Oil Price Shocks and Economic Growth: The Volatility Link*

John M. Maheu[†] Yong Song[‡] Qiao Yang[§]

April 2019

Abstract

This paper shows that oil shocks primarily impact economic growth through the conditional variance of growth. Our comparison of models focuses on density forecasts. Over a range of dynamic models, oil shock measures and data we find a robust link between oil shocks and the volatility of economic growth. A new measure of oil shocks is developed and shown to be superior to existing measures and indicates that the conditional variance of growth increases in response to an indicator of local maximum oil price exceedance. The empirical results uncover a large pronounced asymmetric response of growth volatility to oil price changes. Uncertainty about future growth is considerably lower compared to a benchmark AR(1) model when no oil shocks are present.

key words: Bayes factors, predictive likelihoods, nonlinear dynamics, density forecast

JEL: C53, C32, C11, Q43

*We are grateful for comments from the Editor Michael McCracken, two anonymous referees and Francesco Ravazzolo. Maheu thanks the SSHRC for financial support and Yang thanks ShanghaiTech University for financial support.

[†]DeGroote School of Business, McMaster University and RCEA. Email:maheujm@mcmaster.ca

[‡]University of Melbourne and RCEA. Email:yong.song@unimelb.com.au

[§]School of Entrepreneurship and Management, ShanghaiTech University. Email: yangqiao@shanghaitech.edu.cn

1 Introduction

This paper provides new results to the debate on how oil shocks impact real economic growth. We find no evidence that oil shocks affect the conditional mean of economic growth using a variety of oil shock measures. However, oil shocks display a strong robust impact on the conditional variance of growth. Related to initial studies (Mork 1989, Hamilton 1996) that find oil price increases are relevant when they exceed the maximum oil price we find that they lead to increases in the conditional variance of growth and provide the best density forecasts for future growth.

The importance of oil price movements and their impact on economic growth was raised in Hamilton (1983). However, the subsequent literature is unclear on the role if any that oil plays in predicting economic growth. The initial findings from Mork (1989) and Hamilton (1996) were that oil price increases are relevant when they exceed the maximum oil price and oil price decreases have no significant effects on economic growth. These stylized facts were further confirmed by Hamilton (2003), Hamilton (2011) and Ravazzolo & Rothman (2013) among others.

This asymmetric response to oil shocks was challenged by Kilian & Vigfusson (2011*a*) and Kilian & Vigfusson (2011*b*). These papers focus on impulse response functions and found no significant difference between positive shocks and negative shocks. Hamilton (2011) argued that their results are caused by different data sets, measures of oil price and price adjustment.

The recent study by Kilian & Vigfusson (2013) performs a comprehensive predictive analysis of the effect of oil price shocks on economic growth. Among several economically plausible nonlinear specifications, they find that including negative oil price shocks further improves economic growth forecasting. In addition, the best predictive model preserves symmetry between positive and negative shocks.

One common feature for the majority of the literature is that the predictive models are nonlinear in oil prices, but linear in oil price shocks. First, a measure of oil price shocks is constructed such as net oil price increase (Hamilton 1996) or large oil price change (Kilian & Vigfusson 2013). Then, the constructed variable enters into a linear model as one regressor to have its predictive performance examined in a homoskedastic setting. One exception is Hamilton (2003), who has modelled the nonparametric conditional mean function to study the nonlinear marginal effect of the oil price on economic growth.

Although our paper is the first to explore the impact oil shocks have on economic uncertainty, other papers by Lee et al. (1995) and Elder & Serletis (2010) have investigated second order moments from oil prices. Lee et al. (1995) argue that an oil shock standardized by a GARCH model is more important to the conditional mean of growth. This two step estimation approach is extended to a bivariate GARCH-in-mean specification in Elder & Serletis (2010). The latter paper finds volatility in oil prices have a negative effect on several measures of output.

The most significant contribution of this paper is to demonstrate that oil shocks primarily affect economic growth through a volatility channel. We find little to no gains by including oil shocks into the conditional mean but very significant forecast improvements when oil shocks enter the conditional variance of real growth. Of course, this volatility channel does not show up in point forecasts of the conditional mean as the literature has focused on but becomes readily apparent in density forecasts.

Working with density forecasts have several advantages. First, a density forecast contains a complete description of future outcomes of growth, including the predictive mean. Second, a density forecast can provide a measure of economic uncertainty about the future and will be sensitive to models with different conditional moment specifications. For example, volatility measures and density intervals from the predictive density will be sensitive to heteroskedasticity. Lastly, density-forecasting based model comparison leads to standard Bayesian methods of model comparison based on Bayes factors. Predictive or marginal likelihoods automatically penalize complex models that do not improve predictions. From a classical perspective, density forecasts can be evaluated through scoring rules (Gneiting & Raftery 2007, Elliott & Timmermann 2008) which have a close equivalence to Bayesian predictive likelihoods and Bayes factors.

This volatility channel is shown to be robust to different oil shock measures, real time data and the use of industrial production index for output. We consider five measures of oil price shocks, four from the academic literature and one developed in this paper. The analysis also consider a range of different lag structures for the dependent variable and the oil shock measures.

One oil shock measure, an indicator variable on the net oil price increase, results in the best density forecasts for economic growth. This specification dominates a GARCH model as well as a hybrid GARCH model that includes these shocks and a stochastic volatility specification. This implies that economic uncertainty increases in response to a local maximum oil price exceedance. However, the impact of this exceedance is independent of the shock size and we discuss some possible reasons for this. When an exceedance occurs the standard deviation on real growth innovations, 2 quarters ahead, almost doubles. Thus our empirical results uncover a large pronounced asymmetric response of growth volatility to net oil price increases. An implication of our findings is that the uncertainty about future growth is considerably lower compared to a benchmark AR(1) model when no oil shocks are present.

The remainder of the paper is organized as follows. Section 2 reviews the data. Section 3 explains out-of-sample density forecasts and the computation method. Section 4 defines our benchmark homoskedastic model along with other linear specifications. Oil price shocks are defined in Section 5 while Section 6 introduces the model that allows oil shocks to affect the conditional mean and conditional variance of growth. The empirical results are discussed in Section 7 while Section 8 reports robustness checks. Section 9 concludes and this is followed by an Appendix that provides details on posterior simulation methods and an online Supplementary Material contains additional forecasting results.

2 Data

The paper restricts attention to two popular series: U.S. real GDP growth rate and Refiners' Acquisition Cost composite index (RAC). The first represents economic growth and the latter represents oil price.[1] For the oil price, there is a fair amount of discussion regarding whether using real or nominal oil price data is appropriate. We use the nominal price following Hamilton (2003), because we believe that a nominal price shock is conceptually more related to behavioural responses from the economy.

Define O_t as the U.S. RAC composite index at time t.[2] The change in oil price, denoted by r_t, is defined as the log difference of O_t, scaled by 100. The economic growth rate, denoted by g_t, is defined as the log difference of the real GDP level scaled by 100.[3] The data spans from 1974Q1 to 2018Q3 with 178 observations in total. Figure 1 shows their time series plot.

[1]Section 8 considers other variables.
[2]Obtained from https://www.eia.gov
[3]Data is from https://fred.stlouisfed.org/series/GDPC1.

3 Out-of-Sample Density Forecasts

Almost all existing papers on the predictive relationship between oil prices and economic growth compare point forecasts.[4] In this paper, we evaluate the predictive relationship from a more general perspective. Models can produce better density forecasts due to a more accurate predictive mean, as the literature has focused on, but improvements may come from other higher order moments that affect the shape of the density forecast. Although our focus is on density forecasts we also report the accuracy of models' predictive mean forecasts.

From a Bayesian perspective, density forecasts or the predictive density, integrates out parameter uncertainty. Model comparison is based on the predictive likelihood, which is the evaluation of the predictive density function at the realized data. Predictive likelihoods are the main input to construct Bayes factors. Although we conduct Bayesian inference the predictive likelihoods are equivalent to a log-scoring rule for density forecasts and there are well defined classical methods to compare models (Amisano & Giacomini 2007).

The predictive density at period t is defined as the distribution of the random variable of interest, in this case g_t, conditional on the past information $I_{t-1} = \{g_{1:t-1}, r_{1:t-1}\}$, for model $\mathcal{M}$ as,

$$p(g_t \mid I_{t-1}, \mathcal{M}) = \int p(g_t \mid \theta, I_{t-1}, \mathcal{M}) p(\theta \mid I_{t-1}, \mathcal{M}) d\theta.$$

The first part in the integral $p(g_t \mid \theta, I_{t-1}, \mathcal{M})$ is the data density, the second part $p(\theta \mid I_{t-1}, \mathcal{M})$ is the posterior density given data I_{t-1}. The posterior density is generally of an unknown form but with Markov chain Monte Carlo (MCMC) methods draws can be obtained from this distribution. Given a large sample $\{\theta^{(i)}\}_{i=1}^{M}$ of MCMC draws from the posterior distribution $p(\theta \mid I_{t-1})$ a simulation consistent estimate of the predictive likelihood is calculated as

$$p(g_t \widehat{\mid I_{t-1}}, \mathcal{M}) = \frac{1}{M} \sum_{i=1}^{M} p(g_t \mid \theta^{(i)}, I_{t-1}, \mathcal{M}).$$

In order to evaluate density forecasts, we compute the predictive likelihood which is the value of the predictive density for a model evaluated at the observed data g_t. Models

[4] Ravazzolo & Rothman (2016) is an important exception. They compare models based on density forecasts but do not investigate the impact of oil shocks on the variance of real growth.

can be compared based on predictive likelihood values. A model with a larger value implies that the density forecast is more plausible, while a model with a smaller value indicates that the model has less support from the data. Of course model comparison cannot be based on any single observation but over a range of out-of-sample density forecasts models' predictive likelihoods become informative.

The log-predictive likelihood (LPL) for g_t, $t = t_0, \ldots, T$ can be decomposed into a sequence of the one-period ahead predictive likelihoods

$$LPL_{\mathcal{M}} = \log p(g_{t_0:T} \mid I_{t_0-1}, \mathcal{M}) = \sum_{t=t_0}^{T} \log p(g_t \mid I_{t-1}, \mathcal{M}). \tag{1}$$

The predictive likelihood is an out-of-sample measure and favors parsimony. Complex models only deliver higher predictive likelihood values if their predictions are more accurate and otherwise leads to smaller values.

Log-Bayes factors for model comparison are directly computed from LPL values. For illustration, log-Bayes factor for $\mathcal{M}_0$ versus $\mathcal{M}_1$ is defined as $\log BF_{01} = LPL_0 - LPL_1$, where $LPL_i = \log p(g_{t_0:T} \mid I_{t_0-1}, \mathcal{M}_i)$, $i = 0, 1$. If $\log BF_{01} > 0$, the data support model $\mathcal{M}_0$ and vice versa. Values of $\log BF_{01} > 5$ are considered very strong support for $\mathcal{M}_0$ (Kass & Raftery 1995).

Besides the log-predictive likelihoods we report the continuous rank probability score (CRPS) for a model. For model $\mathcal{M}$'s cumulative predictive density $F_{t-1}(g_t)$ the CRPS is defined for the out-of-sample data as $\text{CRPS}_{\mathcal{M}} = \frac{1}{T-t_0+1} \sum_{t=t_0}^{T} \int_{-\infty}^{\infty} (F_{t-1}(y) - \mathbb{1}\{y \geq g_t\})^2 dy$. This captures the difference beteen the model's cumulative predictive density and that of the data. Smaller values are preferred. We approximate the $\text{CRPS}_{\mathcal{M}}$ for each model from the MCMC draws.

Finally, the root mean squared forecast error (RMSFE) for $\mathcal{M}$ defined as $\text{RMSFE}_{\mathcal{M}} = \sqrt{\frac{1}{T-t_0+1} \sum_{t=t_0}^{T} (g_t - E(g_t \mid I_{t-1}, \mathcal{M}))^2}$, where $E(g_t \mid I_{t-1}, \mathcal{M})$ denotes the predictive mean.

In the following $t_0 = 10$ leaving 168 observations for out-of-sample forecast evaluation. Because there are many models, we report the log-predictive likelihood LPL_i, and for any pair of models, their log-predictive Bayes factors can be inferred. For posterior simulation $M = 20,000$ after discarding $20,000$ burnin samples. Additional details on posterior simulation for the models is found in the Appendix.

4 Benchmark Model ARX

In the ARX model, real GDP growth rate g_t, is modelled as,

$$g_t = \mu + \sum_{j=1}^{q} \alpha_j g_{t-j} + \sum_{j=1}^{p} \beta_j r_{t-j} + \sigma e_t, \qquad e_t \overset{iid}{\sim} N(0,1). \tag{2}$$

There are q and p number of lags for economic growth and oil price change, respectively, and the maximum value is 4.[5] Note that $p = 0$ reduces to an AR(q) model.

Table 1 shows the log-predictive likelihood (LPL), RMSFE and the cumulative rank probability score (CRPS) for the ARX model for various lag length values. The results are from 168 one-period ahead out-of-sample forecasts for each of the model specifications. The first out-of-sample forecast is for 1976Q4. Each model is re-estimated at each time period in the out-of-sample data so that estimation is based on an expanding window. These results indicate that oil price changes do not predict economic growth when oil changes enter the conditional mean in a linear framework. In the following the AR(1) will serve as a benchmark model for comparison.

