Determinants of Oil Price Movements in the Global Market for Crude Oil using Vector Auto Regression

By

Stefan Caputo

An essay submitted to the Department of Economics in partial fulfillment of the requirements for the degree of Master of Arts

Queen's University

Kingston, Ontario, Canada

August 2013

Abstract

There has been much debate regarding the true nature of fluctuations in the price of oil, especially those fluctuations occurring after 2003. The purpose of this study is to align the views of Hamilton (2009) with Killian and Murphy (2012); using a vector auto regression (VAR) to model the oil market while showing that speculative demand plays a significant role in the oil price fluctuations of the mid 2000's. I also seek to develop a model with fewer restrictions that can perform exceptionally well in longer term forecasting exercises. My findings show that by using the year-over-year percentage change of the data, sign restrictions and elasticity bounds are not necessary to develop a model that performs well in long term, out of sample forecasting. Historical shock decompositions of this paper's specification of the model also show that speculative demand has played a significant role in oil price fluctuations occurring over the past decade. In addition, recent literature has argued that search volume intensity (SVI) found through Google search data is a good predictor of all asset prices and holds significant explanatory power towards understanding asset price movements. I test this claim to see if SVI can be used in a multivariate analysis of oil markets and find that SVI is not appropriate when using the methodology of Killian and Murphy (2012) and performs poorly in forecasting exercises. However, historical shock decompositions are consistent with the findings of Da, Engelberg, and Gao (2011a), suggesting that a different type of model may better suit the SVI variable.

Acknowledgements

I would like to thank Dr. Edwin H. Neave for willingly taking on his role as my supervisor, his endless support, insightful advice and helpful discussions; I truly feel that this paper would not meet the level of quality and standard it currently holds without his guidance. I would also like to thank Stephen Snudden for his mentorship throughout the winter semester. I endowed a great wealth of knowledge from him regarding modeling techniques and what it means to do quality research and I am forever grateful. Last but not least, I would like to thank my parents, Luigi and Gisele Caputo, for their unmatched support, love and motivation throughout my academic career. My ambition to complete a Master's Degree in Economics from a highly regarded university such as Queen's would never have been realized if it wasn't for them.

Table of Contents

Introduction

In this paper I propose a vector auto regression (VAR) model of the global market for crude oil, which identifies speculation shocks via the inclusion of above ground oil inventories.[1] This aims to capture the change in forward-looking expectations for future oil supply and demand due to an arbitrage condition that links the oil futures market and the spot market for crude oil, as shown in Hamilton (2009). This arbitrage condition is characterised by taking a long position in a near term futures contract, selling it before expiry, and then using the proceeds to enter another near term futures contract. As the price of oil rises, the price at which one can sell the futures contract is higher than the purchase price, thus allowing for profit without ever taking delivery. Variables for the flow of supply and demand of crude oil are included in the structural estimation, allowing for the simultaneous identification of demand, supply, and speculative shocks.

I also examine retail investor attention measured using Google search data. This explores the work of Da, Engelberg, and Gao (2011a) who suggest that search volume intensity (SVI) is positively correlated with price momentum. Their research also suggests that SVI can be used to predict spot prices in financial markets as search data supposedly captures investor activity and is an indicator of trading volume. Despite their findings, I find little to no use for this particular variable in a VAR of the global market for crude oil. Issues with the time series itself and poor performance in forecasting exercises lead me to these beliefs.

The purpose of this paper is to answer several questions regarding determinants of oil prices and to test the validity of the claim that SVI is a good predictor of asset prices. I hope to

[1] The term speculation in this report refers to the financial positions that investors undertake based on market expectations towards short term capital gains.

answer (if not contribute towards the answer) the question of whether or not oil shocks are demand driven or supply driven or a combination of both factors. In addition, does speculative demand, represented through inventories data, play a role in determining changes in oil prices? If there is a role, to what extent is it significant? The answers come in the forms of various papers such as Hamilton (2009) and Killian and Murphy (2012) and open themselves up to intense debate.

Ultimately, my goal is to align the views of Hamilton (2009) with the methodology of Killian and Murphy (2012) while testing the explanatory power of SVI under more specific conditions. Furthermore, I wish to provide a model useful for those who wish to study oil prices using a longer forecasting horizon. The current literature is mainly focused on shorter forecasting horizons, thus motivating my development of a model for investors concerned with longer term positions.

I then use this model to decompose historical oil price fluctuations into the supply, demand, social media, and speculative components. In contrast to Killian and Murphy (2012), who impose a number of sign and elasticity restrictions, I find that just by using different measures of inventories and demand and by utilizing the year-over-year percentage change in the data give impulse response signs that are consistent with theory. Further, this specification suggests that the supply contractions behind the oil price increases occurring from 2003 to 2007 are more aligned with Hamilton (2009). These supply contractions were caused by a drastic decline in Saudi oil production due to the depletion of the Ghawar oilfield and policy decisions designed to take advantage of the low degree of price sensitivity associated with demand for oil's use in production, as outlined by Hamilton (2009). I also find that speculative demand shocks, represented by inventories data, drive the price fluctuations in early 2008 and throughout 2011, in contrast to Killian and Murphy (2012) who suggest that the fluctuations are demand driven.

I find that slight variations in the specification of the model can drastically change the historical shock decompositions. In particular, the high degree of multicollinearity between the independent variables makes the shocks hard to identify. I illustrate this using a couple of examples, and suggest caution in the recent use of these models' conclusions in policy debate surrounding the regulation of commodity markets. Nevertheless, I feel that this paper makes a significant contribution to the literature regarding models of oil markets and forecasting of oil prices. This paper seeks to explain that policy debate should be resolved through case dependent action. Furthermore, policy should not only account for one driver of changes in oil prices, but rather take individual action depending on the decomposition and nature of the shock.

I also test the forecasting capabilities of this paper's specification of the model and find that the model performs well in out-of-sample forecasting beyond the two year horizon, statistically beating random walk and consensus forecasts. While the model of Alquist, Killian, and Vigfusson (2011) is superior to my model at the one year horizon, the difference is small as the authors obtain a root mean square error of 0.94 while my model is scored at 1.13. In addition, Alquist, Killian, and Vigfusson (2011) do not test their model's capabilities beyond the one year horizon. Therefore, the specification of the VAR found in this paper would be far more appropriate for investors with longer forecasting horizons.

Review of the Literature

Killian and Murphy (2011) propose a structural vector auto regression (VAR) model of the global oil market which identifies speculation shocks via the inclusion of above ground oil inventories. This aims to capture the change in forward-looking expectations for future oil supply and demand due to an arbitrage condition that links the oil futures market and the spot market for crude oil, as shown in Hamilton (2009). Variables for the flow of supply and demand of crude oil are included in the structural estimation, allowing for the simultaneous identification of demand, supply, and speculative shocks. These shocks are identified through the use of historical shock decompositions described below. These experiments are useful, because they explain the degree to which a specific variable can affect changes in oil prices. Multicollinearity can affect the results of the shock decompositions by producing varying results depending on the lag specification of the model, and thus requires concern since the shocks become more difficult to identify. I address this in later sections of the paper.

Killian and Murphy (2011) use these shock estimates to decompose historical oil price fluctuations into the supply, demand, and speculative components. They do not mention the usefulness of the model for out-of-sample forecasting. In addition, they find that oil inventories can identify a speculative demand component, and the demand, supply and speculative components can explain close to 98 percent of oil price fluctuations. They posit that the run up in oil prices since 2003 is demand driven and not due to speculation or supply constraints. Further, they find larger short run price elasticity of demand relative to previous estimates due to the role of inventories in smoothing oil consumption. This is a stark contrast to the work of Hamilton (2009), suggesting that a speculative price bubble where investors seek profit without taking delivery contributes to price increases, since the speed at which these increases take place reflects the instantaneous nature of asset price fluctuations in capital markets. In addition, Hamilton (2009) finds smaller price elasticity of demand at the time of a surge in global demand coupled with rising oil prices.

I find merit in both the Killian and Hamilton arguments and wish to use my model to explain that the nature of the shock depends on the mix of variables and macroeconomic conditions at the time. For example, upward shifts in prices may be attributed to times when developing countries experience radical surges in growth due to sudden expansions in manufacturing, thus highlighting a demand shock. A fine example of this is the growth occurring in China over the past decade, as mentioned by Hamilton (2009). Alternatively, sudden and large price increases may also be attributed to times where the price of oil was perceived to increase in the future. Investors at both the institutional and retail levels may begin to purchase oil related financial instruments, thus explaining a speculative demand shock. In later sections, I explain how my specification of the data makes identifying these shocks easier through reducing restrictions and bias. Da, Engelberg, and Gao (2011a) propose the use of Google SVI as a measure for retail investor attention. SVI captures retail investor attention on a more "timely" basis and is a direct measure of individual interest. Da, Engelberg, and Gao (2011a), henceforth Da et al. (2011a), suggest that SVI holds some explanatory power towards returns on IPOs and price momentum (the latter will be the focus of this paper). The purpose of Google search data is to provide an active and more direct measure of investor attention.

Traditionally, passive measures such as advertising expenditure and media coverage have been used to evaluate investor attention, suggesting that larger firms with a higher degree of media presence will attract more shareholders. However, Ding and Hou (2009) argue that active attention measures deserve attention for several reasons. First, media coverage can vary and can experience high levels for infrequent amounts of time. Second, media coverage is not always a direct cause for positive investor attention. Negative news reports have a tendency to defer investors, leading to lower volumes of shares purchased or higher volumes of shares sold, as outlined by Daniel et al. (1998) who find that investors overreact to news related signals, thus influencing asset price movements. Finally, in order to truly understand the magnitude of retail investor attention, one must disentangle institutional investors from retail investors. While SVI has been shown to proxy retail investor attention, the question regarding institutional investors remains unanswered. Perhaps SVI captures a certain degree of intelligent investment as well as big firms may use this proxy to monitor naive investment in order to determine how to structure their own positions. Identifying and separating the retail from the institutional investor using SVI data is a daunting task to say the least. Solving this matter still requires the development of a foundational theory which is beyond the scope of this paper.

Da et al. (2011a) examine the relationship between changes in SVI and stock prices. They argue that large positive changes in SVI will lead to upward pressure on stock prices. They measure the change in SVI as the difference between the logarithm of SVI and the logarithm of the median of the previous 8 weeks' SVI.

$$\text{SVI\_Change}_t = \log(\text{SVI}_t) - \log[\text{Med}(\text{SVI}_{t-1},\ldots,\text{SVI}_{t-8})]$$

Da et al. (2011a) use the change in SVI to examine the price momentum effect of active investor attention proposed by Barber and Odean (2008). The authors find that high levels of SVI are indicative of high returns in asset prices of IPOs and historical stock data. Empirical investigation on the relation between SVI and stock returns is conducted using Russell 3000[2] stocks. The predictive power of SVI is measured by examining interactions between changes in retail investor attention and actual trading volume. Actual trading volume is measured using the

[2] Russell 3000 is an index of the 3000 largest companies. 90% of total US equity market capitalization is represented using this index thus results are not vulnerable to bid-ask bounce (price fluctuates between the bid ask spread due to impatient traders).

ratio between Dash-5[3] trading volume and the total trading volume of the previous month. This data can be accessed from Market System Incorporated. Da et al. (2011a) find that SVI holds highly significant forecasting capabilities of first week abnormal returns. Thus, their results are consistent with the price pressure hypothesis of Barber and Odean (2008).

Da et al. (2011a) suggest that the price pressure hypothesis of Barber and Odean (2008) applies to all asset prices. However, the research is limited to a single variable analysis of equity and a stock index. How does SVI fare when used in a multivariate framework such as a VAR? In addition, this paper will examine the explanatory power of SVI towards a highly traded commodity, such as oil.

The structural VAR proposed by Killian and Murphy is written as:

$$B_0 y_t = \sum_{i=0}^{24} B_i y_{t-i} + \varepsilon_t,$$

where ε_t is a vector of orthogonal structural innovations, B_i, $i = 0, \ldots, 24$, is the matrix of coefficients and y_t is a vector of endogenous variables. These four variables are described in detail in the next section but include the real USD price of oil, global crude oil production, a global real activity proxy, and a proxy for global above ground crude oil inventories. Two years of lags are used at a monthly frequency for the sample 1973m2 to 2009m8. Seasonal dummies are also included in the VAR model but are not shown above for notational convenience. In addition, the model is estimated using a combination of sign restrictions and bounds on the elasticity of oil supply and demand.

[3] Dash-5 is an abbreviation of the reports of trading volume found in SEC Rule 11AC1-5.

The structural VAR is identified with the use of joint sign restrictions as summarized in Table 1. In particular, the response of the real price of oil and global real activity are set to be negatively correlated with the supply shock for at least twelve months. However, no sign restrictions are imposed on the response of global activity and oil production to speculative shocks.

Table 1: Sign Restriction on Impact Responses in VAR Model

	Flow supply shock	Flow demand shock	Speculative demand shock
Oil Production	+	+	+
Real Activity	+	+	-
Real price of oil	-	+	+
Inventories			+

These sign restrictions for the structural VAR are combined with bounds on the elasticity of oil supply and demand. In particular, the short-run price elasticity of oil supply is bounded between zero and 0.025. The impact price elasticity of oil demand is bounded to be between -0.8 and 0.

The methodology used to estimate the model identifies a range of models consistent with the identifying assumptions. Out of 5 million least-squares reduced-form VAR estimates, Killian and Murphy (2011) find that 14 candidate models satisfy all identifying restrictions. These differences in the estimates of these 14 candidate models primarily concern the relative explanatory power of inventories versus the world industrial production variables. Furthermore, seeking to determine whether speculation or industry based demand can explain more about drastic price shifts in the real price of oil.

