BANK OF CANADA
BANQUE DU CANADA

Working Paper/Document de travail
2012-40

The Effects of Oil Price Uncertainty on the Macroeconomy

by Soojin Jo

Bank of Canada Working Paper 2012-40

December 2012

The Effects of Oil Price Uncertainty on the Macroeconomy

by

Soojin Jo

Canadian Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
sjo@bankofcanada.ca

ISSN 1701-9397

Acknowledgements

I am very grateful to James D. Hamilton for his invaluable guidance and support. I would like to thank Davide Debortoli, John Elder, Marjorie Flavin, Nir Jaimovich, Dimitris Politis, Valerie A. Ramey, Giacomo Rondina and Rossen Valkanov for helpful comments. I also thank Christiane Baumeister for providing the Baumeister and Peersman (2012) data series, code and helpful suggestions.

Abstract

This paper investigates the effect of oil price uncertainty on real economic activity using a quarterly VAR with stochastic volatility in mean. Stochastic volatility allows oil price uncertainty to vary separately from changes in the level of oil prices, and thus the impact of oil price uncertainty can be examined in a more flexible yet tractable way. In addition, this paper substantially improves on the recovery of a historical uncertainty series by incorporating an additional uncertainty indicator, i.e., a realized volatility series from daily oil price data, into the estimation process. The estimation results show that an oil price uncertainty shock alone has negative effects on world industrial production.

JEL classification: E32, C32, Q43
Bank classification: Business fluctuations and cycles; Econometric and statistical methods

Résumé

L'auteure analyse les effets de l'incertitude des cours du pétrole sur l'activité économique réelle à l'aide d'un modèle VAR trimestriel dont les équations de moyenne incluent un terme de volatilité stochastique. La volatilité stochastique permet à l'incertitude entourant les prix du pétrole de varier indépendamment de leur niveau; l'incidence de cette incertitude peut ainsi être étudiée dans un cadre plus flexible qui reste maniable. En outre, l'intégration au processus d'estimation d'un indicateur additionnel de l'incertitude, à savoir une série de volatilités réalisées calculées à partir des cours quotidiens de l'or noir, constitue un réel progrès par rapport à l'extraction d'une série temporelle relative à l'incertitude. Les résultats de l'estimation montrent qu'une hausse inattendue de l'incertitude des prix du pétrole a en soi des effets défavorables sur la production industrielle mondiale.

Classification JEL : E32, C32, Q43
Classification de la Banque : Cycles et fluctuations économiques; Méthodes économétriques et statistiques

1 Introduction

The Great Recession of 2007 to 2009 has brought the importance of understanding the propagation mechanism of uncertainty to the forefront. The unprecedentedly slow recovery to which both economists and policy makers identify uncertainty as one potential contributor has reinforced the significance. Nonetheless, most macroeconomic literature in the past have not focused so primarily on time-varying uncertainty, presuming that it would have rather little effects on the macroeconomy. However, several studies have demonstrated that heightened uncertainty can deteriorate a variety of real economic activity such as hiring, investment, and durable consumption. For instance, Bernanke (1983) shows that firms postpone irreversible investment decisions and wait for more information to arrive under high uncertainty, resulting in cyclical fluctuations in the economy. More recently, Bertola et al. (2005) find that uncertainty of income flows widens the size of the inaction band, making durable goods adjustment less frequent. Finally, Bloom et al. (2011) indicate that an uncertainty shock, defined as an unexpected change in the conditional second moment of a productivity innovation process, results in a sharp and rapid economic decline even though the first moment remains unchanged.

Building on this line of literature, this paper investigates how oil price uncertainty affects real economic activity. Since oil is a salient factor for both households' consumption and firms' production decisions, it is conceivable that changes in oil price uncertainty may also have effects on economic fluctuations, in addition to changes in the oil price level. In other words, it is plausible that not only the actual change in the oil price, but also the variability of the future oil price forecast have a significant impact on economic agents' decision-making process.

To assess how oil price uncertainty affects economic activity, I first define oil price uncertainty as the time-varying standard deviation of the one-quarter ahead oil price forecasting error. The standard deviation of the forecasting error controls the size of an unanticipated oil price change, and likewise most of previous studies on uncertainty also define it as the time-varying second moment of a shock process. Hence, the standard deviation constitutes a good proxy of price uncertainty.[1] In particular, I model oil price uncertainty to evolve as a *stochastic volatility* process, and include the time-varying volatility in the mean equations of a three-variable vector autoregression (VAR) model that can directly measure the impact of uncertainty.

Several empirical studies have demonstrated the significance of oil price uncertainty from various perspectives. Lee et al. (1995) were among the first to emphasize the importance of accounting for the second moment of oil prices in forecasting economic activity. The new oil price shock variable proposed in their paper reflects both the size and the variability of the forecast error, and explains GNP growth much better than e.g., real oil price changes or regular forecast error, which implies that the effect of an oil price change in certain size can differ depending on whether it is an unusual event or just another adjustment. Kellogg (2010) tests for the responsiveness of firms' investment decisions to changes in uncertainty using Texas oil well drilling data and expectations of future oil price volatility. The result is in support of the real option as firms reduce their drilling activity when expected volatility rises. Lastly, Elder and Serletis (2010) and Bredin et al. (2010), which are most closely related to this paper in the sense that they also directly measure the impact of oil price uncertainty, use a two-variable Generalized Autoregressive Conditional Heteroskedasticity (GARCH)-in-Mean VAR with oil price and economic activity for the U.S. and G-7 countries. They find that an increase in oil price uncertainty decreases real economic activity, measured by output, investment, and consumption in the U.S. and four of the G-7 countries. They conclude that the 2003-2008 oil price surge has been rather steady and continuous, keeping oil price uncertainty at a very low level. Hence, the overall change in oil price was less disruptive than previous oil price episodes and did not lead to an immediate economic recession.

[1] Hence, I use the term *oil price uncertainty* and *oil price volatility* interchangeably henceforth.

Most of the approaches taken in the previous literature define oil price uncertainty to evolve as a GARCH process, and as a result, changes in volatility are fully determined by changes in the level. In contrast, this paper proposes a novel framework for investigating the dynamic responses to an unanticipated oil price uncertainty shock which is independent of any changes in the price level. This is done by modeling the oil price uncertainty process to follow stochastic volatility which permits the volatility process to have its own innovation term in addition to that of the the first moment. This, in fact, is not possible in a GARCH model where the volatility process is completely dependent upon changes in the level and the volatility of the past. Time-varying oil price uncertainty is then included in the mean equations of the VAR, from which one can have a direct estimate of the effect of uncertainty on real economic activity. Hence, this framework provides a flexible yet tractable way of examining the dynamic effect of an exogenous oil price uncertainty shock independent of any changes in the price level.[2] Furthermore, the proposed methodology is applicable to a number of other cases where the dynamic effects of time-varying uncertainty of a variable on another are examined, for instance, when investigating the effect of exchange rate uncertainty on trade flows.

Another contribution of this paper is that it substantially improves on the recovery of a time series of historical uncertainty. This again is one of the advantages of modeling oil price uncertainty as stochastic volatility, since it offers room to incorporate an additional oil price uncertainty indicator, i.e., realized volatility. This method is based on Dobrev and Szerszen (2010) who find that there is a significant efficiency gain in estimating time-varying stock return volatility if the realized volatility series constructed from high frequency data is augmented with stochastic volatility of lower frequency. Unfortunately, high frequency price data is not available before 1983 in the oil market. However, this paper resolves the issue by using time-varying Kalman filter, which extends their framework to the oil market.

[2] The model can also be modified for the cases where the oil price shock and the oil price uncertainty shock are correlated to each other, and thus, can nest the GARCH setup.

As a result, the paper estimates the historical oil price uncertainty series with substantial improvement in precision, spanning the long period from 1958Q2 to 2008Q3.

The main finding of this paper is that oil price uncertainty affects global real economic activity in a significantly negative way. The impulse response analysis indicates that an oil price uncertainty shock has the immediate and persistent negative effect of lowering economic activity by about 0.1%-point. While the focus of the previous literature rather lies on measuring the impact of price uncertainty at the national level (e.g. Elder and Serletis (2012) and Bredin et al. (2010)), I use a world industrial production series, extending the analysis to a global context. Related to this issue, Kilian and Vigfusson (2010) suggest that it may not be appropriate to define oil price uncertainty as the time-varying conditional standard deviation when analyzing one specific country since only one side of the price distribution is likely to matter depending on its position in the oil market. For instance, an oil-exporting country would only consider the higher chance of a large oil price drop as risk and make a necessary adjustment, but would not respond to the higher chance of an oil price increase. Hence, looking at global level data has the advantage of providing evidence robust to such possibility.[3] The main result is also robust to changes in the oil price series, to sub-periods of the sample, and holds for advanced economies. Overall, the findings presented in this paper emphasize the importance of accounting for oil price uncertainty since ignoring it can distort the effect of a policy designed under the presumption of linearity in the oil price-economic activity relationship.[4]

The rest of the paper is organized as follows: Section 2 introduces a VAR model with stochastic volatility and further presents the augmentation of the model with realized volatil-

[3] However, according to the theoretical literature that defines uncertainty to be the overall dispersion of the distribution in both directions, the conditional standard deviation can still be a proper proxy of uncertainty for a country (see e.g., Bloom et al (2011)).

[4] It should be noted that this paper examines one specific type of uncertainty channel, i.e., the *linear* effect of oil price uncertainty on real economic activity, among the wide variety of non-linear oil-economy relationships found in Mork (1989), Hamilton (2003) and Baumeister and Peersman (2012, forthcoming) among others.

ity; Section 3 presents the empirical results; and Section 4 concludes.

