

THE ROLE OF OIL PRICES IN THE FORECASTS OF SOUTH AFRICAN INTEREST RATES: A BAYESIAN APPROACH

Rangan Gupta* & Kevin Kotzé†

May 2015

Abstract

This paper considers whether the use of real oil price data can improve upon the forecasts of the interest rate in South Africa. We employ various Bayesian vector autoregressive (BVAR) models that make use of various measures of oil prices and compare the forecasting results of these models with those that do not make use of this data. The real oil price data is also disaggregated into positive and negative components to establish whether this would improve upon the forecasting performance of the model. The full dataset includes quarterly measure of output, consumer prices, exchange rates, interest rates and oil prices, where the initial in-sample extends from 1979q1 to 1997q4. We then perform rolling estimations and forecasts over the out-of-sample period 1998q1 to 2014q4, after the in-sample period is extended to incorporate an additional observation. The results suggest that models that include information relating to oil prices outperform the model that does not include this information, when comparing their out-of-sample forecasts. In addition, the model with the positive component of oil price tends to perform better than other models at the short- to medium-run horizons. Then lastly, the model that includes both the positive and negative components of the oil price, provides superior forecasts at longer horizons, where the improvement is large enough to ensure that it is the best forecasting model on average. Hence, not only do real oil prices matter when forecasting interest rates, but the use of disaggregate oil price data may facilitate additional improvements.

JEL Classifications: C32, C53, E43, E47, Q41.

Keywords: Interest rate, oil price, forecasting, South Africa.

* *Corresponding author:* Department of Economics, University of Pretoria, Pretoria, 0002, South Africa. Email: rangan.gupta@up.ac.za

†School of Economics, University of Cape Town, Rondebosch, 7700, South Africa. Email: kevin.kotze@uct.ac.za

1 Introduction

Since the seminal contribution of Hamilton (1983), which investigated the effects of oil shocks on the business cycles in the United States, a large number of studies across developed and developing economies have analysed the impact of oil price shocks on macroeconomic and financial variables.¹ Within the set of emerging economies considered, South Africa - an oil importing and inflation targeting small open economy - has featured prominently, with a large corresponding literature devoted to studying the impact of oil shocks on macroeconomic, agricultural and financial variables.² The evidence from this literature largely suggests that oil price shocks have a significant effect domestic macroeconomic variables. In addition, the in-sample vector autoregression-based evidence suggests that the South African central bank responds to oil price shocks when setting interest rates (Aye *et al.*, forthcoming; Chisadza *et al.*, forthcoming).

Against this backdrop, the objective of this paper is to test whether the evidence from in-sample explanatory investigations may be extended to test for the predictive power of oil prices in a forecasting exercise. In this case we make use of the quarterly out-of-sample period 1998q1 to 2014q4, which includes the start of the inflation targeting era that begun in February 2000. The initial in-sample period for the estimation of the parameters in the respective models extends from 1979q1 to 1997q4, which is used to generate the parameters that will influence a subsequent forecast that extends for eight quarters. Thereafter, the in-sample period is extended to 1998q1, before the model parameters are estimated once again to generate the second eight-step ahead forecast. This process continues until we have the last of the forecasts at the end of the out-of-sample.

The baseline model that has been used in this exercise employs a vector autoregressive (VAR) structure that is estimated with Bayesian techniques. This model includes measures of output, price level, exchange rate and interest rate to generate the respective rolling forecasts. The out-of-sample ability of this model is then compared to those that make use of the same structure and estimation techniques, but where the variables are supplemented by the addition of a measure of oil prices. The evaluation is conducted with the aid of root-mean squared-error (RMSE) statistics and measures that consider the significance of any potential improvement, to determine whether the models with information relating to oil prices are able to improve upon the accuracy of the predictions for interest rates.

The use of Bayesian vector autoregressive (BVAR) models allow us to work with integrated data in level-form, where the prior is appropriately specified. This procedure ensures that it would not be necessary to apply necessary trans-

¹For a detailed literature review in this regard, the reader is referred to Balcilar *et al.* (forthcoming).

²See for example, Gupta & Hartley (2013); Gupta & Modise (2013); Aye *et al.* (2014, forthcoming); Balcilar *et al.* (2014a,b,c); Gupta & Kanda (2014); Ajmi *et al.* (2015, forthcoming); Chisadza *et al.* (forthcoming); de Bruyn *et al.* (forthcoming), and references cited therein for earlier works.

formations to the underlying variables to induce stationarity in the data, as with a VAR model that employs a classical estimation approach. This may avoid possible misspecification errors that could be due to the identification of a stochastic trend, the removal of deterministic time or seasonal trends, or the extraction of information relating to possible long-run relationships. Hence, in the BVAR an appropriate prior specification would account for features such as unit roots and cointegration, which we control for in this paper. In addition, we also employ the recent developments that follow Giannone *et al.* (2015), which allow for the prior to adjust to the in-sample properties of the data.³

As a part of the forecasting evaluation, the information content from real oil price data is disaggregated into positive and negative components, to consider whether such disaggregation can produce more accurate out-of-sample forecasts, relative to those that use undecomposed oil price data. In addition, we also seek to detect, which component of the oil price matters most. To the best of our knowledge, this is the first attempt that investigates the importance of oil price data and its component in interest rates forecasts for South Africa.