Data

Overview of the Data

There are five variables used in the model: the real USD price of oil, global crude oil production, global industrial production, above ground crude oil inventories, and search volume intensity (SVI) in the OECD. The sample range used in the paper is from 1973m1 to 2012m10. These five series are displayed in Figures 1-4. In addition, I use futures data, analyzing both volume and settle prices to test for correlations with inventories. The purpose of this is to examine the merit behind using inventories to capture the speculative component of the market for crude oil.

The real USD price of oil is measured as the nominal Brent crude oil price from the *International Financial Statistics* database reported by the International Monetary Fund. I choose to use the Brent price of crude oil since it better reflects the global price of oil due to the separation of West Texas Intermediate crude oil prices from global prices, as well as the regulation of the US market for crude oil during the 1970s and early 1980s. This nominal price of oil is deflated by the consumer price index for OECD countries. Since the OECD countries represent over 85% of global GDP, I use this deflator rather than the U.S. deflator to better capture the global price of crude oil.

The measure of global crude oil inventories used in this paper is based on above ground inventories data for total U.S. crude oil inventories provided by the U.S. Energy Information Administration (EIA). Similar to Hamilton (2009), I scale this data by OECD petroleum stocks which only became available after 1988. The series displays seasonality with accumulations of stocks over the summer and decumulation over the winter months. This series lacks data on inventories for non-OECD countries which may bias the results. Specifically, the absence of data on the accumulation of strategic reserves in China may prove to be misleading.

I use inventories to capture the speculative position of commodities as suggested by Alquist and Killian (2010) and Hamilton (2009a). This is because an arbitrage condition links the spot market and the oil futures market. This arbitrage condition is essentially the fact that through the use of derivatives (such as futures and options on futures) oil stocks can be purchased and sold at prices that vary from current spot prices so as to earn a profit. Giannone and Reichlin (2006) use Granger causality tests[4] for futures and inventories, attempting to find if one has more information than the other. They, along with Killian and Murphy (2012) find that there is no additional information in the futures markets other than inventories, supporting the argument that inventories can be used to identify the speculative position. This is also confirmed by Fattouh and others (2012) who find that this relationship has been particularity strong over the last decade.

The measure of world industrial production used in this paper is world industrial production excluding construction provided by *Global Data Services*, see Figure 2. I use this measure over the widely used global activity index (GAI) developed by Killian (2009). The GAI is based on the dry cargo shipping freight rate index and I believe that there are several reasons to cast doubt on the use of the global activity index. First, there has been a remarkable and well recognized decline in shipping rates due to an excess of capacity in shipping. They claim to capture this decline by detrending the data. However, this would not account for changes in the volatility induced by the shipping glut. Further, since this series is detrended, it may fail to

[4] The Granger causality test aids in determining whether the time series data of one variable can produce an accurate forecast of another variable's time series data, as presented by Granger (1969).

capture the increase in the growth rates of emerging market economies over the last decade. In addition, oil is an input in the production of shipping and an increase in the price of oil may drive up the marginal cost of freight rates. This pass through of the cost of oil in shipping rates would imply fictitious correlation which would over identify the role of demand shocks. For these reasons I use the measure of world industrial production excluding construction

The measure of global crude oil supply is the world crude oil production, available in the *Monthly Energy Review* of the EIA, measured in thousands of barrels per day. It is a measure of the flow supply of oil from all countries of the world. This series is widely used to identify flow supply as in Killian and Murphy (2011) and Hamilton (2009). This is a global measure of crude oil supply and takes into account OPEC members, Persian Gulf Nations, and OECD countries.

The measure of SVI is the search phrase “crude oil prices”, available from Google Trends. The data is measured on a scale from 1 to 100; 100 measuring the peak of attention or highest volume of search hits. This series is global but search hits are most concentrated in the US, Canada, India, and Pakistan. Since SVI is a fairly recent discovery regarding economic analysis, Google trends data is the only reliable source for this type of data. SVI is the only variable that is not price data and will be modified into an appropriate metric for this paper. This series will be differenced in order to examine percentage changes along with the other variables.

The measure of oil futures closing prices is the monthly series of ICE Brent crude oil futures available from Quandl. As mentioned previously, I examine both the monthly volumes and settle prices of oil futures. Given that the active life span of a futures contract can vary, I use a continuous contract series that strings together contracts of varying lengths. This results in the ability to model and analyse futures trading over a longer time horizon, without the adverse effects of short term volatility.

Stylized Facts of the Data

From Figures 1-4, it is noticeable that all of the series are not stationary. This is confirmed from the Augmented Dickey Fuller tests for unit root and the Phillips-Perron tests, whose results are not reported here. Autocorrelation tests for all series show near perfect unit root for the first lag. Further, I find that there is a statistically significant break in the mean of the real price of oil between 1986 and 2003. Before differencing the data, I test for cointegrating relationships between the variables but fail to find any.

Table 2: Summary Statistics of the Y/Y % Change in the Data

	Mean	Std. Dev.	Min	Max
Real Price of Oil	10.8577	43.14726	-65.27418	264.9432
Global Oil Production	0.86149	3.877729	-12.23735	17.51737
Global Industrial Prod.	4.17084	5.092993	-13.34764	22.75339
Oil Inventories	1.60751	10.5658	-29.21814	61.17733
SVI	0.02105	23.0288	-76.75	67

I choose to focus on the year-over-year percentage change in the price based variables and the first difference for SVI. The stylized facts of these series are summarized in Table 2. These measures are graphed along with the real price of oil in Figures 5-8. I choose the year-over-year measure for several reasons. First, the measure accounts for the seasonality observed in the time series, especially for inventories which show accumulations of stocks over the summer and decumulation over the winter months. This result contrasts those of Killian and Murphy (2012), who use the month-over-month log difference of inventories and oil supply.

However, I find that the year-over-year series do not significantly differ from white noise, meaning that there is no autocorrelation present within this specification of the series. In addition, the frequencies of the year-over-year series are similar which I believe will provide a more accurate contribution to the change in the oil price movements. Given that SVI is measured on a scale of 1 to 100 and exhibits large jumps, the first difference is the most appropriate measure for this series of data. While year-over-year percentage changes are suitable for inventories, global oil supply, world industrial production and price, SVI data exhibits too much volatility. The year-over-year percentage changes for SVI data can range from tens to hundreds of thousands. In addition, the series would still not be stationary deeming it inappropriate for analysis. Thus, the first difference of SVI is used.

Table 3: Augmented Dickey-Fuller Test for Unit Root

	Lags	Test Statistic Z(t)	1% Critical Value	MacKinnon p-value
Real Price of Oil	13	-4.96	-3.43	0
Global Oil Prod.	14	-5.14	-3.43	0
World Indust. Prod.	27	-5.01	-3.43	0
Inventories	25	-5.06	-3.43	0
SVI	10	-6.97	-3.509	0

The Augmented Dickey-Fuller tests for unit root in the year-over-year percent changes in the series are presented in Table 3. The purpose of this test is to ensure that a time series is stationary and does not possess unit root. If a time series possesses unit root, shocks to variables have permanent effects and variance becomes time dependent. This will pose problems for estimation and forecasting accuracy. Furthermore, I report these results to show that the specification of the data used in this model is most optimal and possesses virtually zero risk of unit root.

I reject the presence of a unit root for all series at the 1% significance level. As shown in Figures 9-13, all series except for SVI display significant autocorrelation at the 99 percent level for up to a year. SVI exhibits significant autocorrelation for no longer than 3 months. The partial autocorrelation tests, Figures 14-17, show that there is an AR process of the variables (except SVI) up to the fourth lag, consistent with the AR(4) ARIMA model used as the benchmark of Alquist, Killian, and Vigfusson (2011), henceforth AKV (2011). Alternatively, Figure 18 displays that SVI has an AR process up to the sixth lag. However, I also find that the there is a role of the 12th and the 24th lag for all of the series. This is due to the year-over-year specification. This is captured in the specifications of the suggested models and is not reflected in the residuals of the estimates.

The pair-wise correlation coefficients for these series are presented in Table 4. All variables, except for SVI, are significantly correlated at the 5 percent significance level. This shows that there may be an issue of multicollinearity in the data suggesting that there would be a loss of asymmetric efficiency in the parameter estimates. Since I have 467 observations I suggest that the models are able to estimate the parameter coefficients effectively. However, this may be an issue when I conduct the historical variance decompositions of the estimates, since the multicollinearity between the variables may reduce the predictive power to identify the role of each of the flow shocks in the estimates. I come back to this in later sections.

Table 4: Pair-wise Correlation Coefficients

	Real Price of Oil	Global Oil Prod.	World Indus. Prod.	Crude Oil Inv.	SVI
Real Price of Oil	1				
Global Oil Prod.	0.12*	1			
World Indus. Prod.	0.19*	0.52*	1		
Crude Oil Inv.	-0.19*	-0.1	-0.19*	1	
SVI	-0.007	-0.0472	-0.0154	0.0492	1

*Significance at 0.001 level

Correlations with Oil Futures

The use of inventories as a proxy for oil futures prices and the speculative component of demand, as suggested by Alquist and Killian (2010) and Hamilton (2009a) are appropriate as confirmed by their research and tests for correlation between both time series, as shown in Table 5. Inventories are negatively correlated with the settle price of futures contracts. This is consistent with theory as more expensive futures contracts would be less enticing to investors, thus meaning that fewer inventories are available for purchase. The real price of oil is also found to have a negative correlation with inventories, meaning that higher spot prices make speculative short term investment less appealing, thus consistent with the decrease in inventories. This relationship only holds for the settle price however, as no significant correlation was found between volume of futures and inventories.

In contrast to Da et al (2011a), no relation between SVI and futures price was found. In addition, a negative correlation between investor attention and the volume of futures traded exists, but at the third lag. This result contrasts the work of Daniel et al (1998) who model a momentum effect where increased attention is causation for increased trading activity. However, oil futures are traded by more sophisticated investors as well. This may be the reason for the lack of correlation between not only futures prices and SVI, but between price movements and SVI as well. Fully understanding the effect of less sophisticated investors on asset prices and the market for crude oil may require separating institutional investors from individual investors. In addition, individual investors may use other derivatives, such as options on oil futures when trading in financial markets. I leave the above analysis to further research and modeling methodology.

Table 5: Correlations with Oil Futures Settle Price

	Correlation
Oil Production	+
Real Activity	+
Real price of oil	+
Inventories	-
SVI	N/A

Modeling Methodology

I propose the VAR model as follows:

$$B_0 y_t = \sum_{i=0}^{j} B_i y_{t-i} + \varepsilon_t,$$

where ε_t is a vector of orthogonal structural innovations, B_i, $j = 0, \ldots, J$, is the matrix of coefficients potentially up to the 24th lag and y_t is a vector of endogenous variables. These variables are described in the previous section and include the year-over-year percent change in the real USD price of oil, global crude oil production, world industrial production, and OECD above ground crude oil inventories.

The model aims to harmonize the views of Killian and Murphy (2011) and Hamilton (2009) while testing the validity of the claim made by Da et al (2011a) that SVI holds explanatory power for all asset prices. The questions posed by the aforementioned papers that I hope to answer include inquiries raised by Blinder (2009): Are oil shocks primarily demand or supply driven? What are the major determinant(s) of the run up of oil prices in the early to mid 2000's? Does speculative demand play a significant role? In addition, I hope to answer questions of my own regarding the formidability of SVI in predicting asset prices such as: Is SVI truly a good indicator for all asset prices? Does it belong in multivariate analysis?

The Killian and Murphy (2011) framework is adopted through the use of VAR modeling of the real USD price of oil, global crude oil supply, global industrial demand, and inventories up to 24 lags. However, I feel that by using a less restrictive measure of the variables and data I can align the views of Hamilton (2009) with the modeling methodology of Killian and Murphy (2011); that speculative demand shocks can play a significant role in determining oil prices using a VAR for the global market for crude oil.

I also aim to test whether or not Killian and Murphy (2011) can be expanded to include additional variables accounting for investor attention.. Furthermore, do the claims of Da et al (2011a) apply to multivariate commodities analysis? While the authors posit that SVI holds a great deal of explanatory power for asset prices, the model seeks to test this claim with the addition of an SVI variable to the recognized model of Killian and Murphy (2011).

Model selection criteria, excluding the SVI variable, suggests the VAR(2) using the Bayesian information criterion (BIC), but a VAR(14) using the Akaike Information Criterion (AIC). A VAR(24) model is also close on the AIC, and is preferable using the root mean squared error (RMSE) estimates. AIC and BIC criterion recommend only a VAR(24) when including the SVI variable. However, Final Prediction Error (FPE) criterion recommends a VAR(15) model when including the SVI variable. The AIC, BIC, adjusted R-squared and the RMSE for these three models are summarized in Table 6.

The findings of the VAR(24) with a SVI variable are troublesome, as shown in Table 6. While a perfect R-squared and RMSE of 0 are statistics that most models would strive for, they may be the result of misspecification.. These results are unattainable, and as a result undesirable. Further, closer examination of the parameters of the model shows that a number of vectors are omitted, compromising the completeness of the model. For these reasons, the VAR(24) with SVI will not be used for further analysis.