In addition to the ARX we investigated an Almon lag version, an ARX with a single lag and an ARX with moving average components. None of these models were competitive with the ARX in terms of out-of-sample forecasts. The definition of the models and forecasting results are contained in the Supplementary Material.

5 Oil Price Shock Measures

Consistent with the literature, the previous section confirms the non-existence of a linear relationship between the oil price changes and economic growth. This section defines a number of *nonlinear oil shock* measures used in the literature as well as a new one. We follow the prevailing papers to adopt four types of oil price shocks: net price increase (Hamilton (1996)), symmetric/asymmetric net price change and large price change (Kilian & Vigfusson (2013)). The new measure we propose uses the sign of the net price increase and provides robustness to the magnitude of oil price changes.

Recall that O_t is the U.S. RAC composite index at time t. The following oil price shocks are constructed.

1. Net price increase

[5] We did not find any forecast improvements for larger values of p and q.

This is probably the most popular way to define an oil price shock, which is developed by Hamilton (1996) as

$$d_t^+ = 100 \max\left\{0, \log \frac{O_t}{O_t^*}\right\},$$

where $O_t^* = \max\{O_{t-1}, ..., O_{t-12}\}$ is the highest oil price in the past three years. Hamilton (1996) used one year history to construct O_t^*. Hamilton (2011) and Kilian & Vigfusson (2013) found that the three-year history is better for prediction. We report the results based on the three year net price increase in this paper.

2. Asymmetric net price change

   Kilian & Vigfusson (2013) showed that a negative oil price shock may also improve prediction, but in an asymmetric way. Therefore, we include both positive and negative shocks to our predictive models in this paper. A positive shock d_t^+ is defined the same as the net price increase. A negative shock is defined as

$$d_t^- = 100 \min\left\{0, \log \frac{O_t}{O_t^{**}}\right\},$$

   where $O_t^{**} = \min\{O_{t-1}, ..., O_{t-12}\}$ is the lowest oil price in the past three years. Notice that there are two shock variables in this setting d_t^+ and d_t^-.

3. Symmetric net price change

   This is the best predictor in Kilian & Vigfusson (2013). They found that restricting a positive and a negative shock to have the same effect can further improve out-of-sample prediction. The new shock measure is

$$d_t^* = d_t^+ + d_t^-,$$

   where d_t^+ is the net price increase and d_t^- is the net price decrease. This variable treats positive and negative shock symmetrically. It is 0 when the price O_t is between the highest (O_t^*) and the lowest (O_t^{**}) historical price.

4. Large price increase

   A shock may impact the economy only when it is unexpected, which is proxied by the "large deviation" in Kilian & Vigfusson (2013). We consider a large price increase as

$$d_t^{large} = r_t I\big(r_t > \text{std}(\{r_{t-1}, ..., r_{t-12}\})\big),$$

where $I(\cdot)$ is the indicator function and equals 1 if its argument is true and 0 otherwise. The shock d_t^{large} is positive if the oil price change r_t is larger than the standard deviation of the oil price change in the past three years. Notice that this measure is asymmetric. If the price decreases, the shock d_t^{large} is zero.

5. Net price increase indicator

   We construct a 0/1 indicator to signal an exceedance of O_t over O_t^* defined as

$$d_t^I = I(d_t^+ > 0),$$

   where d_t^+ is the net price increase. This indicator contains less information than the net price increase d_t^+ but does not suffer from the large magnitudes that d_t^+ can have for *outliers* and may be more robust in capturing an asymmetric response.

6 The Volatility Link

In this section we extend the literature to investigate the transition of oil shocks to the conditional variance of economic growth. Our starting point is the best benchmark model from Section 4 which was an AR(1) without exogenous variables ($q = 1$ and $p = 0$). We augment the AR(1) model by the aforementioned various types of oil price shocks to compare their predictive performance. The general heteroskedastic specification is

$$g_t = \mu + \alpha g_{t-1} + \lambda d_{t-p} + \sigma \exp(\delta d_{t-p}) e_t, \qquad e_t \overset{iid}{\sim} N\big(0, 1\big). \tag{3}$$

The shock d_{t-p} will be replaced by the previously defined measures: r_{t-p}, d_{t-p}^+, d_{t-p}^*, d_{t-p}^{large} or d_{t-p}^I. For the asymmetric net price change, $d_{t-p} = (d_{t-p}^+, d_{t-p}^-)'$ is a vector. The subscript $t-p$ means a p period lag. Model (3) incorporates the lag effects from oil shocks on both the conditional mean and variance. The coefficient λ represents the impact of an oil shock d_{t-p} on the conditional mean of economic growth g_t. The coefficient δ measures the impact of an oil price shock on the conditional variance of g_t. The λ and δ are vectors when the shock d_{t-p} is a vector.

In this model the oil price shocks are transmitted to economic growth through two channels. The first one follows the existing literature to incorporate the shock in the conditional mean function. The second channel is the impact of an oil price shock on the volatility of economic growth. To the best of out knowledge, no one has investigated this issue before.[6] This latter channel will display little to no impact on predictive mean forecasts and it is critical to evaluate model forecasts from the more general metric of density forecasts.

In (3) the oil shock d_{t-p}, affects the mean and variance of economic growth at the same lag time p. Given the weak evidence for oil shocks appearing in the conditional mean we do not consider different lag lengths for d_t in the two moments. Nevertheless our analysis does investigate this possibility indirectly. This is done by considering restricted models. One is to shut down the conditional mean transmission channel by restricting $\lambda = 0$. The other is to restrict $\delta = 0$ to turn off the volatility transmission channel. If we restrict both λ and δ to be zero, we have the benchmark AR(1) model. By comparing the unrestricted and restricted models, we are able to assess oil price shocks impact on the conditional mean and conditional variance and which channel is more relevant.

7 Empirical Results

Table 2 summaries the best models based on out-of-sample density forecasts. This table reports the best models based on LPL values from the more extensive results contained in Tables 3-8. The LPL, RMSFE and CRPS measures of forecast performance are reported. As before each model is re-estimated in the out-of-sample period and the forecast data is the same as in Section 4. Different oil shock measures are included along with restricted versions of the models as well as conditional variance models which are discussed in the next section.

Moving from the LPL of -181.2 for the benchmark AR(1) model in Table 2 we see large improvements from most of the other specifications. Ignoring the model with d_{t-p} set to the large net price increase every other model improves upon the benchmark model. The log-Bayes factors for the new models against the benchmark model range from -2.8 to 19.5. Each of the improved models feature an oil shock measure that enters the conditional variance of real growth. With one exception, the best specification for

[6] Elder & Serletis (2010) study how volatility of oil price shocks affect the mean of economic growth.

each given oil measure occurs with the restriction of $\lambda = 0$. That is, density forecasts are better when the oil shock enters the conditional variance only. The transmission from the oil market to the conditional variance of real growth occurs with 2 or 3 quarters lag, depending on the oil shock measure. Based on the CRPS only the net indicator shock dominates the benchmark AR(1) model. The AR(1) has a CRPS of 0.3914 while the best heteroskedastic model (net price increase indicator) delivers 0.3787.

These results document an important transmission from the oil market to economic growth through the conditional variance of growth with a significant lag effect. According to the LPL values this result is robust to different oil shock measures as well. Consistent with the existing literature, oil shocks in the conditional mean are not generally important nor does allowing heteroskedasticity alter those findings.

The model with the largest LPL value includes the new oil shock, net price increase indicator. The log-Bayes factor for this model against the AR model is 19.5 and is decisive evidence in favor of it. This measure works much better than the other shock measures. The predictive Bayes factor between this model and the best model with the other oil price measures (symmetric net price change) is 10 which is strong evidence in favor of the new measure.

In order to learn more about where the gains from using oil shock measures in the conditional variance of growth come from we plot the cumulative log-predictive Bayes factors. It is calculated as

$$\text{cumlog}BF_t^{01} = \log p(y_{t_0:t} \mid y_{1:t_0-1}, \mathcal{M}_0) - \log p(y_{t_0:t} \mid y_{1:t_0-1}, \mathcal{M}_1),$$

to compare $\mathcal{M}_0$ to $\mathcal{M}_1$. An increase (decrease) in cumlogBF_t^{01} at time t is support for model $\mathcal{M}_0$ ($\mathcal{M}_1$). Cumulative log-Bayes factors appear in Figure 2. The top figure of Figure 2 shows the Bayes factor of the model in (3) with different oil shock measures against the AR(1) benchmark models at each time in the out-of-sample period. Except for the one oil shock all models make regular gains against the AR(1). The improvements these models offer are not due to a few observations but are widespread over the out-of-sample period.

An expanded set of forecasts results are reported in Tables 3-8.[7] These tables report a range of forecast results for different oil shock measures, lag lengths and parameter restrictions. These tables confirm our findings that oil shocks do predict economic

[7] As in the ARX specification we found no forecast improvements from larger maximum values of q and p above 4.

growth through a volatility channel. We can observe that 4 out of 5 shock measures (the exception is the net price increase) have larger LPL when restricting $\lambda = 0$. Incorporating oil price shocks directly into the conditional mean of economic growth is not supported by the data.

Turning to the RMSFE from point forecasts of the conditional mean and comparing the second columns ($\delta \neq 0, \lambda \neq 0$) in Tables 3-8 to the baseline AR(1) model, the best specification in each table does result in a lower RMSFE. However, the gains are small, the AR(1) model has a RMSFE of 0.7078 while the best heteroskedastic model (large net price increase) delivers 0.6995. The CRPS values are consistent with the LPL values but point to a smaller value of p in some cases.

In conclusion, oil price shocks affect the volatility of economic growth. From LPL, RMSFE and CRPS, we can conclude that the best models are always associated with $\delta \neq 0$, which means that an oil price shock predicts the volatility of economic growth. In addition, Table 9 shows the full sample posterior summary of δ from the best models for each type of oil price shock. Except for δ_2 none of their 90% density intervals include 0.

7.1 Volatility of Growth

Our results indicate that oil shocks predict volatility changes in real growth. To further investigate this we estimate an AR(1)-GARCH(1,1) model for heteroskedasticity for comparison,

$$g_t = \mu + \alpha g_{t-1} + e_t, \tag{4a}$$

$$e_t \sim N\big(0, \sigma_t^2\big), \tag{4b}$$

$$\sigma_t^2 = \omega_0 + \omega_1 e_{t-1}^2 + \omega_2 \sigma_{t-1}^2. \tag{4c}$$

See the Appendix for additional details on this specification including estimation. Figure 3 displays the full-sample posterior mean of the standard deviations over time for the best models of each shock type using (3). The black line in each panel is the standard deviation from the GARCH model. The first four panels clearly show that the oil price shock associated with the Gulf War in 1990Q3 exaggerates the volatility of economic growth compared to GARCH. The worst two measures based on the LPL, asymmetric net price change and large price increase, are plotted in the second and fourth panel. It is obvious that the volatilities are distorted relative to the AR(1)-GARCH(1,1) model.

The final panel of this figure shows the model with the net price increase indicator to be the closest to the GARCH implied standard deviation.

In attempt to disentangle the time series effect and the oil shock effect on the volatility of growth, we propose a hybrid model to incorporate both impacts into the second moment. Specifically, the AR-GARCH-Shock model is define as,

$$g_t = \mu + \alpha g_{t-1} + \exp(\delta d_{t-2}) e_t, \tag{5a}$$

$$e_t \sim N\big(0, \sigma_t^2\big), \tag{5b}$$

$$\sigma_t^2 = \omega_0 + \omega_1 e_{t-1}^2 + \omega_2 \sigma_{t-1}^2. \tag{5c}$$

In this model d_{t-2} is the net price increase indicator. The volatility now has two parts: the oil shock effect $\exp(\delta d_{t-2})$ and the GARCH component σ_t^2.

The LPL for the density forecasts from these two GARCH specifications along with the preferred bivariate GARCH-in-mean VAR model from Elder & Serletis (2010) are found in the final entries in Table 2. The univariate GARCH models do improve upon the AR benchmark model but they are still inferior to the best specification which only has oil shocks directing the conditional variance. For instance, the log-Bayes factor for the model with the net price increase indicator entering the conditional variance in (3) against the AR(1)-GARCH(1,1) is 10.8 while it is 3.6 against the AR(1)-GARCH(1,1)-shock. We conclude that the GARCH parameterization is not a proxy for oil shock effects on the conditional variance. Oil price shocks contain additional information value for forecasting output.

These results can be seen from the bottom plot of Figure 2. These are cumulative log-Bayes factors for each specification against the AR(1)-GARCH(1,1) model. The figure shows an upward trend of the log-Bayes factor of the best model with the net price increase indicator. One exceptional period is between 2005-2007, when several oil price shocks are identified but economic growth is tranquil. But the drop in the Bayes factor during that period is quickly compensated during the financial crisis in 2008-2009, when the GARCH model predicts a large volatility but the oil price shocks do not.

The top diagram of Figure 4 displays the cumulative log-Bayes factors for each specification against the GARCH-Shock model. Interestingly, model (3) with the net price increase indicator is still competitive except for the period around 2007.

7.2 Net Price Increase Indicator

This section discusses in more detail the net price increase indicator and the implication of the best forecasting model.