The rest of the paper is organized as follows. Section 2 describes the methodology, while section 3 provides details of the data. The results are discussed in section 4 and the conclusion is contained in section 5.

2 Methodology

In this section, we briefly outline the features of the model that has been used in the forecasting exercise.⁴ The general specification of a multivariate VAR model may be formally represented as follows:

$$\begin{aligned} y_t &= C + B_1 y_{t-1} + \dots + B_p y_{t-p} + \epsilon_t, \\ \epsilon_t &\sim N(0, \Sigma), \end{aligned} \tag{1}$$

In this case, y_t represents an $n \times 1$ vector of endogenous variables, ϵ_t represents an $n \times 1$ vector of exogenous shocks, and C , B_1 , \dots , B_p and Σ represent matrices that contain the unknown parameters.

When applying noninformative priors in a Bayesian framework, the posterior parameter estimates would approximate those that would have been obtained under the classical approach to estimation. This would imply that when conditioning on the initial observations, the posterior distribution of $\beta \equiv \text{vec}([C, B_1, \dots, B_p]')$ would be centred on the ordinary least square estimates. However, the use of flat priors may not be appropriate as it could lead to inadmissible estimators and poor inference, especially when the number of variables in the VAR structure is relatively large.⁵ The use of poor in-sample estimates

³Therefore, this model has features that are commonly associated with hierarchical or multilevel BVAR models.

⁴This section relies on the discussion in Giannone *et al.* (2015)

⁵See, for example, (Stein, 1956); Sims (1980); Litterman (1986); Bańbura *et al.* (2010); and Koop & Korobilis (2010).

would obviously affect the forecasting results and may lead to inaccurate out-of-sample predictions. To improve upon the forecasting performance of VAR models, the Bayesian literature has proposed that researchers should combine the likelihood function with some informative prior distributions. This strategy would often prove to be successful when the use of informative priors effectively reduces the estimation error.

Following Giannone *et al.* (2015), we consider prior distributions for the VAR coefficients that belong to the following Normal-Inverse-Wishart family:

$$\Sigma \sim IW(\Psi; d), \quad (2)$$

$$\beta|\Sigma \sim N(b, \Sigma \otimes \Omega), \quad (3)$$

where the parameters Ψ, d, b and Ω are summarised by γ , which is a vector of hyperparameters of a lower dimensionality. The model parameters, β and Σ , could then be collected in the vector, θ . The reason for focusing on these priors is that this class includes the priors that are most commonly used in the existing literature on BVARs (Giannone *et al.*, 2015).

As in Kadiyala & Karlsson (1997), the degrees of freedom of the Inverse-Wishart distribution is set to $d = n + 2$, which is the minimum value that guarantees the existence of the prior mean, $\Sigma = \Psi/(d - n - 1)$. In addition, Ψ is assumed to be a diagonal matrix with an $n \times 1$ vector ψ on the main diagonal. Note that, following Giannone *et al.* (2015), we treat ψ as a hyperparameter.

As for the conditional Gaussian prior for β , a combination of the three most popular prior densities used in the existing literature for the estimation of BVARs in levels is used. The baseline prior is a version of the so-called Minnesota prior, introduced by Litterman (1979, 1980). This prior assumes that each variable follows a random walk process, possibly with drift. Formally, this prior may be described by the respective first and second moments:

$$\begin{aligned} \mathbb{E} \left[(B_s)_{ij} \mid \Sigma \right] &= \begin{cases} 1 & \text{if } i = j \text{ and } s = 1 \\ 0 & \text{otherwise} \end{cases} \\ cov \left[(B_s)_{ij}, (B_r)_{hm} \mid \Sigma \right] &= \begin{cases} \lambda^2 \frac{1}{s^2} \frac{\Sigma_{ih}}{\psi_j/(d-n-1)} & \text{if } m = j \text{ and } r = s \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

The above moments could then be used for prior in the VAR model, which are considered in equation (3). Note that the variance of this prior is lower for the coefficients associated with more distant lags, and the coefficients associated with the same variable and lag in different equations are allowed to be correlated. The hyperparameter λ is then used to control the scale of all the variances and covariances.

Following the important contribution of Litterman (1979, 1980), a number of refinements have been made to the Minnesota prior to emphasize the role of unit roots and cointegration (Sims & Zha, 1998). Intuitively, the objective of these refinements is to reduce the importance of the deterministic component, when the estimation of the VAR parameters is conditioned on the initial observations (Sims, 1992).