Table 6: Criteria Based Model Selection

Criteria	AIC	BIC	FPE	R-sq	RMSE
VAR(2)	20.2	20.5	4603.4	0.88	15.01
VAR(14)	19	21.1	2045.7	0.91	11.55
VAR(15) with SVI	13.4	24.7	1220.5	0.99	7.46
VAR(24)	19.1	22.7	2302.05	0.92	11.44
VAR(24) with SVI	-251.9	-239.9	NR	1	0

The estimates of the VAR models fit the data reasonably well, as shown in Table 7. The results of the tests on the residuals of all models show that the VAR(14) and VAR(24) are most optimal as they possess the best set of qualities prior to forecasting exercises. All of these characteristics warrant testing of the model’s forecasting capabilities and comparison with the results of AKV (2011). This also begins to raise concern about the use of SVI in multivariate analysis of commodities markets.

To gather more insight into the appropriate AR length of non-SVI models, I conduct a Lagrange Multiplier test. Testing VAR(2) models against the VAR(14) and VAR(24) model, I can reject the null hypothesis that the smaller model is the true model at the 0.001 percent and five percent significance levels, respectively. In addition, I am unable to reject that the VAR(14) is the true model compared to the VAR(24) model. Thus, I find conflicting evidence between the VAR(14) and VAR(24) models.

I continue by testing for possible misspecification by running a battery of tests on the residuals of all estimated models as summarised in Table 7, and the residuals of the VAR(2), VAR(14), and VAR(15) models are graphed in Figures 18-20. The residuals of all models have mean zero and the Augmented Dickey-Fuller tests for unit root for all models are able to reject the null of a unit root below a p-value of one percent. In addition, all non-SVI models display eigenvalue stability while the VAR(15) with the SVI variable does not, which is represented by the unit root circle for the VAR(2) , VAR(14), and VAR(15) models, as shown in Figures 15-17.

The residuals of the VAR(2) models display some degree of persistence, which is confirmed by the presence of an AR(12) in the residuals. These results are consistent with the Lagrange Multiplier test for the VAR(2) described above. This is confirmed by the Portmanteau test for white noise which rejects the null hypothesis of no serial correlation at 0.1 percent significance level for the VAR(2) model. In addition, I find that the Box-Pierce' Q statistic of the VAR(2) models' residuals show conditional heteroskedasticity.

Table 7 shows that the residuals of the VAR(15) with the SVI variable do not display as much persistence, but still warrant concern. Similar to the VAR(2), there is presence of an AR(12) in the residuals. In addition, the Box-Pierce' Q statistic of the VAR(15) models' residuals are increasing over time, thus showing conditional heteroskedasticity. Interestingly, the residuals of the VAR(15) model possess white noise. This is confirmed by the Portmanteau test for white noise which fails to reject the null hypothesis of no serial correlation at the one percent significance level. This result may suggest that the degree of AR persistence in the twelfth lag may not be as severe as the VAR(2).

Table 7: Results of the Tests on the Model Residuals

	VAR (2)	VAR(14)	VAR(15)	VAR(24)
Heteroskedasticity	Heteroskedastic	Homoskedastic	Heteroskedastic	Homoskedastic
Stationarity	No Unit Root	No Unit root	No Unit Root	No Unit root
White Noise	Not White Noise	White noise	White Noise	White noise
Autocorrelation	AR (12)	No AC	AR (12)	No AC

In contrast, the residuals of the VAR(14) and VAR(24) models are better behaved. They do not have AR persistence, consistent with the Lagrange Multiplier test described above. This is confirmed by the Portmanteau test for white noise which fails to reject the null of no serial correlation at the one percent significance level. In addition, I find that the Box-Pierce' Q statistics of these models residuals are not increasing over time and are able to capture the heteroskedasticity in the data. For these reasons, I reject the VAR(2) model in favor of the VAR(12) model. However, I am unable to reject the VAR(24) model over the VAR(14) model.

I test the joint restriction that all of the estimates of a series are significant using Granger causality tests; results are displayed in Table 8. The tests on the estimation of the parameters of the models are consistent with the evidence found in the initial sections. In particular, the VAR(2) model rejects that the price of oil has a significant impact on global industrial production with two lags, but has a significantly positive impact with 14 lags (confirmed by the impulse response functions described below). Interestingly, this becomes insignificant again in the VAR(24). This is consistent with the observation that the use of contracts and real rigidities could cause a slow pass through of the price of oil to the real economy. In addition, the VAR (15) model suggests that the price of oil has a significant impact on the SVI. This is consistent with Ding and Hou (2009) and Da et al. (2011a) who suggest that retail investor attention is stimulated by price changes. Google SVI is also shown to Granger cause all of the other variables, suggesting that retail investor attention may have an impact on the speculative component of oil price movements via oil futures. These results suggest that the VAR(14), VAR(15), and VAR(24) models may be more appropriate for conducting the historical shock decompositions and forecasting exercises in the next sections.

Table 8: Granger Causality Wald Tests

Equation	Excluded	VAR(2)*	VAR(14)*	VAR(15)*	VAR(24)*
Global Oil Prod.	Real Price of Oil	1	0.05	0.003	0.23
	World Indust. Prod.	0.01	0	0	0
	Inventories	0	0	0	0
	Google SVI	n/a	n/a	0	n/a
	ALL	0	0	0	0
Real Price of Oil	Global Oil Prod.	0.44	0.04	0	0.02
	World Indust. Prod.	0	0.02	0	0
	Inventories	0.82	0	0	0
	Google SVI	n/a	n/a	0	n/a
	ALL	0.05	0	0	0
World Indust. Prod.	Global Oil Prod.	0	0	0	0
	Real Price of Oil	0	0	0	0.02
	Inventories	0.27	0.22	0	0.06
	Google SVI	n/a	n/a	0	n/a
	ALL	0	0	0	0
Inventories	Global Oil Prod.	0	0	0	0
	Real Price of Oil	0.08	0.02	0	0
	World Indust. Prod.	0	0	0	0
	Google SVI	n/a	n/a	0.008	n/a
	ALL	0	0	0	0
SVI	Global Oil Prod.	0	0	0	0
	Real Price of Oil	0.08	0.02	0	0
	World Indust. Prod.	0	0	0	0
	Inventories	n/a	n/a	0	n/a
	ALL	0	0	0	0

*Prob > chi^2

Results

Historical Shock Decompositions and IRFs

The impulse responses of the VAR(14) model are summarized in Figure 21, and the types (positive or negative) of impulse responses, if significant at the ninety-five percent level, are summarized in Table 9. The purpose of this test is to determine whether a variable will experience an increase or decrease in response to an increase of another variable. This shows the model's ability to capture the reactionary behavior of the variables and test whether or not the model is consistent with theory and observed fluctuations.

The response of the real price of oil to world industrial production and supply are positively and negatively correlated, respectively. This is consistent with theory and with the observations of the pair-wise correlation coefficients summarized in Table 4. In addition, positive demand shocks increase oil supply, and an increase in oil supply lowers the price of oil which is also consistent with theory.

Table 9: Signs of Impulse Responses in the VAR(14) Model

	Flow supply shock	Flow demand shock	Speculative demand shock
Oil Production	+	+	-
Real Activity	+	+	
Inventories	+	-	+
Real price of oil	-	+	

Interestingly, inventories increase in response to supply shocks, suggesting that the price movements are compounded by the additional speculation when there are supply disruptions. In contrast, the negative sign of the oil inventories in response to a demand shock shows a negative correlation. This suggests that changes in inventories play the role of smoothing changes in the price of oil when the fluctuations are driven by demand shocks. This may be due to higher uncertainly of oil supply due to the role of OPEC in the determination of oil supply.

The impulse responses of the speculative demand shocks on the price of oil and real activity are insignificant. This may reflect a decreased role of speculation in determining these two variables, or a time varying role; having a positive effect in some cases and negative effects at other times. However, the speculative demand shock has a significant negative effect on oil production. This suggests that when oil inventories are drawn upon, the supply tends to respond to smooth price movements. This may suggest that OPEC has been responding to the strategic accumulation in oil reserves over the past few years by responding with increasing supply to match the demand.

The results of the impulse responses of the VAR(15) are less promising. Unlike the VAR(14), the results are not consistent with theory as the real price of oil is negatively correlated with demand and positively correlated with supply. The model also fails to capture the effect of a demand shock on oil supply. If supply were to remain constant while demand increases, theory suggests that a price increase must also take place. However, the VAR(15) suggests that a price decrease would take place, leading me to believe that the VAR(15) produces conflicting results when accounting for active retail investor attention.

The impulse responses of SVI in the VAR(15) are slightly more realistic. Retail investor attention is positively correlated with inventories, suggesting that increased investor confidence and attention results in a larger amount of oil futures purchased. In addition, positive supply shocks result in an increase in SVI. This result suggests that upon news of an increase in supply, investors may become more curious about oil, thus deciding to conduct more research and read news regarding these supply shocks. The one questionable result however, is the response of SVI to demand shocks. An increase in demand leads to an initial decrease in investor attention, but then an increase with the cycle repeating itself over the ten step time frame. I do not find this to be consistent with theory or typical investor behavior as demand and retail investor attention should display a lagged positive correlation. Furthermore, as demand increases, retail investor attention would increase, but at a later time due to the time lapse between when a futures contract is purchased and then exercised.

Da et al. (2011a), suggest that SVI along with "investor overconfidence" results in strong price momentum. While shocks to investor attention and speculative demand result in an increase in price, the changes are not significant. In addition, the small increase in price in response to SVI does not happen immediately, but rather five periods out.

Killian and Murphy (2012) identify demand shocks from speculative demand shocks by the fact that demand shocks should increase real output and speculative demand shocks should not. I find this is not well identified in my model due to the lack of information or the varying effects on the speculative demand shock on real activity. If this is the case, their model may be imposing a potentially incorrect assumption. Alternatively, it may just be that there is not enough of a link between speculation and real output. The evidence thus far points to speculation having a smoothing role in response to demand shocks and an exacerbating role in the presence of supply shocks.

The estimates of the historical shock decompositions for the four VAR models are displayed in Figures 23-26. The models place a large role in the AR process for explaining the price movements. I choose not to distribute this persistence to the observed shocks since there is no reason to assume that all of the shocks have identical persistence effects on the price of oil. Interestingly, I find that the historical shock decompositions vary with each model. Particularly noticeable, the shock decompositions for the VAR(2) model place a large role on demand shocks for explaining price movements, whereas the VAR(14) and VAR(15) are mixed. The VAR(24) model suggests different roles for different time periods. For example, from 2003-2007 price fluctuations were mainly supply driven, but from 2007-2008 speculative demand was the main driver, and then demand driven from 2008 onward.. I focus on VAR(14), VAR(15) and VAR(24) models since the above tests suggest that they are the better identified models. In particular, I focus on major movements in the price of crude oil post 2003.

Although the historical shock decompositions of the VAR(14) and VAR(24) agree in that a demand shock caused the fall in the price of oil in 2008-2009, it disagrees in the other periods. In particular, the large increase in the price of oil over 2003-2007 is attributed to supply shocks in the VAR(24) model, whereas it is mixed in the VAR(14) model. If caused by supply shocks this would lead evidence to the so called ‘peak oil’ theory and as suggested by Hamilton (2009).

The VAR(15) model with the SVI variable attributes the price fall in 2008 to shocks in inventories and SVI. This is consistent with the findings of Da et al. (2011a) who suggest a strong link between investor attention and trading volume. After the financial crisis of the late 2000’s, “less sophisticated” investors gained skepticism towards financial markets, possibly reflected by the negative shock of SVI. Naive investors losing interest may also be reflected by the downward trend of inventories as the volume of futures purchased drops dramatically.

The VAR(15) agrees with the VAR(2) with respect to demand shocks driving the increase in the price of oil over 2003-2007. However, once the SVI variable is introduced, the VAR(15) credits pre crisis price movements to retail investor attention, consistent with Da et al (2011a). Interestingly, the downward pressures on oil from 2007 to 2008 are explained by negative shocks to SVI. Similar to the events during the crisis, lower degrees of investor attention and changes in inventories corresponds to downward pressures on prices, hence less futures purchased in financial markets.

These results are not without controversy however, as the VAR(15) does not attribute strong price momentum to SVI on a consistent basis. Figure 25 shows that there are periods of upward price movement whilst retail investor attention is decreasing and vice versa. Perhaps this may be attributed to the short term goals of naive investors closing out positions prior to maturity for speculative purposes. Questions regarding the degree of impact of independent and less sophisticated investors as opposed to sophisticated and larger institutional investors (banks, insurance companies, hedge funds, etc.) arise; I leave this to further research. Positive price momentum is consistently attributed to demand shocks (except during 2010 where oil prices rose by 33% resulting in a clawback on demand), a feature that is present in the questionable VAR(2) and in contrast to Hamilton (2009). In addition, the VAR(15) does not pass the eigenvalue stability test mentioned in the previous section. Heteroskedastic errors may also add some degree of bias to the model's results.

Killian and Murphy (2012) attributed the movement to demand pressures. However, they bound the price elasticity of oil supply to between zero and 0.025. Without imposing this constraint and their sign restrictions, I estimate the short-run (time span of a single month) supply elasticity to be 0.167. This may suggest that oil supply is actually able to respond in the short run to changes in the price of oil. Again, this may be due to the ability of OPEC to change their supply to fulfill their mandate of price stability. Not taking this supply response into account may push the actual effect of the trend into the other variables thus distorting shock decompositions.

Interestingly, I find that the VAR(14) and VAR(24) models agree that supply responded countercyclically between 2008 and 2011. This is consistent with the objectives of OPEC in smoothing the price of oil, and may further suggest that they are achieving some degree of their price smoothing objective. However, it may also reflect the response of private oil companies to attempt to smooth prices. A more detailed analysis would have to be undertaken to back out the role of each of these players, which I leave to future research.