First, the indicator function is critical to the improved performance of this oil shock measure. Although the net price increase shock does improve the LPL (Table 2) it is nowhere near as good as the indicator version. Figure 5 shows a histogram of the shocks measured by net price increase. The largest shock is associated with the Gulf War in 1990. It is about 10 times larger than the shocks in the left tail. An exponential transformation implies that the variance change associated with this shock is e^{10} times larger than a small shock! Given the heterogeneous nature of net price increase shocks it is not surprising that the indicator function performs better.

The effect of these two oil shocks on the conditional standard deviation can be seen in the top and bottom panels of Figure 3. While the net price increase leads to some extreme measures of volatility the indicator function preserves the direction of the measure but removes the extreme values.

Full sample estimates of the best forecasting model are reported in Table 10 along with the AR(1) model. Adding oil shocks into the conditional variance causes important changes. First, when no oil shock is present the conditional standard deviation is much smaller (0.5285 versus 0.7020), meaning that we are much more certain about growth than the AR model would lead us to conclude. On the other hand, when an oil shock is present the conditional standard deviation doubles (0.5285 to $1.0422 = 0.5285 \exp(0.66790)$). Thus our empirical results uncover a large pronounced asymmetric response of growth volatility to net oil price increases. The impact of this oil shock happens with a two quarter lag.

8 Robustness

This section considers robustness checks on priors, structural stability, a stochastic volatility (SV) model, industrial production data and real time data.

8.1 Prior Sensitivity Check

We did a prior sensitivity check on the best model (3) with net price increase indicator when $\lambda = 0$ and $\delta \neq 0$. By changing the prior parameters relative to the original prior, we propose four alternative settings: loose, tight, very loose and very tight. Details and results are reported in the Supplemental Material. Except for the very tight prior, all other settings shows that the log-predictive likelihoods values are robust to prior changes. Even for the very tight prior case, the log-predictive likelihood strongly supports the heteroskedastic model over the AR(1) benchmark model.

8.2 Structural Instability

We estimated the benchmark linear model AR(1) with a 5-year rolling window to control for structural instability. The log-predictive likelihood, RMSFE and CRPS are -176.9, 0.7480 and 0.4142 respectively. In comparison to the AR(1) model without a rolling window ($LPL = -181.2$, $RMSFE = 0.7078$ and $CRPS = 0.3914$), the gain on the density forecast is not prominent and there is even a loss in precision of the point forecasts.

For the other models we perform a subsample analysis. The data are split into three periods $1976Q4 - 1989Q4$ (before Gulf War), $1990Q1 - 2002Q4$ (before oil price surge) and $2003Q1 - 2018Q3$. Table 11 shows the forecast results in each subsample. Using the net price increase indicator as a shock measure is always competitive in each subsample.

8.3 Stochastic Volatility (SV)

Ravazzolo & Rothman (2016) demonstrate significant gains in density forecasts by adding a SV component. As an alternative to the GARCH model we consider the following SV model.

$$g_t = \mu + \alpha g_{t-1} + \exp(h_t/2)u_t \quad u_t \sim N(0,1), \tag{6a}$$

$$h_t = \omega_0 + \omega_1 h_{t-1} + \sigma_\nu \nu_t \quad \nu_t \sim N(0,1). \tag{6b}$$

The last panel of Table 2 indicates a LPL of -166.1 for the SV model and is very close to the GARCH-Shock model. The log-Bayes factor of the best heteroskedastic oil shock model against the SV is 3.6. The SV model does not appear to offer any advantages over our heteroskedastic model under the specification of net price increase indicator but does better than a number of other specifications. The second plot of Figure 4 shows the SV model to perform well against other alternatives except for our preferred model.

8.4 Data

The shocks are constructed by using the 3-year window according to the standard literature such as Hamilton (2003) and Kilian & Vigfusson (2013). As a robustness check, we also reconstruct these measures by using only a one year window. These results favor the same heteroskedasticity model with the net price increase indicator shock.

By using real industrial production as another proxy for output (https:// fred.stlouisfed.org /series/INDPRO), we carried out the same analysis as we did on real GDP. Industrial production growth is monthly data with a full sample size of 537. Out-of-sample forecasts start at observation 11 for a total of 527 forecasts.

The best of the homoskedastic specifications was the ARX-MA[8] with q=1 and p=0, which has an out-of-sample LPL of -525.0, RMSFE of 0.6599 and CRPS of 0.3571. As in the GDP case, Table 12 shows there is strong evidence that oil shocks impact the conditional variance of industrial production growth. The log-Bayes factor of the best heteroskedastic oil shock specification against the AR(1) is $10 = -515 - (-525)$. This model uses the net price increase indicator.

8.5 Real-time Data

Using real-time data can deliver different forecasting results. For real growth forecasts several papers have addressed this issue[9]. In order to check robustness of our results under real-time, we investigate our results under 1st, 2nd and 3rd released vintage.[10] Table 13 reports the forecasting accuracy for a range of models with the real-time data. In short, there is a substantial difference in forecast accuracy across various vintage, but within each vintage, our net price increase indicator shock model shows competitive and robust forecast performance.

[8]The ARX-MA model is

$$g_t = \mu + \alpha \frac{1}{3} \sum_{j=q}^{q+2} g_{t-j} + \beta \frac{1}{3} \sum_{j=p}^{p+2} r_{t-j} + \sigma e_t, \qquad e_t \stackrel{iid}{\sim} N(0,1).$$

See Supplementary Material for this model and additional forecasting results.

[9]Ravazzolo & Rothman (2013), Clark (2011) and Ravazzolo & Rothman (2016)

[10]The details of data construction is referred to Croushore & Stark (2011). It is available at https://www.philadelphiafed.org/research-and-data/real-time-center/real-time-data

9 Conclusion

This paper shows that the primary channel in which oil shocks affect real growth is through the conditional variance of real growth and not the conditional mean. The paper performs an extensive forecasting analysis using different models, oil shock measures as well as real growth measures to demonstrate the robustness of this volatility link.

Incorporating oil shocks into the conditional variance of real growth leads to large improvements in density forecasts but little to no improvement in conditional mean point forecasts. A new shock measure, net price increase indicator, produces the best density forecasts. An implication of our findings is that the uncertainty about future growth is considerably lower compared to a benchmark AR(1) model when no oil shocks are present.

Table 1: The Log-predictive Likelihood, RMSFE and CRPS of ARX

	$p=0$	$p=1$	$p=2$	$p=3$	$p=4$
	-190.9621	-193.9410	-198.5313	-205.7093	-205.9531
	(0.7463)	(0.7861)	(0.8154)	(0.9419)	(0.9473)
$q=0$	{0.4083}	{0.4313}	{0.4597}	{0.5048}	{0.5294}
	-181.1729	-183.7543	-188.1428	-194.1468	-191.1639
	(0.7078)	(0.7356)	(0.7639)	(0.8908)	(0.9207)
$q=1$	**{0.3914}**	{0.4118}	{0.4372}	{0.4833}	{0.5195}
	-181.1741	-186.2719	-189.0330	-194.4778	-197.6055
	(0.7131)	(0.7635)	(0.7722)	(0.9011)	(0.9329)
$q=2$	{0.3968}	{0.4278}	{0.4452}	{0.4902}	{0.5284}
	-183.8563	-189.7557	-193.6134	-196.8827	-200.0005
	(0.7367)	(0.8114)	(0.8622)	(0.9258)	(0.9570)
$q=3$	{0.4116}	{0.4560}	{0.4910}	{0.5116}	{0.5503}
	-187.6794	-192.0069	-196.6988	-199.0541	-202.5128
	(0.7622)	(0.8005)	(0.9030)	(0.9595)	(0.9752)
$q=4$	{0.4284}	{0.4623}	{0.5079}	{0.5471}	{0.5726}

This table reports log-predictive likelihood values, root mean squared forecast errors (RMSFE) in parentheses, and cumulative rank probability score (CRPS) in braces for the 168 out-of-sample observations. A bold number indicates the best model with respect to each measurement.ARX: $g_t = \mu + \sum_{i=1}^{q} \alpha_i g_{t-i} + \sum_{i=1}^{p} \beta_i r_{t-i} + \sigma e_t$

Table 2: Summary of the Best Forecasting Models

Benchmark Homoskedastic Model

$$g_t = \mu + \sum_{i=1}^{q} \alpha_i g_{t-i} + \sum_{i=1}^{p} \beta_i r_{t-i} + \sigma e_t,\ e_t \overset{iid}{\sim} N(0,1).$$

Model	q	p	LPL	RMSFE	CRPS
ARX	1	0	-181.2	0.7079	0.3914

Heteroskedastic Model: Oil Shock Enters Conditional Mean and Variance

$$g_t = \mu + \alpha g_{t-1} + \lambda d_{t-p} + \sigma \exp(\delta d_{t-p}) e_t,\ e_t \overset{iid}{\sim} N(0,1).$$

Oil Shock, d_{t-p}	Shock-lag, p	Restriction	LPL	RMSFE	CRPS
Net Price Change	3	$\delta \neq 0, \lambda \neq 0$	-171.7	0.7439	0.4157
Asymmetric Net Price Change	3	$\delta_{1:2} \neq 0, \lambda_{1:2} = 0$	-175.9	0.7146	0.4114
Symmetric Net Price change	3	$\delta \neq 0, \lambda = 0$	-175.2	0.7142	0.4030
Large Net Price Increase	3	$\delta \neq 0, \lambda = 0$	-184.0	**0.6995**	0.4119
Net Price Increase Indicator	2	$\delta \neq 0, \lambda = 0$	**-161.7**	0.7044	**0.3787**
Oil Price Log-difference	3	$\delta \neq 0, \lambda = 0$	-178.1	0.7030	0.3974

GARCH and SV Models

Model Type	Shock Type	Shock-lag	LPL	RMSFE	CRPS
GARCH	N/A	N/A	-172.5	0.7177	0.3932
GARCH-Shock	Net Indicator	4	-165.3	0.7319	0.3965
SV	N/A	N/A	-166.1	0.7148	0.4166
VAR4-GARCH-M	N/A	N/A	-184.3	0.8201	0.4064

This table reports the log-predictive likelihood (LPL), RMSFE and CRPS for a range of models and oil shock measures. The top entry is the best homoskedastic model from Table 1. The second set of results are the bold entries for LPL from Tables 3-8. The GARCH models are defined in Section 7.1 and SV in Section 8.3. VAR4-GARCH-M is the preferred bivariate GARCH-in-mean VAR model from Elder & Serletis (2010). A bold entry denotes the largest LPL value, and smallest RMSFE and CRPS values.

Table 3: Forecasts from Heteroskedastic model. Shock: oil price log-difference

	$\delta \neq 0,\ \lambda \neq 0$	$\delta \neq 0,\ \lambda = 0$	$\delta = 0,\ \lambda \neq 0$
$p = 1$	-186.9 (0.7429) {0.4225}	-181.1 (0.7092) {0.3976}	-183.7 (0.7353) {0.4116}
$p = 2$	-180.9 (0.7199) {0.4219}	**-178.1 (0.7030) {0.3974}**	-183.3 (0.7170) {0.4062}
$p = 3$	-182.8 (0.7453) {0.4279}	-178.4 (0.7194) {0.4061}	-185.4 (0.7815) {0.4190}
$p = 4$	-184.8 (0.7675) {0.4379}	-181.0 (0.7149) {0.4062}	-184.7 (0.7370) {0.4161}

This table reports log-predictive likelihood values, root mean squared forecast errors (RMSFE) in parentheses, and cumulative rank probability score (CRPS) in braces for the 168 out-of-sample observations. A bold number indicates the best model with respect to each measurement. Shock Definition: $d_{t-p} = r_{t-p} = \ln \frac{O_{t-p}}{O_{t-p-1}}$. Model: $g_t = \mu + \alpha g_{t-1} + \lambda d_{t-p} + \sigma \exp(\delta d_{t-p}) e_t$.

Table 4: Forecasts from Heteroskedastic model. Shock: net price increase

	$\delta \neq 0$, $\lambda \neq 0$	$\delta \neq 0$, $\lambda = 0$	$\delta = 0$, $\lambda \neq 0$
$p = 1$	-181.2 (0.7152) {0.4080}	-177.8 (0.7105) **{0.3958}**	-181.8 (0.7409) {0.4108}
$p = 2$	-178.2 (0.7827) {0.4386}	-175.7 **(0.7030)** {0.4162}	-183.8 (0.7387) {0.4145}
$p = 3$	**-171.7** (0.7439) {0.4157}	-172.3 (0.7158) {0.4049}	-178.7 (0.7271) {0.4050}
$p = 4$	-175.9 (0.8048) {0.4322}	-176.4 (0.7179) {0.4170}	-182.2 (0.7924) {0.4201}

This table reports log-predictive likelihood values, root mean squared forecast errors (RMSFE) in parentheses, and cumulative rank probability score (CRPS) in braces for the 168 out-of-sample observations. A bold number indicates the best model with respect to each measurement. Shock Definition: $d_{t-p} = \max\left\{0, \ln \frac{O_{t-p}}{O^*_{t-p}}\right\}$. Model: $g_t = \mu + \alpha g_{t-1} + \lambda d_{t-p} + \sigma \exp(\delta d_{t-p}) e_t$.

