁰Restrictions on the responsiveness of one variable relative to another have been applied previously by, among others, Farrant and Peersman (2006), Eickmeier and Ng (2015) and Aastveit et al. (2015), albeit in different empirical contexts. The paper most similar to the present one is that of Aastveit et al. (2015) who use this strategy to identify the relative effects of an emerging-country demand shock and a developed-country demand shock in the global market for crude oil.

¹¹An exception is the extreme case in which the price elasticity of oil demand is zero (Kilian and Murphy, 2014).

tionary demand shocks elicits a positive response to real output is in line with theoretical arguments in [Alquist and Kilian \(2010\)](#) and empirical results in [Jo \(2014\)](#).

3.3.3 Narrative Sign Restrictions

In addition to the aforementioned identification assumptions, [Zhou \(2019\)](#) postulated the use of five additional narrative sign restrictions on the historical decomposition of the real price of oil. In the first instance, storage demand shocks are assumed to (cumulatively) raise the log real price of oil by at least 0.2 (or approximately 20%) between May and December of 1979. This is consistent with anecdotal evidence of a dramatic surge of inventory building in the oil market during that time. Second, following the collapse of OPEC in December of 1985, storage demand cumulatively lowered the log real price of oil by at least 0.15 up until December 1986. Third, in line with the established belief that Iraq would invade its neighbors, storage demand shocks raised the log real price of oil by at least 0.1 cumulatively between June 1990 and October 1990. Fourth, following the invasion of Kuwait and the cessation of Iraqi and Kuwaiti oil production in early August of 1990, flow supply shocks are assumed to have raised the log real price of oil cumulatively by at least 0.1 between July and October of 1990. Fifth and final, the cumulative effect of flow demand shocks on the log real price of oil between June and October of 1990 is bounded by 0.1, given that the oil price spike of 1990 was not associated with the global business cycle.

One difficulty in directly applying these narrative restrictions in our framework is that it remains unclear which component of storage demand, i.e. speculative and precautionary motives, is associated with the first three of the above restrictions. Since imposing a restriction on the wrong component would bias our results, we instead impose the narrative restrictions on the sum of their cumulative effects. For instance, the first restriction translates to imposing that the log real price of oil increased by at least 0.2 as a result of precautionary and speculative shocks between May and December of 1979.

4 Oil Market Results

The SVAR model is estimated using Bayesian methods as in [Zhou \(2019\)](#). The reported point estimates correspond to the most likely structural model from the posterior distribution for which computational details can be found in [Inoue and Kilian \(2013, 2019\)](#). When reporting the impulse response functions, we follow recent advice in [Baumeister and Hamilton \(2021\)](#) and present a full set of posterior draws as opposed a credible set.

4.1 Responses of Variables to Oil Market Shocks

The impulse response functions (IRFs) for a one standard deviation shock from each of the identified demand and supply shocks are shown in [Figure 5](#). The boldface (black) line corresponds to the most likely structural model, and the thin (red) lines correspond to distinct structural models from the posterior distribution. To aid in the interpretation of the dynamics, we report numerical results for the modal structural model on impact, 6 months and 1 year after the shock in [Table 2](#). Following convention in the literature, all shocks have been normalized such that they increase in the real price of oil on impact. In particular, the flow supply shock refers to an unanticipated flow supply disruption.

Since we are the first to jointly identify the precautionary and speculative demand components underlying storage demand, the IRFs in rows three and four of [Figure 5](#) are new. First, focusing on the third row, we observe that the most likely response to a positive speculative demand shock is an immediate and persistent jump in the real price of oil, which is accompanied by a large persistent increase in above-ground inventories. Such shocks also generate a gradual decline in oil production and a temporary reduction in real activity. These responses are in line with the proposed mechanism outlined in [Kilian and Murphy \(2014\)](#). Specifically, when speculators purchase a large quantity of inventories, they signal to oil producers that they expect higher prices in the future. This causes oil producers to withhold oil from the market in order to sell the stored oil at the higher price. Such withholding can take place by either increasing holdings of above-ground inventories or reducing the number of barrels pumped out of the ground. Our result that a speculative demand shock elicits an immediate increase in above-ground inventories and a gradual decline in oil production, suggests that both types of behavior are at play.

