Targeted Filtration with a Difference for Inference and Long-Horizon Forecasts: Application to the Real Price of Crude Oil

Stephen Snudden
Department of Economics
Dunning Hall, Room 209
Queen's University
94 University Avenue
Kingston, ON K7L 3N, Canada
Office: (613) 533-6660
snudden@econ.queensu.ca

February 2016

Filtration by growth rates with targeted lag selection can influence inference and improve long-horizon forecast accuracy. This paper proposes the method of targeted filtration. The method targets lower-frequencies of the data which correspond to respective forecast horizons. Frequency-dependent structural relationships can be targeted to influence inference and dynamic properties. The method is applied to auto-regressive models of the global market for crude oil. The frequencies targeted vary structural shock persistence and the corresponding oil market elasticities. Targeted filtration can significantly improve forecast precision at horizons of up to four years.

JEL classification: C1, C53, Q43
Keywords: Forecasting and Prediction Methods, Oil Prices, Filters, Spectral Analysis
This research was supported by the Social Sciences and Humanities Research Council of Canada. SSHRC Award Number: 767-2013-2637.

Econometric forecasting methods have difficulty achieving high degrees of accuracy at long horizons in small samples, particularly in the markets for crude oil (Alquist *et al.* 2013; Baumeister and Kilian, 2011). This paper proposes the method of targeted filtration, a modification to the standard auto-regressive model construction method. Lags in growth rate filtration are chosen to target select lower-data frequencies. Low frequency oil price movements are of primary interest to physical investors in the oil industry and policy decisions at central banks (Hamilton, 2009; Kilian, 2009). The method of targeted filtration is found to significantly improve long-horizon forecast mean-squared precision and directional accuracy in both univariate and multivariate auto-regressive methods.

The method of targeted filtration can be understood using spectral analysis to explain how selecting an appropriate lag in growth rates can target specific frequencies of the data. The method removes select high frequencies and emphasises select low frequencies which correspond to respective forecast horizons. Parametric estimates of the spectral density using univariate models show that targeted filtration can improve model identification of frequencies corresponding to horizons beyond one year.

In particular, targeted filtration is studied and applied to auto-regressive models of the global market for crude oil. The application to univariate models serves as an intuitive way to illustrate the modification to the standard univariate Box-Jenkins methodology (Box *et al.* 2008). Moreover, Alquist *et al.* (2013) have shown that these methods beat no-change and futures forecasts for the real price of crude oil for horizons less than one year. When targeted filtration is combined with univariate methods, model forecasts can achieve the same degree of accuracy at two and three years that these model have previously only been able to achieve at the one-month horizon and can produce significant mean-squared precision and directional accuracy of forecasts for up to four years.

In addition, targeted filtration can modify inference in SVAR models. By targeting select frequencies, structural shocks correspond to the preserved lower frequencies. With frequency-dependent structural relationships, targeting specific frequencies may change model inference and can have implications for the persistence, signs, and magnitudes of impulse response functions. The method is applied to SVAR models of the global market for crude oil (Kilian, 2009; Kilian and Murphy, 2014). The elasticities are found to be increasing in shock persistence.

A common finding is that VAR models of the global market for crude oil are unable to beat no-change forecasts at horizons beyond one year (Alquist *et al.* 2013; Baumeister and Kilian, 2011). Moreover, the accuracy of forecasts can evolve over time as has been shown by Baumeister *et al.* (2014) and Baumeister and Kilian (2015). When targeted filtration is applied to VAR forecasts, the method can improve long-run forecasts. The method exhibits some robustness over sample periods and alternative oil price series. Targeted filtration for VAR forecasts is applied to the West Texas Intermediate (WTI) and US Refiners' import prices of crude oil.

A systematic investigation of oil demand and oil inventory series is examined for VAR forecasts at longer horizons. An interesting finding is that the use of world industrial production rather than the global activity index produces superior forecast at longer horizons. Moreover, U.S. oil inventory are found to produce superior forecasts of U.S. oil price series. The extra predictive power of U.S. crude oil inventories for U.S. oil price series is consistent with similar findings for short-run forecasts by Baumeister *et al.* (2015).

The paper is structured as follows. Section 2 introduces the intuition for targeted filtration methods for model construction using spectral analysis. Section 3 analyses these methods when applied to univariate methods and compares parametric estimates of the spectral density. Section 4 extends the analysis to structural multivariate models and discusses the implication for inference. Section 5 examines the ability of the multivariate models with targeted filtration to forecast alternative price series of crude oil. Section 6 concludes.

2. Filtration Method

Filtration using growth rates is well documented, see for example Hamilton (1994). Higher lags in growth rates are often applied to produce more intuitive scales. For example, macroeconomic data is often announced as year-over-year growth rates due to its intuitive format. This section provides the intuition behind targeted filtration and why departing from the standard first difference may be desirable for some applications.

Let Y be a covariance-stationary series with absolutely summable autocovariances. Let $s_Y(\omega)$ be the population spectrum and $g_Y(\kappa)$ be the autocovariance generating function of Y, where

$$s_Y(\omega) = (2\pi)^{-1} g_Y(e^{-i\omega}), \tag{1}$$

and ω is the frequency, $\omega \in (0, \pi)$. Let X be a transformation of Y given by $X = h(L)Y$, with $\sum_{-\infty}^{\infty} |h_j| < \infty$. The autocovariance generating function of X is known to be calculated from Y by:

$$g_X(\kappa) = h(\kappa)h(\kappa^{-1})g_Y(\kappa) \tag{2}$$

The population spectrum of X is:

$$s_X(\omega) = (2\pi)^{-1} h(e^{-i\omega})h(e^{i\omega})g_Y(\omega) \tag{3}$$

Hence the population spectrum of X is related to that of Y by:

$$s_X(\omega) = h(e^{-i\omega})h(e^{i\omega})s_Y(\omega) \tag{4}$$

Thus, applying the $h(L)$ filter to Y is the same as multiplying the spectrum of Y by $h(e^{-i\omega})h(e^{i\omega})$. The original series must be co-variance stationary to that $s_Y(\omega)$ exists otherwise the population spectrum is not zero at frequency zero.

The growth rate of Y can be approximated by applying the first difference filter to the log of Y. Hence, $h(e^{-\omega}) = 1 - L$, and $h(e^{-i\omega})h(e^{i\omega})$ is given by $(1 - e^{i\omega})(1 - e^{-i\omega}) = 2 - 2\cos(\omega)$. Define the operator, L^Z, such that $L^Z x_t = x_{t-Z}$ for $Z \in \mathbb{R}$. More generally, filtration of the logged variable Y using the growth rate on the Z^{th} lag, $(1 - L^Z)$, to produce the spectrum of X, is equivalent to filtering the spectrum of Y by

$$h(e^{-iz\omega})h(e^{iz\omega}) = 2 - 2\cos(Z\omega) \tag{5}$$

Figure 1 graphs the growth rate filter applied to the spectral density of Y using $Z = 1,\ 2,\ 3$ and 6, respectively. The growth rate filtration on the first lag, $Z = 1$, preserves high frequencies, maximized at $\omega = \pi$, but removes the lowest frequencies, minimized at $\omega = 0$. Hence the first difference filter is an example of a high-pass filter as it retains higher frequencies. Applying growth rates on the third lag, (q-o-q for monthly data) places zero weights at frequencies $\omega \in \{0, 2/3\pi\}$, and maximum weights at frequencies $\omega \in \{1/3\pi, \pi\}$.