2 Model and Estimation

2.1 A VAR model with time-varying volatility

In order to measure the effect of oil price uncertainty on the economy, I develop a modified version of the VAR model with time-varying stochastic volatilities in the spirit of Cogley and Sargent (2005), Primiceri (2005), and Baumeister and Peersman (2012, forthcoming). The main difference is the inclusion of the term that captures the effect of time-varying oil price uncertainty on the dynamics of economic activity. Another difference of the model lies in the time-invariance of the VAR coefficients while the correlation and variance parameters' time-dependence is preserved. Together, this specification permits to focus on the specific form of non-linearity, namely, the linear effect of the conditional second moment of oil price.[5] Hence, the VAR can be written as,

$$y_t = B_0 + B_1 y_{t-1} + ... + B_p y_{t-p} + \Lambda \log h_t + u_t, \tag{1}$$

where y_t is a 3×1 vector consisting of quarterly world crude oil production, the real price of crude oil and global real economic activity, covering 1958Q2 - 2008Q3.[6] All variables are in first differenced logs multiplied by 100, and represent the quarterly growth rate. The 3×1 vector B_0 is an intercept, B_i for $i = 1, \ldots, p$ are 3×3 coefficient matrices with the number of lags p set at 4 to allow for sufficient dynamics of the system. The reduced-form innovation vector u_t is defined to have the conditional mean zero and the conditional time-varying

[5] Not much variation over time is observed for the coefficient of interest (λ) when the coefficients are permitted to be time-varying, and hence, time-invariant coefficients seem to capture the intended relationship sufficiently well in a parsimonious way.

[6] The real crude oil price is the U.S. refiners' acquisition cost of imported crude oil (IRAC) deflated by the U.S. consumer price index. The real economic activity is measured by the industrial production index series of the global economy. Details on data can be found in the Appendix A.

variance-covariance matrix given by Ω_t such that $\Omega_t = A_t^{-1}\Sigma_t\Sigma_t'(A_t^{-1})'$ where

$$A_t = \begin{bmatrix} 1 & 0 & 0 \\ a_{21,t} & 1 & 0 \\ a_{31,t} & a_{32,t} & 1 \end{bmatrix}, \qquad \Sigma_t = \begin{bmatrix} h_{1,t} & 0 & 0 \\ 0 & h_{2,t} & 0 \\ 0 & 0 & h_{3,t} \end{bmatrix}. \tag{2}$$

By letting the lower triangular elements of A_t and the diagonals of Σ_t be time-varying, any change in the correlations between variables will be captured, reflecting possible structural changes in the oil market.

Let $a_t \equiv [a_{21,t}\ a_{31,t}\ a_{32,t}]'$ and $\log h_t \equiv [\log h_{1,t}\ \log h_{2,t}\ \log h_{3,t}]'$. Then the dynamics of the volatilities are modeled as follows:

$$a_t = a_{t-1} + e_t, \tag{3}$$

$$\log h_t = \mu + \rho \log h_{t-1} + \eta_t, \tag{4}$$

where $e_t \sim N(0, S)$ and $\eta_t \sim N(0, W)$. That is, a_t evolves as a random walk process, and the logarithms of h_t, as a first-order autoregressive process. Hence, h_t falls into the category of stochastic volatility models. Here, ρ is a diagonal matrix with AR(1) coefficients on the diagonal and μ is a 3×1 vector of intercepts. Instead of defining $\log h_t$ as a unit root process as has been done in the previous literature, I let the AR(1) coefficients be determined by data.

In equation (1), the term $\Lambda \log h_t$ captures the effect of time-varying oil price uncertainty on global economic activity. Given that the focus is on oil price uncertainty and the economy, the structure of Λ is set to be,[7]

$$\Lambda = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & \lambda & 0 \end{bmatrix}. \tag{5}$$

In sum, oil price uncertainty is modeled as stochastic volatility, and included in the mean equations of the VAR. Stochastic volatility is one common and flexible way of modeling time-varying volatilities. One popular alternative would be the GARCH model developed by Engle (1982), as in Elder and Serletis (2010). However, the two classes of models differ crucially in that the former has a free driving variable (η_t) in the volatility process and the latter does not. Thus, in the GARCH model, the shock that changes the oil price level is the same shock that increases volatility. In contrast, in the stochastic volatility framework, the volatility can in principle evolve independently of any changes in the level. As a result, stochastic volatility allows to separate the dynamic impact of an exogenous oil price uncertainty shock from the shock to the level of the oil price. This is not possible with the GARCH setup where the variation of the volatility is tied to the changes in level. In sum, the stochastic volatility model offers flexibility, and yet it keeps the model fairly parsimonious and hence computationally tractable.

In matrix form, the VAR can be rewritten as,

$$\begin{aligned} y_t &= B_0 + B_1 y_{t-1} + ... + B_p y_{t-p} + \Lambda \log h_t + u_t \\ &= [X_t \ \log h_t]'[\beta \ \lambda] + A_t^{-1} \Sigma_t \epsilon_t. \end{aligned}$$

where X_t denotes $I_3 \otimes x_t$ with x_t being a vector containing the constant and all four lags

[7] It serves the purpose of this paper sufficiently well to define the Λ matrix as above, but it would be also interesting to look at other elements of Λ in future studies.

of y_t,[8] β denotes the vector of parameters with $3 \times (3p+1)$ elements with each row of the coefficient matrices (B) stacked, and the error term ϵ_t follows a conditional multivariate standard normal distribution.

The conditional error terms of the whole system ϵ_t, e_t, η_t are assumed to follow a Normal distribution and to be uncorrelated to each other given the history up to $t-1$:

$$\begin{pmatrix} \epsilon_t \\ e_t \\ \eta_t \end{pmatrix} \sim N \left(0, \begin{pmatrix} I_3 & 0 & 0 \\ 0 & S & 0 \\ 0 & 0 & W \end{pmatrix} \right).$$

Here, S is a block diagonal matrix as,

$$S \equiv Var(e_t) = \begin{bmatrix} S_1 & 0 \\ 0 & S_2 \end{bmatrix},$$

where $S_1 \equiv Var(e_{21,t})$ and $S_2 \equiv Var([e_{31,t} e_{32,t}]')$, and W is assumed to be a diagonal matrix with the error terms being independent of each other.

2.2 Augmented Model

In this section, I present the benchmark VAR model augmented with an additional oil price uncertainty indicator, realized volatility. The idea of using an extra indicator for oil price uncertainty in addition to the stochastic volatility component of the VAR follows Dobrev and Szerszen (2010). They show that there is a substantial efficiency gain in estimating daily stock market volatility when the information content of realized volatility is added to stochastic volatility. Here, the realized volatility series is constructed from high frequency, i.e., intra-day stock return data. To be more specific, the efficiency gain stems from adding a measurement equation based on the asymptotic distribution of the realized volatility estimator to the

[8] Hence $[X_t \ \log h_t]$ is a (40×3) matrix.

state-space model.

In the context of my paper, where the VAR models quarterly dynamics, high-frequency data would refer to daily oil price changes. Yet the daily oil price data to construct the realized volatility series is not available for about the first half of the sample period. The framework originally designed by Dobrev and Szerszen (2010) is only appropriate when both low and high frequency data are attainable for the full sample period, which is not the case for the oil price. Hence, I extend the framework using a time-varying Kalman filter to account for different sample sizes. In this way, the new indicator derived from the high-frequency oil price data can still improve on the inference of unobservable volatility states substantially, as will be shown in Section 3.

In particular, I construct the following state-space model for volatilities with three measurement and one state equations:

$$\textit{Measurement equations} \qquad y^*_{2,t} = 2 \times \log(h_{2,t}) + \log(\epsilon^2_{2,t}), \tag{6}$$

$$\tilde{y}_{3,t} = \lambda \log(h_{2,t}) + h_{3,t}\epsilon_{3,t}, \tag{7}$$

$$\log(\widehat{RV}_{t,M}) = 2 \times \log(h_{2,t}) + \sqrt{\frac{\nu}{M}\frac{\widehat{IQ}_{t,M}}{\widehat{RV}^2_{t,M}}}\xi_t, \tag{8}$$

$$\textit{State equation} \qquad \log h_{2,t} = \mu_2 + \rho_2 \log h_{2,t-1} + \eta_{2,t}. \tag{9}$$

Equation (6) results from transforming the real oil price ($y_{2,t}$) equation in the VAR based on the mixture Normal treatment of Kim et al. (1998), where $y^*_{2,t} \equiv \log(\{A_t(y_t - [X_t \;\; \log h_t]'[\beta \;\; \lambda])\}^2 + c)$ after squaring and taking logarithms.[9] Equation (7) is the economic activity equation in the VAR after subtracting lags of endogenous variables, i.e., the third equation of $\tilde{y}_t \equiv A_t(y_t - X_t'\beta) = A_t\Lambda\log(h_t) + \Sigma\epsilon_t$, and is included since economic activity is modeled to

[9] A small offset constant c is added to avoid the case that $A_t(y_t - [X_t \;\; \log h_t]'[\beta \;\; \lambda])^2$ is too small and thus a logarithm is not well-defined. It is set to be 0.001.

be affected by oil price uncertainty ($\log h_{2,t}$).

Observation equation (8) is derived from the asymptotic distribution of a general class of consistent realized volatility estimators ($\widehat{RV}_{t,M}$) of the volatility state, $h_{2,t}$, i.e., $\sqrt{M}(\widehat{RV}_{t,M} - h^2_{2,t}) \to^D N(0, \nu{\cdot}IQ_t)$, where M is the number of days in each quarter, ν is a known asymptotic variance factor, and IQ_t is the asymptotic variance of the realized volatility estimator. In particular, I use the jump-robust median realized volatility estimator (MedRV) of Andersen et al. (2009) as $\widehat{RV}_{t,M}$.

Then applying the Delta method yields,

$$\sqrt{M}\frac{log(\widehat{RV}_{t,M}) - log(h_{2,t})^2}{\sqrt{\nu\frac{\widehat{IQ}_{t,M}}{\widehat{RV}^2_{t,M}}}} \to^D N(0,1).$$

where $\widehat{IQ}_{t,M}$ is a consistent estimator of IQ_t.