One such refinement resulted in what is known as the “sum-of-coefficients” prior, which was originally proposed by Doan *et al.* (1984). The estimation under this prior is implemented with the aid of Theil’s (1971) mixed estimation, where a set of n artificial observations are used to suggest that there is no-change in the forecasted values at the beginning of the sample. To achieve this objective, this approach makes use of the following set of dummy observations:

$$\begin{aligned} y_{n \times n}^+ &= \text{diag} \left(\frac{\bar{y}_0}{\mu} \right) \\ x_{n \times (1+np)}^+ &= \begin{bmatrix} 0 \\ n \times 1, y^+, \dots, y^+ \end{bmatrix}, \end{aligned}$$

where \bar{y}_0 is an $n \times 1$ vector containing the average of the first p observations for each variable, and the expression $\text{diag}(v)$ denotes the diagonal matrix that has a vector v on the main diagonal. These artificial observations are added to the dataset, such that the matrices would take the form $y \equiv [y_{p+1}, \dots, y_T]'$ and $x \equiv [x_{p+1}, \dots, x_T]'$. These matrices may then be used for inference, where the prior implied by these dummy observations is centred at one for the sum of coefficients on own lags for each variable, and at zero for the sum of coefficients on the lags of the other variables. It also allows for correlation among the coefficients on each variable in each equation. The hyperparameter μ is then used to control the variance that is associated with the prior beliefs. For example, if $\mu \rightarrow \infty$ the prior becomes uninformative, while if $\mu \rightarrow 0$ it implies the presence of a unit root in each equation, which rules out the existence of cointegration.

Since the sum-of-coefficients prior is not consistent with cointegration at the limit, Sims (1993) developed an additional prior, which is known as the “dummy-initial-observation” prior. Formally, this prior may be implemented with the aid of the following dummy observations:

$$\begin{aligned} y_{1 \times n}^{++} &= \frac{\bar{y}_0'}{\delta} \\ x_{n \times (1+np)}^{++} &= \begin{bmatrix} 1 \\ \frac{1}{\delta}, y^{++}, \dots, y^{++} \end{bmatrix}. \end{aligned}$$

This expression states that a no-change forecast for all variables is a good forecast at the beginning of the sample. In this case, the hyperparameter δ controls the tightness of the prior, which is implied after including the additional external observation. As $\delta \rightarrow \infty$, the prior becomes uninformative, while as $\delta \rightarrow 0$, all the variables of the VAR are forced to be at their unconditional mean. This would be the case where the system is characterized by the presence of an unspecified number of unit roots without drift, which would imply that the dummy-initial-observation prior is consistent with cointegration.

In essence, the settings for these priors depends on the hyperparameters λ, μ, δ and ψ . As this general specification of hyperpriors for λ, μ and δ , embodies the work of Sims & Zha (1998), we follow their choice and make use of Gamma densities with modes equal to 0.2, 1 and 1, while the standard deviations are set equal to 0.4, 1 and 1, respectively. As in Giannone *et al.* (2015), we then

set the hyperprior for $\psi/(d - n - 1)$, which characterizes the prior mean of the main diagonal of Σ , such that it takes an Inverse-Gamma distribution with scale and shape equal to 0.02, which is proper, but quite disperse, as it has neither a variance nor a mean.

We are then able to apply Bayes' law to derive more appropriate values for the hyperparameters, which are contained in γ . Hence, the posterior of the hyperparameters may be derived as

$$p(\gamma|y) \propto p(y|\gamma) \cdot p(\gamma),$$

where $p(\gamma)$ is the prior for the hyperparameters and $p(y|\gamma)$ is the marginal likelihood. In this particular case, inference from the marginal likelihood may be attributed to both the hyperparameters and the model parameters, θ . Hence, the marginal likelihood function may be expressed as,

$$p(y|\gamma) = \int p(y|\theta, \gamma) p(\theta|\gamma) d\theta,$$

This procedure allows for the optimal selection of the model hyperparameters, based on the in-sample fit of model. As a result, the properties of the data would influence the degree to which the prior informs the posterior, where both a dogmatic and a flat prior are nested within the possible outcomes.

3 Data

The data used in this study covers the quarterly period of 1979q1 to 2014q4, with the start and end date being purely driven by data availability. The variables that were used include real gross domestic product (GDP, seasonally adjusted with a base year of 2010), consumer price index (CPI, seasonally adjusted with a base year of 2012), three-month Treasury bill rate, real effective exchange rate, and real oil price. The real oil price is based on the nominal Western Texas Intermediate (WTI) crude oil price (Cushing, Oklahoma) quoted in U.S. dollars per barrel deflated by the U.S. CPI (seasonally adjusted with base years of 1982 to 1984).