As shown in Figure 2, the frequency, ω, can be converted into the period of a cyclical function, T, using $T = 2\pi/\omega$. Thus, for a monthly time series, q-o-q growth rates preserves cycles two and six periods in length. Models that use growth rates on the first lag do not preserve all information on cycles beyond one year since these frequencies are removed in filtration. Conversely, taking the growth rate beyond the first lag preserves select lower frequencies.

The method of targeted filtration can be defined as follows. Target the lag in the growth rate, Z^*, to maximize the variance at frequencies related to the targeted horizon, K.

First consider the use of targeted filtration for forecasting. In practice, a forecaster estimates a model for each forecast horizon K, using a selected Z^*. Importantly, the filter not only keeps information on the select frequencies but also excludes higher frequencies. For integer values of $Z > 1$, there are $Z+1$ extrema. Hence, the longer lags in differencing imply that zero weights are applied more often. The exclusion of higher frequencies is advantageous. Simply put, the identification of lower frequencies is a small-sample problem, and the elimination of some higher frequencies allows the model estimation to identify the lower frequencies.

The filtered series depends on the filter as well as the spectral density of the original series. The choice of the cycle to target with Z^* depends on the available sample size and the forecast horizon. When the forecast horizon is not too long given the amount of data, the forecaster can choose Z^* to target a frequency that is a maximized at the forecast horizon K. Given less data or lower frequencies, one can choose Z^* to target a fraction of the full cycle. For example, targeting a frequency which corresponds to half forecast horizon K, means that the model forecasts horizon K using the second cycle of targeted frequency.

Despite the non-linearity of the spectral filter, converting between frequency and time-domain, and the choice of cycle targeted, the method can be summarized by a simple rule. Let $K, C \in \mathbb{N}$ and $Z \in \mathbb{R}^+$ so that the frequency corresponding to horizon K/C is targeted. Then the optimal lag in the growth rate filtration, Z^*, is given by $Z^* = K/2C$. For example, to target the full cycle ($C = 1$) or a third of the full cycle ($C = 3$), the rule is to select $Z^*_{C=1} = K/2$ and $Z^*_{C=3} = K/6$.

Now consider the use of targeted filtration for inference in multivariate structural VARs (SVARs). First, it should be clear that applying a growth rate or moving average is an explicit filtration which has implications for inference. When some series are growth rate filtered with $Z = 1$ and other series are kept in levels, inference at lower frequencies explicitly excludes some information from the filtered series. Hence, the choice of filtration in multivariate models may influence inference at lower-frequencies.

When lower frequencies are preserved via growth rate filtration beyond the first lag, the structural shocks correspond to the preserved lower frequencies. With frequency dependent structural relationships, targeting specific frequencies may change model inference and can have implications for the persistence, signs, magnitudes of impulse response functions. This is explored in more detail in section 4.

Despite the simplicity of the filter, targeted filtration has nice qualities for forecasting. First, it is a backward looking filter, with no end-point bias. The method is easily implemented, and the level of a series can always be recovered. Moreover, as some higher frequencies are preserved, estimates are more stable in small samples. In contrast, applying a low-pass or band-pass filter to a series prior to estimation can also target low frequencies, but is known to result in over-fitting and introduces end point bias (see Azevedo and Pereira, 2013).

Growth rates, g_t, can be calculated using per cent change, $g_t = (y_t - y_{t-1})/y_{t-1}$ or approximated using log differences, $\tilde{g}_t = (1 - L)ln(y_t)$. Since, $(1 - L)lny_t = g_t - \frac{1}{2}g_t^2 +$

$\frac{1}{3}g_t^3 \ldots$, it follows that $g_t = \tilde{g}_t - e_t$, where $e_t = \sum_{n=2}^{\infty}(-1)^{n+1}g_t^n/n$. As e_t is a lower order of magnitude than g_t it is a close approximation for low values of g_t. Percent change is a non-linear filter, so it is not possible to derive exact simple rules for targeted filtration as for the log-difference filter. However, the rules for targeted filtration derived above are still a first-order approximation of the growth rate filter. Consider when the two growth rate filters differ.

First, if $\frac{\partial g}{\partial Z} > 0$ then $|\frac{\partial e}{\partial Z}| > 0$. This implies that taking larger lags when log-differencing will create a larger discrepancy between these methods. This suggests that the rules for targeted filtering may only be a good approximation at low values of Z when applied to percent change. Second, suppose we want to calculate dynamic forecasts for $ln\hat{y}_{T+1|T}$ with a estimated value of $\hat{\tilde{g}}_{T+1|T}$. Then generally $ln\hat{y}_{T+k|T} = lny_T + \sum_{i=1}^{k}(\hat{g}_{t+i|T} + \hat{e}_{t+i|T})$. Thus, the approximation error between the two methods can compound in forecasting at long forecasting horizons. The trade-offs of these two filtration methods are explored further in section 3.

3. Application to Univariate Methods

This section applies targeted filtration to univariate forecast models of the global market for crude oil. This illustrates the modification to the standard univariate Box-Jenkins methodology (Box *et al.* 2008). Moreover, these methods are a well documented benchmark for forecasting of real oil prices and have been shown to beat no-change forecasts for the real price of crude oil for horizons of six months and less (Alquist *et al.*, 2013).

The real U.S. dollar price of oil is measured as the monthly nominal U.S refiners' acquisition cost of crude oil imports deflated by the consumer price level. Both series are from the U.S. Energy Information Administration (EIA). Figure 3 graphs the log level as well as the filtered series with $Z = 1$ and $Z = 6$. All raw series are found to be weakly stationary. Estimation in log levels is standard in this literature under the assumption that the real oil price series is co-variance stationary (Alquist *et al.*, 2013; Kilian, 2009).

The model is estimated using maximum likelihood. Let y be a stationary process, ϕ_l denote the k autoregressive parameter, c a constant, and ε_t the error. The AR representation takes the following form:

$$y_t = c + \sum_{l=1}^{p} \phi_l y_{t-l} + \varepsilon_t, \tag{6}$$

The parametric spectral density can be calculated as shown by Box *et al.* (2008) as:

$$f(\omega; \phi, \sigma_\varepsilon^2, \gamma_0) = \frac{\sigma_\varepsilon^2}{2\pi\gamma_0} \frac{1}{1 - \phi_1 e^{-i\omega} - \phi_2 e^{-i2\omega} - \cdots - \phi_p e^{-ip\omega}} \quad (7)$$

Where $\omega \in [0, \pi]$, and γ_0 is the variance of the variable, and σ_ε^2 is the variance of the error.

Figure 4 illustrates the sample parametric spectral density estimation of the AR models for the real price of crude oil estimated in log-levels and when y_t is filtered in growth rates, where Z^* targets frequencies corresponding to two year horizon with $C \in \{1, 2, 3\}$. Lags are selected using the AIC criterion, (Akaike, 1974). The parametric sample density is very similar for, models estimated with 12 lags so is not shown. The area under the sample periodogram represents the portion of the variance of the series attributable to cycles of different frequencies, ω. The sample periodograms show that the targeted growth rate filtration increases the portion of the variance of the series attributable to the frequencies that is 24 periods in length.