Table 5: Forecasts from Heteroskedastic model. Shock: asymmetric net price change

	$\delta_{1:2} \neq 0$, $\lambda_{1:2} \neq 0$	$\delta_{1:2} \neq 0$, $\lambda_{1:2} = 0$	$\delta_{1:2} = 0$, $\lambda_{1:2} \neq 0$
$p = 1$	-193.6 (0.7109) {0.6887}	-180.5 (0.7173) **{0.4066}**	-200.1 (0.7436) {0.6752}
$p = 2$	-193.0 (0.7887) {0.6593}	-179.5 (**0.7036**) {0.4278}	-193.5 (0.7401) {0.6645}
$p = 3$	-189.9 (0.7478) {0.6696}	**-175.9** (0.7146) {0.4114}	-190.5 (0.7273) {0.6596}
$p = 4$	-188.6 (0.8010) {0.6810}	-179.7 (0.7286) {0.4285}	-193.6 (0.7901) {0.6730}

This table reports log-predictive likelihood values, root mean squared forecast errors (RMSFE) in parentheses, and cumulative rank probability score (CRPS) in braces for the 168 out-of-sample observations. A bold number indicates the best model with respect to each measurement. Shock Definition: $d^+_{t-p} = \max\left\{0, \ln \frac{O_{t-p}}{O^*_{t-p}}\right\}$, $d^-_{t-p} = \min\left\{0, \ln \frac{O_{t-p}}{O^{**}_{t-p}}\right\}$. Model: $g_t = \mu + \alpha g_{t-1} + \lambda_1 d^+_{t-p} + \lambda_2 d^-_{t-p} + \sigma \exp(\delta_1 d^+_{t-p} + \delta_2 d^-_{t-p}) e_t$.

Table 6: Forecasts from Heteroskedastic model. Shock: symmetric net price change

	$\delta \neq 0$, $\lambda \neq 0$	$\delta \neq 0$, $\lambda = 0$	$\delta = 0$, $\lambda \neq 0$
$p = 1$	-179.7 (0.7239) {0.4171}	-176.0 (0.7180) **{0.3954}**	-184.1 (0.7680) {0.4216}
$p = 2$	-179.5 (0.8184) {0.4546}	-177.7 (**0.7021**) {0.4093}	-184.9 (0.7428) {0.4203}
$p = 3$	-180.0 (0.7434) {0.4370}	**-175.2** (0.7142) {0.4030}	-182.6 (0.7834) {0.4217}
$p = 4$	-182.1 (0.8180) {0.4513}	-178.1 (0.7294) {0.4199}	-184.7 (0.8175) {0.4323}

This table reports log-predictive likelihood values, root mean squared forecast errors (RMSFE) in parentheses, and cumulative rank probability score (CRPS) in braces for the 168 out-of-sample observations. A bold number indicates the best model with respect to each measurement. Shock Definition: $d_{t-p} = \max\left\{0, \ln \frac{O_{t-p}}{O^*_{t-p}}\right\} + \min\left\{0, \ln \frac{O_{t-p}}{O^{**}_{t-p}}\right\}$. Model: $g_t = \mu + \alpha g_{t-1} + \lambda d_{t-p} + \sigma \exp(\delta d_{t-p}) e_t$.

Table 7: Forecasts from Heteroskedastic model. Shock: large net price increase

	$\delta \neq 0$, $\lambda \neq 0$	$\delta \neq 0$, $\lambda = 0$	$\delta = 0$, $\lambda \neq 0$
$p = 1$	-187.4 (0.7337) {0.4442}	-184.5 (0.7043) {0.4091}	-183.9 (0.7271) {0.4328}
$p = 2$	-185.7 (0.7016) {0.4296}	-185.7 (0.7016) **{0.4007}**	-185.7 (0.7152) {0.4286}
$p = 3$	-187.2 (0.7344) {0.4445}	**-184.0** (**0.6995**) {0.4119}	-184.3 (0.7359) {0.4299}
$p = 4$	-191.9 (0.9102) {0.4793}	-186.7 (0.7145) {0.4274}	-186.7 (0.8503) {0.4487}

This table reports log-predictive likelihood values, root mean squared forecast errors (RMSFE) in parentheses, and cumulative rank probability score (CRPS) in braces for the 168 out-of-sample observations. A bold number indicates the best model with respect to each measurement. Shock Definition: $d_{t-p} = r_{t-p} I\{r_{t-p} > \text{std}\{r_{t-p-1}, ..., r_{t-p-12}\}\}$, where $I(\cdot)$ is the indicator function and the symbol std represents one standard deviation of past three years oil price change rates. Model: $g_t = \mu + \alpha g_{t-1} + \lambda d_{t-p} + \sigma \exp(\delta d_{t-p}) e_t$.

Table 8: Forecasts from Heteroskedastic model. Shock: net price increase indicator

	$\delta \neq 0$, $\lambda \neq 0$	$\delta \neq 0$, $\lambda = 0$	$\delta = 0$, $\lambda \neq 0$
$p = 1$	-175.7 (0.7396) {0.4049}	-171.5 (0.7150) {0.3924}	-186.3 (0.7330) {0.4069}
$p = 2$	-162.5 (0.7064) {0.3849}	**-161.7 (0.7044) {0.3787}**	-181.2 (0.7109) {0.3975}
$p = 3$	-164.3 (0.7156) {0.3909}	-163.5 (0.7163) {0.3857}	-181.3 (0.7105) {0.3984}
$p = 4$	-162.9 (0.7066) {0.3839}	-163.3 (0.7252) {0.3859}	-180.3 (0.7067) {0.3946}

This table reports log-predictive likelihood values, root mean squared forecast errors (RMSFE) in parentheses, and cumulative rank probability score (CRPS) in braces for the 168 out-of-sample observations. A bold number indicates the best model with respect to each measurement. Shock Definition: $d_{t-p} = I(d_t^+ > 0)$. $I(\cdot)$ is the indicator function. Model: $g_t = \mu + \alpha g_{t-1} + \lambda d_{t-p} + \sigma \exp(\delta d_{t-p}) e_t$.

Table 9: Posterior Summary of δ

Shock, d_{t-p}	Mean	90% DI
Oil Price Log-difference	0.0132	$(0.007696, 0.01836)$
Net Price Increase	0.0495	$(0.029796, 0.07194)$
Asymmetric Net Price Change	δ_1 0.0477	$(0.02777, 0.06965)$
	δ_2 0.0102	$(-0.00238, 0.0206)$
Symmetric Net Price Change	0.0198	$(0.01299, 0.02625)$
Large Net Price Increase	0.0146	$(0.00311, 0.02711)$
Net Price Increase Indicator	0.6790	$(0.48120, 0.88160)$

This table reports full sample posterior statistics of δ from the best models according to LPL values from Tables 3–8. The model is $g_t = \mu + \alpha g_{t-1} + \lambda d_{t-p} + \sigma \exp(\delta d_{t-p}) e_t$ except for Asymmetric Net Price Change it is $g_t = \mu + \alpha g_{t-1} + \lambda_1 d_{t-p}^+ + \lambda_2 d_{t-p}^- + \sigma \exp(\delta_1 d_{t-p}^+ + \delta_2 d_{t-p}^-) e_t$. 90% DI is the 90 % density interval.

Table 10: Posterior Summary of the Best Forecasting Model

	Mean	90% DI		Mean	90% DI
	Best Shock Model			AR(1)	
μ	0.4764	(0.3664,0.5864)	μ	0.4234	(0.3070,0.5380)
α	0.3634	(0.2485,0.4780)	α	0.3680	(0.2529,0.4841)
σ	0.5285	(0.4765,0.5860)	σ	0.7020	(0.6448,0.7647)
δ	0.6790	(0.4812,0.8816)			

This table reports posterior estimates for $g_t = \mu + \alpha g_{t-1} + \sigma \exp(\delta d_{t-2}) e_t$, $e_t \sim N(0,1)$ with d_{t-2} equal to net price increase indicator and the benchmark AR(1), ($\delta = 0$).

Table 11: Out-of-Sample Forecasts for Subsamples

Model	LPL	RMSFE	CRPS
	1976Q4 to 1989Q4		
Net Price Increase	-68.89	0.9659	0.5416
Asymmetric Net Price Change	-69.243	0.9652	0.5634
Symmetric Net Price Change	**-66.47**	0.9654	0.5371
Large Net Price Increase	-77.72	0.9575	0.5857
Oil Price Log Difference	-74.85	0.9586	0.5667
Net Price Increase Indicator	-67.17	**0.9505**	**0.5245**
ARX-1 Benchmark	-74.16	1.0089	0.5480
GARCH	-74.07	0.9698	0.5579
GARCH-Shock	-67.07	0.9839	0.5555
SV	-70.85	0.9631	0.5926
VAR4-GARCH-M	-82.28	1.1888	0.5921
	1990Q1 to 2002Q4		
Net Price Increase	-49.23	0.5565	0.3732
Asymmetric Net Price Change	-49.85	0.5498	0.3672
Symmetric Net Price Change	-49.12	0.5487	0.3672
Large Net Price Increase	-51.37	**0.5456**	0.3488
Oil Price Log Difference	-49.85	0.5468	0.3363
Net Price Increase Indicator	-45.10	0.5565	**0.3226**
ARX-1 Benchmark	-50.10	0.5553	0.3307
GARCH	-46.73	0.5540	0.3281
GARCH-Shock	–**44.93**	0.5607	0.3267
SV	-46.52	0.5479	0.3475
VAR4-GARCH-M	-47.31	1.1947	0.3270
	2003Q1 to 2018Q3		
Net Price Increase	-54.14	0.5778	0.3160
Asymmetric Net Price Change	-56.80	0.5801	0.3200
Symmetric Net Price Change	-59.62	0.5796	0.3199
Large Net Price Increase	-55.43	0.5599	0.3102
Oil Price Log Difference	-53.43	**0.5531**	0.3052
Net Price Increase Indicator	-49.42	0.5614	**0.3023**
ARX-1 Benchmark	-56.90	0.5617	0.3095
GARCH	-51.69	0.5737	0.3083
GARCH-Shock	-53.34	0.5955	0.3203
SV	**- 48.71**	0.5782	0.3255
VAR4-GARCH-M	-55.24	0.5813	0.3157

This table reports the log-predictive likelihood (LPL), root-mean squared forecasts error (RMSFE) and continuous rank probability score (CRPS) for different subsamples. Bold denotes the largest LPL (smallest RMSFE and CRPS) in a subsample. The heteroskedastic model with different oil shocks ($\lambda = 0$ and $\delta \neq 0$) is included along with a homoskedastic AR(1) (benchmark), AR(1)-GARCH(1,1), AR(1)-GARCH(1,1) with shock and SV.

Table 12: Robustness Check using Industrial Production Index

Shock Type, d_{t-p}	$\delta \neq 0, \lambda = 0$	$\delta = 0, \lambda \neq 0$
Net Price Increase	-516 (0.6587) {0.3576}	-526 (0.6596) {0.3615}
Asymmetric Net Price Change	-520 (0.6585) {0.3589}	-553 (0.6585) {0.4874}
Symmetric Net Price Change	-520 (0.6593) {0.3581}	-529 (0.6636) {0.3643}
Large Net Price Increase	-531 (0.6563) {0.3609}	-531 (0.6641) {0.3699}
Oil Price Log-Difference	-529 (0.6714) {0.3614}	-531 (0.6676) {0.3638}
Net Price Increase Indicator	-515 (0.6585) {0.3572}	-528 (0.6633) {0.3611}

The table displays the log-predictive likelihood, RMSFE (parenthesis), and continuous rank probability score (braces) for the best heteroskedastic shock models using industrial production growth as economic growth. The full sample size is 537 with the out-of-sample size of 527. The oil shock is measured according to the past 36 months. The best homoskedastic model (ARX-MA) delivers a LPL of -525.0, RMSFE of 0.6599 and CRPS of 0.3571.