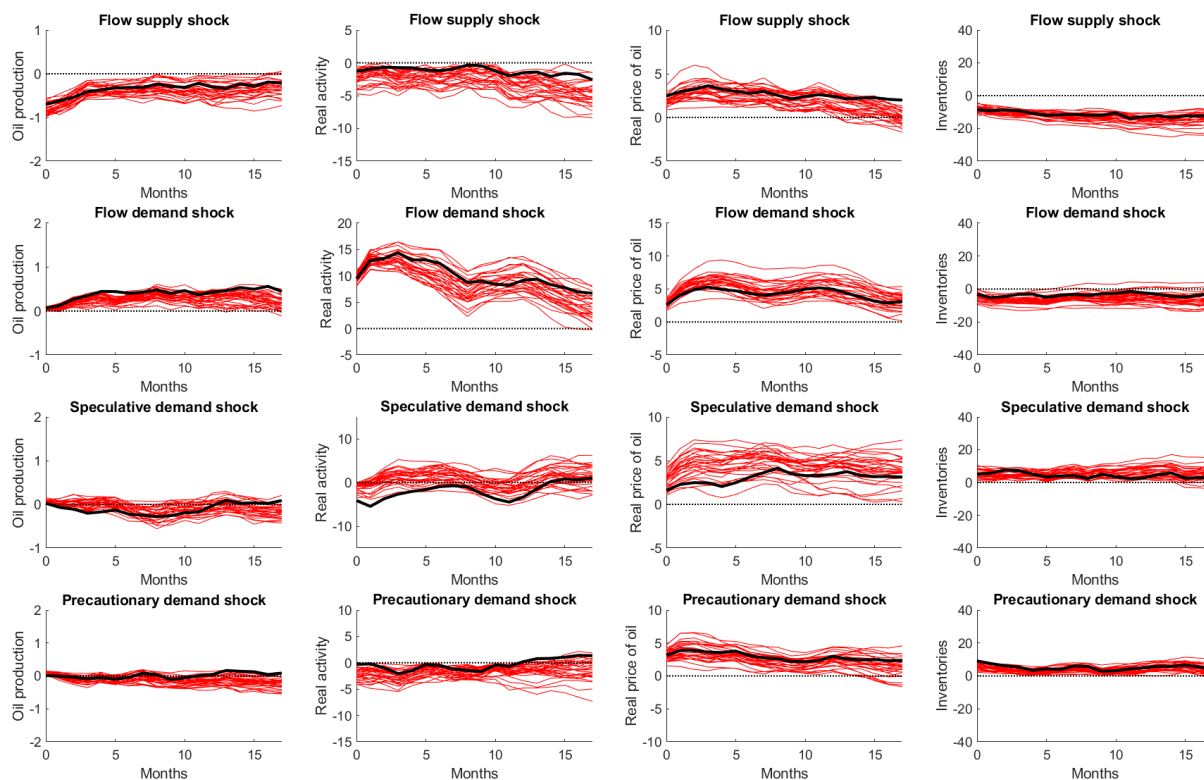


Figure 5: Structural impulse response functions.

Notes: Impulse response functions for a one standard deviation shock from each of the identified demand and supply shocks. The response in boldface represents the most likely structural model computed with methods in Inoue and Kilian (2013, 2019) and the remaining responses are from 100 distinct structural models, drawn from the posterior distribution. Oil production and inventories are expressed as the cumulative percent change of the respective impulse responses.

The results in row four show that an unanticipated increase in uncertainty induced precautionary demand also elicits an immediate and persistent positive effect on the real price of oil. The relatively large magnitude of the price response suggests that such shocks highlights the importance of disentangling such shocks from speculative demand shocks. In line with the theoretical mechanism proposed in Alquist and Kilian (2010), we also observe that the oil price increase is associated with an increased accumulation in above-ground oil inventories and a decrease in real economic activity.