Approximating the above distribution results in equation (8). In short, the observed realized volatility is considered as a function of the unobserved stochastic volatility, and provides additional information that otherwise would not have been available, and consequently improves on the efficiency.[10]

The key feature due to the incorporation of the realized volatility series is the appearance of the new measurement equation (8). In particular, for the earlier period during which daily data is not available, the state-space model consists of equations (6), (7), and (9), and in the later period, the model also includes equation (8). In this way, it is possible to make use of the extra information content of realized volatility as soon as the high frequency data become available.

To show the improvement in efficiency more clearly, I will estimate both the benchmark VAR with equation (8) and the VAR without equation (8), and present both results for

[10]More detailed explanations about the Kalman filter setup are provided in Appendix B and C along with notes on the estimation algorithm.

comparison; as a brief preview, the coefficients in the VAR exhibit higher precision, not to mention oil price volatility. Together with realized volatility, this paper estimates a reliable world oil price uncertainty series for an extended time period starting from 1958Q2.

2.3 Bayesian estimation strategy

The benchmark VAR model is estimated using Bayesian methods to assess the joint posterior distribution of the parameters of interest, unobserved states and hyperparameters. Bayesian estimation is a natural choice as the model has several state variables appearing non-linearly in the measurement equation. As pointed out by Primiceri (2005, pp.826), classical likelihood methods may not be the optimal choice when high dimensionality and nonlinearity exist in the model due to the danger of having multiple local peaks in the likelihood function. Bayesian methods, on the other hand, deal with this type of problem particularly well by separating the parameter space into several blocks, which simplifies the estimation process to a great extent. Thus the Markov Chain Monte Carlo (MCMC) algorithm, particularly the Gibbs sampling procedure, is applied to draw from a series of conditional posterior distributions of parameter blocks. Furthermore, this algorithm is expedited by imposing conjugate prior distributions. Each step of the algorithm can be found in Appendix C and the details on priors are in Appendix D.

3 Results

3.1 The estimated effect of oil price uncertainty

Figure 1 displays the posterior distribution of λ and column (1) of Table 1 reports the summary statistics for the benchmark model. Most of the posterior draws of λ range below zero; the mean of the coefficient λ is -0.1136, the standard deviation is 0.0514, and 98.89% of λ draws are negative. In fact, the high probability of a negative λ suggests that oil price uncertainty alone can hamper economic activity when increased to a higher level. This result further confirms that not only oil price movements but also changes in oil price uncertainty matter, supporting the non-linearity in the oil price - macroeconomy relationship.

To better understand this result, it is instructive to study a specific episode in 1980s. The mid-1980s is recorded as the first period when the non-linear relationship between oil prices and real economic activity became apparent. At that time, world oil consumption declined as a consequence of the oil crises in the 1970's that led to improved energy efficiency. On the other hand, there was added production particularly from Iran and Iraq to finance their lingering war. In an effort to keep the prices from falling further, Saudi Arabia cut back its production during the early 1980s; however, in 1986Q1, Saudi Arabia reversed its decision and started pumping more oil. Consequently, the over-production and the reduced demand for oil resulted in a huge price drop; the real price of oil fell from $24.5 in 1985Q4 to $17.4 in 1986Q1.[11] This unexpected oil price decrease resulted in a drastic jump of oil price uncertainty by 100%, i.e., from 14.11%-points to 31.79%-points, according to the median of the posterior draws of oil price volatility.

While the mid-1980s price collapse would have been expected to stimulate oil consumption, and thus, output, the global economy in general did not experience the anticipated expansion. Based on the empirical analysis of this paper, one likely reason for the moderate growth was the unusually rapid and severe increase in oil price uncertainty.

Figure 2 shows the histogram of possible realizations of economic growth in 1986Q1 implied by the posterior distribution of λ, assuming that the oil price uncertainty had remained at the same median level as in 1985Q4, i.e., 14.11%-points instead of 31.79%-points. Had this been the case, the model would have predicted the median quarterly global industrial production growth rate to be 0.9348%, which is about 0.1%-point higher than the predicted rate with the uncertainty surge, 0.8443%.[12]

[11] For a more detailed history of 1980's, see Downey (2009)(pp.19) and Hamilton (2011)(pp.17-18).
[12] The actual growth rate of global industrial production in 1986Q1 was 0.7788%.

3.2 The importance of including realized volatility

Next, I re-estimate the VAR, but without the extra information from the additional uncertainty indicator to illustrate clearly the contribution of augmenting the VAR by realized volatility. That is, equation (8) is omitted from the state-space model for the entire sample period to compare the effect of including the realized volatility series.[13] Figure 3 plots the times series of oil price uncertainty obtained (i) from the benchmark VAR estimation and (ii) from the condensed model without equation (8) with 95% error bands. The difference in the size of the error bands is striking. In the upper panel where two oil price uncertainty indicators are used, the error band becomes much narrower from 1983 onwards when additional information content becomes available.[14] By contrast, the lower panel exhibits much larger error bands for the whole sample period, though the median (solid line) does not differ very much. Therefore, augmenting the model by including realized volatility substantially improves the inference of the posterior distribution of oil price uncertainty.

Next, column (2) of Table 1 presents summary statistics of the posterior distribution of λ when equation (8) is omitted. The point estimate (-0.1995) implies a negative impact of oil price uncertainty, which is consistent with the baseline result, although the mean is larger in size (in absolute value). However, it is estimated with less precision, i.e., the standard deviation is 0.0709, compared to 0.0514 in the baseline case. Thus, another benefit of having the additional information from the high frequency data is that the posterior distribution of λ has become more centered.

Figure 4 shows the median of the oil price uncertainty time series recovered from the benchmark model. Oil price uncertainty jumps up during the periods of the first oil shock

[13] As noted above, the new measurement equation (8) derived from realized volatility plays a role of being an additional uncertainty indicator only when drawing oil price volatility state in the Gibbs sampler algorithm. Hence, it is perfectly feasible to estimate the VAR without having the equation; one can simply omit equation (8) and run Kalman filter with equations (6), (7), and (9) for the entire sample period.

[14] In addition, one can also observe the improvement in precision in the period before than 1983, as a result of the smoothed draws from the forward and backward Kalman filter recursions.

and Iranian revolution followed by Iran-Iraq War. Next, as seen in the above illustration, uncertainty doubles in 1986Q1, accompanied by the modest growth rate in industrial production. Furthermore, the last two quarters in 1991 exhibit the highest oil price uncertainty, which coincide with the First Persian Gulf War. Price volatility in general remains elevated in the 1990's and 2000's, peaking up during the historical episodes such as East Asian Crisis and the Second Persian Gulf War.

3.3 Dynamic effects of an oil price uncertainty shock

Suppose oil price uncertainty increases unexpectedly and substantially without affecting the actual price series. This may be the case when, for example, economic agents fear a much higher oil demand and/or oil depletion in the future that have not yet led to any noticeable changes in the current oil market. Thus there is no significant variation in the first moment, but the underlying distribution of the oil price has become more dispersed and hence, uncertain. Broadly speaking, the uncertainty shock shares some similarity with the *oil-specific demand shock* in Kilian (2009) and the *speculative demand shock* in Kilian and Murphy (2010) in the sense that it does not reflect any current changes in fundamentals, e.g., oil supply disruptions due to geopolitical turmoils and/or demand shifts due to economic expansions, but are related to expectations about the future. However, while the above shocks triggers an increase in the oil price, the uncertainty shock refers to a more exogenous change in uncertainty without implying any variation in the actual price level. It can also reflect cases of small market disruptions which are not associated with any distinct movements in the oil production and/or price. In sum, the uncertainty shock mainly describes unexpectedly heightened chance of facing an extreme price change.

For example, during the first half of 2003 that lies in between the two red vertical lines in Figure 5, the Middle East has undergone the strikes in Venezuela and the Second Persian Gulf War. Although both oil production and oil price were only modestly affected for a relatively short period of time compared to previous unrests, the uncertainty measure indicates that these episodes were associated with a doubling of oil price uncertainty, i.e., from 12.18%-points (2002Q4) to 25.36%-points (2003Q1). This period is a good example of the uncertainty measure moving differently from the oil price level, which can be easily modeled using the stochastic volatility, emphasizing the advantage of the framework developed in this paper. At this time, the global industrial index exhibits a short-lived dip (2003Q2) for which oil price uncertainty may have played an important role. Hence, I further investigate impulse responses to the oil price uncertainty shock.

The impulse responses implied by this kind of uncertainty shock have not been explored so far in the related literature, mainly because the statistical models in the previous papers are based on the GARCH framework. In that case, uncertainty cannot change alone without assuming specific variations in the level of oil prices. In fact, GARCH models by definition do not allow any free driving variable in the volatility generating process and thus, one has to specify changes in the level first, which consequently increase volatility, in order to study the effect of an uncertainty shock.

In contrast, stochastic volatility enables the investigation of the consequences of an unanticipated oil price uncertainty increase itself on economic activity, independent of any other change in the oil price level. This is due to the volatility generating process which has its own free parameter (η_t in equation (5)) that permits exogenous innovations to uncertainty.

The specifics of this exercise are as follows: I generate a one-time oil price uncertainty increase by 100 percent. Since oil price uncertainty is included in the mean equation of the VAR in logs, this is equivalent to having the log oil price uncertainty increase by *one unit*. As an illustration, the doubling of uncertainty is comparable to what happened during the first oil shock in 1973Q4 and 1986Q1, according to the posterior draws of volatility. This shock will be highly persistent over time, as reflected in the posterior distribution of ρ_2,[15] and the

[15] The point estimate of the oil price volatility AR(1) coefficient ρ_2 is 0.9590 (with the standard deviation 0.0221), confirming the prior belief that oil price volatility follows a process close to unit root. In case of oil

uncertainty series comes back to its normal path very slowly. Again, it is worth noting that it is not necessary to consider any innovations in the oil price level as the uncertainty shock does not imply any first moment movements.