Table 13: Summary of the Best Forecasting Models under Real-time Data

		1st Released Data			
Model	q	p	LPL	RMSFE	CRPS
ARX-1	1	4	-156.8	0.6349	0.3550
ARX-MA	1	4	-164.9	0.6397	0.3607
ARX	4	0	-154.1	0.6053	0.3428
Model	Lags	Restrictions	LPL	RMSFE	CRPS
Net Price Increase	$p=3$	$\delta \neq 0, \lambda = 0$	-148.0	**0.5710**	0.3311
Asymmetric Net Price Change	$p=4$	$\delta \neq 0, \lambda = 0$	-151.8	0.5942	0.3399
Symmetric Net Price Change	$p=4$	$\delta \neq 0, \lambda = 0$	-148.9	0.5932	0.3486
Large Net Price Increase	$p=4$	$\delta = 0, \lambda \neq 0$	-157.0	0.5898	0.3631
Oil Price Log-Difference	$p=4$	$\delta = 0, \lambda \neq 0$	-154.8	0.5945	0.3447
Net Price Increase Indicator	$p=4$	$\delta \neq 0, \lambda = 0$	**-136.4**	0.5978	**0.3218**
GARCH	N/A	ARX-1	-141.3	0.5850	0.3272
GARCH-Shock	N/A	ARX-1	-139.1	0.5972	0.3262
SV	N/A	ARX-1	-139.4	0.5870	0.3526
		2nd Released Data			
Model	q	p	LPL	RMSFE	CRPS
ARX-1	1	0	-170.0	0.6159	0.3580
ARX-MA	0	4	-175.1	0.6791	0.3841
ARX	4	0	-166.8	0.6517	0.3684
Model	Lags	Restrictions	LPL	RMSFE	CRPS
Net Price Increase	$p=4$	$\delta \neq 0, \lambda = 0$	-160.0	0.6420	0.3803
Asymmetric Net Price Change	$p=4$	$\delta \neq 0, \lambda = 0$	-163.8	0.6426	0.3881
Symmetric Net Price Change	$p=4$	$\delta \neq 0, \lambda = 0$	-161.0	0.6395	0.3720
Large Net Price Increase	$p=4$	$\delta = 0, \lambda \neq 0$	-168.9	0.6379	0.3888
Oil Price Log-Difference	$p=4$	$\delta = 0, \lambda \neq 0$	-166.6	0.6381	0.3699
Net Price Increase Indicator	$p=4$	$\delta \neq 0, \lambda = 0$	**-149.2**	0.6448	**0.3477**
GARCH	N/A	ARX-1	-151.9	**0.6243**	0.3787
GARCH-Shock	N/A	ARX-1	-151.4	0.6244	0.3485
SV	N/A	ARX-1	-150.5	0.6408	0.3500
		3rd Released Data			
Model	q	p	LPL	RMSFE	CRPS
ARX-1	1	4	-174.8	0.6994	0.3940
ARX-MA	1	4	-179.6	0.6961	0.3939
ARX	1	0	-173.8	0.6775	0.3829
Model	Lags	Restrictions	LPL	RMSFE	CRPS
Net Price Increase	$p=4$	$\delta \neq 0, \lambda = 0$	-166.8	0.6644	0.3942
Asymmetric Net Price Change	$p=4$	$\delta \neq 0, \lambda = 0$	-170.3	0.6641	0.4022
Symmetric Net Price Change	$p=4$	$\delta \neq 0, \lambda = 0$	-167.0	0.6633	0.3869
Large Net Price Increase	$p=4$	$\delta = 0, \lambda \neq 0$	-175.6	0.6601	0.4046
Oil Price Log-Difference	$p=4$	$\delta = 0, \lambda \neq 0$	-172.8	0.6613	0.3835
Net Price Increase Indicator	$p=4$	$\delta \neq 0, \lambda = 0$	**-157.0**	0.6680	**0.3608**
GARCH	N/A	ARX-1	-163.6	0.6495	0.3655
GARCH-Shock	N/A	ARX-1	-163.1	0.6609	0.3659
SV	N/A	ARX-1	-157.7	**0.6490**	0.3907

This table reports the log-predictive likelihoods (LPL), RMSFE and CRPS for a range of models and oil shock measures using real-time data. Bold entries are for the largest LPL and smallest RMSFE and CRPS within a data vintage.

Figure 1: RAC Composite Quarterly Oil Price (Top)
RAC Composite Quarterly Percentage Change of Oil Price (Middle)
U.S Quarterly Real GDP Growth Rates (Bottom)

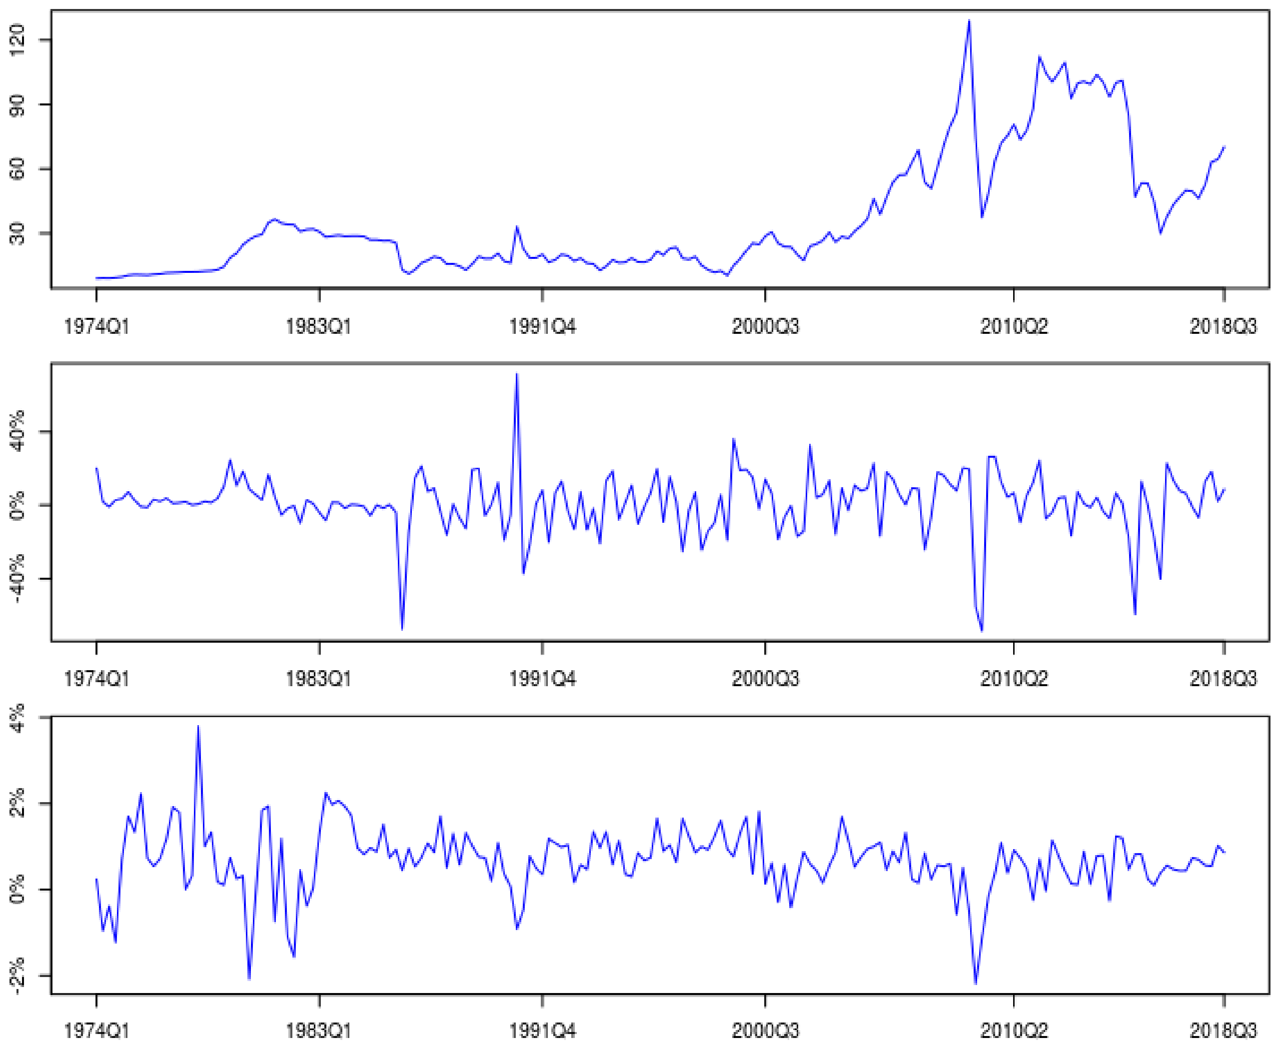

Figure 2: Cumulative Predictive Bayes Factors of Heteroskedastic Model with Oil Shock against Benchmark AR(1) (Top) and GARCH (Bottom)

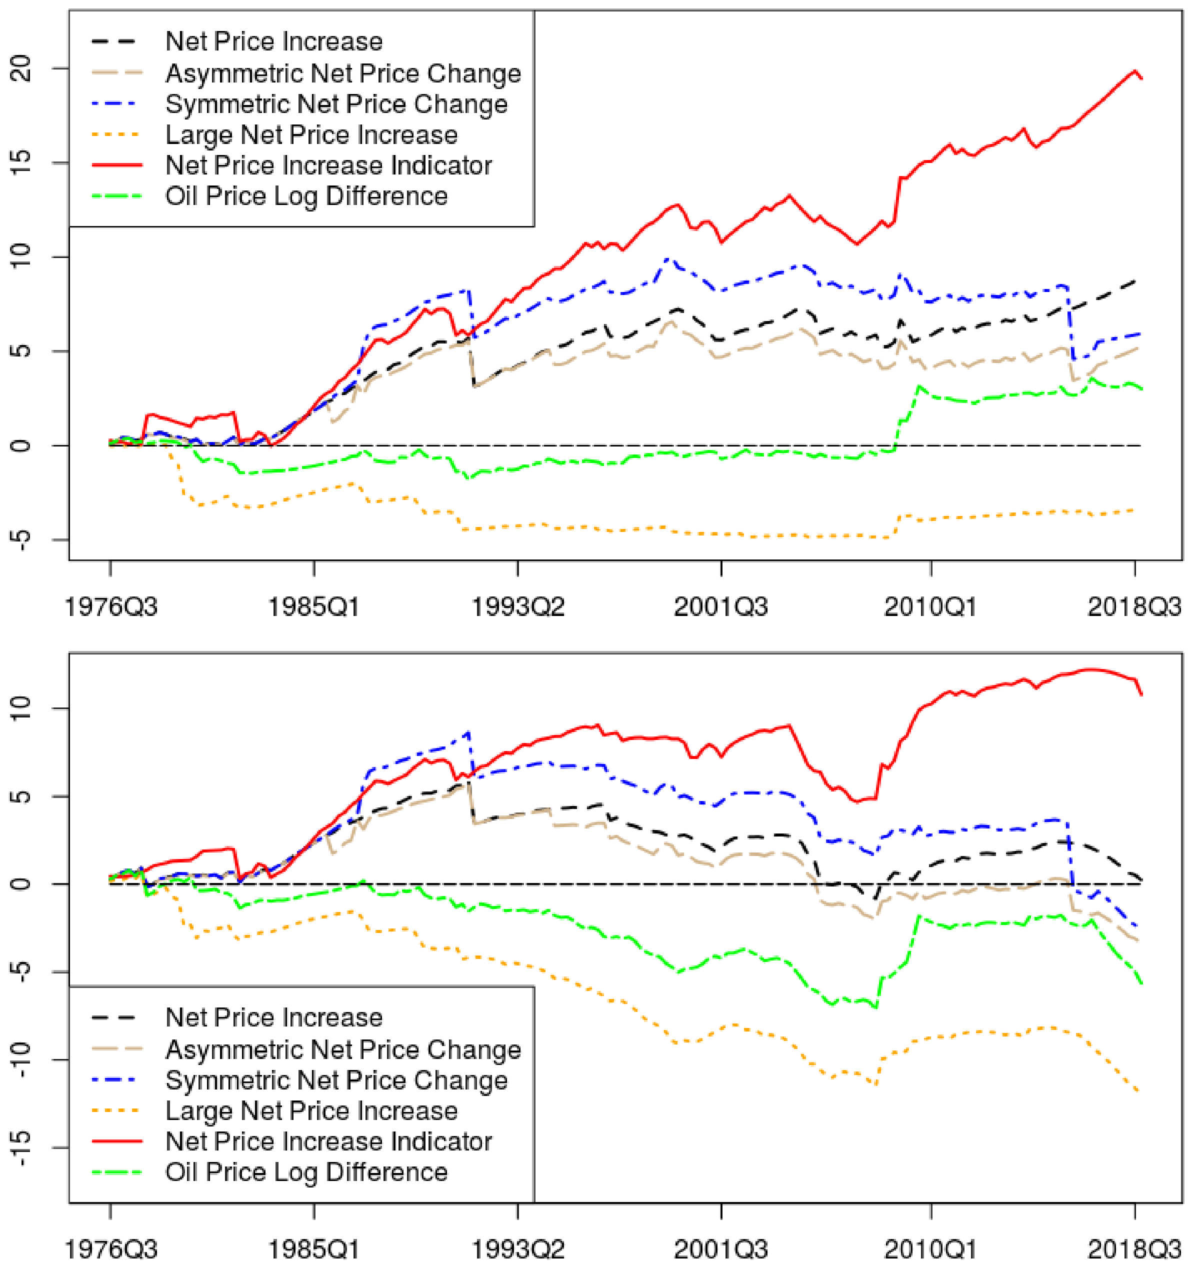

This figure plots the cumulative Bayes factor at each point in time in favour of a model against the benchmark AR(1) model (top plot) and GARCH (bottom plot). Various lines represent model (3) with corresponding shock variables.