Table 2: Structural impulse response functions at selected horizons

	Flow supply shock			Flow demand shock			Speculative demand shock			Precautionary demand shock		
	h=0	h=6	h=12	h=0	h=6	h=12	h=0	h=6	h=12	h=0	h=6	h=12
Oil production	-0.70	-0.32	-0.22	0.06	0.44	0.37	0.02	-0.13	-0.17	0.01	-0.01	0.01
Real economic activity	-1.31	-1.10	-1.99	9.46	13.02	8.18	-4.12	-1.64	-4.48	-0.32	-0.33	-0.68
Real price of oil	2.45	2.97	2.61	2.59	4.73	5.18	1.60	2.48	3.29	3.34	5.36	4.34
Inventories	-8.63	-11.96	-13.99	-0.99	-3.29	-4.99	5.08	3.42	2.27	9.12	4.31	4.23

Notes: The numbers correspond to the impulse response functions to one-standard-deviation structural shocks at horizons $h = 0, 6, 12$ as shown in Figure 5. Oil production and inventories are expressed as the cumulative percent change of the respective impulse responses. Real economic activity is expressed as percentage deviation from trend. The real price of oil is expressed in terms of real US Dollars per barrel.

While it is not directly related to our primary research questions, the result that an unexpected increase in oil price uncertainty is most likely to have a negligible effect on the production of crude oil is relevant for a related literature on real options theory. Real options theory posits that an increase in oil price uncertainty may impact the decision-making process of irreversible firm-level investments (Bernanke, 1983). The key mechanism is that an increase in oil price uncertainty causes firms to postpone major purchases of capital goods and wait-and-see what happens to the oil price. If such a real options channel is important in the market for crude oil, then a positive precautionary demand shock should result in a sustained increase in inventories, and associated reductions in oil production and real economic activity until the oil price situation manifests. Following a precautionary demand shock, our results from the full posterior distribution suggest that firms tend to reduce production and increase their inventories resulting in a decline of real economic activity. This corroborates results in Elder and Serletis (2010) and Jo (2014) who find evidence of a real options effect in global oil market models that abstract from speculative demand. That being said, the magnitudes from the most likely structural model reported in Table 2 suggests that the quantitative strength of this theoretical mechanism within the global market for crude oil is somewhat weak.

Finally, we note that the size and qualitative patterns of the responses presented in the first two rows of Figure 5 are comparable to those in Kilian and Murphy (2014) and Zhou (2019). For instance, a negative flow supply shock is associated with a reduction

in oil production, global real activity and oil inventories while increasing the real price of oil. We also observe that the price of oil rises only temporarily, peaking after about three months and declining below its starting value after about one year, due to the drop in real activity. In contrast, a positive shock to the flow demand for crude oil, is associated with a slight increase in oil production as firms meet the persistent demand associated with the increased real activity. Such shocks also cause a persistent hump-shaped increase in the real price of oil and have a negligible effect on inventories.

4.2 What Drives Fluctuations in the Real Price of Oil?

Which shock is the most important driver of the real price of oil? We address this question using a forecast error variance decomposition (FEVD) for which the results are shown in Figure 6.¹² The results show that, on impact, around 36% of the the variation in the real price of oil is driven by precautionary demand shocks, followed by flow demand shocks with 24%, flow supply shocks with 22% and precautionary demand shocks with 10%. The relative importance of these shocks changes over time, however. In the short run, we find that the overall contribution of flow demand shocks and precautionary demand shocks tends to increase, while the relative contribution of precautionary demand shocks and flow supply shocks declines. In the long run, we observe that 40% of the variation in the real price of oil can be attributed to flow demand shocks, compared with 26% due to precautionary demand shocks, 15% due to speculative demand shocks and 15% due to flow supply shocks. This evidence suggests that while flow demand shocks are the primary driver of long run variations in the real price of oil, our newly identified precautionary and speculative demand shocks have a significant effect on the real price of oil, particularly in the short run to medium run.

¹²We do not include the contribution of the residual demand shock because that shock has no structural interpretation and makes relatively little systematic contributions to real price of oil dynamics.




Figure 6: Forecast Error Variance Decomposition.

Notes: The results are for the most likely structural model from 100 distinct structural models drawn from the posterior distribution. The residual variance contribution is attributed to the unidentified component of the structural VAR model.