The VAR(24) suggests that accumulations of inventories were a major cause of the oil price fluctuations in early 2008 and late 2011, whereas the VAR(14) is ambiguous. The findings of the VAR(24) model are in contrast with that of Killian and Murphy (2012) who attribute the increase in the price of oil in 2008 to demand pressures, particularly in developing countries. The model lends some evidence to the popular perception that speculation was behind the increases in the price of oil during these periods. The estimated short-run price elasticity of demand and inventories are 0.145, and 0.064, respectively. This price elasticity of oil with respect to output is in line with the estimates found in Hamilton (2008).

As mentioned earlier, there may be a trade-off between the explanatory power of the shocks due to the multicollinearity between the variables of the VAR (14) and VAR(24) models. Without the sign restriction and bounds used in Killian and Murphy (2012), my findings place more emphasis on the role of speculative demand and more closely align the results in Hamilton (2009). The absence of these restrictions and the ability to lose bias found in demand and inventories using the year-over-year percentage change specification lead me to believe that the degree to which multicollinearity may affect the model is less severe than that of Killian and Murphy (2012). Either way, these results suggest caution in the use of these models to identify the historical shock decompositions.

In summary, impulse response functions and historical variance decompositions suggest that the VAR(14) and VAR(24) models may be the most useful towards analyzing the determinants of oil price fluctuations in the global market for crude oil. The results of these models are most consistent with theory and suggest a role for speculative demand in explaining oil price fluctuations. These tests also contribute to the ability of this paper to align the views of Hamilton (2009) with the modeling methodology of Killian and Murphy (2012), whilst improving and adding to the results of AKV (2011). In addition, I find further evidence that SVI may not be appropriate for multivariate analysis of commodities markets as impulse responses are not consistent with theory and do not capture the effects of all variables. However, the results of the historical variance decompositions of the VAR(15) are aligned with the findings of Da et al. (2011a) and suggest that SVI may be hold more significant explanatory power under a different modeling procedure.

Forecasting

I consider the VAR(14) model as a candidate, but report the root mean square error (RMSE) for the VAR(2), VAR(15), and VAR(24) models when examining forecasting performance for the sake of comparison. I developed a procedure that estimates the one-year-ahead to five-year-ahead dynamic, rolling, out-of-sample forecasts for the VAR(14) and one-year-ahead to three-year-ahead for the VAR(15). In addition, I test the prediction of the VAR(14) model in the 1990m1-2012m10 sample period and the prediction of the VAR(15) in the 2009m2-2012m10 sample. The year-over-year growth rates of the model are converted back into levels when calculating the RMSE. I also estimate the results for other sample periods but find that the performance ordering of the models do not change. These estimates, as well as the random walk RMSE are summarized in Table 10.

Table 10: Dynamic, Out-of-Sample, Rolling, Forecast Root Mean Squared Errors

Years-Ahead	One	Two	Three	Four	Five
VAR(2)	8.83	8.39	7.88	**7.74**	**7.72**
VAR(14)	8.75	**8**	**7.69**	7.79	7.77
VAR(15)	578.41	61229.63	50672	n/a	n/a
VAR(24)	8.66	9.39	8.6	8.47	8.29
Random Walk	**7.72**	8.72	9.24	10	10.95

The model forecasts of non-SVI VARs perform well in out-of-sample for forecast horizons of two years or more. The VAR(14) model has the lowest RMSE at the two- and three-year–ahead horizon statistically beating random walk estimates. The VAR(2) model's estimates are comparable to the VAR(14) and slightly lower beyond the four-year-ahead forecast horizon. Throughout, I am unable reject that the RMSE of the VAR(2) models are significantly different than the VAR(14) models. In contrast, the VAR(24) model is unable to beat random walk until after the three year forecast horizon. In addition, the other models have significantly lower RMSE than the VAR(24) model at the five percent level.

The forecasting results of the VAR(15) indicate that search volume intensity combined with other macroeconomic variables for oil holds little to no predictive power. This contrasts the work of Da et al. (2011a) who suggest that online investor attention may be able to better predict spot prices. Problems with the VAR(15), namely instability and heteroskedastic errors, appear to be the main reason for the model's weakness in forecasting. While retail investor attention does display correlation with price fluctuations and changes in inventories and futures (in so far as the historical shock decomposition), its use in a VAR model for forecasting purposes is not warranted.

There has been extensive work on evaluating the forecasts for the price of oil. AKV (2011) evaluate over 20 models for forecasting the price of oil up to one year, as well as for the oil futures price and the survey forecasts. The RMSE of these forecasts divided by the RMSE of the random walk estimate is summarized in Table 10. Their best model, of those surveyed, can beat random walk by a ratio of 0.94 after one year, and is unable to significantly beat random walk beyond that horizon. In comparison, oil futures have a relative RMSE of 0.93 relative to random walk at the one year horizon, but fail to beat random walk past the one year horizon. The authors suggest that the cost of oil inventories break down the no-arbitrage condition beyond the one year horizon.

Table 11: RMSE of Out of Sample Forecasts Relative to Random Walk

Years-Ahead	One	Two	Three	Four	Five
VAR(2)	1.14	0.96	0.85	0.77	0.71
VAR(14)	1.13	0.92	0.83	0.78	0.71
AKV (2011)	0.94	-	-	-	-
Oil Futures*	0.93	1.16	1.17	1.21	1.28
Survey*	1.02	-	-	-	0.86

*As reported in Alquist, Kilian, and Vigfusson (2011), 1990m1-2012m10

Beyond the one-year-ahead horizon, I find that the non-SVI models are able to beat oil futures. Further, I find that the estimates beat the survey estimates for the five-year-ahead horizon. The survey forecasts are obtained from Consensus Economics Inc., based on private sector forecasts in a variety of countries as summarized in AKV (2011). The survey provides the arithmetic average for each survey month beginning in 1989m10 and ending in 2009m12. AKV (2011), find that these survey forecasts are worse at predicting oil prices than no-change models at one year but much better at the five year horizon. The VAR(2) and VAR(14) models are able to significantly beat the consensus forecasts for the five year horizon.

Concluding Remarks

In summary, I use percentage changes in inventories to identify the speculative demand component of oil price movements, as suggested by Killian and Murphy (2011). I conduct the analysis using a VAR for the global market of crude oil. In contrast to the large set of sign restrictions and elasticity bounds imposed by Killian and Murphy (2011), I am able to match the theoretical movements in the price of crude oil. I do this using more appropriate measures of demand, oil prices, and inventories data and the specification of year-over-year data.

By utilizing the year-over-year specification of the data I am able to correct for the unit root in the price of oil, and any potential bias imposed by global demand and OECD inventory proxy in Killian and Murphy (2011). Further, since the role of inventories is not well understood and I do not impose the large set of sign restrictions, the data was able to speak regarding its role. I find that supply tends to respond countercyclically to the oil price movements and inventories respond countercyclically to global demand. This suggests some degree of forward looking behaviour of the markets and of OPEC or producers in smoothing oil price movements over time.

In addition, I find that speculation may exacerbate the oil price movements when the there is volatility and uncertainty in the supply of oil. This may be the reason for the massive array of sign restrictions imposed by Killian and Murphy (2011). The role of speculation in influencing price movement may be a symptom of the volatility "hump", a stylized fact observed in trading markets where volatility is higher for assets with shorter maturities (Hull, 2012). The overreaction of traders to price shifts and short rates tends to increase spot volatility and can have large effects, such as regular and significant fluctuations of oil prices.

The VAR(14) and VAR(24) models' results are consistent with the theory . I prefer the VAR(24) model when analysing the market for oil and understanding shifts in the real price since these yield more sensible results in terms of the persistence of the shocks. The model is able to identify the demand shocks driving the price of oil in late 2008 to 2010, thus consistent with expectations. However, in contrast to Killian and Murphy (2011) I find that supply contractions were a driving factor in the run up of the price of oil in 2003 to 2007. This aligns the results more closely with Hamilton (2009) and the theories of peak oil.

The VAR(15) model with an SVI variables poses several issues and its results indicate that the findings of Da et al. (2011a) do not apply to all spot prices and investment mechanisms. Using the proxy for retail investor attention in addition to inventories, world industrial production, global crude oil supply and the real price of oil in a VAR holds insignificant predictive power, displayed through the poor performance in forecasting. The specification of search volume data and brief historical database may be the reason that including such a variable in multivariate econometric models leads to eigenvalue instability and heteroskedastic errors.

However, historical shock decompositions give results that are consistent with Da et al. (2011a) and Barber and Odean (2008). Increases in retail investor attention are positively correlated with increases in inventories. This is a result of more oil futures being purchased at the retail level, thus creating an upward pressure on the spot price of oil. SVI found on Google trends may be able to give some insight on retail investor behaviour in oil markets, thus providing more information on the determinants of oil price fluctuations. In addition, SVI as a metric for retail investor attention may be appropriate for forecasting and other forms of economic analysis when not used in a VAR or for analysis of commodities markets.

Perhaps simple regression such as ordinary least squares analysis on highly traded equities may be a more appropriate foundation and further support the findings of Da et al (2011a). Unfortunately, with regards to the global market for crude oil, the data exhibit questionable issues, such as the 0-100 scale and brief history (9 years); thus using SVI as a forecasting tool in this realm is inappropriate.

I also find a larger role of inventories in affecting the price of crude oil, consistent with the Granger causality tests. In contrast to Killian and Murphy (2012), I find that the popular assumption that speculation was behind the increases in the price of oil in 2008 has some degree of merit. Further, I find some role of speculative demand in the increases in the price of oil in 2011. Interestingly, at other times, I find that the supply of oil tends to respond in a countercyclical fashion over the oil price cycles, especially if demand driven. The countercyclical interaction is simply an increase in the supply of oil that corresponds with a downward shift in real prices and vice versa. This is consistent with the objectives of oil producers to smooth the price of oil without the use of inventories within the OECD countries.

I add to the investigation of AKV (2011), and evaluate the model in forecasting the price of oil over longer time horizons. I find that the VAR(2) and VAR(14) models perform well in forecasts over the one-year-ahead horizons, being able to beat random walk estimates. Beyond the one-year-ahead horizon, I find that the models developed in this paper are able to beat oil futures and consensus survey estimates up to the five-year-ahead horizon, suggesting that the model may be appropriate for forecasting experts, such as futures traders, central banks, and oil producers, whose horizons are above one year. Consistency in forecast performance and parameter stability suggests that these models may be useful for understanding the crude oil market in the future.

References

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

Alquist, R., Kilian, L., and Vigfusson, R., (2011). “Forecasting the price of oil.” CEPR Discussion Papers 8388.

Barber, Brad M., and Terrance Odean, 2008, All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors, Review of Financial Studies 21(2), 785-818.

Barber, Brad M., Terrance Odean, and Ning Zhu, 2009, Do Retail Trades Move Markets? , Review of Financial Studies 21(1), 151-186.

Barsky, R.B., and L. Kilian, (2002). “Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative,” in: *NBER Macroeconomics Annual 2001*, B.S. Bernanke and K. Rogoff (eds.), MIT Press: Cambridge, MA, 137-183.

Benes, J., Chauvet, M., Kamenik, O., Kumhof, M., Laxton, D., Mursula, S., and Selody, J., (2012). “The Future of Oil: Geology versus Technology” (No. 12/109). IMF Working Paper.

Blinder, Alan S. and Kilian, Lutz. "Comments and Discussion." *Brookings Papers on Economic Activity* 2009.1 (2009): 262-78. Print.

Da, Zi, Joseph, Engelberg, and Pengjie Gao, 2011, In Search of Attention. Journal of Finance 66, 1461-1499.

Daniel, Kent D., Mark Grinblatt, Sheridan Titman and Russ Wermers, 1997, Measuring Mutual Fund Performance with Characteristic-Based Benchmarks, Journal of Finance 52, 1035-1058.

Daniel, Kent D., David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor Psychology and Security Market Under- and Overreactions, Journal of Finance 53(5), 1839-1885.

Ding, Rong and Hou, Wenxuan, 2009, Retail Investor Attention and Stock Liquidity. Available at SSRN, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1786762

Fattouh, B., Kilian, L., and Mahadeva, L., (2012). "The Role of Speculation in Oil Markets: What Have We Learned So Far?" CEPR Discussion Papers 8916.

Giannone, D., and L. Reichlin (2006), "Does Information Help Recover Structural Shocks from Past Observations," Journal of the European Economic Association, 4, 455-465.

Granger, C. W. J. (1969). "Investigating Causal Relations by Econometric Models and Cross-spectral Methods". *Econometrica* **37** (3): 424–438.

Hamilton, J.D., (2009), "Causes and Consequences of the Oil Shock of 2007-08," Brookings Papers on Economic Activity, 1 (Spring), 215-261.

Hull, John. "Interest Rate Derivatives: The Standard Market Models." *Options, Futures, and Other Derivatives*. Boston: Prentice Hall, 2012. 656-57. Print.

Hyndman, R., Koehler, A. (2006). "Another look at measures of forecast accuracy" International Journal of Forecasting: 679–688.

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., and Murphy, D., (2011). "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," forthcoming: *Journal of Applied Econometrics*.