The forecast performance is expected to be superior at horizons that correspond to frequencies where the maximum portion of the variance is explained. Autoregressive models with fixed lag lengths of 12 are evaluated at several forecast horizons to provide a direct comparison with Alquist *et al.* (2013). The models are estimated over the 1974.1-2009.9 sample period. The parameters of the model are estimated recursively in each time period, and the out-of-sample dynamic forecasts are evaluated against observed data for the 1991.12-2009.9 period.

Forecast criteria reported include the mean squared prediction error (MSPE) and success ratios. Model MSPE is expressed as a ratio relative to the no change forecast. Success ratios are calculated to quantify the accuracy of the forecast direction and represent the fraction of times the forecast correctly predicts the direction of change in the real price of oil. All forecast criterion are evaluated in the real levels of the price of oil. The p-values are calculated by bootstrapping the model under the null following Pesaran and Timmermann (2009). Models include those estimated in log-levels, $Z = 1$, and using targeted filtration with $C \in \{1, 2, 3\}$ and Z^* given by $Z^* = K/2C$.

As shown in Table 2, the model estimated in log-levels does best at horizons less than 6 months in both precision and directional accuracy. At horizons between 6 months and three years targeted filtration with $C = 1$ produces superior forecasts to the model in log-levels or in growth rates with $Z = 1$. At horizons of two years and beyond, lower values of C become desirable as the small sample makes identification of lower-frequencies more difficult. This also holds when the constant is dropped from the estimations. Overall, targeted filtration produces superior forecasts for up to four years when compared to models estimated in log levels or with $Z = 1$.

Interestingly, directional accuracy of the forecasts using targeted filtration is quite high for all horizons beyond one year relative to log levels. Success ratios greater than 0.6 are found for horizons of 12 months and beyond. These are high values when compared to the empirical finance literature (Pesaran and Timmermann 1995) and are on par with short-horizon success ratios of the multivariate modeling method of Baumeister and Kilian (2013). The identification of the lack of trend is guaranteed only asymptotically. The directional accuracy of the targeted forecasts falls when the constant is dropped, particularly at horizons of three years and beyond. Overall, the use of targeted filtration robustly achieves superior forecasts to models estimated in levels and with $Z = 1$ at horizons greater than six months and less than three years.

Table 3 reports the auto-regressive model estimated using targeted filtration in log-differences. The results are very similar to the model filtering in per cent change, particularly for horizons of three years and less. The use of log-differences provide slightly superior MSPE accuracy compared to per cent change. One downside of using log-differences in practice is that if a non-zero constant is estimated in sample, it is applied linearly. Thus, at very long horizons (ten years or more), the forecast can became unbounded and the real price may become less than zero. This suggests some advantage to calculating growth rates using percent change rather than log differences at very long-forecast horizons for small samples.

Forecasts based on real-time data sets rather than historical data are preferable for model evaluation. To keep the exposition of the targeted filtration method as simple as possible, the analysis is limited to historical data. The evaluation of the method of targeted filtration using real-time data sets is left to future research. For a discussion of real-time data in forecasting the real price of crude oil, see Baumeister and Kilian (2011) and Baumeister *et al.* (2014).

The insight that data transformations are explicit filtrations is a useful insight. Previous efforts to forecast the real price of oil have focused on log-level or log-differenced data, so it should be expected that these models are designed to forecast well at shorter horizons. Moreover, other forms of filters including moving average filters could also be beneficial for targeting select lower frequencies. For example, using data at lower observation frequencies, such as annual data, is a form of a moving average filter. This may explain why there is some evidence that time series models that explicitly model trends and estimated at annual frequency can outperform random walk forecasts in the long run (see, for example, Bernard, *et al.*, 2015)

4. Inference in SVARs

A shortcoming of univariate models is that they lack an interpretation of the economic shocks driving the oil price dynamics. Understanding the source of price movements is of primary interest to market observers. This section shows how targeted filtration applied to structural vector auto regression (SVAR) can influence inference. The SVAR takes the following form:

$$B_0 y_t = c + \sum_{l=1}^{p} B_l y_{t-l} + \alpha D_t + \epsilon_t \tag{8}$$

where ϵ_t is a vector of orthogonal structural innovations, c is the vector of constants, D_t is a matrix of seasonal dummy variables and α is the corresponding coefficient matrix, B_l, $l = 1, \ldots, p$ is the matrix of autoregressive coefficients, and y_t is a vector of endogenous variables.

The three endogenous variables include global crude oil production, world industrial production, and the real price of crude oil. The structural errors are identified using a recursive assumption exploiting institutional knowledge of the global market for crude oil as motivated by Kilian (2009). Hence, structural shocks to oil supply, oil demand, and other oil demand shocks are identified. Oil inventories are excluded for simplicity in exposition, since identification of speculative, demand-driven price movements requires sign restrictions (Baumeister and Peersman, 2013; Kilian and Murphy, 2014).

As in the last section, the real price of crude oil is the U.S refiners' acquisition cost of crude oil imports deflated by the consumer price index as reported by the EIA. Global demand is proxied using world industrial production (WIP) from the International Monetary Fund's International Financial Statistics database. Global crude oil supply is that reported by the EIA. All series are monthly. Alternative measures of oil demand and prices are considered in the next section but do not substantially change the qualitative results. As in the univariate cases, the SVAR is estimated over the 1974.1-2009.8 sample period.

The recursive identification method is valid only for small values of Z and is examined up to $Z = 3$ consistent with the recursive identification motivated by Kilian (2009). It would be difficult to motivate the recursive structure and assume that WIP would not respond to oil price movements greater than one quarter ago. Table 1 reports the elasticities of the SVAR when the price equation is estimated in log-levels, $Z = 1$ and $Z = 3$. The oil supply and oil demand series are estimated in growth rates with $Z = 1$ for the models where the real price of oil is estimated using log-levels or with growth rates using $Z = 1$, but are in growth rates with $Z = 3$ for the model where the price equation is estimated in growth rates with $Z = 3$.

Unlike the model estimated in log levels, the model with the price equation estimated using $Z = 1$ and $Z = 3$ can significantly identify the sign of structural supply shocks. In particular, as supply rises, prices fall and production rise. The demand shock is identified at the 10 per cent significance level for all models consistent with the definition of Kilian (2009). In particular, as demand rises, prices and supply rise. The significance of the demand and supply response increases in Z.

Table 1: Structural Shock Persistence and Elasticities

	Supply Shock			Demand Shock		
	Persistence	Ave. Elas.	Impact Elas.	Persistence	Ave. Elas.	Impact Elas.
Level	6	-0.001	0.000	15	0.110*	0.027*
Z=1	6	-0.165	-0.109	15	3.146*	2.566*
Z=3	12	-0.387	-0.311	16	2.663*	2.620*

Order from Kilian (2009): global supply, world industrial production, U.S. refiners crude oil import price. Persistance is the number of periods the shock is signifiganty different than zero at 90% signifigance. Ave. Elas. is the average price elasticity for the period the shock is signifigant. Impact Elas. is the impact elasticity calculated from the first period. * indicates significant price response

The structural shocks in the estimated SVARs exhibit more persistence for higher values of Z, consistent with the preservation of lower frequencies in data transformation. This provides a method to examine the persistent dependency of elasticities. For example, the impact price elasticity of supply is over two times larger when the supply shock is twice as persistent, as seen from models estimated with $Z = 1$ and $Z = 3$. The persistence dependency of the elasticities is an often overlooked factor when comparing elasticities as shock persistence can vary across models. We leave a complete investigation of the persistence dependency of the elasticities to further research.