Note that, following the studies by Chisadza *et al.* (forthcoming) and Aye *et al.* (forthcoming), and the recent trend in the international oil price shock literature, as discussed in Baumeister *et al.* (2010) (and references cited therein), we do not convert the oil price into domestic currency, i.e. rand values (using the rand-dollar exchange rate). This practice facilitates the identification of the effects of (exogenous) oil shocks, when investigating the respective forecasts of the interest rate. The real gross domestic product data was obtained from the South African Reserve Bank, while the consumer price index, interest rate and the real effective exchange rate was obtained from the International Financial Statistics database of the International Monetary Fund.⁶ Data on the WTI oil price and the U.S. CPI was obtained from the Global Financial database.

⁶This choice of the variables is largely consistent with the literature on open-economy interest-rate rules, as considered in Ball (1999) and several others.

Sims *et al.* (1990) show that when adopting a Bayesian approach that is based entirely on the likelihood function, the associated inference does not need to specifically account for nonstationarity. This is due to the fact that the likelihood function will continue to adopt a Gaussian shape, despite the presence of nonstationarity. Hence, barring the interest rate, all variables have been specified in their natural logarithmic form, and have been plotted in Figure 1.⁷

To decompose the real oil price into its positive and negative components, we make use of the method of Granger & Yoon (2002), which may be applied to data that follows the behaviour of a unit root process. To confirm that this particular time series is integrated of the first order, we perform the Augmented Dickey-Fuller (ADF, 1979) and Phillips and Perron (PP, 1988) tests. The results of these statistics are reported in Table 2, which is in the appendix of the paper.

Following Granger & Yoon (2002), we then define the real oil price (OP) process as a random walk with the appropriate moving average representation,

$$\begin{aligned} OP_t &= OP_{t-1} + u_{1t} \\ &= OP_0 + \sum_{i=1}^t u_{1i} \\ &= OP_0 + \sum_{i=1}^t u_{1i}^+ + \sum_{i=1}^t u_{1i}^-, \end{aligned} \tag{4}$$

where $t = 1, 2, \dots, T$. The initial value of oil price is OP_0 and u_{1i} indicates a white noise error term, which is defined as the sum of positive and negative shocks, $u_{1i} = u_{1i}^+ + u_{1i}^-$. In this case the positive shocks may be represented by $u_{1i}^+ = \max(u_{1i}, 0)$, and the negative shocks may be represented by $u_{1i}^- = \min(u_{1i}, 0)$. The cumulative form of the positive oil price shocks is then $OP^+ = \sum_{i=1}^t u_{1i}^+$. Similarly, the cumulative form of the negative oil price shocks is $OP^- = \sum_{i=1}^t u_{1i}^-$. This provides us with two additional measures of real oil prices that represent positive and negative real oil price shocks.

The advantage of this procedure is that we are not only able to test whether those forecasts that include real oil prices as an explanatory variable are able to provide more accurate forecasts of the South African interest rate, but we are also able to ascertain whether positive and negative oil price shocks provide unique information when forecasting the interest rate. This would be of particular relevance to cases where the central bank is faced with nominal downside rigidities in the pricing mechanism. Under such conditions the central bank would be more inclined to react to positive oil price shocks over the short-term. However, over a longer time period the pricing mechanism would have more time to adjust and the effect of both positive and negative shocks would be of importance (were the long-term prices could possibly be approximated by the flexible-price level).⁸

⁷These plots of the data are contained in the appendix.

⁸Ball & Mankiw (1994) and Cabral & Fishman (2012) consider the theoretical foundations of an asymmetric price adjustment mechanism, while Peltzman (2000) provides empirical support.

4 Results

Given that our objective is to analyse the importance of the oil price when generating forecasts for the interest rate, our benchmark model is the model that does not include any measure of oil price, which we denote BVAR1. The forecasts from this model are then generated from parameters that have been estimated on data for the interest rate, output, consumer prices and the real effective exchange rate.

We then compare the four models that include various measures of the real oil price with the benchmark model. The additional models are organised as follows: (i) BVAR2 denotes the model with real oil price data, in addition to the four variables that are included in BVAR1; (ii) BVAR3 denotes the model with the positive component of the real oil price, in addition to the variables in BVAR1; (iii) BVAR4 denotes the model with the negative component of the real oil price, in addition to the variables in BVAR1; and (iv) BVAR5 denotes the model with the positive and negative components of real oil prices, in addition to the variables in BVAR1. Hence, the BVAR1 model is nested within the respective BVAR2, BVAR3, BVAR4 and BVAR5 models.

To conduct our out-of-sample forecasting exercise, we divide the total sample period into an in-sample period (1979q1 to 1997q4) and an out-of-sample period (1998q1 to 2014q4), with the models being estimated recursively over the latter period. Since we produce one- to eight-quarter-ahead forecasts, this choice implies that the evaluation period starts exactly two-years (i.e. eight-quarters) before the point at which the South African central bank formally adopted the inflation targeting regime (which was during the first quarter of 2000). In addition to this particular structural break, the use of the Bai & Perron (2003) tests of multiple structural breaks also indentifies three other potential breaks (1985q2, 1998q4 and 1999q3) in the out-of-sample period, which were obtained from the interest rate equation with the real oil price included.