Appendix

Figure 1: Real Crude Oil Price and Inventories

1970m1 - 2012m10

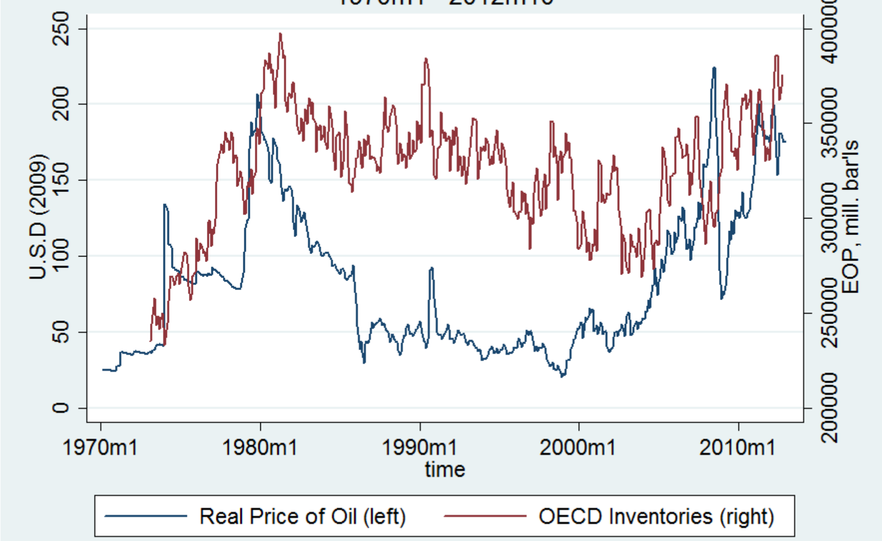

Figure 2: Real Crude Oil Price and World Industrial Production

1970m1 - 2012m10

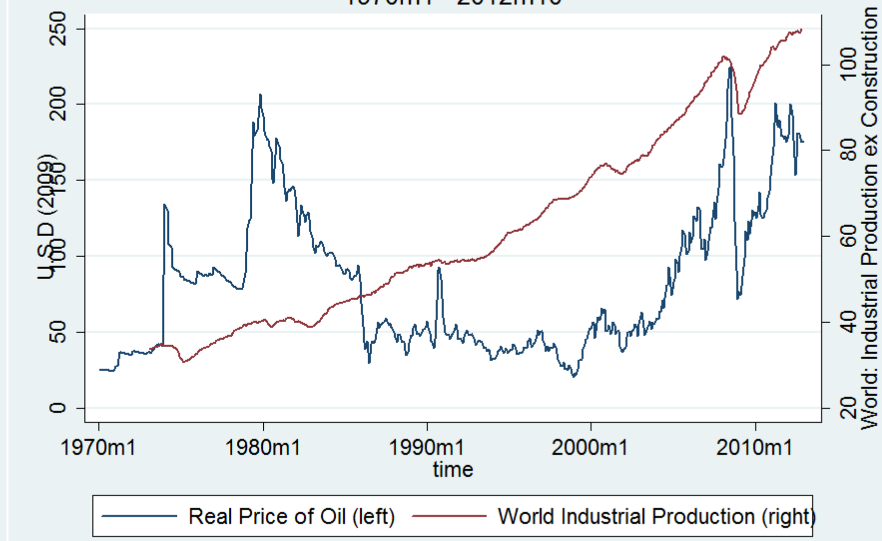

Figure 3 : Real Crude Oil Price and Global Crude Oil Production

1970m1 - 2012m10

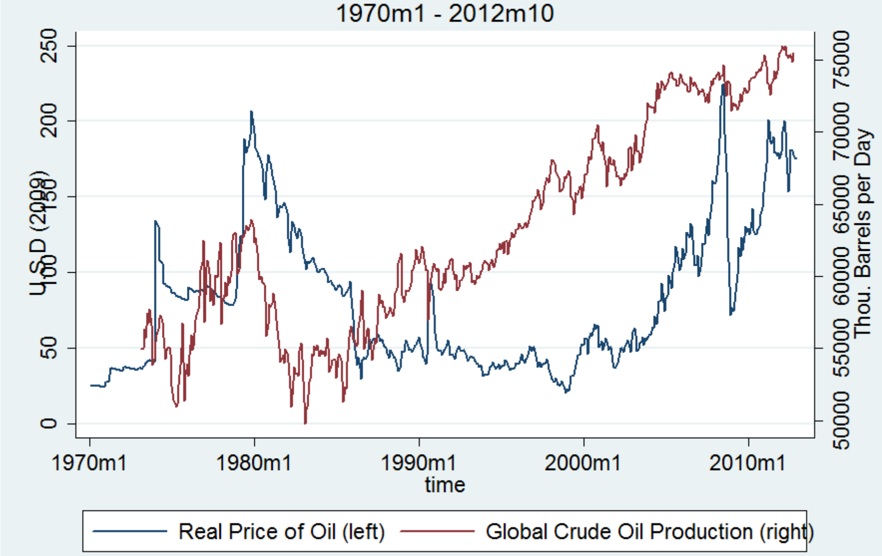

Figure 4 : Real Crude Oil Price and Google SVI

2004m1 - 2013m7

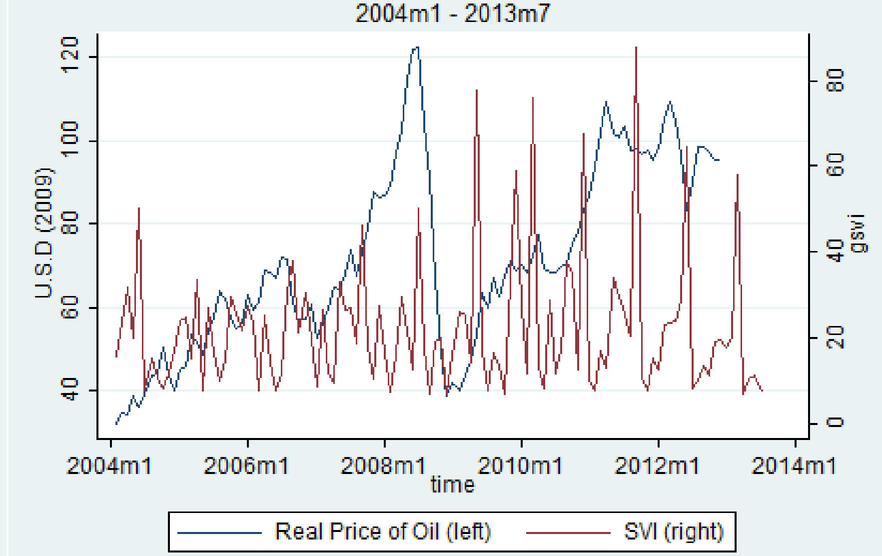

Figure 5 : Real Crude Oil Price and Inventories

1970m1 - 2012m10

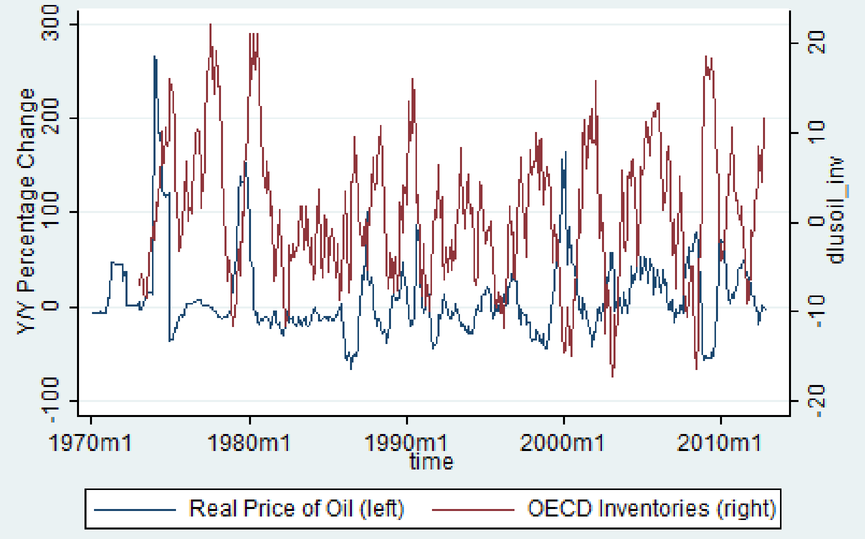

Figure 6 : Real Crude Oil Price and World Industrial Production

1970m1 - 2012m10

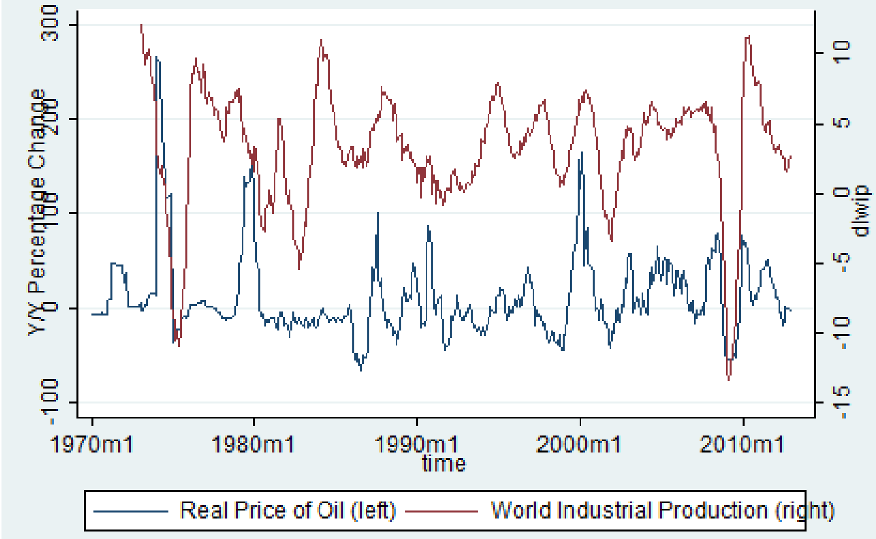

Figure 7 : Real Crude Oil Price and Global Crude Oil Production

1970m1 - 2012m10

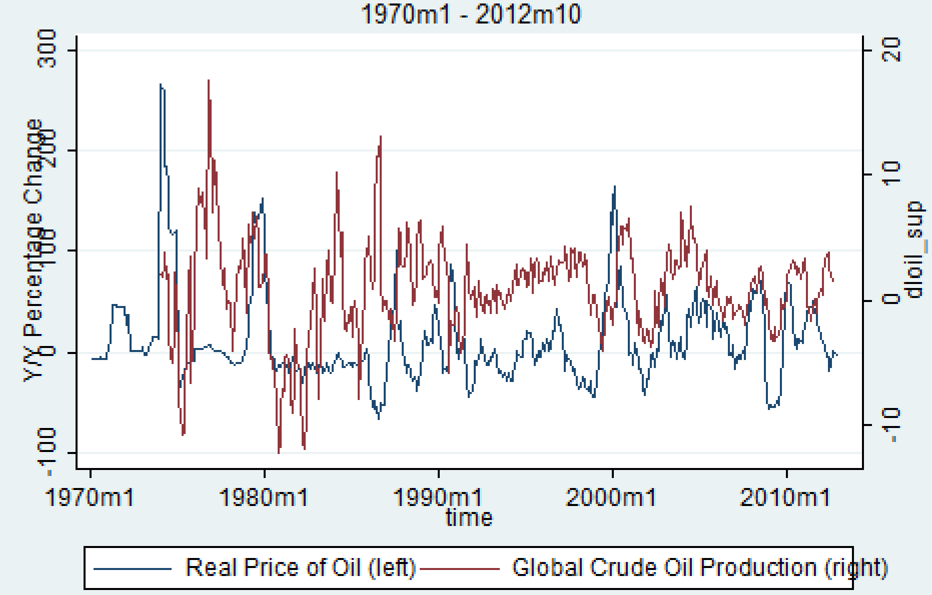

Figure 8 : Real Crude Oil Price and Google SVI

2004m1 - 2013m7

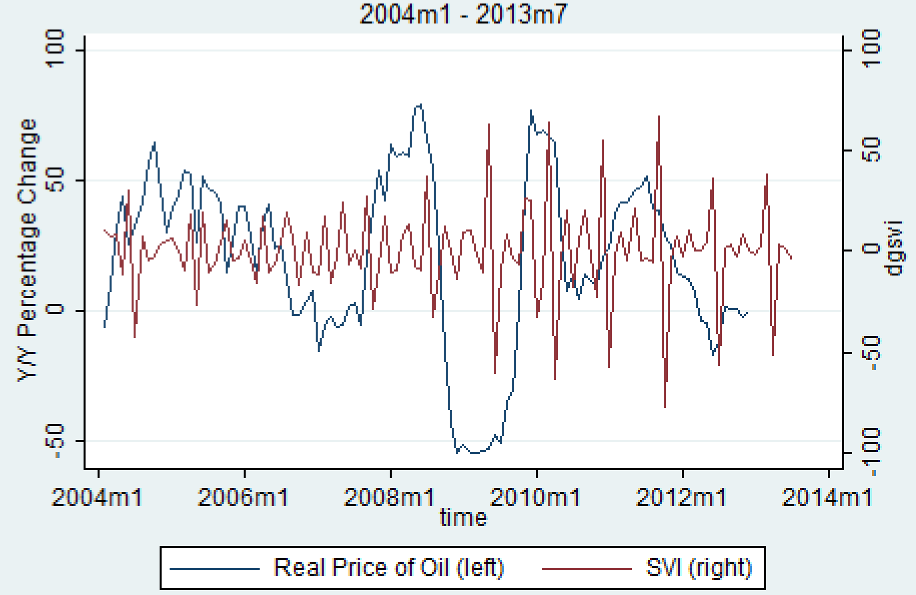

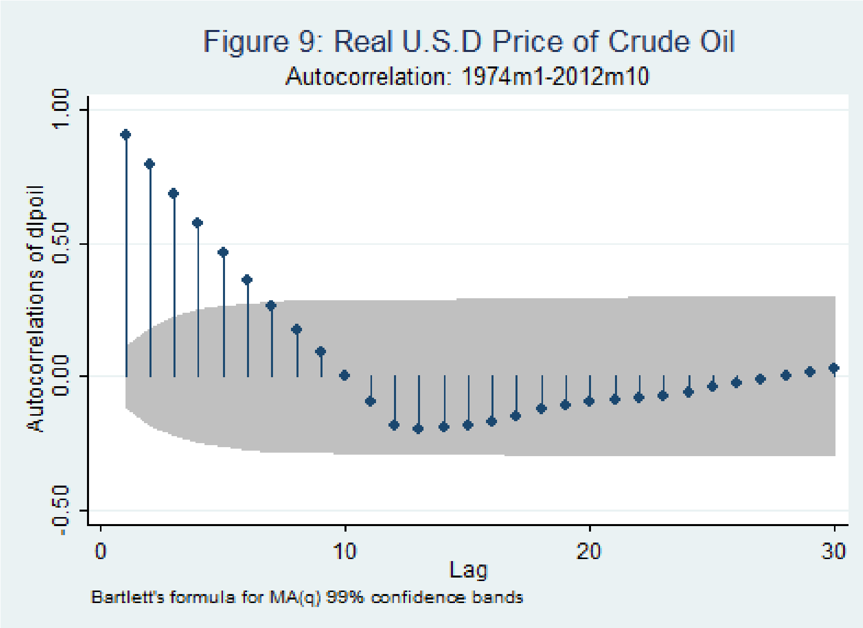

Figure 9: Real U.S.D Price of Crude Oil