The perspective of data transformations as filtration raises questions when series in multivariate structural models are filtered to different frequencies. In particular, combining various data transformations may not preserve equal opportunity of all series to explain the dependent variable at all frequencies. This may be a concern in previous models of the oil market (e.g. Kilian, 2009; Kilian and Murphy, 2014) that combine the log-level of prices and a global activity index based on global shipping rates in Kilian (2009), both of which exhibits low frequency adjustment, with the first difference of supply and inventories. A similar frequency filtration for supply and inventories could be considered to provide these series an equal opportunity to explain lower frequency prices movements.

To formally test the explanatory power of targeted filtration for the price of oil we employ the Breitung and Candelon (2006) Granger causation tests in the frequency domain. The null that the variables in the VAR do not cause the price of oil at frequency ω is tested for the log-level, $Z = 1$, and targeted filtered price of crude oil at all frequency horizons considered in sections 3 and 5. Is is found that the VAR variables significantly Granger Cause oil prices (jointly and individually) at almost all considered horizons beyond one year when the data is targeted filtered but not when in levels or first difference. This Granger causation at longer horizons in the frequency domain suggests improved explanatory power and potential improvements in forecasting performance at these frequencies when targeted filtration is employed.

5. Application to VAR Forecasts

This section evaluates the forecast performance of targeted filtration when applied to standard VAR models. To begin, a systematic investigation of alternative oil demand and oil inventory series are conducted to evaluate forecast performance at longer horizons. The method of target filtration is then applied and found to further improve forecast performance at longer horizons. The method exhibits some robustness over sample periods and holds generally across alternative oil price series.

The VAR model of the previous section is augmented by the inclusion of crude oil inventories. Two oil inventory series are evaluated: U.S. crude oil inventories and the OCED crude oil inventories (Kilian and Murphy, 2014). Moreover, two oil demand series are evaluated: world industrial production (WIP), and the global activity index (GAI) of Kilian (2009). Recursive, dynamic, out-of-sample forecasts are conducted over the 1991.8–2009.9 evaluation period. Models are estimated using 24 lags. This is close to the lags selected by the LR criterion (Lutkepohl, 2005, 143–144) and ensures that residuals of the models pass the Portmanteau (Q) test for white noise (Box and Pierce, 1970; Ljung and Box, 1978). Moreover, using a fixed lag length allows for clearer comparison across the choice of Z as lags are identical across models. The results are found to be qualitatively similar given other lag selection criterion.

The MSPE and success ratios of the model forecasts, with real oil prices in log levels or with $Z = 1$ are shown in Table 4. U.S. oil inventories are generally found to produce superior forecasts of U.S. oil price series. The extra predictive power of U.S. crude oil inventories for U.S. oil price series is consistent with the improvement in predictive power for short run-forecasts by Baumeister *et al.* (2015). However, the largest improvements in forecasts are made by using WIP rather than the GAI. Forecasts from models with real prices in log-levels and with $Z = 1$ estimated with WIP and U.S. inventories can significantly beat the random walk in MSPE precision at the two year horizon.

The method of targeted filtration is applied to the models using WIP and U.S. inventories to see if forecast performance can be further improved at longer horizons. As shown in Table 5, the application of targeted filtration improves the MSPE ratios compared to models where the price of oil is estimated in log-levels or with $Z = 1$. The MSPE ratios of the model with $C = 3$ are 0.74 and 0.83 at the two and three year horizon, respectively. These are lower than the results previously found at this horizon in the literature by Alquist *et al.* (2013) and Kilian and Baumeister (2011). The success ratios of the model with $Z = 3$ also improve over the models where the real price of oil is estimated in log-levels or with $Z = 1$. At the three year horizon the success ratio of the model with $Z = 1$ is slightly better than with $C = 3$. As often found with VAR forecasts, there exists some trade-off between mean-square precision and directional forecast accuracy.

Tables 6 report the MSPE and success ratios for the model estimated using WTI prices and U.S. inventories. The estimation begins in 1981.11 as the WTI series is not back-casted due to potential failures to preserve lower-frequency information. Over the 2000.1-2009.9 evaluation period, the model estimated with $C = 3$ produces superior MSPE and success ratios at the one- to three-year horizon compared to the models estimated with $Z = 1$ and in log levels. Overall, for the 1991.11 to 2009.9 period, application of targeted filtration to VAR forecasts of the real price of crude oil can help improve long-horizon forecasts.

Applications of VAR forecast techniques have found that no individual model performs well over all sample periods (Baumeister *et al.*, 2014; Baumeister and Kilian, 2015; Manescu and Robays, 2014). Tables 7 and 8 report MSPE ratios and success ratios for the sub-periods considered in Manescu and Robays (2014). Targeted filtration with $C = 3$ exhibits robust MSPE performance at the two year horizon, consistently beating the model estimated with $Z = 1$ and in log levels for all sample periods apart from 2008.1 to 2009.12. At the three year horizon targeted filtration with $C = 3$ consistently beats the model estimated with $Z = 1$ apart from the 2008.1 to 2009.12 sub period, but is outperformed by the model estimated in log levels. The success ratios with $C = 3$ also consistently exceed those of the models estimated with $Z = 1$ and in log levels for all sample periods at the two-year horizon.

6. Conclusion

This paper suggests an amendment to traditional auto-regressive model construction methods when conducting forecasts at longer horizons. Lags in growth rate filtration are chosen to target select lower-data frequencies. This is in contrast to the standard approach of estimating in levels which makes lower-frequencies hard to identify in small samples or which takes first differences and explicitly filters lower frequencies of the data. These lower frequency price movements are of primary interest for investment decisions in the oil industry and policy decisions at central banks.

Targeted filtration has an advantage as a backward-looking filter. Hence, the end-point bias that is often associated with centered moving average filtration is not present. Moreover, the method is easily implemented, and the level of a series can always be recovered. When applied to the global market for crude oil, targeted filtration can improve long-horizon forecast precision and directional accuracy in both univariate and multivariate methods. Targeted filtration can provide similar degrees of precision at horizons up to three years that have previously been found at shorter horizons. Moreover, this method exhibits robustness over model choice, variables, and sub-periods.

Targeted frequencies can provide insight into the time dependency of short-run elasticities. In particular, targeted filtration can also influence the significance, signs, and persistence of structural dynamics. The discussion also highlights potential complications from applying different filtrations to variables in a VAR model. The combining of filtrations may not provide equal opportunity to explain the dependent variable at all frequencies.

Application of targeted filtration can provide improvements in longer-horizon forecast accuracy. Due to the compromises of forecast evaluation criterion, this method may be combined with forecast combinations as in Baumeister *et al.*, (2014) or Baumeister and Kilian (2015) to achieve more robust long-term forecasts. The innate time dependency of elasticities also warrants further investigation.

References

Akaike H. 1974. A new look at the statistical model identification. *IEEE Transactions on Automatic Control* **19**: 716–723.