4.3 Reassessing the Historical Narrative

Our results so far suggest that, on average, the precautionary and speculative demand components underlying conventional storage demand shocks have different effects on the real price of oil. In light of this evidence, our objective in this section is to reassess the underlying historical narrative of what caused the ups and downs in the real price of oil since the late 1970s. The historical decomposition in Table 3 enables us to draw inference on the causal underlying dynamics of the real price of oil and above-ground inventories during the 1979 oil crisis, the Iran-Iraq War of 1980, the collapse of OPEC in 1986, Iraq’s invasion of Kuwait in 1990, the early millennium surge in the real price of oil between 2003 and mid-2008, the price drop in the Great Recession of 2008 and the oil price collapse of 2014-15.

Table 3: Cumulative effects on the real price of oil (percent)

	1979 oil crisis	Iran-Iraq War	Collapse of OPEC	Persian Gulf War	2003-08 Price Surge	Great Recession	2014-15 Price Drop
	1979:1-1980:1	1980:9-1980:12	1985:12-1986:12	1990:5-1990:10	2002:7-2008:6	2008:6-2008:12	2014:6-2015:12
Flow Supply Shocks	-2	8	7	30	-4	-11	-46
Flow Demand Shocks	29	-3	-20	5	135	-74	-24
Speculative Demand Shocks	8	-3	-43	13	5	0	-25
Precautionary Demand Shocks	38	5	-15	9	8	-26	4

Notes: Cumulative effects on the real price of oil (percent) from the workhorse oil market model.

The results in the first column reveal that the rise in the real price of oil in late 1979 associated with the Iranian Revolution was primarily driven by a sharp increase in precautionary demand associated with uncertainty around future supply shortfalls. This result supports the hypothesis in [Kilian \(2009\)](#) that the increased importance of “oil market-specific demand shocks” starting in 1979 is consistent with an increase in precautionary demand. As noted in that paper, this period was plagued by various sociopolitical events, including Khomeini’s arrival in Iran, the Iranian hostage crisis and the Soviet invasion of Afghanistan. These geopolitical events in 1979/80 occurred against the backdrop of expectations associated with persistent fears of a regional war and the destruction of oil fields in Iran and Saudi Arabia, or strong demand for oil. Importantly, without that backdrop, these events would have had no effect on the real price of oil. We also observe that such shocks played a key role in shaping the real oil price dynamics during the two wars of 1980 and 1990, however it is also noted that supply disruptions also played a significant role during these periods. Finally, we find that precautionary demand shocks also played a major role in the oil price decline during the Great Recession, but was relatively less important in the preceding oil price surge of 2003-08 and 2014 price collapse.

While uncertainty driven precautionary motives are important for explaining the real oil price dynamics during these adverse sociopolitical and economic events, our results reveal that oil price decline following OPEC’s collapse in late 1985 was largely the result of poor global economic conditions and speculative demand. In line with existing research, however, we observe no evidence of rising speculative demand during the early millennium surge in the real price of oil between 2003 and mid 2008 ([Fattouh et al., 2013](#); [Kilian and Murphy, 2014](#); [Kilian and Lee, 2014](#)). Instead, our results suggest that the surge was primarily due a sustained global economic expansion, which is generally attributed to unexpectedly high growth from emerging Asia ([Kilian, 2009](#); [Kilian and Hicks, 2013](#); [Aastveit et al., 2015](#)), with much smaller contributions stemming from speculative demand. A similar result is observed for the oil price decline during the Great Recession, however such shocks provided a relatively large contribution towards the 2014-15 price decline.

5 Conclusion

The workhorse oil market model allows researchers to examine the effects of *storage demand* shocks in addition to more conventional *flow demand* and *flow supply* shocks. The key idea underlying the identification of storage demand shocks is that latent expectation shifts about future oil market conditions are reflected by observable changes in above-ground crude oil inventories. Implicit in this assumption, however, are two very different types of economic behavior. On the one hand, *speculative demand* for oil occurs because buyers are anticipating future demand or supply conditions. In contrast, *precautionary demand* for oil occurs in response to heightened uncertainty about the future price of oil. Despite this distinction, the fact that these underlying motives are latent and share similar transmission mechanisms renders the joint identification of these two distinct shocks difficult in practice.