The initial evaluation of the forecasts is performed after calculating the root mean square errors (RMSEs) for the one- to eight-quarter-ahead forecasts. When reporting the results we provide the actual RMSEs for the BVAR1 model and the relative RMSEs for the models that include oil prices. Therefore, a value of less one for the relative RMSE of the BVAR2 model would suggest that it provides a lower RMSE, when compared to the BVAR1 model.

Thereafter, we also analyse whether the differences in the forecasts that are generated from the respective models (that either include or exclude measures of the real oil price) are statistically significant. For this purpose, we use the *MSE-F* test statistic proposed by McCracken (2007), which investigates whether or not two models have equal population-level predictive ability. Hence, this statistic is used to test the null hypothesis that the restricted BVAR1 model and the unrestricted BVAR2, BVAR3, BVAR4 and BVAR5 models have equal forecasting ability. In addition, as the BVAR3 model is nested within the BVAR5 model we also apply the test to these forecasts. The null is then tested against the one-sided alternative hypothesis, that the *MSE* for the unrestricted model

forecasts, (\hat{MSE}_1) , is less than the MSE for the restricted model forecasts, (\hat{MSE}_0) . Formally, the statistic is given as:

$$MSE-F = (T - R - h + 1) \cdot \bar{d} / \hat{MSE}_1 \quad (5)$$

where T is the total sample, R is the number of observations used in estimation of the model from which the first forecast is derived (i.e. the in-sample portion of the total number of observations). In this case, $\hat{MSE}_i = (T - R - h + 1)^{-1} \sum_{t=R}^{T-h} (u_{i,t+1})^2$, where $i = 1, 0$, and $\bar{d} = MSE_0 - MSE_1$, with u_i being the forecast error. Note that $h = (1, 2, \dots, 8)$ is the forecasting horizon. A positive and significant $MSE-F$ statistic indicates that the unrestricted model forecasts are statistically superior to those of the restricted model.

The RMSEs obtained from the BVAR1 through BVAR5 models for the one-to eight-quarter-ahead forecasts for the interest rate, as well as the results from the $MSE-F$ tests, are reported in Table 1. We make the following observations: (i) The RMSEs for models BVAR2 to BVAR5 relative to the BVAR1 model (which does not include real oil price), is always below unity, implying that the inclusion of the real oil price in the respective models would result an improved forecasting performance; (ii) More importantly, based on the $MSE-F$ test statistic, these improvements are statistically significant, at least at the 5% level of significance; (iii) Within the group of models that include measures of the real oil price in its various forms, the BVAR3 (i.e. the model with the positive component of the real oil price) is the best performing model for horizons 1 to 5. However, beyond that, the BVAR5 model (i.e., the model with both positive and negative components of the real oil price) provides more impressive forecasting results. The gains at longer horizons, that are derived from the BVAR5 model, causes it to produce the lowest relative RMSE on average, when comparing the results from all the models. The BVAR3 finishes a close second following its superior performance at short- to medium-run horizons; (iv) Finally, since the BVAR5 nests the BVAR3 model, the one sided $MSE-F$ test reveals that the former significantly outperforms the latter at the 5 percent level of significance for horizons 6 till 8.⁹

This suggests that not only does the real oil price matter when forecasting the South African three-month Treasury bill rate, but there are added forecasting gains to be derived from disaggregating oil price data into its positive and negative components. This could be due to a number of reasons, as the central bank may react more strongly to positive shocks over the short- to medium-term horizon, as is possibly the case in South Africa.

⁹The two-sided Diebold & Mariano (2002) test has a non-standard distribution under nested models. However, though not ideal, when this test is applied to compare the BVAR3 and the BVAR5 models, we find that the BVAR3 outperforms the BVAR5 significantly (at conventional levels of significance) for horizons 1 to 5, while the BVAR5 outperforms the BVAR3 significantly (at conventional levels of significance) for the remaining forecasting horizons. Complete details of these results are available upon request from the authors.

5 Conclusion

Against the backdrop of some in-sample evidence that oil price shocks affect South African interest rates, this paper investigates whether real oil price movements could assist in the (out-of-sample) forecasting of the three-month Treasury bill rate in South Africa. A number of BVAR models were used to generate the respective forecasts, where variants of the model would either exclude or include various measure of oil prices. The use of such Bayesian models has allowed us to work with the data in level-form, without any need to induce stationarity with the aid a particular transformation. In addition, the choice of priors is essentially determined by the in-sample properties of the data as a part of a hierarchical Bayesian approach. Together, these features of the model may avoid important instances of potential model misspecification.

We further disaggregate the real oil price into its positive and negative component, to analyse whether such disaggregation can produce more accurate out-of-sample forecasts relative to the real oil price that has not been decomposed. The models are applied to a dataset that comprise of interest rates, real output, consumer prices and the real effective exchange rate; and in certain cases, measures of oil price and its components. This dataset is based on an out-of-sample period of 1998q1 to 2014q4, with an initial in-sample of 1979q1 to 1997q4.