Alquist R, Kilian L. 2010. What do we learn from the price of crude oil futures? *Journal of Applied Econometrics*, **25(4)**: 539–573.

Alquist R, Kilian L, Vigfusson R. 2013. Forecasting the price of oil. In *Handbook of economic forecasting* Elliott G, Timmermann A (eds) Vol. 2. Newnes.

Azevedo J, Pereira A. 2013. Macroeconomic Forecasting using low-frequency filters. Banco de Portugal Working Paper No. 2003-01.

Baumeister C, Kilian L. 2011. Real-Time Forecasts of the Real Price of Oil. *Journal of Business and Economic Statistics* **30(2)**: 326–336.

Baumeister C, Kilian L. 2015 Forecasting the Real Price of Oil in a Changing World: A Forecast Combination Approach, *Journal of Business and Economic Statistics* **33(3)**, 338–351.

Baumeister C, Kilian L, Lee T. 2014. Are There Gains from Pooling Real-Time Oil Price Forecasts? *Energy Economics* **46(S1)**: S33–S43.

Baumeister C, Peersman G. 2013. The Role Of Time Varying Price Elasticities In Accounting For Volatility Changes In The Crude Oil Market. *Journal of Applied Econometrics* **28(7)**: 1087–1109. DOI: 10.1002/jae.2283

Bernard, J. T., Khalaf, L., Kichian, M., and Yelou, C. 2015. Oil Price Forecasts for the Long-Term: Expert Outlooks, Models, or Both?. Cahier de recherche/Working Paper, 3.

Box G, Jenkins G, Reinsel G. 2008. *Time Series Analysis: Forecasting and Control.* 4th ed. Hoboken, NJ: Wiley.

Box G, Pierce D. 1970. Distribution of residual autocorrelations in autoregressive-integrated moving average time series models. *Journal of the American Statistical Association* **65**: 1509–1526.

Breitung, J., and Candelon, B. 2006. Testing for short-and long-run causality: A frequency-domain approach. *Journal of Econometrics*, **132(2)**: 363-378.

Hamilton J. 1994. *Time series analysis* Vol. 2. Princeton: Princeton university press.

Hamilton J. 2009. Causes and Consequences of the Oil Shock of 2007-08. *Brookings Papers on Economic Activity* **1**: 215–261.

Kilian L. 2009. Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market. *American Economic Review* **99(3)**: 1053-69.

Kilian L, Murphy D. 2014. The Role of Inventories and Speculative Trading in the Global Market for Crude Oil. *Journal of Applied Econometrics* **29(3)**: 454–78. DOI: 10.1002/jae.2322

Ljung G, Box G. 1978. On a Measure of Lack of Fit in Time Series Models. *Biometrika* **65**: 297–303.

Lutkepohl H. 1993. *Introduction to Multiple Time Series Analysis*. 2nd ed. New York: Springer.

Manescu C, Robays I. 2014. Forecasting the Brent Oil Price: Addressing Time-Variation in the Forecast Performance. European Central Bank Working Paper No. 1735.

Pesaran M, Timmermann A. 1995. Predictability of Stock Returns: Robustness and Economic Significance. *Journal of Finance* **50**: 1201–1228.

Pesaran M, Timmermann A. 2009. Testing Dependence Among Serially Correlated Multicategory Variables. *Journal of the American Statistical Association* **104**: 325–337.