Figure 3: Volatility of Real Growth

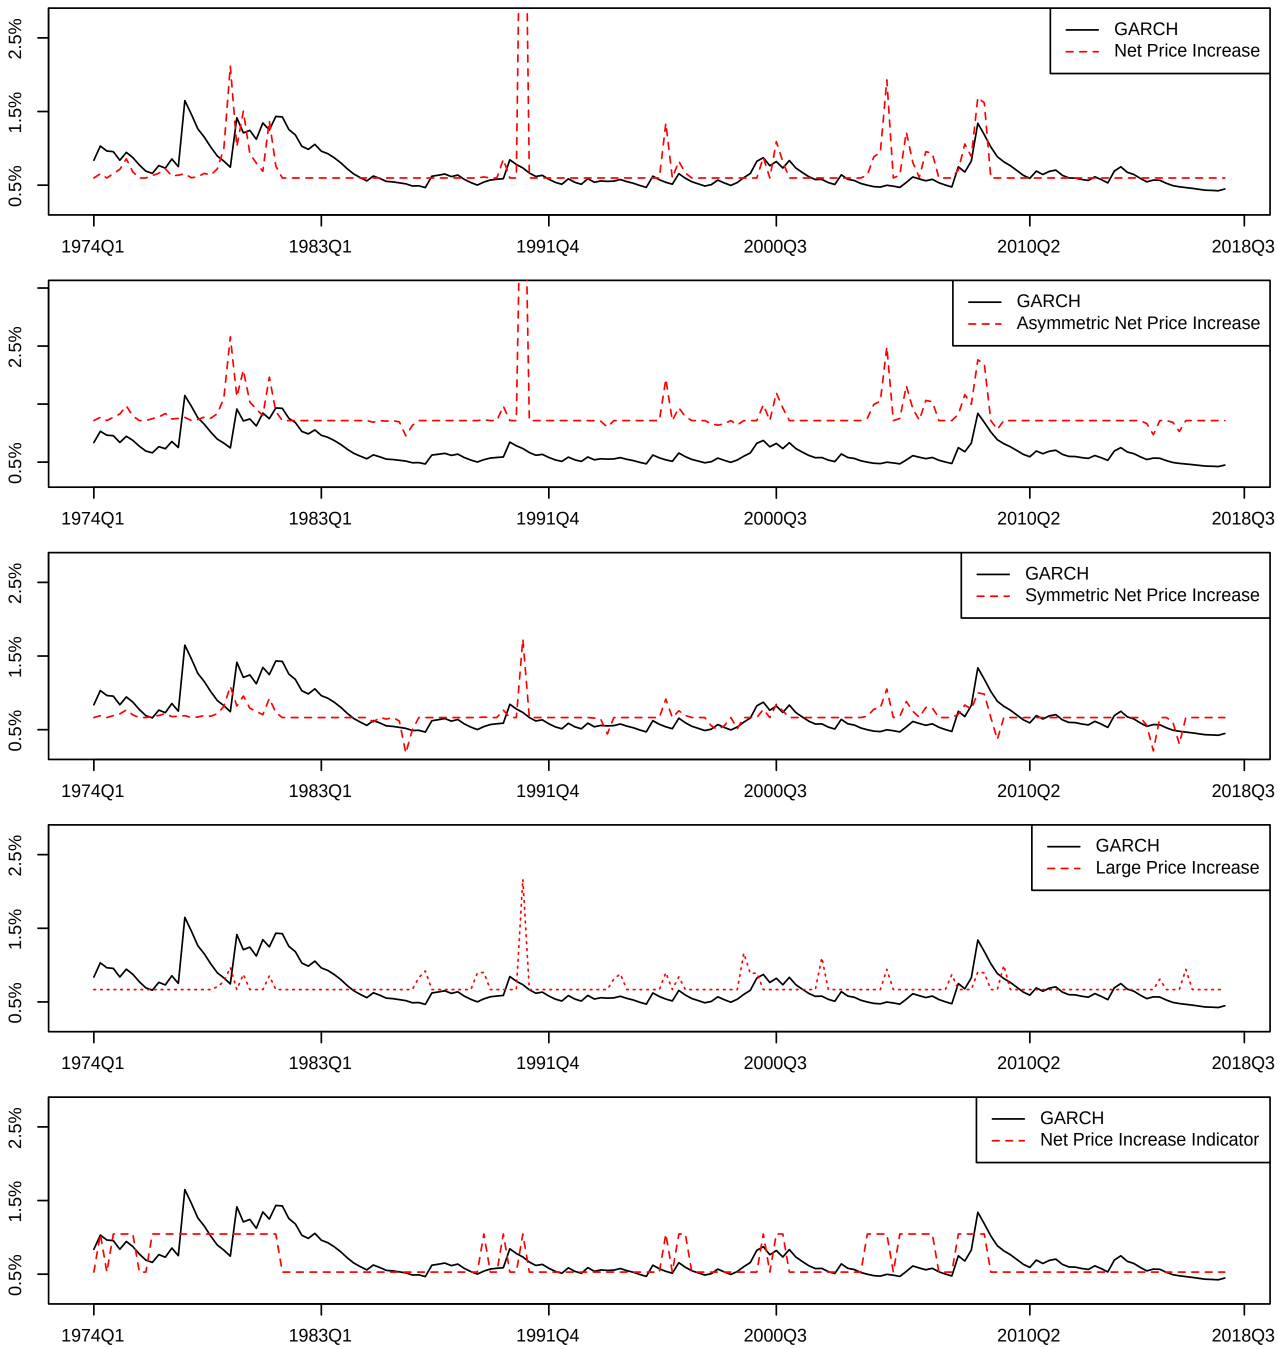

Each line of each sub-figure represents the posterior mean of $\sigma \exp(\delta d_{t-p})$ corresponding to different heteroskedastic shock models along with the AR(1)-GARCH model. For the case of asymmetric net price change, it is the posterior mean of $\sigma \exp(\delta_1 d^+_{t-p} + \delta d^-_{t-p})$. The heteroskedastic shock model in each sub-figure represents the best specification according to Tables 4 - 8.

Figure 4: Cumulative Predictive Bayes Factors of Heteroskedastic Model with Oil Shock against Benchmark GARCH-Shock (Top) and SV (Bottom)

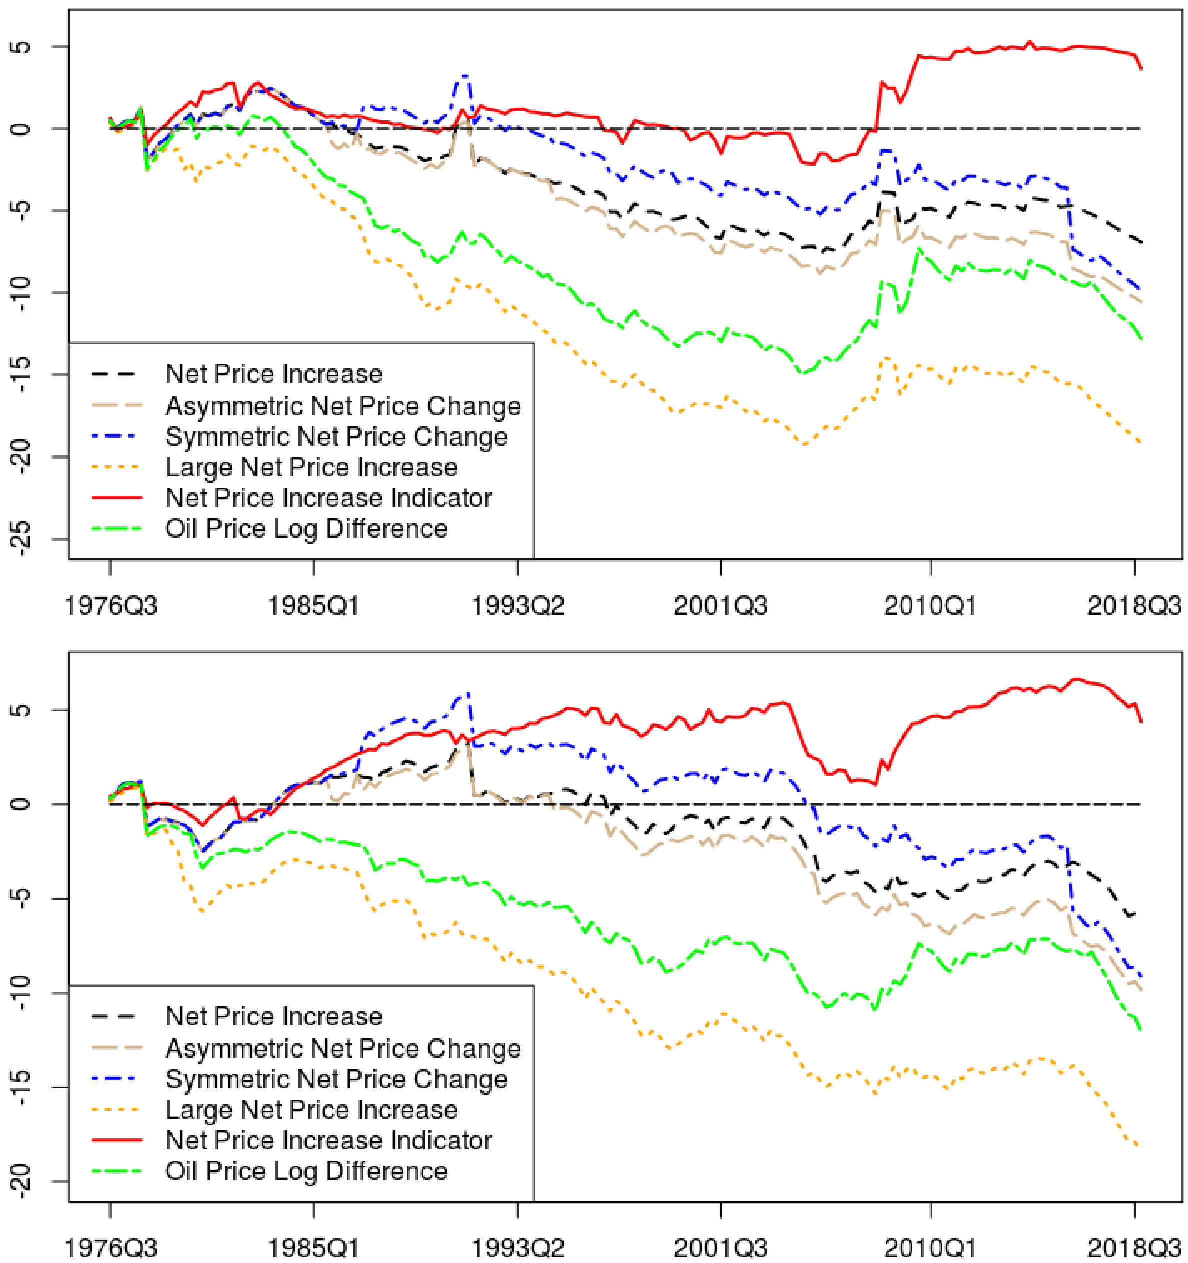

This figure plots the cumulative Bayes factor at each point in time in favour of a model against the AR-GARCH-Shock model of (equation 5) (top plot) and SV (equation 6) model (bottom plot). Various lines represent model (3) with corresponding shock variables.

Figure 5: Histogram of Net Price Increase Shocks, d_t^+

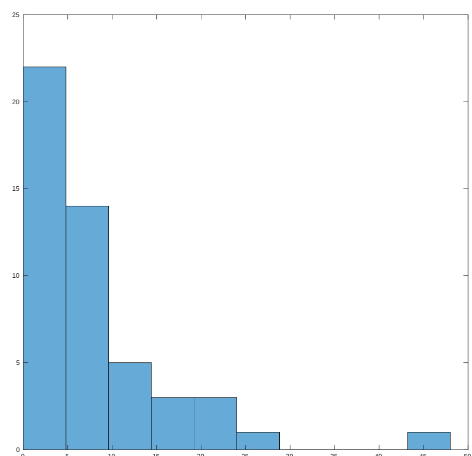

References

Amisano, G. & Giacomini, R. (2007), 'Comparing density forecasts via weighted likelihood ratio tests', *Journal of Business & Economic Statistics* **25**, 177–190.

Clark, T. (2011), 'Real-time density forecasts from Bayesian vector autoregreesions with stochastic volatility', *Journal of Business & Economic Statistics* pp. 327–341.

Croushore, D. & Stark, T. (2011), 'A real-time data set for macroeconomists', *Journal of Econometrics* pp. 111–130.

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

Elliott, G. & Timmermann, A. (2008), 'Economic forecasting', *Journal of Economic Literature* **46**(1), 3–56.

Gneiting, T. & Raftery, A. E. (2007), 'Strictly proper scoring rules, prediction, and estimation', *Journal of the American Statistical Association* **102**(477), 359–378.

Hamilton, J. D. (1983), 'The oil and the macroeconomy since World War II', *Journal of Monetary Economics* **91**, 228–248.

Hamilton, J. D. (1996), 'This is what happened to the oil price-macroeconomy relationship', *Journal of Monetary Economics* **38**, 215–220.

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

Hamilton, J. D. (2011), 'Nonlinearities and the macroeconomic effects of oil prices', *Macroeconomic Dynamics* **15**, 364–378.

Kass, R. E. & Raftery, A. E. (1995), 'Bayes factors', *Journal of the American Statistical Association* **90**, 773–795.

Kilian, L. & Vigfusson, R. J. (2011*a*), 'Are the responses of the U.S. economy asymmetric in energy price increases and decreases', *Quantitative Economics* **2**, 419–453.

Kilian, L. & Vigfusson, R. J. (2011*b*), 'Nonlinearities in the oil price-output relationship', *Macroeconomic Dynamics* **15**, 337–363.

Kilian, L. & Vigfusson, R. J. (2013), 'Do oil prices help forecast US real GDP? the role of nonlinearities and asymmetries', *Journal of Business and Economic Statistics* **31**, 78–93.

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

Mork, K. (1989), 'Oil and the macroeconomy when prices go up and down: an extension of Hamilton's results', *Journal of Political Economy* **97**, 228–248.

Ravazzolo, F. & Rothman, P. (2013), 'Oil and U.S. GDP: A real-time out-of-sample examination', *Jounral of Money, Credit and Banking* **45**, 449–463.

Ravazzolo, F. & Rothman, P. (2016), 'Oil-price density forecasts of US GDP', *Studies in Nonlinear Dynamics and Econometrics* pp. 441–453.