The results suggest that models that include oil prices in various forms are able to producing forecasts that outperform those model that only make use of data for the interest rate, output, consumer prices and the real effective exchange rate. In addition, we note that the BVAR model with the positive component of oil price tends to perform better relative to all other models at short- to medium-run horizons, but the model that includes both the positive and negative components of oil price performs provides superior forecasts over the six- to eight-quarter-ahead horizons. The gains at the longer horizon from the latter model is large enough to ensure that the BVAR model with both positive and negative components of the oil price outperforms the other models on average.

Hence, not only does the real oil price matter when forecasting the three-month Treasury bills rate in South Africa, but additional performance gains may be derived from disaggregating the oil price into its positive component at shorter horizons, and from both positive and negative components at longer horizons. Even though we analyse the role of oil price asymmetry in forecasting the interest rate, as part of future research, it would be interesting to compare our results from the linear BVAR models with time-varying (and hence nonlinear) BVAR models. This would allow us to not only account for nonlinearity in the relationship between the respective variables, but it would also for the accommodation of possible breaks (which we show exist) within the in-sample.

References

- AJMI, A.N., GUPTA, R., BABALOS, V. & HEFER, R. 2015. Oil price and consumer price nexus in South Africa revisited: A novel asymmetric causality approach. *Energy, Exploration & Exploitation*, 33(1):63–74.
- AJMI, A.N., GUPTA, R., KRUGER, M., SCHOEMAN, N. & WALTERS, L. forthcoming. The Nonparametric Relationship between Oil and South African Agricultural Prices. *Energy Economics*.
- AYE, G.C., DADAM, V., GUPTA, R. & MAMBA, B. 2014. Oil price uncertainty and manufacturing production. *Energy Economics*, 43:41–47.
- AYE, G.C., GADINABOKAO, O. & GUPTA, R. forthcoming. Does the South African Reserve Bank (SARB) Respond to Oil Price Movements? Historical Evidence from the Frequency Domain. *Energy Sources, Part B: Economics, Planning, and Policy*.
- BAI, J. & PERRON, P. 2003. Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, 18(1):1–22.
- BALCILAR, M., CHANG, S., GUPTA, R., KASONGO, V. & KYEI, C. 2014a. The Relationship between Oil and Agricultural Commodity Prices: A Quantile Causality Approach. Working paper no. 201468, University of Pretoria, Department of Economics.
- BALCILAR, M., GUPTA, R. & MILLER, S.M. 2014b. Regime Switching Model of US Crude Oil and Stock Market Prices: 1859 to 2013. *Energy Economics*, 49:317–327.
- BALCILAR, M., UWILINGIYE, J. & GUPTA, R. forthcoming. Dynamic Relationship between Oil Price and Inflation in South Africa. *Journal of Developing Areas*.
- BALCILAR, M., VAN EYDEN, R., UWILINGIYE, J. & GUPTA, R. 2014c. The impact of oil price on South African GDP growth: A Bayesian Markov Switching-VAR analysis. Working paper no. 201470, University of Pretoria, Department of Economics.
- BALL, L. & MANKIW, G.N. 1994. Assymmetric Price Adjustments and Economic Fluctuations. *The Economic Journal*, 104:247–261.
- BALL, L.M. 1999. Policy Rules for Open Economies. In *Monetary Policy Rules*, NBER Chapters, pages 127–156. National Bureau of Economic Research, Inc.
- BAÑBURA, M., GIANNONE, D. & REICHLIN, L. 2010. Large Bayesian vector auto regressions. *Journal of Applied Econometrics*, 25(1):71–92.
- BAUMEISTER, C., PEERSMAN, G. & VAN ROBAYS, I. 2010. The economic consequences of oil shocks: Differences across countries and time. In *Inflation in an era of relative price shocks*, pages 91–128. Reserve Bank of Australia.
- CABRAL, L. & FISHMAN, A. 2012. Business as Usual: A Consumer Search Theory of Sticky Prices and Asymmetric Price Adjustment. *International Journal of Industrial Organisation*, 30(4):371–376.
- CHISADZA, C., DLAMINI, J., GUPTA, R. & MODISE, M.P. forthcoming. *The impact of oil shocks on the South African economy*.