Our contribution in this paper was to resolve this identification problem and examine the relative effects of these distinct shocks on the real price of oil. To that end, we built upon the workhorse oil market model in two ways. First, we constructed a newly observable monthly measure of real oil price uncertainty. Second, we exploited the fact that precautionary motives are necessarily associated with a high state of oil market uncertainty, to disentangle the relative effects of precautionary and speculative demand shocks.

Our analysis provided important new insights on the relative roles of precautionary and speculative behavior in driving the real price of crude oil. Overall, we found that precautionary demand for crude oil is, on average, the primary driver of real oil price fluctuations that were previously associated with storage demand shocks. We also provided new insights on the relative roles of uncertainty induced precautionary motives and pure speculation in various episodes of historical significance. First, precautionary demand for oil was shown to be the primary driver of the oil price spike in the 1979 oil crisis, the second most important driver of the price decline in the Great Recession, and had significant effects on the price dynamics during the wars of 1980 and 1990, and the collapse of OPEC in 1985. Next, speculative demand for oil largely shaped the oil price dynamics around the collapse of OPEC, and contributed towards the oil price spike in both the Persian Gulf War and the oil price decline of 2014. Finally, in line with existing studies, such shocks were found to have had relatively no impact during the early millennium oil price surge of mid-2003-08.

References

- Aastveit, K. A., Bjørnland, H. C., and Thorsrud, L. A. (2015). What drives oil prices? emerging versus developed economies. *Journal of Applied Econometrics*, 30(7).
- Alquist, R. and Kilian, L. (2010). What do we learn from the price of crude oil futures? *Journal of Applied Econometrics*, 25(4):539–573.
- Alquist, R., Kilian, L., and Vigfusson, R. J. (2013). Forecasting the price of oil. In *Handbook of economic forecasting*, volume 2, pages 427–507. Elsevier.
- Antolín-Díaz, J. and Rubio-Ramírez, J. F. (2018). Narrative sign restrictions for SVARs. *American Economic Review*, 108(10):2802–29.
- Anzuini, A., Pagano, P., and Pisani, M. (2015). Macroeconomic effects of precautionary demand for oil. *Journal of Applied Econometrics*, 30(6):968–986.
- Bai, J. and Ng, S. (2008). Forecasting economic time series using targeted predictors. *Journal of Econometrics*, 146(2):304–317.
- Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *The quarterly journal of economics*, 131(4):1593–1636.
- Barsky, R. B. and Kilian, L. (2004). Oil and the Macroeconomy since the 1970s. *Journal of Economic Perspectives*, 18(4):115–134.
- Baumeister, C. and Hamilton, J. D. (2019). Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks. *American Economic Review*, 109(5):1873–1910.
- Baumeister, C. and Hamilton, J. D. (2021). Structural vector autoregressions with imperfect identifying information. Technical report, Technical report, University of Notre Dame.
- Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The Quarterly Journal of Economics*, 98(1):85–106.

- Bonaparte, Y. (2015). Oil Price Uncertainty Index: Capturing Media Information. *Available at SSRN 2641297*.
- Caldara, D. and Iacoviello, M. (2018). Measuring geopolitical risk. *FRB International Finance Discussion Paper*, (1222).
- Castelnuovo, E. (2019). Domestic and global uncertainty: A survey and some new results.
- Diebold, F. X. and Kilian, L. (2001). Measuring predictability: theory and macroeconomic applications. *Journal of Applied Econometrics*, 16(6):657–669.
- Eickmeier, S. and Ng, T. (2015). How do US credit supply shocks propagate internationally? A GVAR approach. *European Economic Review*, 74:128–145.
- Elder, J. and Serletis, A. (2010). Oil price uncertainty. *Journal of Money, Credit and Banking*, 42(6):1137–1159.
- Farrant, K. and Peersman, G. (2006). Is the exchange rate a shock absorber or a source of shocks? new empirical evidence. *Journal of Money, Credit and Banking*, pages 939–961.
- Fattouh, B., Kilian, L., and Mahadeva, L. (2013). The role of speculation in oil markets: What have we learned so far? *The Energy Journal*, pages 7–33.
- Hamilton, J. D. (2003). What is an oil shock? *Journal of Econometrics*, 113(2):363–398.
- Hamilton, J. D. (2009). Causes and consequences of the oil shock of 2007-08. *Brookings Papers on Economic Activity*, page 215.
- Herrera, A. M., Karaki, M. B., and Rangaraju, S. K. (2019). Oil price shocks and U.S. economic activity. *Energy Policy*, 129:89 – 99.
- Herrera, A. M. and Rangaraju, S. K. (2020). The effect of oil supply shocks on US economic activity: What have we learned? *Journal of Applied Econometrics*, 35(2):141–159.
- Inoue, A. and Kilian, L. (2013). Inference on impulse response functions in structural VAR models. *Journal of Econometrics*, 177(1):1–13.