Autocorrelation: 1974m1-2012m10

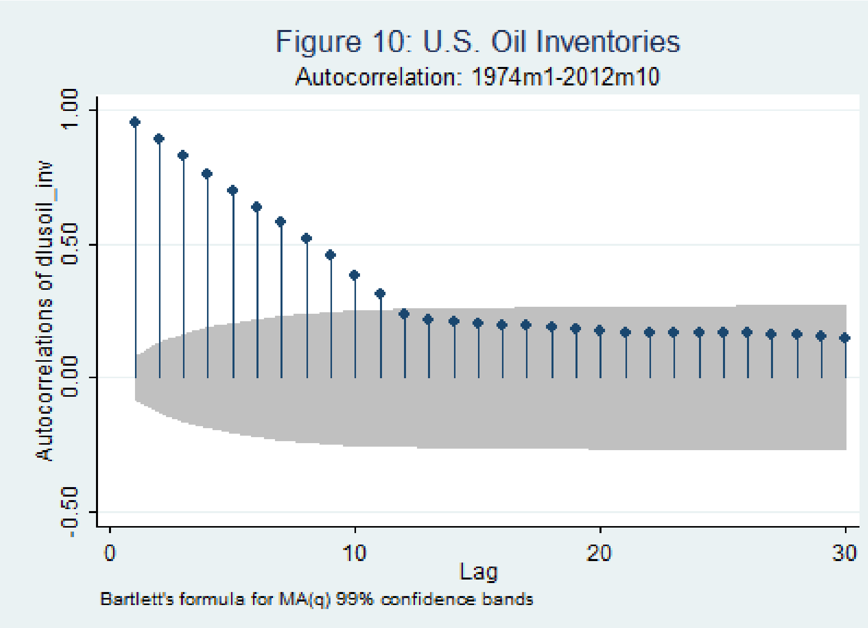

Figure 10: U.S. Oil Inventories
Autocorrelation: 1974m1-2012m10

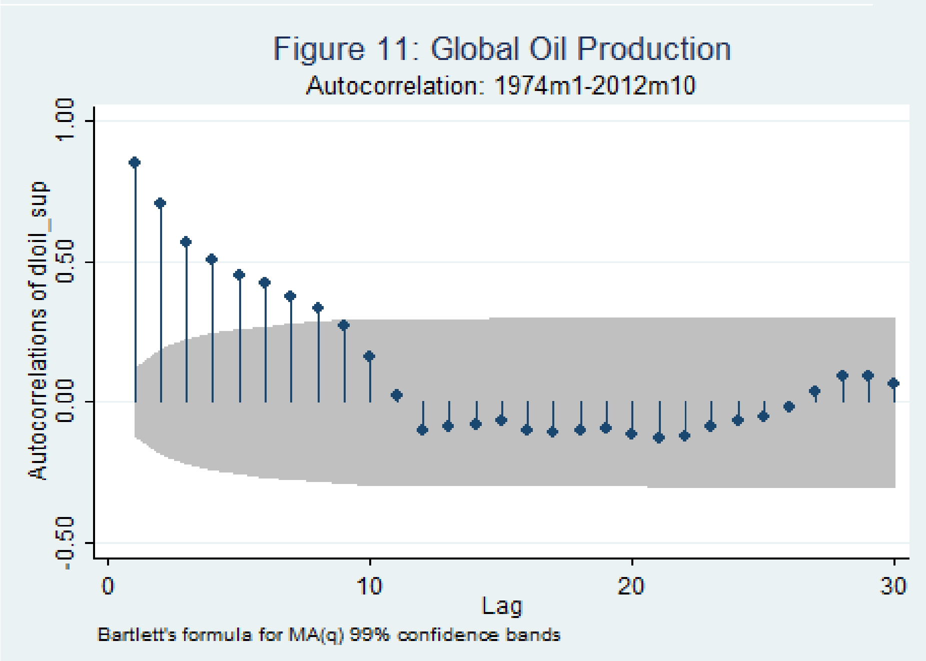

Figure 11: Global Oil Production
Autocorrelation: 1974m1-2012m10

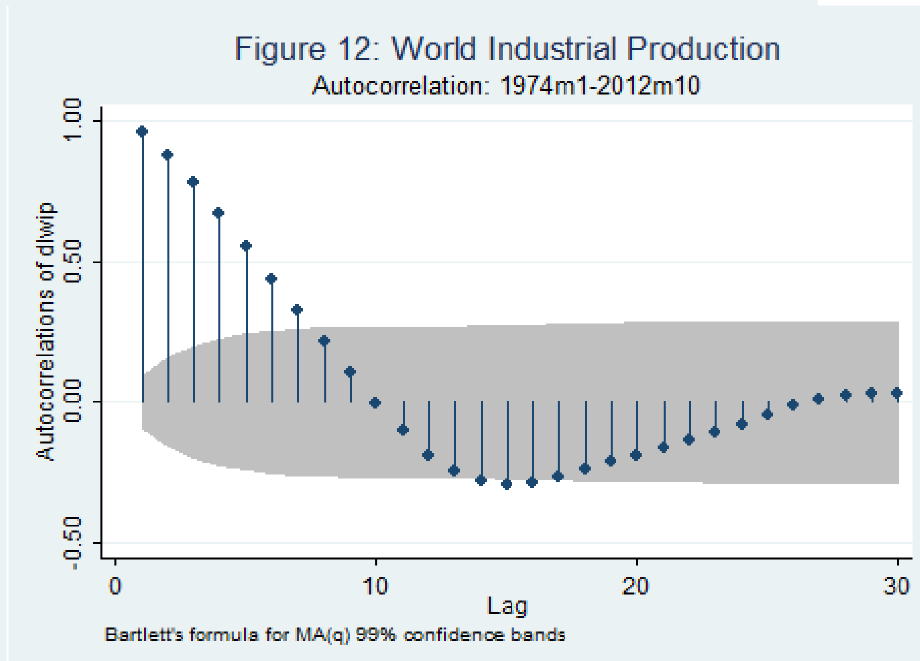

Figure 12: World Industrial Production
Autocorrelation: 1974m1-2012m10

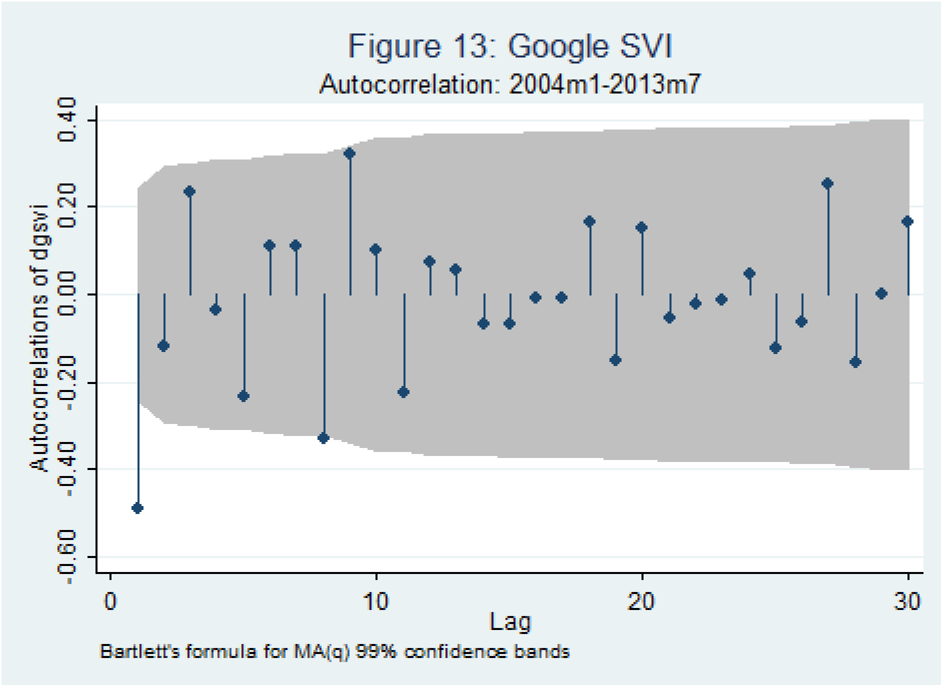
Figure 13: Google SVI
Autocorrelation: 2004m1-2013m7
Autocorrelations of dgsvi
-0.60
-0.40
-0.20
0.00
0.20
0.40
0
10
20
30
Lag
Bartlett's formula for MA(q) 99% confidence bands

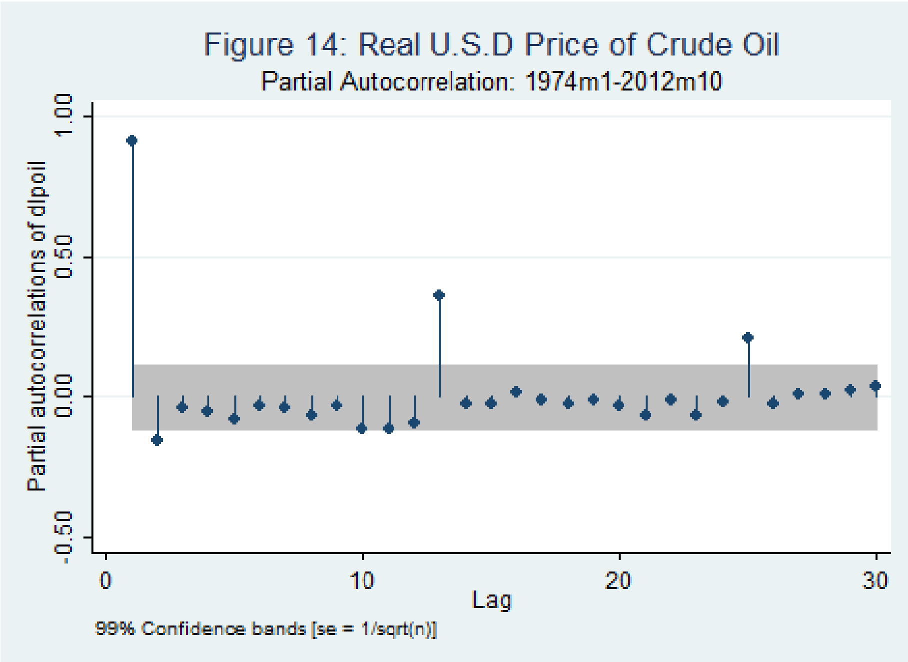
Figure 14: Real U.S.D Price of Crude Oil
Partial Autocorrelation: 1974m1-2012m10
Partial autocorrelations of dlpoil
-0.50
0.00
0.50
1.00
0
10
20
30
Lag
99% Confidence bands [se = 1/sqrt(n)]

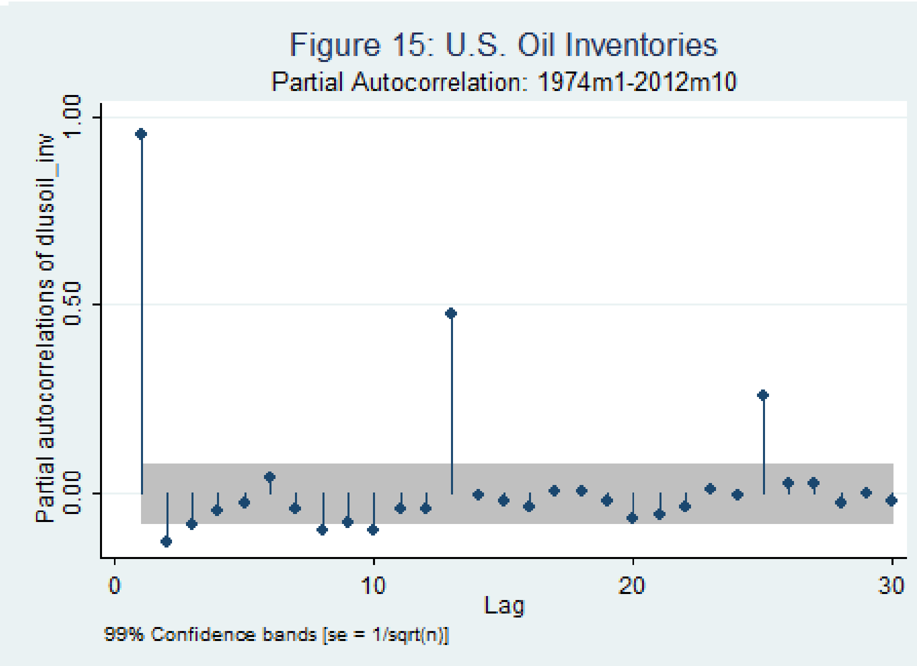
Figure 15: U.S. Oil Inventories
Partial Autocorrelation: 1974m1-2012m10
Partial autocorrelations of dlusoil_inv
0.00
0.50
1.00
0
10
20
30
Lag
99% Confidence bands [se = 1/sqrt(n)]