Appendix

The following gives the posterior sampling steps for the models. Let I denote the information set and $I_x(\cdot)$ the indicator function for x.

Sampling Steps for ARX, ARX-A, ARX-1 and ARX-MA

The sampling method is straightforward when we confront linear models like Given conditionally conjugate priors the posterior distribution is conjugate and we apply Gibbs sampler. For simplicity, we use the following matrix form to represent linear models.

$$g = X\overline{\beta} + u \quad u \sim NID(0, \sigma^2 \mathrm{I}),$$

where g is a vector growth rates with dimension of T, which is total number of observations. Let $\overline{\beta}$ denote the parameter vector we are interested. Let $X = [X_1, \ldots, X_t]'$ and the input of each element various according to the following models.

ARX: Let $\overline{\beta}$ =$[\mu, \alpha_{1:q}, \beta_{1:p}]'$ with dimension of $p + q + 1$, and corresponding $X_t = [1, g_{t-1}, \ldots, g_{t-q}, r_{t-1}, \ldots, r_{t-p}]$ for $t = 1, \ldots, T$.

ARX-A: Let $\overline{\beta}$ =$[\mu, a_{1:f}, b_{1:j}]'$ with dimension of $f + h + 1$. Let $X_t = [1, z(q, 0), \ldots, z(q, f), s(p, 0), \ldots, s(p, h)]$ for $t = 1, \ldots, T$, and the X has T rows and $f+h+1$ columns. Please refer to Supplementary Material for details of the polynomial construction.

ARX-1: Let $\overline{\beta}$ =$[\mu, \alpha, \beta]'$. The X has a dimension of T by 3, where each $X_t = [1, g_{t-1}, r_{t-1}]$ for $t = 1, \ldots, T$.

ARX-MA: Let $\overline{\beta}$ =$[\mu, \alpha, \beta]'$. The X has a dimension of T by 3, where each $X_t = [1, \frac{1}{3}\sum_{j=q}^{q+2} g_{t-j}, \frac{1}{3}\sum_{j=p}^{p+2} r_{t-j}]$ for $t = 1, \ldots, T$.

The priors are $\overline{\beta} \sim MN(b, B)$ and $\sigma^{-2} \sim Gamma(\chi, \nu)$. The MN denotes multivariate normal distribution. We set b and B be respectively a vector of zeros and an identity matrix. We set $\chi = 3$ and $\nu = 1$ correspondingly for the prior of σ^{-2}. The conditional posterior distribution for $\overline{\beta}$ and σ^{-2} are the following:

$$\overline{\beta}|\sigma^2, I \sim MN(M, V^{-1}) \quad V = (\sigma^{-2}X'X + B^{-1}) \quad M = V^{-1}(\sigma^{-1}X'g + B^{-1}b)$$

$$\sigma^{-2}|\overline{\beta}, I \sim Gamma\left(\chi + \frac{T}{2}, \nu + \frac{1}{2}u'u\right).$$

Sampling Steps for the Shock Model

Net Price Increase, Symmetric Net Price Change, Large Price Increase, Net Price Increase Indicator

The sampling steps on $\mu, \beta, \lambda, \sigma$ of shock model[11] are similar as they are sampled in previous section. Besides, we apply the Metropolis-Hasting (MH) algorithm on δ due to its non-conjugate feature.

$$g_t = \mu + \beta g_{t-1} + \lambda d_{t-p} + \sigma \exp(\delta d_{t-p}) e_t \qquad e_t \overset{iid}{\sim} N\big(0, 1\big)$$
$$(\mu, \beta, \lambda) \sim MN(b, B), \qquad \delta \sim N(a, A), \qquad \sigma^{-2} \sim Gamma(\chi, \nu)$$

Let b be a vector of zeros and B be an identity matrix. We set $a = 0, A = 1, \chi = 3$ and $\nu = 1$. At each ith MCMC draw, the parameter space is sampled by the following conditional posterior density:

1. Draw $\sigma^{-2(i)} \sim p(\sigma^{-2}|\mu^{(i-1)}, \beta^{(i-1)}, \lambda^{(i-1)}, \delta^{(i-1)}, I)$
2. Jointly Draw $\mu^{(i)}, \beta^{(i)}, \lambda^{(i)} \sim p(\mu|\beta^{(i-1)}, \lambda^{(i-1)}, \delta^{(i-1)}, \sigma^{-2(i)}, I)$
3. Draw $\delta^{(i)} \sim p(\delta|\mu^{(i)}, \beta^{(i)}, \lambda^{(i)}, \sigma^{-2(i)}, I)$

Step 1 and 2 are sampled through the following transformations:

$$\frac{g_t}{\exp\left(\delta d_{t-p}\right)} = \frac{\mu}{\exp\left(\delta d_{t-p}\right)} + \frac{\beta g_{t-1}}{\exp\left(\delta d_{t-p}\right)} + \frac{\lambda d_{t-p}}{\exp\left(\delta d_{t-p}\right)} + \sigma e_t \tag{8a}$$
$$g_t^* = \mu x_{t-1}^* + \beta g_{t-1}^* + \lambda d_{t-p}^* + \sigma e_t \tag{8b}$$

where the x_{t-1}^* is $\frac{1}{\exp\left(\delta d_{t-p}^n\right)}$. The equation (8b) is derived with given δ. Then, we can easily sample the $\{\mu, \beta, \lambda, \sigma^2\}$ under conjugacy. For simplicity, the equation (8b) is rewritten in the following matrix form:

$$g = X\overline{\beta} + u \quad u \sim NID(0, \sigma^2 \mathrm{I})$$

The g and X are respectively a vector of $g_{1:T}$ and T by 3 matrix. Let $X = [X_1, \ldots, X_T]'$ and $X_t = [x_{t-1}^*, g_{t-1}^*, d_{t-p}^*]$ for $t = 1, \ldots, T$. We set $\overline{\beta} = [\mu, \beta, \lambda]'$. The conditional posterior distribution of $\overline{\beta}$ and σ^{-2} are the following:

$$\overline{\beta}|\sigma^{-2}, I \sim MN(M, V^{-1}) \quad V = (\sigma^{-2} X^T X + B^{-1}) \quad M = V^{-1}(\sigma^{-1} X^T g + B^{-1} b)$$

[11] Equation (3)

$$\sigma^{-2}|\overline{\beta}, I \sim Gamma\left(\chi + \frac{T}{2}, \nu + \frac{1}{2}u'u\right)$$

As mentioned above, the δ is sampled through the Metropolis-Hastings random walk algorithm. To simplify the notations, we set $m_t(\mu, \beta, \delta) = \mu + \beta g_{t-1} + \lambda d^n_{t-p}$. The step 3 is then sampled by the following density function,

$$p(\delta|\sigma^2, \mu, \beta, \lambda, x) \propto \exp\big(-\frac{(\delta - a)^2}{2A}\big) \exp\big(-\delta \sum_{t=1}^{T} d_{t-p}\big) \exp\Big(-\frac{1}{2\sigma^2} \sum_{t=1}^{T} \frac{\big(g_t - m_t(\mu, \beta, \delta)\big)^2}{\exp(2\delta d_{t-p})}\Big)$$

with given $\mu^{(i)}, \beta^{(i)}, \lambda^{(i)}$ and $\sigma^{-2(i)}$, the $\delta^{(i)}$ at ith MCMC iteration is sampled such as: We first draw $\delta^{new} = \delta^{(i-1)} + N(0, s)$, where s is the tuning parameter for adjusting the acceptance probability and set $\delta^{old} = \delta^{(i-1)}$. Then, we decide on accepting δ^{new} or keeping δ^{old} according to the following rule,

$$\theta = \min\Big[\frac{p(\delta^{new}|\mu^{(i)}, \beta^{(i)}, \lambda^{(i)}, \sigma^{-2(i)}, I)}{p(\delta^{old}|\mu^{(i)}, \beta^{(i)}, \lambda^{(i)}, \sigma^{-2(i)}, I)}, 1\Big] \tag{9}$$

Next, we draw $u \sim Uniform(0, 1)$, if $u \leq \theta$, set $\delta^{(i)} = \delta^{new}$, otherwise set $\delta^{(i)} = \delta^{old}$.

Asymmetric Net Price Change

$$g_t = \mu + \alpha g_{t-1} + \lambda_1 d^+_{t-p} + \lambda_2 d^-_{t-p} + \sigma \exp(\delta_1 d^+_{t-p} + \delta_2 d^-_{t-p}) e_t \qquad e_t \overset{iid}{\sim} N\big(0, 1\big) \tag{10a}$$

$$(\mu, \beta, \lambda_1, \lambda_2) \sim MN(b, B), \quad \delta_1 \sim N(a_1, A_1), \quad \delta_2 \sim N(a_2, A_2), \quad \sigma^{-2} \sim Gamma(\chi, \nu) \tag{10b}$$

Let b and B be respectively a vector of zeros and an identity matrix. We set $a_1 = a_2 = 0$ and $A_1 = A_2 = 1$ correspondingly. At each ith MCMC draw, we apply the Gibbs sampler to the following conditional posterior density:

1. Draw $\sigma^{-2(i)} \sim p(\sigma^{-2}|\mu^{(i-1)}, \beta^{(i-1)}, \lambda_1^{(i-1)}, \lambda_2^{(i-1)}, \delta_1^{(i-1)}, \delta_2^{(i-1)}, I)$
2. Jointly Draw $\mu^{(i)}, \beta^{(i)}, \lambda_1^{(i)}.\lambda_2^{(i)} \sim p(\mu, \beta, \lambda_1, \lambda_2|\delta_1^{(i-1)}, \delta_2^{(i-1)}, \sigma^{-2(i)}, I)$
3. Draw $\delta_1^{(i)} \sim p(\delta_1|\mu^{(i)}, \beta^{(i)}, \lambda_1^{(i)}.\lambda_2^{(i)}, \sigma^{-2(i)}, \delta_2^{(i-1)}, I)$
4. Draw $\delta_2^{(i)} \sim p(\delta_2|\mu^{(i)}, \beta^{(i)}, \lambda_1^{(i)}.\lambda_2^{(i)}, \sigma^{-2(i)}, \delta_1^{(i)}, I)$

Step 1 and 2 can be sampled under perfect conjugacy if all variables of equation (10a) are divided by $\exp(\delta_1 d^+_{t-p} + \delta_2 d^-_{t-p})$ such as:

$$g^*_t = \mu x^*_{t-1} + \beta g^*_{t-1} + \lambda_1 d^{+*}_{t-p} + \lambda_2 d^{-*}_{t-p} + \sigma e_t \tag{11}$$

The x^*_{t-1} is denoted as $\frac{1}{\exp(\delta_1 d^+_{t-p}+\delta_2 d^-_{t-p})}$. Again, for simplicity, we rewrite the equation (11) as the following matrix form:

$$g = X\overline{\beta} + u \quad u \sim NID(0, \sigma^2 \mathrm{I})$$

Let the g and X denote respectively a vector of $g_{1:T}$ and $X = [X_1, \ldots, X_t]'$, a T by 4 matrix such as $X_t = [x^*_{t-1}, g^*_{t-1}, d^{+*}_{t-p}, d^{-*}_{t-p}]$ for $t = 1, \ldots, T$. Let $\overline{\beta} = [\mu, \beta, \lambda_1, \lambda_2]'$ as the parameter space.

$$\overline{\beta}|\sigma^{-2}, I \sim MN(M, V^{-1}) \quad V = (\sigma^{-2}X^T X + B^{-1}) \quad M = V^{-1}(\sigma^{-1}X^T g + B^{-1}b)$$

$$\sigma^{-2}|\overline{\beta}, I \sim Gamma\left(\chi + \frac{T}{2}, \nu + \frac{1}{2}u'u\right)$$

Due to a lack of conjugacy, the δ_1 and δ_2 are sampled through the Metropolis-Hastings algorithm of random walk. To simplify the notations, we set $m_t = \mu + \beta g_{t-1} + \lambda_1 d^+_{t-p} + \lambda_2 d^-_{t-p}$. The step 3 is sampled with the following joint density,

$$\begin{aligned}&p(\delta_1|\sigma^2, \mu, \beta, \lambda_1, \lambda_2, \delta_2, I)\\ &\propto \exp\Big(-\frac{(\delta_1 - a_1)^2}{2A_1}\Big)\exp\big(-\delta_1 \sum_1^T d^+_{t-p}\big)\exp\Big(-\frac{1}{2\sigma^2}\sum_{t=1}^T \frac{(g_t - m_t)^2}{\exp(2\delta_1 d^+_{t-p} + 2\delta_2 d^-_{t-p})}\Big)\end{aligned}$$

The following example shows posterior sampling steps of $\delta_1^{(i)}$: given $\mu^{(i)}, \beta^{(i)}, \lambda_1^{(i)}, \lambda_2^{(i)}$, $\delta_2^{(i-1)}$ and $\sigma^{-2(i)}$, we first draw $\delta_1^{new} = \delta_1^{(i-1)} + N(0, s)$, where s is the tuning parameter for adjusting the acceptance probability. Let $\delta_1^{old} = \delta_1^{(i-1)}$. Then, we decide on accepting δ_1^{new} or keeping δ_1^{old} according to the following rule,

$$\theta = \min\Big[\frac{p(\delta_1^{new}|\mu^{(i)}, \beta^{(i)}, \lambda_1^{(i)}, \lambda_2^{(i)}, \delta_2^{(i-1)}\sigma^{-2(i)}, I)}{p(\delta_1^{old}|\mu^{(i)}, \beta^{(i)}, \lambda_1^{(i)}, \lambda_2^{(i)}, \delta_2^{(i-1)}\sigma^{-2(i)}, I)}, 1\Big]$$

Next, we draw $u \sim Uniform(0, 1)$, if $u \leq \theta$, set $\delta_1^{(i)} = \delta_1^{new}$, otherwise set $\delta_1^{(i)} = \delta_1^{old}$.