Figure 6 shows the median impulse responses to a 100-percent increase in oil price uncertainty over a 12-quarter horizon. As oil price uncertainty is negatively correlated with economic activity, a shock which unexpectedly increases oil price uncertainty would result in a drop in the industrial production growth rate. Confirming this view, a doubling of oil price uncertainty yields an immediate drop of approximately 0.11%-point in the global industrial production growth rate in the same quarter. In other words, the exogenous doubling of oil price uncertainty alone can decrease real economic activity growth by almost 0.4%-point annually. This negative response remains very persistent over the 12-quarter horizon due to the close-to-unit root characteristics of the oil price volatility process, although the oil price level responds little. Finally, oil production decreases slightly along with the drop in industrial production through the oil price uncertainty channel.

3.4 Sensitivity Check

Structural break in the oil market

In this section, I first look at whether the *structural break in the oil market* detected in the mid-1980's would change the way oil price uncertainty affects real economic activity. A number of recent studies asserts that oil price shocks affect economic activity far less than they did in the past, e.g., Blanchard and Galí (2007), although no consensus on the existence of a break has emerged so far(e.g., Ramey and Vine (2010)). If there was a structural break that the model did not account for, then the negative oil price uncertainty coefficient might have resulted from averaging out the weakened price effects throughout the whole sample

---

production volatility, the posterior distribution of ρ_1 is much less persistent with mean 0.5287 and a standard deviation of 0.1670. Finally, the point estimate of the economic activity volatility AR(1) coefficient is 0.2807 and the standard deviation is 0.0988.

period. Or, it is also possible that the structural change may have been so fundamental that the oil price uncertainty effect might also have been moderated after the structural break as did the oil price level impact.

Hence, in order to control for possible changes in the oil market, I first add an indicator variable for 1984Q1 in the right hand side of the main VAR, which is the quarter with the structural break presented in Blanchard and Galí (2007). Second, I split the sample period into two, one until 1983Q4, the other from 1984Q1 onwards, and repeat the VAR estimation for different sample periods.[16]

Tables 2 and 3 report the summary statistics of the posterior λ draws of the above specifications. All three posterior distributions of λ show high probabilities of λ being negative, and furthermore, the means are in line with the baseline result. Although standard deviations, particularly during the first subsample period, are generally larger due to the smaller number of observations, it is apparent that oil price uncertainty has had consistent negative effects on world economic activity throughout the whole sample period.

Effects on advanced economies

Second, I use the industrial production series of *advanced economies* (1957Q1-2010Q1) instead of that of the global economy. The first column of Table 4 reports the summary statistics of the posterior λ draws. The point estimate, -0.1154, is very similar to that of the global economy, and the posterior distribution is in a similar range with a very high chance that λ is negative (95.09%). This result is somewhat predicted considering the size of advanced economies in world economy. It also confirms and extends the results in Elder and Serletis (2010) and Bredin, Elder and Fountas (2010) obtained using the U.S and G-7 countries' various real economic activity data, since all of the countries examined in their

[16]For this exercise, the main VAR is run without including equation (8) of realized volatility. The reason for excluding realized volatility is that the realized volatility series starts only from 1983Q1, and thus only four observations are available for the first subsample. Hence, to treat the two subsamples equally, I exclude equation (8) for both periods.

papers are advanced economies.

In the second column of Table 4, I present the advanced economies' result without incorporating realized oil price volatility data, which is comparable to the result in column (2) in Table 1. This is to confirm that the efficiency gain from including realized volatility can be found consistently across different country coverages. The point estimate of λ shows weaker association, but is well within the range. More importantly, the posterior distribution appears to be more dispersed with lower precision without realized volatility, similar to the global economy's case. This highlights the efficiency gain achieved by the additional oil price uncertainty indicator. Moreover, the realized volatility estimator also helps recover more reliable historical oil price uncertainty series as it does for global economy, though the posterior distributions are not presented here.

Changes in oil price series

Next, I rerun the estimation procedure using *different oil price series* than the Imported Refiners' Acquisition Cost (IRAC) of crude oil as the world oil price series. Columns (1) and (3) of Table 5 report the benchmark model estimation results obtained from using West Texas Intermediate (WTI) as a world oil price series. Affirming the baseline result, oil price uncertainty reduces economic growth both in global and advanced economies in all estimations. Furthermore, comparing columns (1) and (2), and (3) and (4), one finds the extra information obtained from daily oil price series again helps obtaining more centered posterior distributions of λ with the smaller standard deviations, although the efficiency gain is not as evident as the cases when IRAC is used in the VAR.

Measurement error of the realized volatility estimator

Finally, I re-estimate the benchmark VAR, *multiplying a constant* m to the theory-predicted measurement error (i.e., $\sqrt{\nu \widehat{IQ}_{t,M} / M \widehat{RV}^2_{t,M}}$) in equation (8), giving data a chance to determine the actual size of the error. In other words, since the new measurement equation (8) is based on the approximation of the asymptotic distribution, the drastic improvement in oil price uncertainty estimation might have potentially come from forcing the measurement error to be infinitesimal although it may not be the case in a small sample. Hence, equation (8) is substituted by

$$\log(\widehat{RV}_{t,M}) \quad = \quad 2 \times \log(h_{2,t}) + \mathbf{m} \times \sqrt{\frac{\nu}{M} \frac{\widehat{IQ}_{t,M}}{\widehat{RV}^2_{t,M}}} \xi_t.$$

If the approximation of the generic asymptotic distribution is not imposing too strong a restriction, m should be arbitrarily close to one as the data sampling frequency becomes infinitesimally small. Table 6 presents the summary statistics of m, and the posterior draws have the mean 1.18 and the standard deviation 0.59, confirming the model's robustness.

In sum, the baseline result that increases in oil price uncertainty are detrimental to economic activity is consistent even when potential structural changes in the oil market are accounted for. Moreover, it is also applicable to the advanced economies' case, and is insensitive to different combinations of oil price series. Furthermore, the statistical improvement achieved by having realized volatility conforms to the baseline result, consistently shown across different data sets. In addition, the improvement does not seem to be forcefully driven by the small measurement error predicted by the asymptotic distribution of the realized volatility estimator.

4 Conclusion

This paper investigates how oil price uncertainty affects global real economic activity during 1958Q2 - 2008Q3 using a VAR model with time-varying stochastic volatility in mean; oil price uncertainty has significant negative effects on real economic activity measured by the industrial production index, extending the previous studies' findings at the national level (e.g., Elder and Serletis (2010) and Bredin et al. (2010)). In contrast to the earlier studies, this paper shows that high oil price uncertainty *alone* can significantly reduce the industrial production growth, independent of actual price level changes. This result is robust to the use of the different oil price series, when considering the possible structural break in mid-1980's, and holds for advanced economies as well.

The main contribution of this paper is to develop a framework that can separate the first and second moment effects. This is done by modeling oil price volatility as stochastic volatility and including it directly in the mean equations of the VAR. Stochastic volatility by definition allows an innovation term for volatility itself, and thus, makes it possible to examine the effect of volatility more flexibly without having to presume any changes in the level. This is not feasible in the approach using a GARCH model as in many of previous studies since the GARCH process models changes in volatility to be completely determined by changes in the level. The proposed framework with stochastic volatility can also be applied to a number of other cases to examine the dynamic effects of time-varying uncertainty, such as the investigation of the effect of exchange rate uncertainty on trade flows, the relationship between inflation volatility and output, and so on.

Another methodological contribution of this paper is to extend the framework by Dobrev and Szerszen (2010) who find that there is a significant efficiency gain in estimating time-varying volatility when the realized volatility series constructed from high frequency data is augmented with stochastic volatility of lower frequency. This paper extends their framework to be applicable to the case where high frequency data is available only for a limited sub-period of the sample by implementing a time-varying Kalman filter. As a result, the historical oil price uncertainty series is estimated with substantial improvement in precision, and the overall statistical inference becomes more significant. This improvement is observed consistently in case of advanced economies, and robust to the use of the different oil price series.

With the novel framework, this paper explores another channel of non-linearity in the propagation mechanism of the oil price shocks with the specific focus on the effect of oil price uncertainty. Overall, the findings of this paper highlights the importance of tracking and accounting for oil price uncertainty in the oil price-economic activity relationship.

A Data

I use the real U.S. refiners' acquisition cost of imported crude oil (IRAC) as in Baumeister and Peersman (2012) as a measure of the world oil price series. This nominal series is a volume-weighted average price of all crude oils imported into the United States over a specified period, and was provided by Baumeister. Since the U.S. imports more types of crude oil than any other country, this series is often regarded as a good proxy of the world crude oil price. One caveat of this series for this paper is that it is provided at a monthly frequency at most. This means that the same series cannot be used when constructing the realized volatility series that requires daily price data. Thus, daily WTI is used to estimate realized volatility of each quarter starting from 1983. The IRAC is originally taken from the Department of Energy (DoE) and deflated by the U.S. CPI, the WTI series is retrieved from the Global Financial Database after adjusting for inflation, and hence, the real price of crude oil is obtained.

With respect to world economic activity, I use the world index of industrial production series spanning from 1947Q1 to 2008Q3, of which the first 40 observations are used as a training period to construct prior distributions.[17] This index covers global industrial activities in mining and quarrying, manufacturing and electricity, gas and water supply. The advanced economies' industrial production index series used later in the robustness check section is taken from International Financial Statistics of International Monetary Funds from 1957Q1 to 2010Q1.[18] Again, the first 40 observations are used as training period.

[17] This data series is also provided by Christiane Baumeister and is used in Baumeister and Peersman (forthcoming) The original source of world index of industrial production is the United Nations Monthly Bulletin of Statistics, from which a coherent series is constructed by Baumeister by re-weighing and seasonally adjusting the raw data. The series can be obtained by contacting Baumeister at CBaumeister@bank-banque-canada.ca. For more detailed explanations on each series, see the not-for-publication Appendix of Baumeister and Peersman.