- DE BRUYN, R., GUPTA, R. & VAN EYDEN, R. forthcoming. Forecasting the rand-dollar and rand-pound exchange rates using dynamic model averaging. *Emerging Markets Finance and Trade*.
- DICKEY, D.A. & FULLER, W.A. 1979. Distribution of the estimatoes for autoregressive time series with a unit root. *Journal of American Statistical Association*, 74(366):427–431.
- DIEBOLD, F.X. & MARIANO, R.S. 2002. Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 20(1):253–263.
- DOAN, T., LITTERMAN, R. & SIMS, C. 1984. Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews*, 3(1):1–100.
- GIANNONE, D., REICHLIN, L. & PRIMICERI, G. 2015. Prior Selection for Vector Autoregressions. *The Review of Economics and Statistics*, 2(97)(4):436–451.
- GRANGER, C.W. & YOON, G. 2002. Hidden cointegration. Working paper, (2002-02), University of California, Economics.
- GUPTA, R. & HARTLEY, F. 2013. The role of asset prices in forecasting inflation and output in South Africa. *Journal of Emerging Market Finance*, 12(3):239–291.
- GUPTA, R. & KANDA, P. 2014. Does the Price of Oil Help Predict Inflation in South Africa? Historical Evidence Using a Frequency Domain Approach. Working paper no. 201401, University of Pretoria, Department of Economics.
- GUPTA, R. & MODISE, M.P. 2013. Does the source of oil price shocks matter for South African stock returns? A structural VAR approach. *Energy Economics*, 40:825–831.
- HAMILTON, J.D. 1983. Oil and the macroeconomy since World War II. *The Journal of Political Economy*, pages 228–248.
- KADIYALA, K. & KARLSSON, S. 1997. Numerical methods for estimation and inference in Bayesian VAR-models. *Journal of Applied Econometrics*, 12(2):99–132.
- KOOP, G. & KOROBILIS, D. 2010. Bayesian Multivariate Time Series Methods for Empirical Macroeconomics. *Foundations and Trends in Econometrics*, 3:267–358.
- LITTERMAN, R.B. 1979. Techniques of forecasting using vector autoregressions. Working paper 115, Federal Reserve of Minneapolis.
- LITTERMAN, R.B. 1980. A Bayesian Procedure for Forecasting with Vector Autoregression. Technical Report, Massachusetts Institute of Technology, Department of Economics.
- LITTERMAN, R.B. 1986. Forecasting with Bayesian Vector Autoregressions: Five Years of Experience. *Journal of Business & Economic Statistics*, 4(1):25–38.
- MCCRACKEN, M.W. 2007. Asymptotics for out of sample tests of Granger causality. *Journal of Econometrics*, 140(2):719–752.
- PELTZMAN, S. 2000. Prices Rise Faster Than They Fall. *Journal of Political Economy*, 108:466–502.

- PHILLIPS, P.C. & PERRON, P. 1988. Testing for a unit root in time series regression. *Biometrika*, 75(2):335–346.
- SIMS, C.A. 1980. Comparison of Interwar and Postwar Business Cycles. *American Economic Review*, 70(2):250–257.
- SIMS, C.A. 1992. Interpreting the macroeconomic time series facts: The effect of monetary policy. *European Economic Review*, 36:975–1011.
- SIMS, C.A. 1993. A nine-variable probabilistic macroeconomic forecasting model. In *Business Cycles, Indicators and Forecasting*. University of Chicago Press.
- SIMS, C.A., STOCK, J.H. & WATSON, M.W. 1990. Inference in linear time series models with some unit roots. *Econometrica*, 58(1):113–144.
- SIMS, C.A. & ZHA, T. 1998. Bayesian methods for dynamic multivariate models. *International Economic Review*, 39(4):949–968.
- STEIN, C. 1956. Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. In *Proceedings of the Third Berkeley symposium on mathematical statistics and probability*, volume 1(399), pages 197–206.
- THEIL, H. 1971. *Principles of Econometrics*. New York: John Wiley.

Model	Forecast Horizons									
	1 step	2 step	3 step	4 step	5 step	6 step	7 step	8 step	Average	
BVAR1	0.0117	0.019	0.0226	0.0258	0.0286	0.0308	0.0322	0.0335	0.0255	
BVAR2	0.9954**	0.9886**	0.9827**	0.9783**	0.9750**	0.9719***	0.9700***	0.9689***	0.9789	
BVAR3	0.9916**	0.9779**	0.9586***	0.9393***	0.9197***	0.9000***	0.8817***	0.8677***	0.9296	
BVAR4	0.9981**	0.9901**	0.9736**	0.9550***	0.9360***	0.9170***	0.8948***	0.8771***	0.9427	
BVAR5	0.9929**	0.9809**	0.9619***	0.9422***	0.9211***	0.8983*** [**]	0.8738*** [**]	0.8547*** [**]	0.9282	

Table 1: Out-of-Sample (1998q1-2014q2) Root-Mean Square Errors (RMSEs)

Notes: BVAR1 includes GDP, CPI, interest rate and real effective exchange rate; BVAR2: BVAR1 + real oil price; BVAR3: BVAR1 + positive component of real oil price; BVAR4: BVAR1 + negative component of real oil price; BVAR5: BVAR4 + negative component of real oil price; Entries corresponding to BVAR1 is the absolute RMSEs, while entries corresponding to the other BVARs are RMSEs relative to BVAR1. Bold entries indicate the minimum RMSEs. **[**] and *** indicates significance of the MSE-F test statistic at 5 percent [5 percent] level of significance, and 1 percent level of significance respectively, when comparing BVAR2, BVAR3, BVAR4, or BVAR5 with BVAR1 [BVAR5 with BVAR3]