Figure 1: Spectral Density of Growth Rate Filter at Z: Frequency Domain

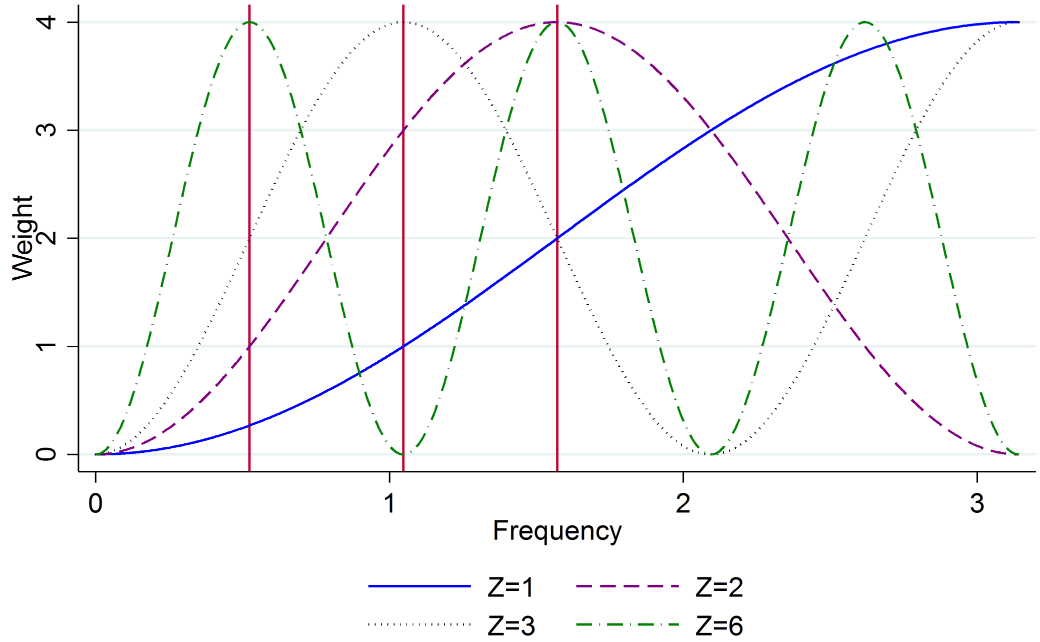

Red vertical line corresponding to frequencies 12, 6, and 4 periods in length

Figure 2: Spectral Density of Growth Rate Filter at Z: Time Domain

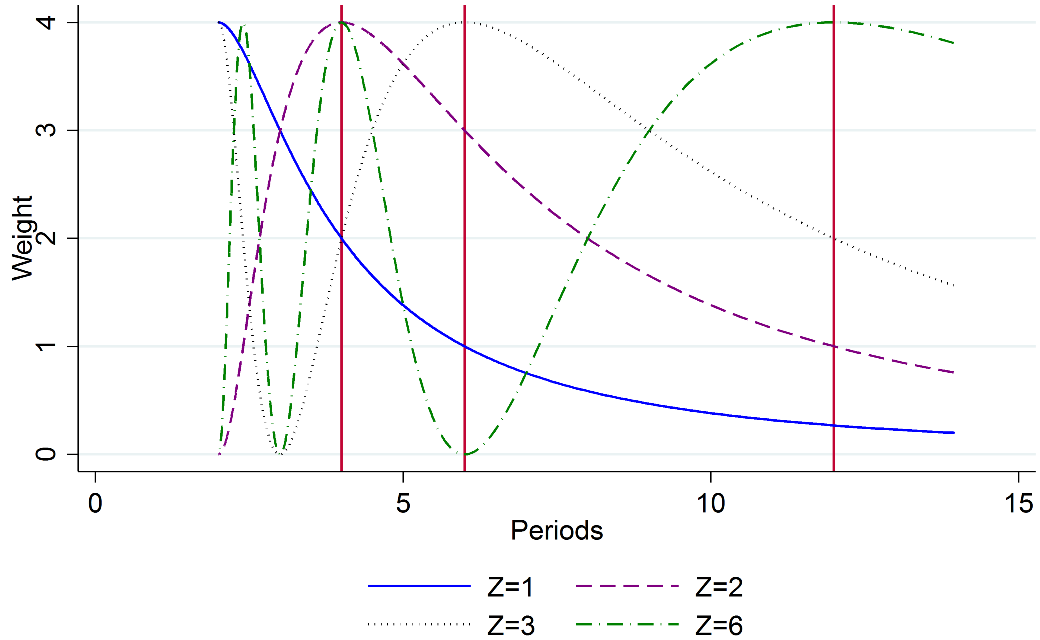

Red vertical line at periods 4, 6, and 12

FIGURE 3: U.S. Refiners Import Real Oil Price at Z

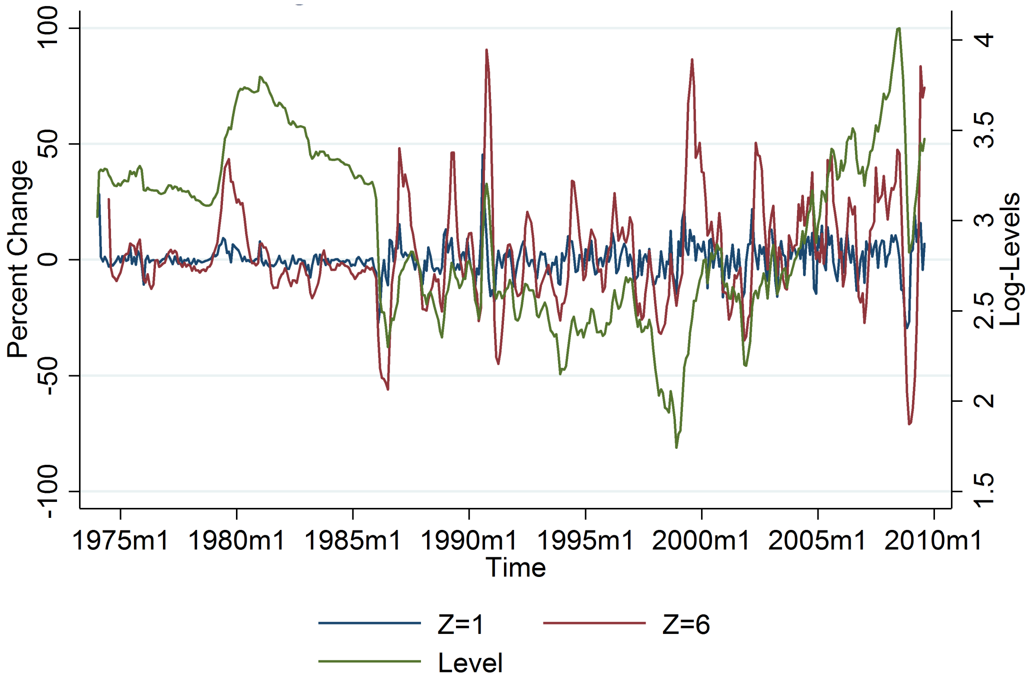

FIGURE 4: Univariate AR Parametric Spectral Density Estimate

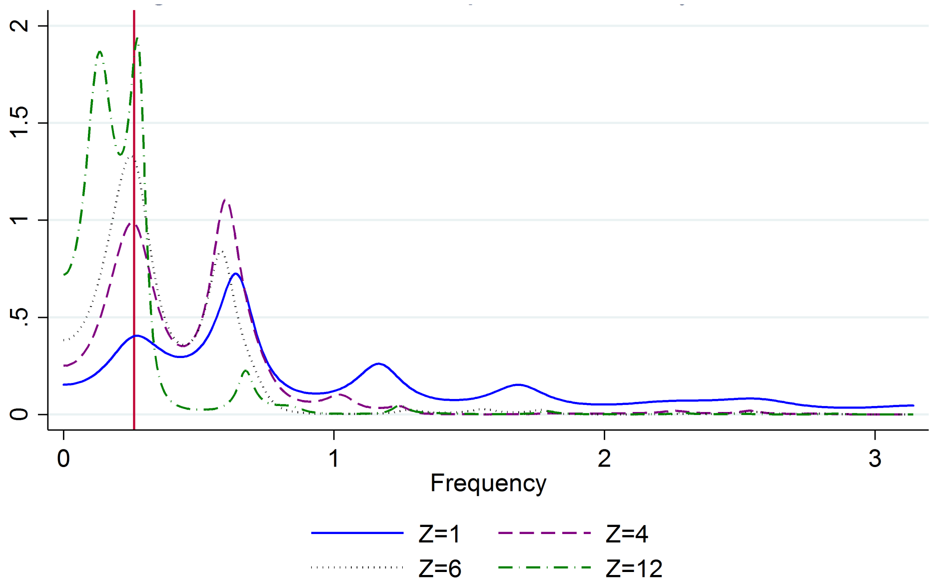

Red vertical line corresponding to frequencies 24 periods in length

Table 2: U.S. Refiners Import Real Oil Price AR Forecasts, 12 Lags, Percent Change

Months-Ahead	1m	3m	6m	9m	12m	24m	36m	48m
				MSPE relative to Random Walk				
Level	**0.74***	**0.88***	0.98	1.04	1.03	1.13	1.32	1.41
Z=1	0.75*	0.90*	1.03	1.12	1.13	1.15	1.13	1.18
Z=K/6	-	-	1.03	1.04	0.99	0.97**	0.97	0.93*
Z=K/4	-	-	0.99	1.10	0.97	0.89*	**0.93***	**0.92***
Z=K/2	**-**	0.88*	**0.97**	**0.91***	**0.86***	**0.80***	0.96**	0.98
Z=1, exc. c	**0.75***	0.89*	1.01	1.08	1.07	1.06	1.03	1.03
Z=K/6, exc. c	-	-	1.01	1.01	0.95*	0.94*	**0.98**	**0.93***
Z=K/4, exc. c	-	-	0.97	1.06	0.94*	0.89*	1.22	0.96**
Z=K/2, exc. c	-	**0.87***	**0.96****	**0.90***	**0.85***	**0.85***	1.03	1.01
				Success Ratio				
Level	**0.60***	0.55	0.49	0.49	0.51	0.53	0.56	0.59*
Z=1	0.59*	0.57*	0.54	0.49	0.51	0.50	0.56	**0.70***
Z=K/6	-	-	0.54	0.49	0.56	0.62*	**0.67***	0.66*
Z=K/4	-	-	0.55	0.50	0.56**	**0.63***	0.55	0.64*
Z=K/2	-	**0.57***	**0.55**	**0.58***	**0.63***	0.53	0.57*	0.46
Z=1, exc. c	**0.59***	0.57*	0.52	0.48	0.51	0.48	**0.53**	**0.55**
Z=K/6, exc. c	-	-	0.52	0.47	0.52	0.55	0.52	0.39
Z=K/4, exc. c	-	-	0.53	0.49	0.56	0.55	0.50	0.43
Z=K/2, exc. c	-	**0.58***	**0.57****	**0.54**	**0.56****	**0.56**	0.51	0.38

Evaluation period: 1991.11-2009.8; Estimation starts 1974.1, model estimated with 12 lags, growth rates in percent change; Recursive, dynamic, out-of-sample-forecasts; Success ratios are the relative frequency that the predictive regression model correctly predicts the sign of the change in the oil price; p-values are calculated by bootstrapping the VAR model under the null; Bold values are the best filtration at that forecast horizon; * and ** indicates 5 and 10 per cent significance.