[18] The coverage of this index is similar to that of world industrial production index, i.e., the index comprises mining and quarrying, manufacturing and electricity, and gas and water, according to the UN international Standard Industrial Classification (ISIC) and is compiled using the Laspeyres formula. The list of "advanced countries" are : Australia, Austria, Belgium, Canada, Cyprus, Czech Republic, Denmark, Finland, France, Germany, Greece, Hong Kong SAR, Iceland, Ireland, Israel, Italy, Japan, Luxembourg, Malta, Netherlands,

The world oil production data provided by Baumeister is originally obtained from the DoE starting January 1973, from the Oil & Gas Journal for the period April 1953 to December 1972, and the earlier series from January 1947 is interpolated from yearly oil production data.[19] Later for the sensitivity check with advanced economies, I add quarterly observations during the period October 2008 to March 2010 from the DoE. Quarterly data are averages of monthly observations.

B Realized Volatility

As shown in Dobrev and Szerszen (2010), there is a high efficiency gain when the information content from high-frequency data is additionally used in estimating the oil price volatility state. In the context of my paper, high-frequency data refers to daily oil price changes compared to the quarterly price series used in the main VAR. However, since a daily oil price price series is not available for the entire sample period, I use a time-varying state space model to solve the problem of the data shortage in the early period.

I begin by estimating realized volatility of each quarter by using the jump-robust median realized volatility estimator (MedRV) of Andersen et al. (2009) for the period when the daily price variation is observable as:

$$\begin{aligned}\widehat{RV}_{t,M} &= MedRV_{t,M} \\ &= \frac{\pi}{6-4\sqrt{3}+\pi}\left(\frac{M}{M-2}\right)\sum_{i=2}^{M-1} med(|\Delta OP_{i-1}|, |\Delta OP_i|, |\Delta OP_{i+1}|)^2.\end{aligned}$$

Here, M is the number of days in each quarter that oil is traded, and ΔOP_i denotes the

New Zealand, Norway, Portugal, Singapore, Slovak Republic, Slovenia, South Korea, Spain, Sweden, Switzerland, Taiwan Province of China, United Kingdom and the United States. I apply the X-12-ARIMA of the U.S Census to adjust for the seasonality of the series.

[19] As noted by Baumeister and Peersman (2012), the use of interpolated data in the earlier periods is of minor importance, as this part of the data set is mostly used in the training sample to construct priors that is dominated in the algorithm fairly quickly by the data.

observed daily change of the logarithmic oil price in quarter t with $i = 1, 2, ..., M$. MedRV is a consistent estimator of the population variance which is equivalent to $h^2_{2,t}$ in this paper.[20] Hence, MedRV can become an extra oil price uncertainty indicator.

Then the central limit theorem implies the following generic asymptotic results:

$$\sqrt{M}(\widehat{RV}_{t,M} - h^2_{2,t}) \rightarrow^D N(0, \nu \cdot IQ_t)$$

where ν is a known asymptotic variance factor, i.e., 2.96 for MedRV, and IQ_t is the integrated quarticity controlling the precision of the realized volatility estimators. By applying the Delta method with a consistent jump-robust estimator $\widehat{IQ}_{t,M}$ of IQ_t,[21] we get

$$\sqrt{M}\frac{log(\widehat{RV}_{t,M}) - log(h_{2,t})^2}{\sqrt{\nu\frac{\widehat{IQ}_{t,M}}{\widehat{RV}^2_{t,M}}}} \rightarrow^D N(0, 1).$$

When obtaining the above asymptotic results, the log transformation is conducted which makes it very simple to augment the main VAR with realized volatility. With this result, I write a new measurement equation (i.e., equation (8) in Section 2.2) which is added into the step in the Gibbs sampler algorithm to estimate oil price uncertainty :

$$log(\widehat{RV}_{t,M}) = log(h^2_{2,t}) + \sqrt{\frac{\nu}{M}\frac{\widehat{IQ}_{t,M}}{\widehat{RV}^2_{t,M}}}\xi_t,$$

where ξ_t follows the standard Normal distribution and is independent of the underlying process. Note that $log(\widehat{RV}_{t,M})$ and $\sqrt{\nu\widehat{IQ}_{t,M}/M\widehat{RV}^2_{t,M}}$ can be readily computed provided

[20] See Huang and Tauchen (2005); Dobrev and Szerszen (2010) for more detail.

[21] I estimate IQ_t by the median realized quarticity estimator (MedRQ) for each quarter such that

$$\begin{aligned}\widehat{IQ}_{t,M} &= MedRQ_{t,M} \\ &= \frac{3\pi M}{9\pi + 72 - 52\sqrt{3}}\left(\frac{M}{M-2}\right)\sum_{i=2}^{M-1} med(|\Delta OP_{i-1}|, |\Delta OP_i|, |\Delta OP_{i+1}|)^4.\end{aligned}$$

a high frequency data series. Therefore, in addition to the measurement equations (6) and (7), I use the above measurement equation for the period during which daily oil price data is observable. In sum, a time-varying state-space model is constructed to generate smoothed draws of volatility states. A more detailed description on how to obtain draws from the posterior distribution of $\{h_{2,t}\}_{t=1}^{T}$ inside the Gibbs sampler algorithm is provided in Appendix C.[22]

C Gibbs sampler algorithm

When it is not feasible to analytically derive the joint posterior distribution due to the model's high dimensionality and non-linearity, the Gibbs sampler algorithm provides a computationally tractable way of simulating the posterior distributions. The Gibbs sampler algorithm repeatedly draws the parameters after separating them into several blocks whose conditional posterior distributions are known, and after iterating the chains for a sufficiently long time, the draws will be equivalent to those from the joint posterior distribution. The first 30,000 "burn-in" draws are discarded in order to eliminate the possible impact of initial values and to ensure the chain mixes well. Then the following 15,000 draws are collected, and thus, 45,000 iterations in total are conducted for each analysis.

Step 1: Drawing coefficients of lags (β) and of uncertainty (λ)

Given a^T, h^T, y^T and other hyperparameters, this step is equivalent to regressing y_t on $[X_t \;\; \log h_t]$. This step is further expedited with the independent conjugate Normal priors, whose parameters are based on the homoskedastic OLS regression result of the training period. The error term of this step is heteroskedastic; however, given all the values of a^T and h^T, the error covariance and variance matrices are completely known in this step as they

[22] I use MedRV and MedRQ estimators by Andersen et al. (2009), but the result does not change when different realized volatility estimators are used.

consist Ω^T.

Step 2: Drawing covariance states (a) and hyperparameter S

Conditional on β, λ, h^T, Y^T and other hyperparameters, the following provides the measurement equations for covariance states a^T with the observable heteroskedastic innovation u_t :

$$A_t(y_t - [X_t \ \log h_t]'[\beta \ \lambda]) = A_t u_t = \Sigma_t \epsilon_t.$$

This transforms the problem to a standard linear Gaussian state space model with state equation (3) shown in Section 2.1. Now, the elements of a^T are divided into two sub-groups, $a_1 \equiv \{a_{21,t}\}_{t=1}^T$ and $a_2 \equiv \{a_{31,t}, a_{32,t}\}_{t=1}^T$, and drawn in turn from the following two state-space models, i.e.,:

$$\begin{aligned} a_{21,t} &= a_{21,t-1} + e_{21,t} & e_{21,t} &\sim N(0, S_1) \\ u_{2,t} &= -u_{1,t}a_{21,t} + \epsilon_{2,t} & \epsilon_{2,t} &\sim N(0, h_{2,t}^2) \end{aligned}$$

and

$$\begin{aligned} \begin{pmatrix} a_{31,t} \\ a_{32,t} \end{pmatrix} &= \begin{pmatrix} a_{31,t-1} \\ a_{32,t-1} \end{pmatrix} + \begin{pmatrix} e_{31,t} \\ e_{32,t} \end{pmatrix} & \begin{pmatrix} e_{31,t} \\ e_{32,t} \end{pmatrix} &\sim N(0, S_2) \\ u_{3,t} &= -u_{1,t}a_{31,t} - u_{2,t}a_{32,t} + \epsilon_{3,t} & \epsilon_{3,t} &\sim N(0, h_{3,t}^2). \end{aligned}$$

After the transformation, one can now obtain the conditional mean $a_{i,t|t-1}$ and the conditional variance $P_{i,t|t-1}$ using the standard forward Kalman filter method up to T for each sub-group of a. Then the backward recursion is conducted which considers the information content of the entire sample in order to collect smoothed draws. Note that the last iteration of the forward Kalman filter yields $a_{i,T|T}$ and $P_{i,T|T}$. One can draw a last state, $a_{i,T|T}$, from $N(a_{i,T|T}, P_{i,T|T})$, which then feeds back into the backward recursion algorithm to get $a_{i,T-1|T}$ and $P_{i,T-1|T}$. This is repeated until $t = 1$ by updating the conditional mean and the conditional variance as following[23] :

$$\begin{aligned} a_{i,t|t+1} &= a_{i,t|t} + P_{i,t|t} P^{-1}_{i,t+1|t} (a_{i,t+1} - a_{i,t}), \\ P_{i,t|t+1} &= P_{i,t|t} - P_{i,t|t} P^{-1}_{i,t+1|t} P_{i,t|t}. \end{aligned}$$

After drawing a^T, the elements of hyperparmeter S are sampled from the inverse-Wishart posterior distributions by updating the innovations which are observable given the new a^T draws.

Step 3: Drawing volatility states (h) and hyperparameter W

Conditional on other parameter values, drawing the volatility state h^T becomes a non-linear and non-Gaussian state-space problem. Since h_t is modeled to follow a log Normal distribution, it is not possible to use the standard linear state-space model applied in the previous step. Moreover, $\log h_{2,t}$ appearing in the third mean equation of the main VAR multiplied by λ complicates this stage further. Therefore, I apply a log transformation to linearize the system and then use the mixture Normal treatment by Kim et al. (1998). In particular, the entire history of each element of the vector h_i^T is drawn one after another starting from h_1^T. After h_1^T is sampled, sampling h_2^T is in order, which requires the use of a time-varying Kalman filter and smoother depending on the data availability of the high-frequency oil price series. Finally, h_3^T can be drawn based on the updated value of h_2^T. This procedure requires the variance covariance matrix W to be diagonal since we implicitly disregard the possible effect of a correlation between h_i^T's.