A Appendix

Tests		Real Oil Price	(Real Oil Price)
ADF	C	-1.69	-9.56***
	C+T	-1.93	-9.62***
	N	-0.37	-9.61***
PP	C	-1.72	-9.03***
	C+T	-1.73	-9.06***
	N	-0.05	-9.09***

Table 2: Unit Root Tests

Notes: C(C+T)[N] stands for tests with a constant (constant and trend) [none] in the equation specification.*** indicates rejection of the null at 1 percent level of significance; ADF (Dickey and Fuller, 1979), PP (Phillips and Perron, 1988).

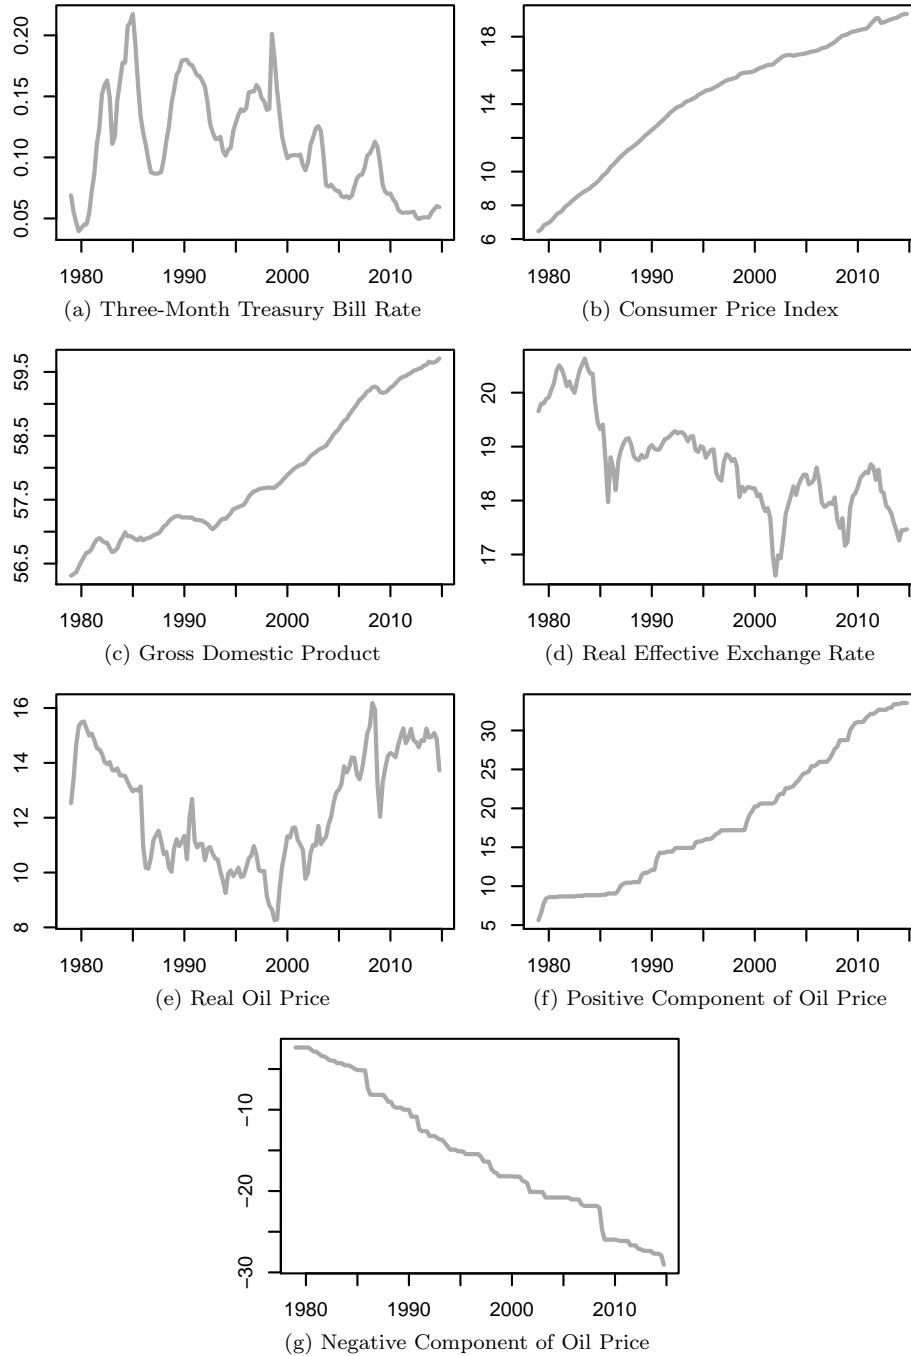


Figure 1: Data Plots