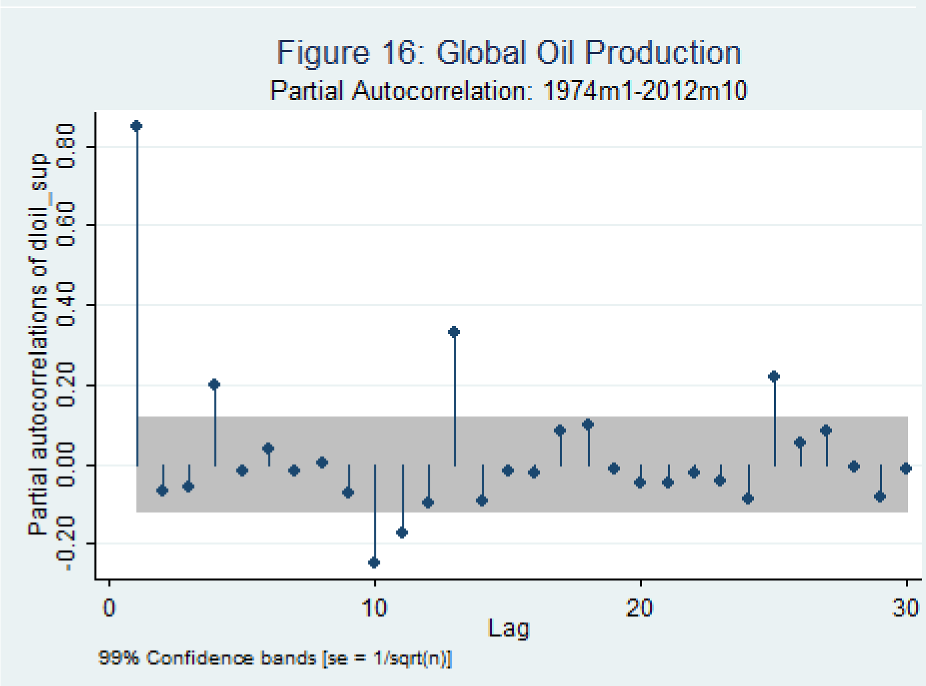
Figure 16: Global Oil Production
Partial Autocorrelation: 1974m1-2012m10
Partial autocorrelations of dloil_sup
-0.20
0.00
0.20
0.40
0.60
0.80
0
10
20
30
Lag
99% Confidence bands [se = 1/sqrt(n)]

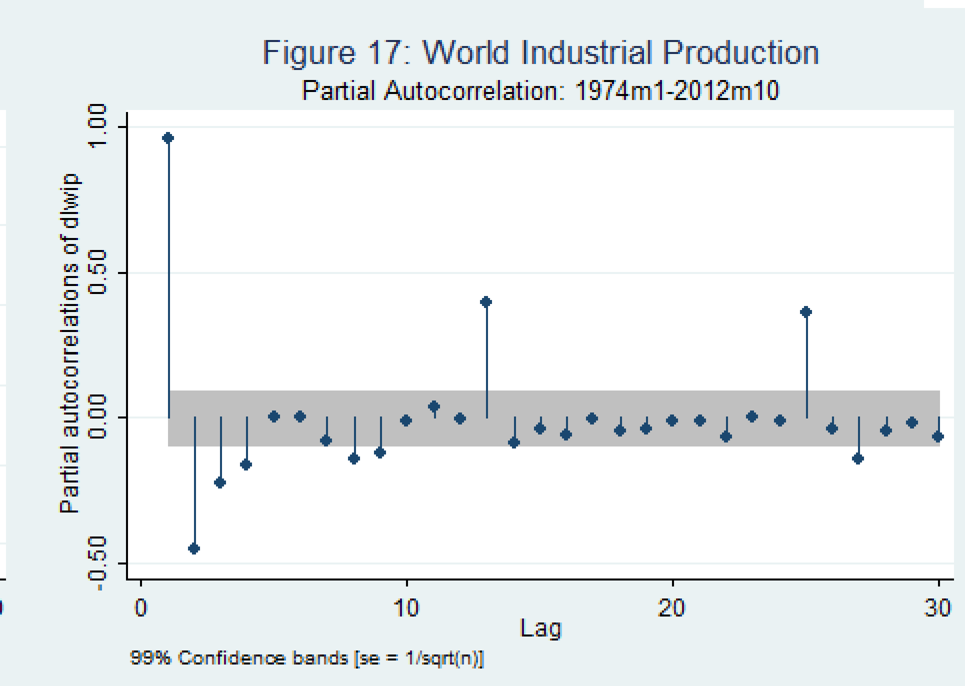
Figure 17: World Industrial Production
Partial Autocorrelation: 1974m1-2012m10
Partial autocorrelations of dlwip
-0.50
0.00
0.50
1.00
0
10
20
30
Lag
99% Confidence bands [se = 1/sqrt(n)]

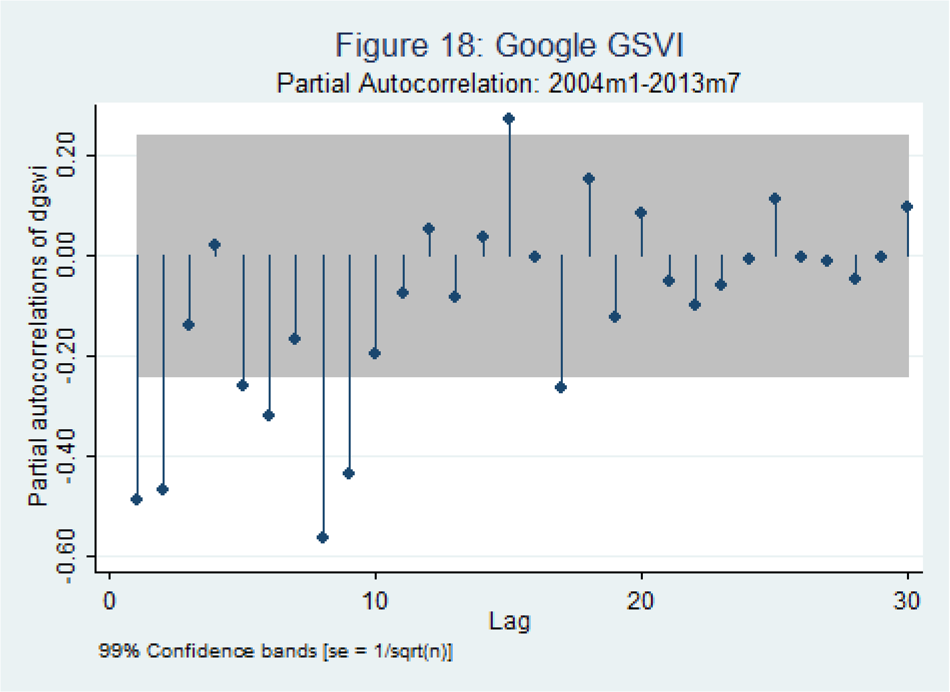
Figure 18: Google GSVI
Partial Autocorrelation: 2004m1-2013m7
Partial autocorrelations of dgsvi
-0.60
-0.40
-0.20
0.00
0.20
0
10
20
30
Lag
99% Confidence bands [se = 1/sqrt(n)]

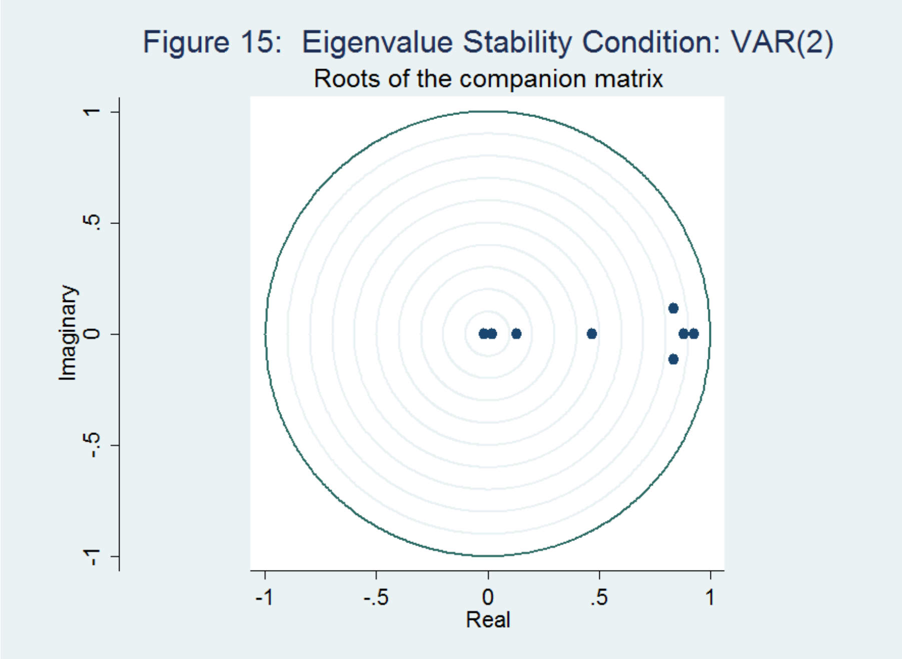
Figure 15: Eigenvalue Stability Condition: VAR(2)
Roots of the companion matrix
Imaginary
-1
-.5
0
.5
1
Real

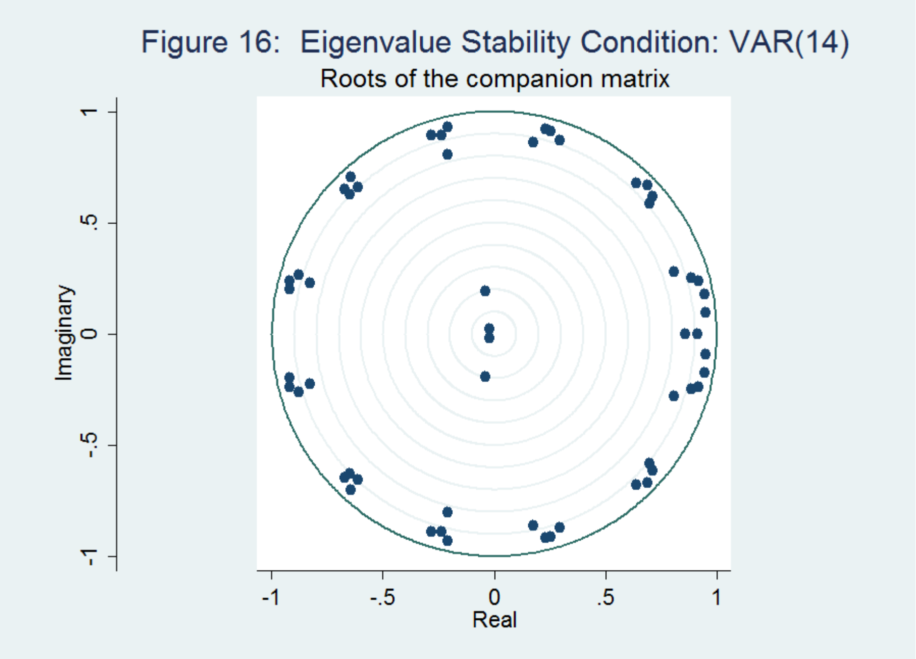
Figure 16: Eigenvalue Stability Condition: VAR(14)
Roots of the companion matrix
Imaginary
-1
-.5
0
.5
1
Real

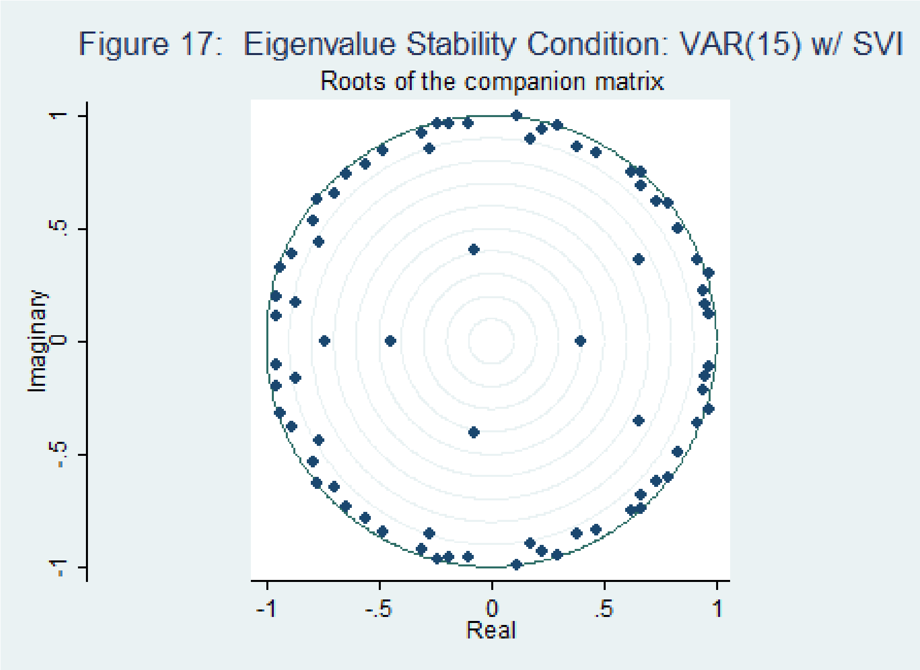
Figure 17: Eigenvalue Stability Condition: VAR(15) w/ SVI
Roots of the companion matrix
Imaginary
Real
-1
-.5
0
.5
1