The procedure in common for h_1^T and h_2^T starts by obtaining the first two elements of

[23]See Carter and Kohn (1994) for more details on the use of Gibbs sampler for a state space model.

the orthogonalized innovation, i.e., $A_t(y_t - [X_t \;\; \log h_t]'[\beta \;\; \lambda]) = \Sigma_t \epsilon_t$. Conditional on a^T draws from the previous step and all other values, the elements are now observable. Then to linearize the equations, I take logarithms after squaring both sides and adding an offset constant c to the left hand sides.[24] Then, the following state-space models are obtained for h_1^T and h_2^T, respectively, i.e., :

$$\begin{aligned} \log h_{1,t+1} &= \mu_1 + \rho_1 \log h_{1,t} + \eta_{1,t+1} \\ y^*_{1,t} &= 2 \log h_{1,t} + \log(\epsilon^2_{1,t}) \end{aligned}$$

and

$$\begin{aligned} \log h_{2,t+1} &= \mu_2 + \rho_2 \log h_{2,t} + \eta_{2,t+1} \\ y^*_{2,t} &= 2 \log h_{2,t} + \log(\epsilon^2_{2,t}) \\ \tilde{y}_{3,t} &= \lambda \log(h_{2,t}) + h_{3,t}\epsilon_{3,t} \end{aligned}$$

where $y^*_{i,t}$ is the first two elements of the vector $\log(\{A_t(y_t - [X_t \;\; \log h_t]'[\beta \;\; \lambda])\}^2 + c)$ after the transformations. Next, economic activity equation is included in the Kalman filter step after subtracting the effect of lags of the endogenous variables from $y_{3,t}$, i.e., the third equation of $\tilde{y}_t \equiv A_t(y_t - X_t'\beta) = A_t \Lambda \log(h_t) + \Sigma \epsilon_t$, since oil price uncertainty ($h_{2,t}$) appears in the mean equation of $y_{3,t}$. The number of measurement equation increases to three during the period when the daily oil price series is available by adding equation (8) as described in detail in Section 2.2 and Appendix B.

The above linear system is still non-Gaussian as the distribution of $\log(\epsilon^2_{i,t})$ follows $\log \chi^2(1)$, and thus approximated by mixing seven different Normal distributions as Kim

[24] An offset constant is added since squared value of the right hand side can be infinitesimal, and thus log transformation may not be well defined. Following the previous literature, I set c to be 0.001.

Table A1: Seven Normal distributions for $\log(\epsilon_{i,t}^2)$

s	$q_j = Pr(s = j)$	m_j	v_j^2
1	0.00730	-10.12999	5.79596
2	0.10556	-3.97281	2.61369
3	0.00002	-8.56686	5.17950
4	0.04395	2.77785	0.16735
5	0.34001	0.61942	0.64009
6	0.24566	1.79518	0.34023
7	0.25750	-1.08819	1.26261

et al. (1998) and Primiceri (2005). In doing so, an indicator variable $s_{i,t}$ is assigned for each time period t, which determines the particular Normal distribution by which the distribution of $\log(\epsilon_{i,t}^2)$ is approximated. More specifically, to start the Gibbs sampling algorithm, one first assigns any number from 1 to 7 randomly to $s_{i,t}$ for all i and t for the initial iteration of the Gibbs sampler algorithm, and applies the forward and backward Kalman filter (as demonstrated in Step 2) to obtain the first set set of the h_1^T and h_2^T draws. Then with new sets of h_1^T and h_2^T draws, one can update the indicator variables $s_{1,t}$ and $s_{2,t}$. In particular, each $s_{i,t}$ is independently decided from the discrete density function defined as

$$Pr(s_{i,t} = j | y_{i,t}^*, h_{i,t}) \propto q_j f_N(y_{i,t}^* | \, 2h_{i,t} + m_j - 1.2704, v_j^2), \qquad j = 1, \ldots, 7, \;\; i = 1, 2. \quad (10)$$

where the means $(m_j - 1.2704)$ and the variances (v_j^2) of the seven Normal distributions are given in Table A1 (see Kim, Shephard and Chib (1998) and the Appendix in Primiceri (2005)).

With the updated value of h_2^T, we are ready to draw the series of h_3^T, since the third element of $A_t(y_t - [X_t \;\; \log h_t]'[\beta \;\; \lambda])$ is now observable. Denote this element as $\hat{y}_{3,t}$. Then, the state-space model becomes as follows :

$$\begin{aligned} \log h_{3,t} &= \mu_3 + \rho_3 \log h_{3,t-1} + \eta_{3,t} \\ y_{3,t}^* &= 2 \log h_{3,t} + \log(\epsilon_{3,t}^2). \end{aligned}$$

where $y^*_{3,t} = \log((\hat{y}_{3,t})^2 + c)$. This step is also completed by sampling the new values of $s_{3,t}$ according to (10), which will be used in approximating the distribution of $\zeta_{3,t}$ in the next iteration.

Finally, the diagonal elements of the hyperparameter W are drawn one at a time by updating the differences using the new set of h^T draws, as each element of the matrix is considered to be distributed following the inverse Gamma distribution, and the innovation is perfectly observable.

D Prior distributions

The conditional priors for the VAR coefficients, time-varying covariance and log standard deviations are assumed to follow Normal distributions. When obtaining parameter values that define the prior distributions, the first 40 observations are used as a training period, i.e., 1947Q2 - 1958Q1. In particular, an OLS regression is run assuming a time-invariant error structure, from which the resulting point estimates are used to construct prior distributions. When running the OLS regression for the training period, I use the logarithms of 5-quarter rolling standard deviation as a proxy for time-varying uncertainty in order to have an estimate of λ, since it will cause a multicollinearity problem if a constant term and homoskedastic volatility are included at the same time on the left hand side of the VAR.

The mean of $[\beta\ \lambda]$ prior comes from the OLS estimates and the variance-covariance matrix of the prior is obtained by multiplying a constant, 4, to the variance of OLS coefficient estimates, following the specification of Baumeister and Peersman (2012). With respect to the prior of α and $\log h$, I follow the previous literature (i.e., Primiceri (2005), Benati and Mumtaz (2007), and Baumeister and Peersman) by applying the Cholesky decomposition to the variance-covariance matrix and using the diagonal and the lower triangular elements after standardization.

Hyperparameters S and W, which govern the variability of α and $\log h$, respectively, follow Inverse Wishart and Inverse Gamma distributions that belong to conjugate prior family. The prior distributions of μ and ρ reflect the belief that the log of volatility is so persistent that the process is close to a random walk, which would be one way of reflecting the modeling conventions of the previous literature but still giving a chance to the data to determine the posterior distributions.

Later, I multiply a constant m to $\sqrt{\nu\ \widehat{IQ}_{t,M}/M\ \widehat{RV}^2_{t,M}}$ in equation (8) for a sensitivity check to determine how large the measurement error of the realized volatility actually is relative to what is predicted by the generic asymptotic distribution. When doing so, I impose the inverse-Gamma(2,2) prior with prior average of 1, reflecting the belief that the measurement error of the sample will be indeed close to what the asymptotic distribution predicts on average.

Table A2 summarizes all of the prior distributions used in the estimation.

Table A2: Prior distributions

$[\beta\ \lambda]$	$N([\hat{\beta}_{OLS}, [\hat{\lambda}_{OLS}], 4 \cdot V([\hat{\beta}_{OLS}, \hat{\lambda}_{OLS}]))$
a	$N(\hat{a}_{OLS}, 10 \cdot V(\hat{a}_{OLS}))$
$logh_0$	$N(\log \hat{h}_{OLS}, \log \hat{h}_{OLS}/\hat{h}^2_{OLS})$
$[\mu_i\ \rho_i]$	N([0 1], $0.05 \cdot I_2$)
S_1	$IW(V(\hat{a}_{1,OLS}), 2)$
S_2	$IW(V(\hat{a}_{2,OLS}), 3)$
w	$IG((0.01)^2/2, 1/2)$
m	$IG(2, 2)$

References

T.G. Andersen, D. Dobrev, and E. Schaumburg. Jump-robust volatility estimation using nearest neighbor truncation. *NBER Working Paper*, 2009.

C. Baumeister and G. Peersman. Time-varying effects of oil supply shocks on the US economy. *Bank of Canada Working Papers*, 2012.

C. Baumeister and G. Peersman. The role of time-vayring price elasticities in accounting for volatility changes in the crude oil market. *Journal of Applied Econometrics*, forthcoming.

L. Benati and H. Mumtaz. US evolving macroeconomic dynamics: a structural investigation. *Working Paper, European Central Bank*, (746), 2007.

B.S. Bernanke. Irreversibility, uncertainty, and cyclical investment. *The Quarterly Journal of Economics*, 98(1):85–106, 1983.

G. Bertola, L. Guiso, and L. Pistaferri. Uncertainty and consumer durables adjustment. *Review of Economic Studies*, 72(4):973–1007, 2005.

O.J. Blanchard and J. Galí. The macroeconomic effects of oil shocks: Why are the 2000s so different from the 1970s? *NBER Working Paper*, 2007.

N. Bloom, M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. Terry. Really uncertain business cycles. *Working Paper, Stanford University*, 2011.

D. Bredin, J. Elder, and S. Fountas. Oil volatility and the option value of waiting: An analysis of the G-7. *Journal of Futures Markets*, 2010. ISSN 1096-9934.

C.K. Carter and R. Kohn. On Gibbs sampling for state space models. *Biometrika*, 81(3): 541, 1994. ISSN 0006-3444.

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

D. Dobrev and P. Szerszen. The information content of high-frequency data for estimating equity return models and forecasting risk. *FRB International Finance Discussion Papers No. 1005*, 2010.

M. Downey. *Oil 101.* Wooden Table Press, 2009.

J. Elder and A. Serletis. Oil price uncertainty. *Journal of Money, Credit and Banking*, 42 (6):1137–1159, 2010.

R.F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4):987–1007, 1982.