Table 3: U.S. Refiners Import Real Oil Price AR Forecasts, 12 Lags, Log-Difference

Months-Ahead	1m	3m	6m	9m	12m	24m	36m	48m
				MSPE relative to Random Walk				
Level	**0.74***	0.88*	0.98	1.04	1.03	1.13	1.32	1.41
Z=1	0.74*	0.88*	1.01	1.09	1.10	1.10	1.07	1.09
Z=K/6	-	-	1.01	1.02	0.92*	0.87*	**0.91***	**0.88***
Z=K/4	-	-	0.97	1.03	0.90*	0.83*	0.95*	0.91*
Z=K/2	-	**0.86***	**0.93***	**0.87***	**0.84***	**0.80***	0.95*	0.98
Z=1, exc. c	**0.74***	0.87*	0.99	1.06	1.05	1.04	1.01	1.02
Z=K/6, exc. c	-	-	0.99	0.99	0.89*	0.85*	**0.92***	**0.89***
Z=K/4, exc. c	-	-	0.96*	1.00	0.87*	0.83*	0.93*	0.92*
Z=K/2, exc. c	-	**0.85***	**0.91***	**0.85***	**0.81***	**0.79***	0.97	0.98
				Success Ratio				
Level	**0.60***	0.55	0.49	0.49	0.51	0.53	**0.56**	**0.59***
Z=1	0.59*	0.57**	0.52	0.50	0.50	0.45	0.47	0.48
Z=K/6	-	-	0.52	0.50	0.57*	0.48	0.43	0.36
Z=K/4	-	-	0.56**	0.57*	0.57*	**0.45**	0.44	0.32
Z=K/2	-	**0.59***	**0.58***	**0.57***	**0.58***	0.46	0.43	0.26
Z=1, exc. c	**0.58***	0.55	0.55	0.47	0.50	0.46	0.50	**0.56****
Z=K/6, exc. c	-	-	0.55	0.50	0.58*	**0.58***	**0.57***	0.43
Z=K/4, exc. c	-	-	**0.56****	0.52	0.57*	0.57*	0.57*	0.44
Z=K/2, exc. c	-	**0.58***	**0.61***	**0.57***	**0.59***	0.57**	0.49	0.41

Evaluation period: 1991.11-2009.8; Estimation starts 1974.1, model estimated with 12 lags, growth rates in log-difference; Recursive, dynamic, out-of-sample-forecasts; Success ratios are the relative frequency that the predictive regression model correctly predicts the sign of the change in the oil price; p-values are calculated by bootstrapping the VAR model under the null; Bold values are the best filtration at that forecast horizon; * and ** indicates 5 and 10 per cent significance.

Table 4: U.S. Refiners Import Real Oil Price VAR Forecasts

Months-Ahead	1m	3m	6m	9m	12m	24m	36m	48m
MSPE ratio relative to Random Walk								
Oil Price in Log-Levels								
GAI & OECD Inv.	0.87*	0.97	1.08	1.24	1.28	1.50	1.32	1.11
GAI & U.S. Inv.	0.87*	0.93*	1.05	1.18	1.14	1.29	1.18	**1.07**
WIP & OECD Inv.	0.85*	0.88*	0.94*	1.03	1.03	1.01	1.06	1.11
WIP & U.S. Inv.	**0.85***	**0.87***	**0.94***	**1.03**	**0.99**	**0.95***	**1.02**	1.09
Oil Price with Z=1								
GAI & OECD Inv.	0.87*	1.02	1.27	1.43	1.45	1.73	1.63	1.47
GAI & U.S. Inv.	0.90*	1.00	1.17	1.35	1.43	1.73	1.63	1.47
WIP & OECD Inv.	0.88*	0.95*	1.09	1.18	1.14	1.14	1.02	1.01
WIP & U.S. Inv.	**0.86***	**0.88***	**0.95***	**1.04**	**1.03**	**0.94***	**0.93***	**0.97**
Success Ratio								
Oil Price in Log-Levels								
GAI & OECD Inv.	0.51	0.55	0.52	0.52	0.58*	0.50	0.51	0.40
GAI & U.S. Inv.	0.52	**0.56****	**0.56****	**0.54**	**0.62***	**0.57****	**0.56****	0.35
WIP & OECD Inv.	0.53	0.54	0.50	0.43	0.46	0.45	0.48	0.51
WIP & U.S. Inv.	**0.53**	0.54	0.52	0.48	0.52	0.48	0.47	**0.52**
Oil Price with Z=1								
GAI & OECD Inv.	0.55	0.52	0.50	0.53	0.56	0.46	0.49	0.41
GAI & U.S. Inv.	0.52	0.54	**0.54**	**0.56****	**0.59***	0.50	0.48	0.36
WIP & OECD Inv.	0.52	**0.56**	0.50	0.42	0.43	0.43	0.47	**0.59***
WIP & U.S. Inv.	**0.55**	0.54	0.53	0.50	0.53	**0.55**	**0.53**	0.55

Evaluation period 1991.11 to 2009.9. Estimation starts 1974.1. Recursive, dynamic, out-of-sample-forecasts; p-values are calculated by bootstrapping the VAR model under the null; Bold values are the best at a forecast horizon; * and ** Indicates 5 and 10 per cent significance. GAI is the level of the global activity index, WIP is world industrial production with Z=1, OECD inv. is OECD crude inventoires with Z=1, U.S. Inv. is U.S. inventories with Z=1.

Table 5: Targeted Filtration U.S. Refiners Import Real Oil Price VAR Forecasts

Months-Ahead	1m	3m	6m	9m	12m	24m	36m	48m
MSPE ratio relative to Random Walk								
Level[1]	**0.85***	**0.87***	0.94*	1.03	0.99	0.95*	1.02	1.09
Z=1	0.86*	0.88*	0.95*	1.04	1.03	0.94*	0.93*	**0.97**
Z=K/6	-	-	0.95*	1.00	**0.97**	**0.74***	**0.83***	1.01
Z=K/4	-	-	**0.92***	**1.00**	1.02	0.78*	0.90*	1.08
Z=K/2	-	0.87*	0.97	1.06	0.98	0.89*	1.50	1.80
Success Ratio								
Level[1]	0.53	0.54	0.52	0.48	0.52	0.48	0.47	0.52
Z=1	**0.55**	0.54	0.53	0.50	0.53	0.55	**0.53**	**0.55**
Z=K/6	-	-	0.53	0.52	0.57*	0.61*	0.51	0.43
Z=K/4	-	-	**0.54**	0.54	0.57**	**0.62***	0.50	0.43
Z=K/2	-	**0.56****	0.53	**0.56**	**0.58***	0.53	0.49	0.39

[1]Oil price in levels, all other variables in percent change. All other models estimated with variables filtered given Z. Evaluation period 1991.11 to 2009.9. Estimation starts 1974.1. Recursive, dynamic, out-of-sample-forecasts; Success ratios are the relative frequency that the predictive regression model correctly predicts the sign of the change in the oil price; p-values are calculated by bootstrapping the VAR model under the null; Bold values are the best filtration at that forecast horizon; * and ** indicates 5 and 10 per cent significance.