- Inoue, A. and Kilian, L. (2019). Corrigendum to “Inference on impulse response functions in structural VAR models” [J. Econometrics 177 (2013) 1–13]. *Journal of Econometrics*, 209(1):139–143.
- Jo, S. (2014). The effects of oil price uncertainty on global real economic activity. *Journal of Money, Credit and Banking*, 46(6):1113–1135.
- Jurado, K., Ludvigson, S. C., and Ng, S. (2015). Measuring uncertainty. *American Economic Review*, 105(3):1177–1216.
- Juvenal, L. and Petrella, I. (2015). Speculation in the oil market. *Journal of Applied Econometrics*, 30(4):621–649.
- Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from opec announcements. *American Economic Review*, 111(4):1092–1125.
- Kastner, G. (2016). Dealing with stochastic volatility in time series using the R package stochvol. *Journal of Statistical Software*, 69(5):1–30.
- Kilian, L. (2008). Exogenous oil supply shocks: how big are they and how much do they matter for the us economy? *The Review of Economics and Statistics*, 90(2):216–240.
- Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069.
- Kilian, L. (2019). Measuring global real economic activity: Do recent critiques hold up to scrutiny? *Economics Letters*, 178:106–110.
- Kilian, L. and Hicks, B. (2013). Did unexpectedly strong economic growth cause the oil price shock of 2003–2008? *Journal of Forecasting*, 32(5):385–394.
- Kilian, L. and Lee, T. K. (2014). Quantifying the speculative component in the real price of oil: The role of global oil inventories. *Journal of International Money and Finance*, 42:71–87.
- Kilian, L. and Lütkepohl, H. (2017). *Structural vector autoregressive analysis*. Cambridge University Press.

- Kilian, L. and Murphy, D. P. (2012). Why agnostic sign restrictions are not enough: understanding the dynamics of oil market var models. *Journal of the European Economic Association*, 10(5):1166–1188.
- Kilian, L. and Murphy, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29(3):454–478.
- Kilian, L. and Zhou, X. (2020a). Does drawing down the US Strategic Petroleum Reserve help stabilize oil prices? *Journal of Applied Econometrics*, 35(6):673–691.
- Kilian, L. and Zhou, X. (2020b). The Econometrics of Oil Market VAR Models. *Working paper*.
- Nguyen, B. H., Okimoto, T., and Tran, T. D. (2021). Uncertainty and sign-dependent effects of oil market shocks. *Journal of Commodity Markets, Forthcomming*.
- Plante, M. and Traum, N. (2012). Time-varying oil price volatility and macroeconomic aggregates. *Center for Applied Economics and Policy Research Working Paper*.
- Rubio-Ramirez, J. F., Waggoner, D. F., and Zha, T. (2010). Structural vector autoregressions: Theory of identification and algorithms for inference. *The Review of Economic Studies*, 77(2):665–696.
- Tran, T. (2020). The macroeconomic effects of commodity price uncertainty. *mimeo*.
- Zhou, X. (2019). Refining the workhorse oil market model. *Journal of Applied Econometrics*, 35(1):130–140.

Centre for Applied Macroeconomics and Commodity Prices (CAMP)

will bring together economists working on applied macroeconomic issues, with special emphasis on petroleum economics.

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

www.bi.no/camp