J.D. Hamilton. What is an oil shock? *Journal of Econometrics*, 113(2):363–398, 2003.

J.D. Hamilton. Historical oil shocks. *NBER Working Paper*, 2011.

X. Huang and G. Tauchen. The relative contribution of jumps to total price variance. *Journal of Financial Econometrics*, 3(4):456, 2005. ISSN 1479-8409.

R. Kellogg. The effect of uncertainty on investment: Evidence from Texas oil drilling. *NBER Working Paper*, 2010.

L. Kilian. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069, 2009.

L. Kilian and D. Murphy. The role of inventories and speculative trading in the global market for crude oil. *CEPR Discussion Paper*, (DP7753), 2010.

L. Kilian and R.J. Vigfusson. Nonlinearities in the oil price-output relationship. *Macroeconomic Dynamics*, 15:337–363, 2010.

S. Kim, N. Shephard, and S. Chib. Stochastic volatility: Likelihood inference and comparison with ARCH models. *Review of Economic Studies*, 65(3):361–393, 1998.

K. Lee, S. Ni, and R.A. Ratti. Oil shocks and the macroeconomy: The role of price variability. *The Energy Journal*, 16(4):39–56, 1995.

K.A. Mork. Oil and the macroeconomy when prices go up and down: An extension of Hamilton's results. *Journal of Political Economy*, 97(3):740–744, 1989.

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

V.A. Ramey and D.J. Vine. Oil, automobiles, and the US economy: How much have things really changed? *NBER Macroeconomics Annual*, 2010.

Table 1: Summary statistics of posterior λ draws - Global economy

	(1)	(2)
Main OP series	IRAC	IRAC
RV series	WTI	n/a
mean	-0.1136	-0.1995
std. dev.	0.0514	0.0709
95% interval	[$-$.216,$-$.014]	[$-$.341,$-$.064]
$P(\lambda < 0)$	98.89%	99.99%

This table shows the summary statistics of the 30,000 posterior λ draws. Column (1) is the result from the benchmark VAR model that makes use of the additional measurement equation (8) obtained from the asymptotic distribution of the realized volatility estimator. Column (2) is the result from the VAR model after excluding equation (8).

Table 2: Robustness check - 1984Q1 dummy

λ	
mean	-0.1495
std. dev.	0.0763
95% interval	[-0.3191, -0.0202]
P($\lambda < 0$)	99.15%

This table reports the summary statistics of the posterior λ draws when the dummy variable for 1984Q1 is included in the right hand side of the main VAR. The inclusion of 1984 dummy variable is to see whether the effect of uncertainty remains when considering the structural change in 1984 as noted in Blanchard and Gali (2007).

Table 3: Robustness check - Split sample VAR

λ	∼1983Q4	1984Q1∼
mean	-0.1115	-0.1310
std. dev.	0.1508	0.1036
P($\lambda < 0$)	78.8%	89.6%

This table reports the summary statistics of the posterior λ draws for two subsample periods. The subsample periods is split before- and after- 1984Q1 and the main VAR is run without including equation (8) of realized volatility. The exclusion of realized volatility is because the realized volatility series starts only from 1983Q1, and thus there is only four observation available for the first subsample period. Hence, to treat two subsample period as equal as possible, I exclude equation (8) for both periods. Again, this exercise is to see whether the effect of uncertainty differs when the structural change in 1984 noted in Blanchard and Gali (2007) is considered.

Table 4: Robustness check - Advanced economies

	(1)	(2)
Main OP series	IRAC	IRAC
RV series	WTI	n/a
mean	−0.1154	−0.0768
std. dev.	0.0963	0.0838
$P(\lambda < 0)$	95.09%	81.38%

This table reports the summary statistics of the posterior λ draws of advanced economies. As in the global economy's case presented in Table 1, column (1) shows the result from the benchmark VAR model augmented with realized volatility, and column (2) is the result from the VAR model after excluding equation (8).

Table 5: Robustness check - Different oil price series

	Global economy		Advanced economies	
	(1)	(2)	(3)	(4)
Main OP series	WTI	WTI	WTI	WTI
RV series	WTI	n/a	WTI	n/a
mean	−0.1087	−0.1887	−0.1711	−0.1903
std. dev.	0.0644	0.0742	0.0734	0.0758
$P(\lambda < 0)$	96.27%	98.81%	99.17%	99.35%

This table reports the robustness check results that changes the quarterly real price of crude oil for the VAR from IRAC to WTI. Columns (1) and (2) are for global economy, and (3) and (4) are for advanced economy. Columns (2) and (4) repeat the exercise of re-estimating the VAR without including realized volatility.

Table 6: Robustness check - Multiplication of the parameter m

m	
mean	1.1807
std. dev.	0.5932
95% interval	[0.4004 2.6825]

This table reports the summary statistics of posterior m draws, that is the parameter multiplied to the asymptotic measurement error of equation (8) in order to give the data set a chance to determine the actual size of the measurement error.

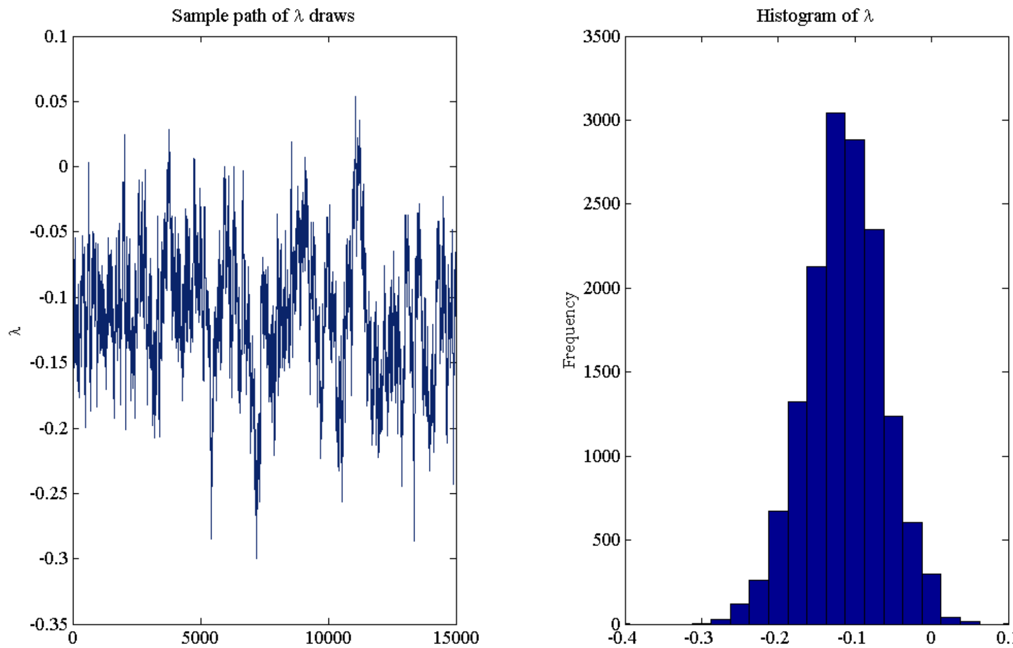

Figure 1: Posterior distribution of λ

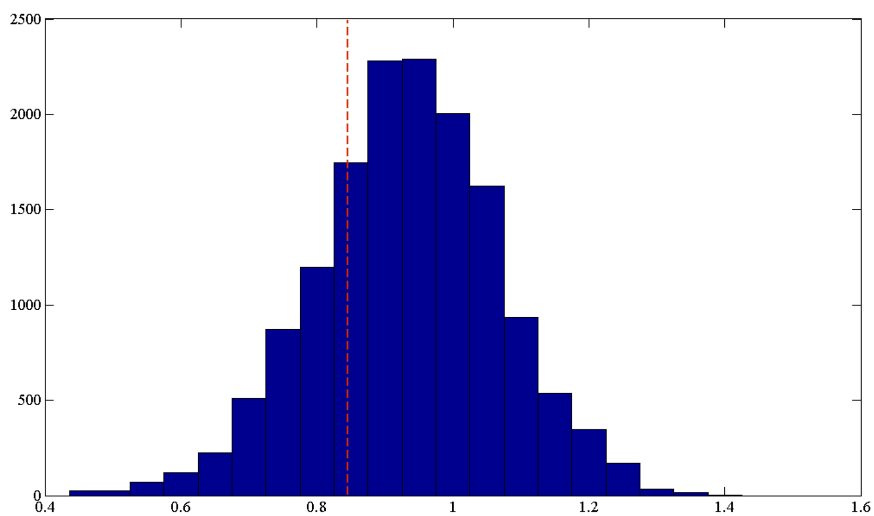

Figure 2: 1986Q1 illustration

This panel is a histogram of possible realization of the global industrial production growth rate in 1986Q1 if oil price uncertainty had not increased at that time. The dotted red line represents the predicted industrial production growth rate, 0.8423%.