Table 6: Targeted Filtration WTI Real Oil Price VAR Forecasts

Months-Ahead	1m	3m	6m	9m	12m	24m	36m	48m
MSPE relative to Random Walk								
Level	**0.99**	**0.88***	0.93*	0.98	0.93*	0.92*	0.97	1.02
Z=1	1.00	0.90*	0.95*	0.98	0.94*	0.87*	0.92*	**0.97**
Z=K/6	-	-	0.95*	**0.95***	**0.90***	**0.71***	**0.77***	1.06
Z=K/4	-	-	**0.93***	0.97**	0.94*	0.74*	0.90*	1.16
Z=K/2	-	0.91*	1.00	0.99	0.91*	0.93*	1.20	1.42
Success Ratio								
Level	0.55	0.51	0.51	0.47	0.52	0.50	0.48	0.53
Z=1	**0.56**	0.51	0.51	0.50	0.51	0.57**	0.55	**0.61***
Z=K/6	-	-	0.51	**0.57****	0.58*	**0.64***	**0.55**	0.44
Z=K/4	-	-	**0.53**	0.55	0.59*	0.64*	0.54	0.39
Z=K/2	-	**0.52**	0.52	0.56**	**0.61***	0.54	0.49	0.39

Evaluation period 1991m11to 2009m9. Estimation starts 1981m.11; Recursive, dynamic, out-of-sample-forecasts; Success ratios are the relative frequency that the predictive regression model correctly predicts the sign of the change in the oil price; p-values are calculated by bootstrapping the VAR model under the null; Bold values are the best filtration at that forecast horizon; * and ** indicates 5 and 10 per cent significance.

Table 7: Targeted Filtration Forecasts: MSPE Relative to RW for Sub Periods

K	1m	3m	6m	9m	12m	24m	36m	48m
			Evaluation Period: 1995.1 to 2001.12					
Level	1.01	**1.07**	1.25	1.25	1.21	1.21	**1.08**	0.90*
Z=1	**1.00**	1.13	1.32	1.26	1.20	1.21	1.16	**0.88***
Z=K/6	-	-	1.32	1.11	1.00	0.91*	1.10	1.36
Z=K/4	-	-	1.21	1.02	0.91*	**0.86***	1.36	2.04
Z=K/2	-	1.08	**1.08**	**0.82***	**0.80***	1.51	4.14	2.86
			Evaluation Period: 2002.1 to 2007.12					
Level	**1.23**	1.33	1.35	1.52	1.23	0.97	**0.70***	**0.80***
Z=1	1.28	**1.32**	**1.19**	**1.33**	**1.15**	0.93**	0.86*	1.02
Z=K/6	-	-	1.19	1.37	1.24	**0.88***	0.84*	1.30
Z=K/4	-	-	1.20	1.46	1.32	0.91*	0.95	1.49
Z=K/2	-	1.32	1.36	1.66	1.37	1.12	1.42	1.93
			Evaluation Period: 2008.1 to 2009.12					
Level	**0.66***	0.74*	0.82*	0.83*	0.71*	**0.45***	**0.54***	**0.41***
Z=1	0.66*	0.74*	0.83*	0.85*	0.77*	0.67*	0.97	0.89**
Z=K/6	-	-	0.83*	0.82*	0.70*	0.54*	0.83*	1.04
Z=K/4	-	-	0.79*	**0.78***	0.70*	0.61*	1.26	0.94
Z=K/2	-	**0.73***	**0.79***	0.79*	**0.68***	0.96	1.95	1.36
			Evaluation Period: 1995.1 to 2009.12					
Level	**0.83***	0.84*	0.92*	1.03	0.99	0.88*	0.94*	**0.97**
Z=1	0.84*	0.84*	0.93*	1.05	1.03	0.88*	0.93*	1.02
Z=K/6	-	-	0.93*	1.00	0.97	**0.72***	**0.83***	1.07
Z=K/4	-	-	**0.89***	**1.00**	1.01	0.73*	1.03	1.16
Z=K/2	-	**0.83***	0.92*	1.04	**0.96**	1.00	1.57	1.53

Model estimated using U.S. Refiners Import Price, U.S. inventories. Estimation starts 1974.1; Recursive, dynamic, out-of-sample-forecasts, MSPE ratios are defined as the MSPE of the predictive regression model reative to a no-change forecast; Bold values are the best filtration at that forecast horizon; * and ** indicates significance at the 5 percent and 10 per cent level, respectively.

Table 8: Targeted Filtration Forecasts: Success Ratio for Sub Periods

K	1m	3m	6m	9m	12m	24m	36m	48m
			Evaluation Period: 1995.1 to 2001.12					
Level	0.55	0.62*	0.50	0.44	0.42	0.43	0.46	**0.65***
Z=1	**0.58**	0.62*	0.50	0.49	0.46	0.49	0.49	0.57
Z=K/6	-	-	0.50	0.57	0.57	0.56	0.54	0.32
Z=K/4	-	-	0.51	0.60**	0.57	**0.58**	**0.58**	0.27
Z=K/2	-	**0.64***	**0.56**	**0.64***	**0.64***	0.56	0.55	0.25
			Evaluation Period: 2002.1 to 2007.12					
Level	0.51	0.51	0.58	0.49	**0.53**	0.51	0.49	0.31
Z=1	**0.53**	0.53	**0.63***	**0.49**	0.46	**0.64***	**0.58**	0.43
Z=K/6	-	-	0.63*	0.46	0.46	0.63*	0.44	0.43
Z=K/4	-	-	0.63*	0.43	0.47	0.57	0.43	**0.50**
Z=K/2	-	**0.54**	0.58	0.42	0.46	0.39	0.43	0.49
			Evaluation Period: 2008.1 to 2009.12					
Level	0.58	**0.67**	0.67	0.79*	0.79*	0.88*	0.54	0.63
Z=1	**0.67**	0.63	0.58	0.79*	0.83*	0.96*	**0.67**	**0.71***
Z=K/6	-	-	0.58	0.83*	**0.83***	**0.96***	0.54	0.42
Z=K/4	-	-	**0.67**	**0.83***	0.83*	0.96*	0.33	0.46
Z=K/2	-	0.63	0.63	0.79*	0.67	0.71*	0.42	0.42
			Evaluation Period: 1995.1 to 2009.12					
Level	0.54	0.58*	0.56	0.51	0.51	0.52	0.48	0.51
Z=1	**0.57****	0.58*	0.56	0.53	0.51	0.61*	**0.55**	0.53
Z=K/6	-	-	0.56	0.56	0.56	**0.64***	0.50	0.38
Z=K/4	-	-	0.58*	0.56	**0.57****	0.63*	0.49	0.39
Z=K/2	-	**0.60***	**0.58***	**0.57****	0.57**	0.51	0.48	0.37

Model estimated using U.S. Refiners Import Price, U.S. inventories. Estimation starts 1974.1; Recursive, dynamic, out-of-sample-forecasts, Success ratio is defined as the relative frequency with which the predictive regression model is able to correctly predict the sign of the change in the oil price; Bold values are the best filtration at that forecast horizon; * and ** indicates significance at the 5 percent and 10 per cent level.