Sampling δ_2 is exactly same manner as sampling the δ_1 with the following joint density:

$$
\begin{aligned}
&p(\delta_2|\sigma^2, \mu, \beta, \lambda_1, \lambda_2, \delta_1, I) \\
&\propto \exp\Big(-\frac{(\delta_2 - a_2)^2}{2A_2}\Big) \exp\big(-\delta_2 \sum_{t=1}^{T} d^-_{t-p}\big) \exp\Bigg(-\frac{1}{2\sigma^2}\sum_{t=1}^{T}\frac{(g_t - m_t)^2}{\exp(2\delta_1 d^+_{t-p} + 2\delta_2 d^-_{t-p})}\Bigg)
\end{aligned}
$$

AR(1)-GARCH(1,1)

The AR(1)-GARCH(1,1) is introduced in section 7.1.

$$g_t = \mu + \beta g_{t-1} + e_t \tag{14a}$$

$$e_t \overset{iid}{\sim} N(0, \sigma_t^2) \tag{14b}$$

$$\sigma_t^2 = \omega_0 + \omega_1 e_{t-1}^2 + \omega_2 \sigma_{t-1}^2 \tag{14c}$$

The sampling approach is the standard Metropolis-Hasting (MH) algorithm with random walk. Each step is sampled through MH. The prior each of the parameter $(\mu, \beta, \omega_0, \omega_1, \omega_2)$ follows a standard normal distribution under the constraints of $\omega_0 > 0$, $\omega_1 > 0$, $\omega_2 > 0$ and $\omega_1 + \omega_2 < 1$. The joint posterior density is:

$$
\begin{aligned}
&p(\mu, \beta, \omega_0, \omega_1, \omega_2) \propto p(\mu)p(\beta)p(\omega_0)p(\omega_1)p(\omega_2)\times \\
&\prod_{t=1}^{T}\frac{1}{\sqrt{2\pi\sigma_t^2}}\exp\Bigg(-\frac{(g_t - \mu - \beta g_{t-1})^2}{2\sigma_t^2}\Bigg) I_{\omega_0>0}(\omega_0) I_{\omega_1>0}(\omega_1) I_{\omega_2>0}(\omega_2) I_{\omega_1+\omega_2<1}(\omega_1, \omega_2)
\end{aligned}
$$

The $p(\mu), p(\beta), p(\omega_0), p(\omega_1)$ and $p(\omega_2)$ are the prior densities. A single move random walk sampler is applied and it is iterated in the following steps:

1. Draw $\mu^{(i)} \sim p(\mu|\beta^{(i-1)}, \omega_0^{(i-1)}, \omega_1^{(i-1)}, \omega_2^{(i-1)}, I)$
2. Draw $\beta^{(i)} \sim p(\beta|\mu^{(i)}, \omega_0^{(i-1)}, \omega_1^{(i-1)}, \omega_2^{(i-1)}, I)$
3. Draw $\omega_0^{(i)} \sim p(\omega_0|\mu^{(i)}, \beta^{(i)}, \omega_1^{(i-1)}, \omega_2^{(i-1)}, I)$
4. Draw $\omega_1^{(i)} \sim p(\omega_1|\mu^{(i)}, \beta^{(i)}, \omega_0^{(i)}, \omega_2^{(i-1)}, I)$
5. Draw $\omega_2^{(i)} \sim p(\omega_2|\mu^{(i)}, \beta^{(i)}, \omega_0^{(i)}, \omega_1^{(i)}, I)$

For the ith MCMC draw, μ can be sampled as: draw $\mu^{new} = \mu^{(i-1)} + N(0, s)$, where s is a tuning parameter used for adjusting the acceptance probability. We set $\mu^{old} = \mu^{(i-1)}$.

Next evaluate the following,

$$\theta = \min\left[1, \frac{p(\mu^{new}|\beta^{(i-1)}, \omega_0^{(i-1)}, \omega_1^{(i-1)}, \omega_2^{(i-1)}, I)}{p(\mu^{old}|\beta^{(i-1)}, \omega_0^{(i-1)}, \omega_1^{(i-1)}, \omega_2^{(i-1)}, I)}\right]$$

Draw a $u \sim Uniform(0,1)$. If $u \leq \theta$ set $\mu^{(i)} = \mu^{new}$, and otherwise set $\mu^{(i)} = \mu^{old}$. The other parameters are sampled in the exactly the same manner.

AR(1)-GARCH(1,1)-Shock

This model is a hybrid model which incorporate both oil shocks and GARCH model. It is introduced in section 7.1.

$$g_t = \mu + \beta g_{t-1} + \exp(\delta d_{t-p}) e_t \tag{16a}$$

$$e_t \overset{iid}{\sim} N\big(0, \sigma_t^2\big) \tag{16b}$$

$$\sigma_t^2 = \omega_0 + \omega_1 e_{t-1}^2 + \omega_2 \sigma_{t-1}^2. \tag{16c}$$

The sampling approach is exactly the same as AR(1)-GARCH(1,1) with one extra step on δ. The prior for AR(1)-GARCH(1,1)-Shock is the same as AR(1)-GARCH(1,1) with $\delta \sim N(0,1)$. The joint posterior density becomes:

$$p(\mu, \beta, \delta, \omega_0, \omega_1, \omega_2) \propto p(\mu)p(\delta)p(\beta)p(\omega_0)p(\omega_1)p(\omega_2)\times$$
$$\prod_{t=1}^{T} \frac{1}{\sqrt{2\pi\sigma_t^2 \exp(2\delta d_{t-p})}} \exp\left(-\frac{(g_t - \mu - \beta g_{t-1})^2}{2\sigma_t^2 \exp(2\delta d_{t-p})}\right) I_{\omega_0>0}(\omega_0) I_{\omega_1>0}(\omega_1) I_{\omega_2>0}(\omega_2) I_{\omega_1+\omega_2<1}(\omega_1, \omega_2)$$

The $p(\mu), p(\beta), p(\omega_0), p(\omega_1)$ $p(\omega_2)$ and $p(\delta)$ represent prior densities. Some restrictions on the sampling: $\omega_0 > 0$, $\omega_1 > 0$, $\omega_2 > 0$ and $\omega_1 + \omega_2 < 1$. The sampling steps are iterated in the following ways:

1. Draw $\mu^{(i)} \sim p(\mu|\beta^{(i-1)}, \omega_0^{(i-1)}, \omega_1^{(i-1)}, \omega_2^{(i-1)}, \delta^{(i-1)}, I)$
2. Draw $\beta^{(i)} \sim p(\beta|\mu^{(i)}, \omega_0^{(i-1)}, \omega_1^{(i-1)}, \omega_2^{(i-1)}, \delta^{(i-1)}, I)$
3. Draw $\omega_0^{(i)} \sim p(\omega_0|\mu^{(i)}, \beta^{(i)}, \omega_1^{(i-1)}, \omega_2^{(i-1)}, \delta^{(i-1)}, I)$
4. Draw $\omega_1^{(i)} \sim p(\omega_1|\mu^{(i)}, \beta^{(i)}, \omega_0^{(i)}, \omega_2^{(i-1)}, \delta^{(i-1)}, I)$
5. Draw $\omega_2^{(i)} \sim p(\omega_2|\mu^{(i)}, \beta^{(i)}, \omega_0^{(i)}, \omega_1^{(i)}, \delta^{(i-1)}, I)$

6. Draw $\delta^{(i)} \sim p(\delta|\mu^{(i)}, \beta^{(i)}, \omega_0^{(i)}, \omega_1^{(i)}, \omega_2^{(i)}, I)$

For ith MCMC draw, the δ can be sampled such that, we draw $\delta^{new} = \delta^{(i-1)} + N(0, s)$, where s is a tuning parameter for adjusting the acceptance probability. Let $\delta^{old} = \delta^{(i-1)}$. Next we evaluate the following,

$$\theta = \min\left[1, \frac{p(\delta^{new}|\mu^{(i)}, \beta^{(i)}, \omega_0^{(i)}, \omega_1^{(i)}, \omega_2^{(i)}, I)}{p(\delta^{old}|\mu^{(i)}, \beta^{(i)}, \omega_0^{(i)}, \omega_1^{(i)}, \omega_2^{(i)}, I)}\right]$$

Draw a $u \sim Uniform(0, 1)$. If $u \leq \theta$ set $\delta^{(i)} = \delta^{new}$, and otherwise set $\delta^{(i)} = \delta^{old}$. The other parameters are sampled in the exactly the same manner.

SV

The AR(1)-SV is introduced in section 8.3.

$$g_t = \mu + \alpha g_{t-1} + \exp(h_t/2)u_t \quad u_t \sim N(0, 1) \tag{18a}$$

$$h_t = \omega_0 + \omega_1 h_{t-1} + \sigma_\nu \nu_t \quad \nu_t \sim N(0, 1) \tag{18b}$$

The sampling approach iterates on the following steps:

1. Draw $\mu^{(i)}, \alpha^{(i)} \sim p(\mu, \alpha|\omega_0^{(i-1)}, \omega_1^{(i-1)}, \sigma_\nu^{(i-1)}, h_{1:T}, I)$
2. Draw $\omega_0^{(i)}, \omega_1^{(i)} \sim p(\omega_0, \omega_1|\mu^{(i)}, \alpha^{(i)}, \sigma_\nu^{(i-1)}, h_{1:T}, I)$
3. Draw $\sigma_\nu^{(i)} \sim p(\sigma_\nu|\mu^{(i)}, \alpha^{(i)}, \omega_0^{(i)}, \omega_1^{(i)}, h_{1:T}, I)$
4. Draw $h_t \sim p(h_t|\mu^{(i)}, \alpha^{(i)}, \omega_0^{(i)}, \omega_1^{(i)}, h_{-t}, I)$, $t = 1, \ldots, T$.

Sampling μ and α can be done jointly. Given $h_{1:T}$, the above model can be rewritten as

$$g_t \exp(-h_t/2) = \mu \exp(-h_t/2) + \alpha g_{t-1} \exp(-h_t/2) + u_t \quad t = 1, \ldots, T$$

Then we can treat above equation as linear regression model and we can write it as matrix form such as $y = X\beta + u$, where $X_t = [\exp(-h_t/2), g_{t-1}\exp(-h_t/2)]$ and $X = [X_1, \ldots, X_T]$, and $\beta = [\mu, \alpha]^T$, and $u \sim MN(0, I)$, where MU denotes multivariate normal distribution. Then with the prior of $\beta \sim MN(b_0, B_0)$, where b_0 is a vector of 0 with dimension of 2 and B_0 is an identity matrix with dimension of 2 by 2.

$$p(\beta|Y,x) \sim MN(M, V^{-1}) \tag{19a}$$
$$\text{where } M = V^{-1}(X^T y + B_0^{-1} b_0),\ V = (X^T X + B_0^{-1}) \tag{19b}$$

Jointly sampling ω_0 and ω_1 can be done similarly as we did above. We write equations as matrix form[12] such as $y = [h_2, \ldots, h_T]$, $\beta = [\omega_0, \omega_1]^T$, $u \sim MN(0, \sigma_\nu I)$ and $X = [X_2, \ldots, X_T]$, where $X_t = [1, h_{t-1}]$. The prior for $\beta \sim MN(b_0, B_0)$,

$$p(\beta|Y,x) \sim MN(M, V^{-1}) \tag{20a}$$
$$\text{where } M = V^{-1}(\sigma_\nu^{-2} X^T y + B_0^{-1} b_0),\ V = (\sigma_\nu^{-2} X^T X + B_0^{-1}) \tag{20b}$$

Sampling σ_ν follows inverse gamma. Using the linear regression form from previous step with $\sigma_\nu^2 \sim IG(\nu/2, s/2)$ we have:

$$\sigma_\nu^2 \sim IG\Big(\frac{T+\nu}{2}, \frac{(y - X\beta)^T(y - X\beta) + s}{2}\Big)$$

Sampling $h_{1:T}$ relies on the data augmentation. Given the proposal $h' \sim N(\overline{\mu}_t, \sigma^2)$, with $\overline{\mu}_t = \mu_t + \frac{1}{\sigma^2}[(r_t - \alpha r_{t-1})^2 \exp(-\mu_t) - 1]$ and

$$\mu_t = \frac{\omega_0(1-\omega_1) + \omega_1(h_{t-1} + h_{t+1})}{1+\omega_1^2}, \quad \sigma^2 = \frac{\sigma_\nu}{1+\omega_1^2}$$

the proposal is accepted with probability

$$\theta = \min\Big[1, \frac{p(h_t'|h_{t-1}, h_{t+1}, \omega_0, \omega_1)/q(h'|h_{t-1}, h_{t+1}, \omega_0, \omega_1)}{p(h_t|h_{t-1}, h_{t+1}, \omega_0, \omega_1)/q(h_t|h_{t-1}, h_{t+1}, \omega_0, \omega_1)}\Big].$$

[12] $y = X\beta + u$