Oil price uncertainty *with* realized volatility

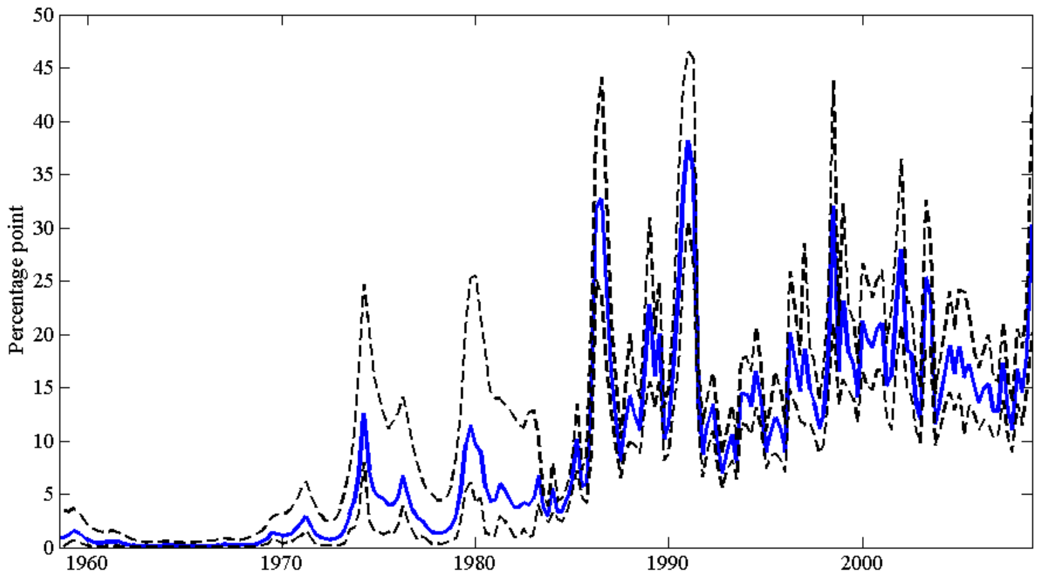

Oil price uncertainty *without* realized volatility

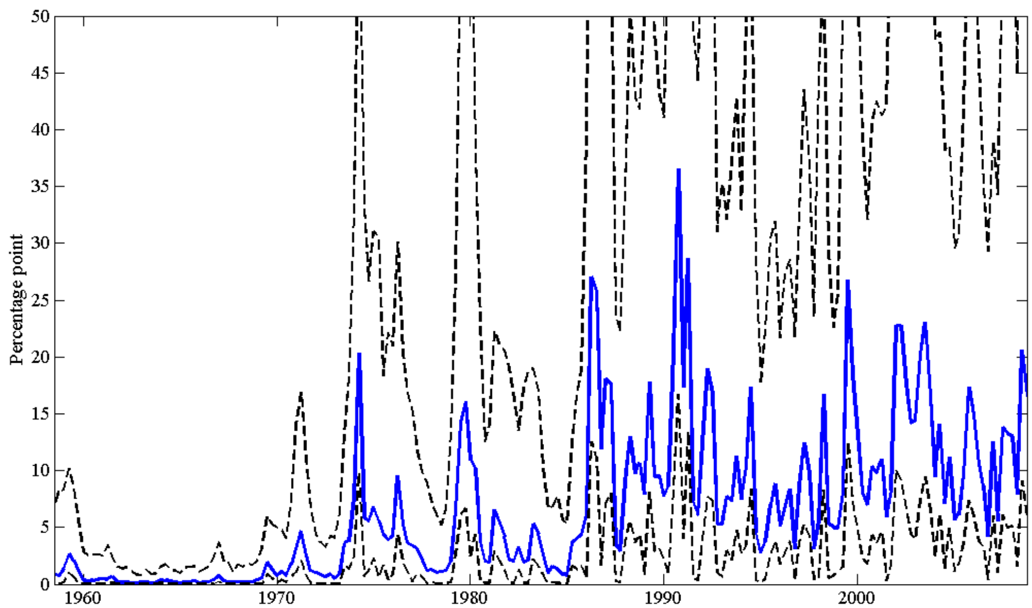

Figure 3: The effect of the augmentation of the extra uncertainty indicator

The posterior distributions of oil price uncertainty from the statistical model of this paper. The upper panel plots the distribution from the benchmark VAR when the information content of high-frequency data is incorporated through realized volatility. The lower panel shows oil price uncertainty from the VAR without using the additional price volatility indicator.

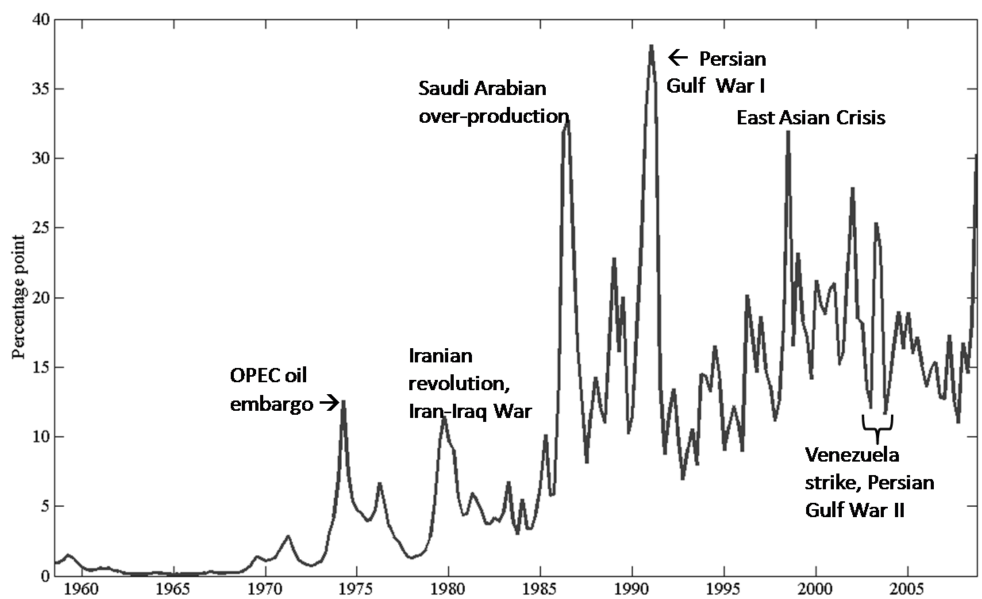

Figure 4: Historical oil price uncertainty jumps

This figure plots the median of the posterior draws of oil price uncertainty from 1958Q2 to 2008Q3.

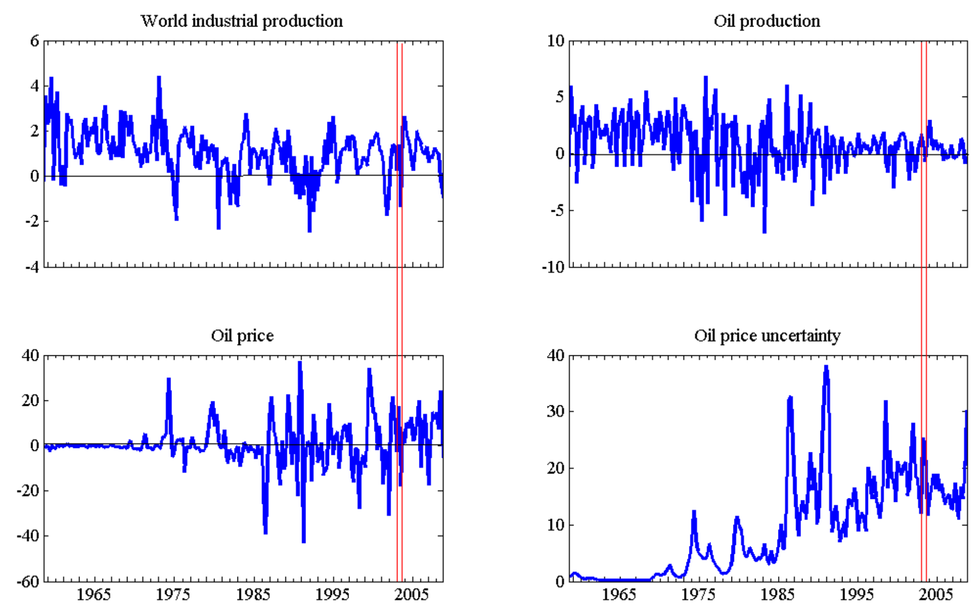

Figure 5: Industrial production, oil production, oil price and oil price volatility

This figure plots the data series used in the main VAR along with the median oil price uncertainty from the posterior draws recovered from the model. The red vertical box denotes 2003Q1 and 2003Q2.

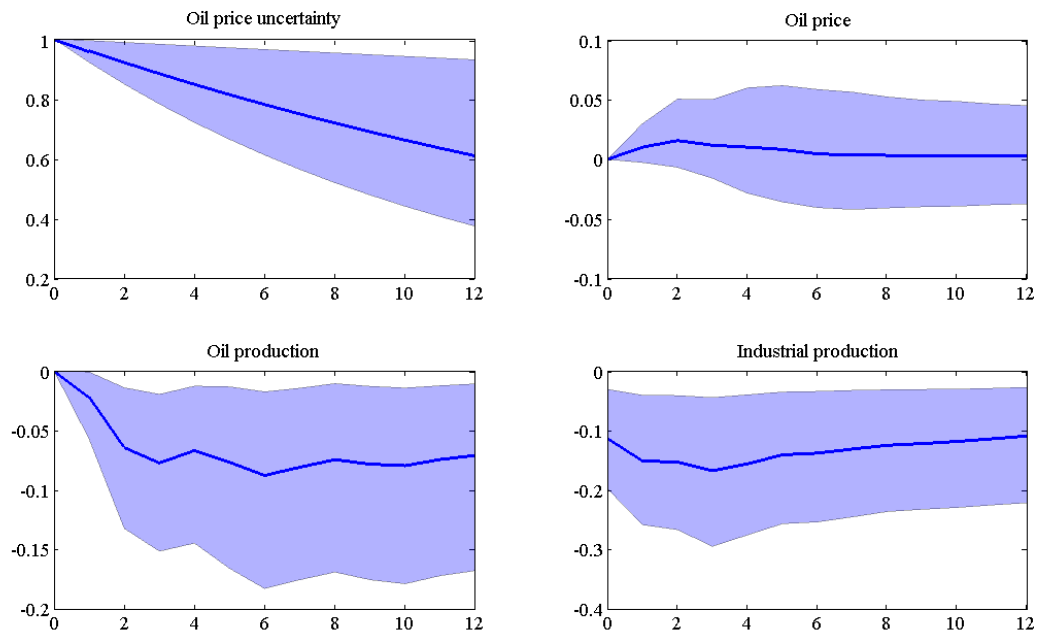

Figure 6: Impulse responses to a 100% uncertainty shock

Impulse responses to the shock that doubles the level of oil price uncertainty. Since oil price uncertainty is included in the mean equation of VAR in logarithm, this means that the log oil price uncertainty increases by a unit. Thus, the left top panel shows the dynamics of log of oil price uncertainty over 12-quarter horizon.