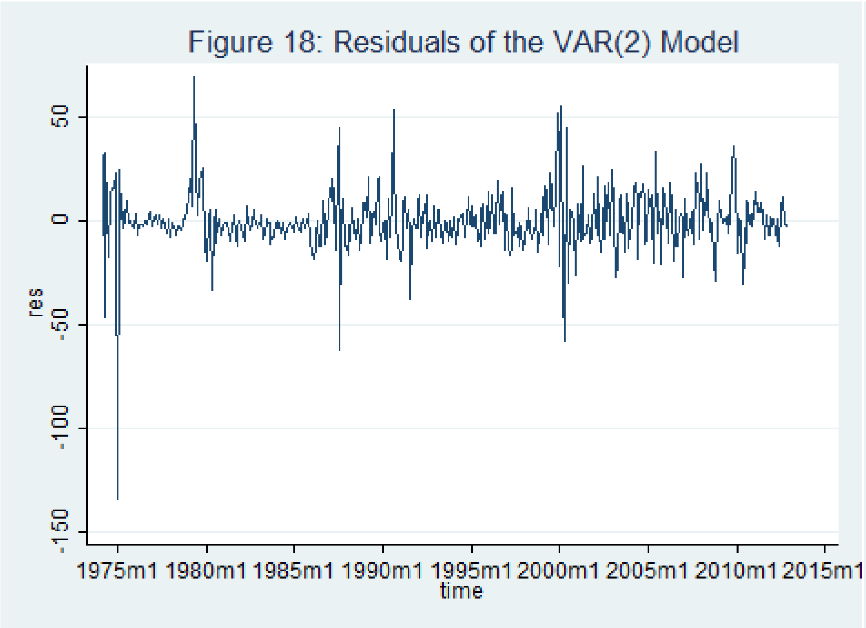
Figure 18: Residuals of the VAR(2) Model
res
time
50
0
-50
-100
-150
1975m1
1980m1
1985m1
1990m1
1995m1
2000m1
2005m1
2010m1
2015m1

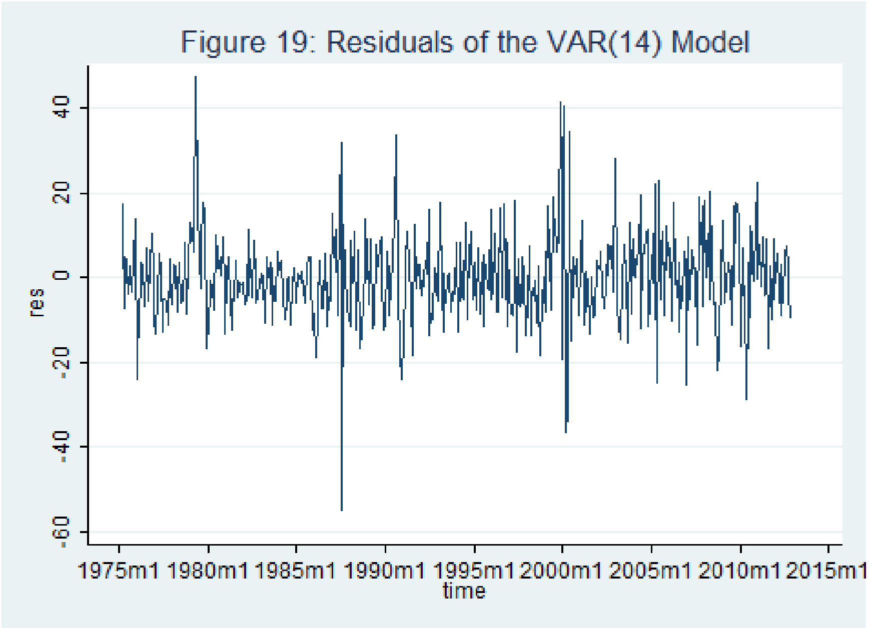
Figure 19: Residuals of the VAR(14) Model
res
time
40
20
0
-20
-40
-60
1975m1
1980m1
1985m1
1990m1
1995m1
2000m1
2005m1
2010m1
2015m1

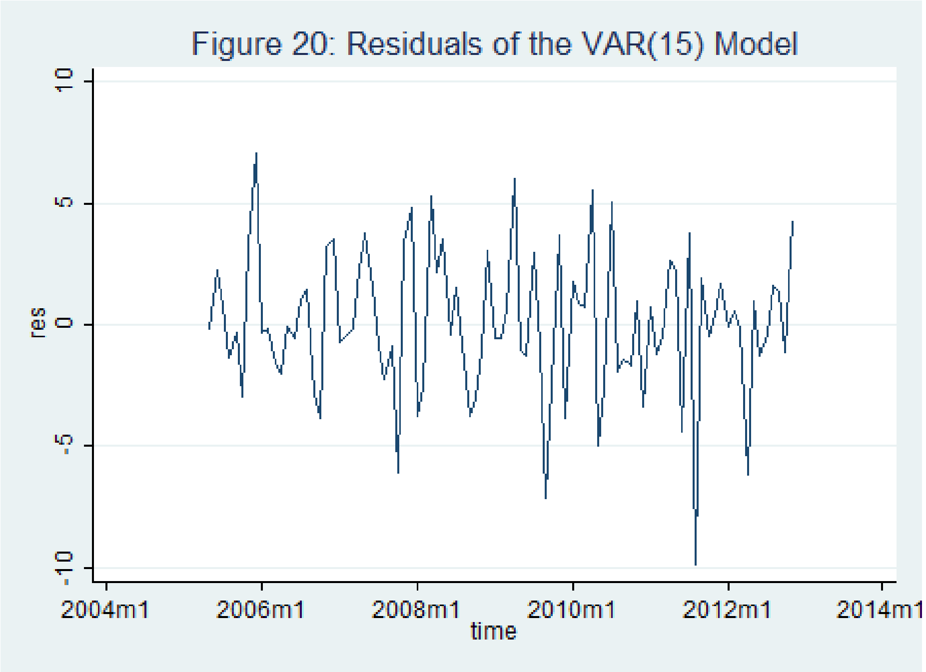
Figure 20: Residuals of the VAR(15) Model
res
-10
-5
0
5
10
2004m1
2006m1
2008m1
2010m1
2012m1
2014m1
time

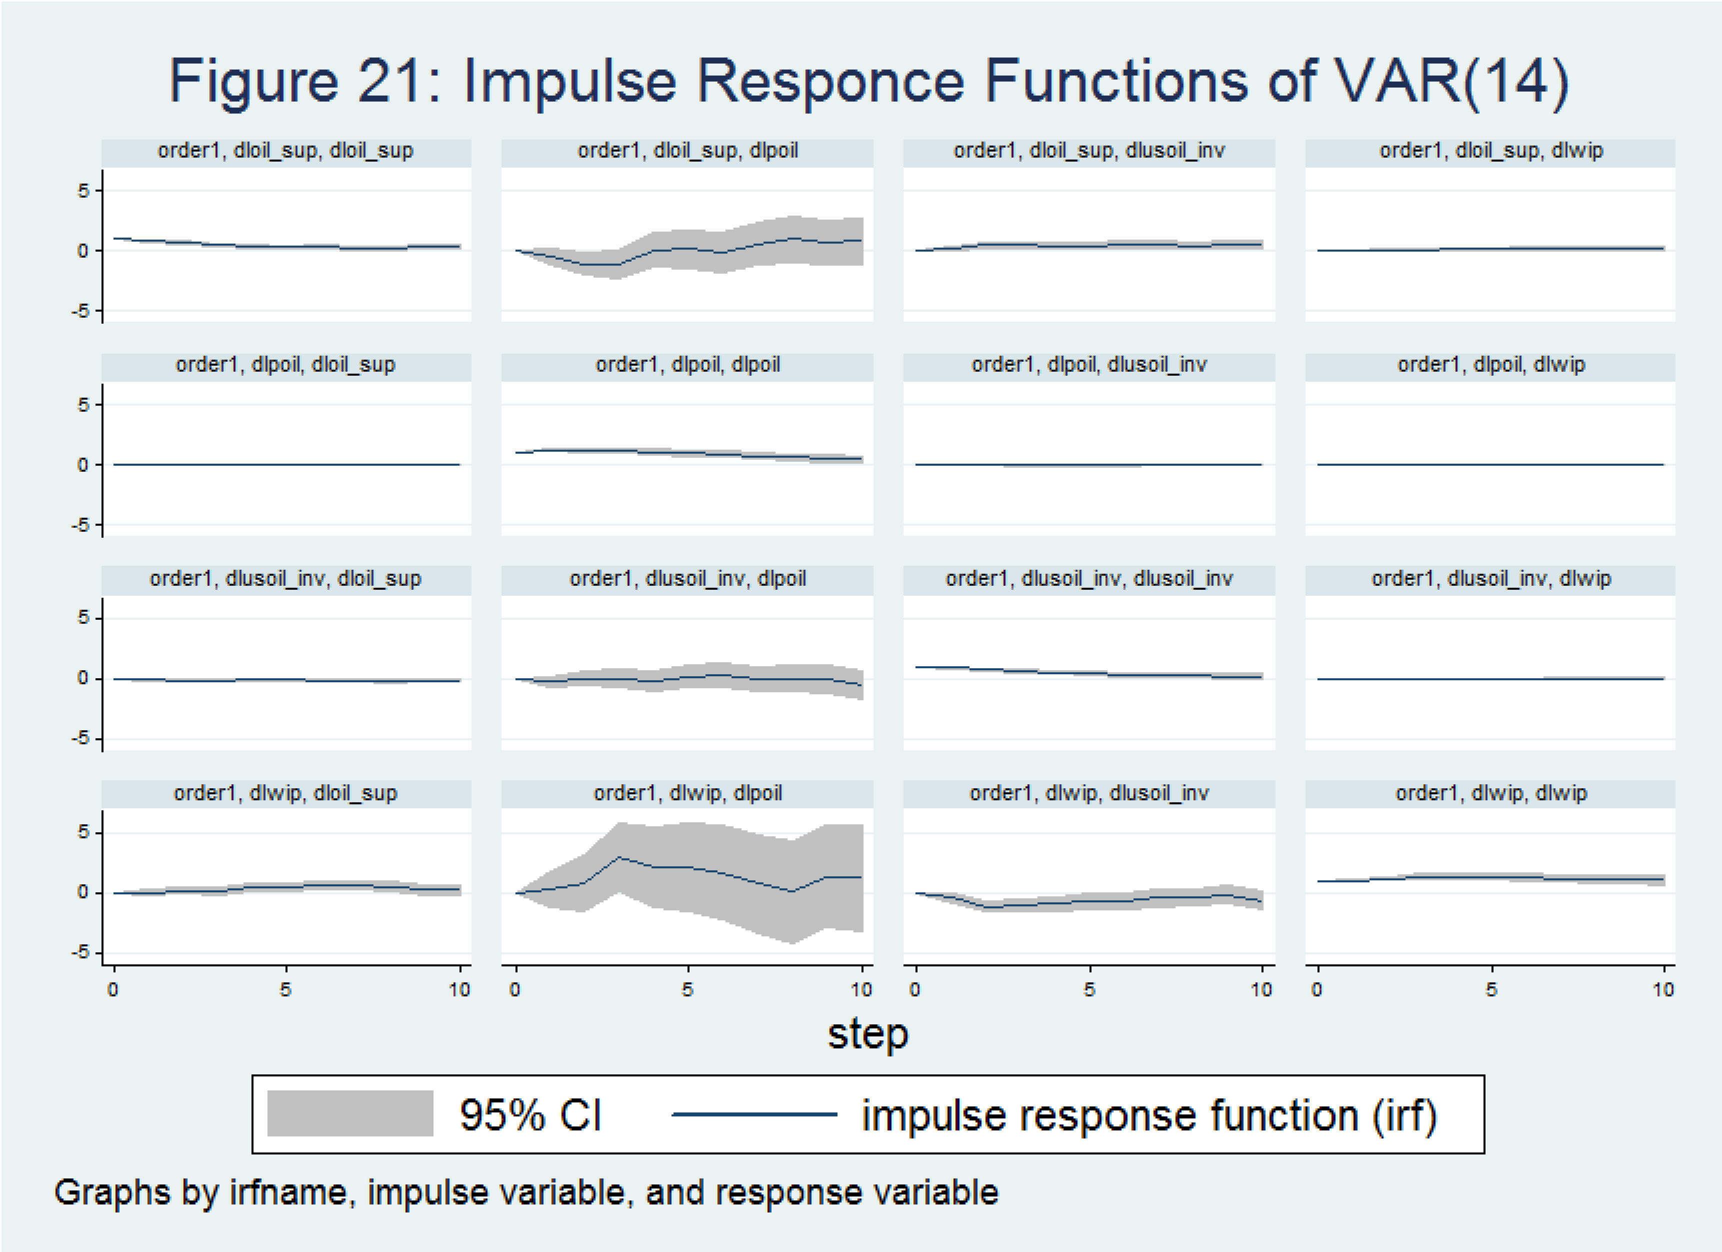
Figure 21: Impulse Responce Functions of VAR(14)
order1, dloil_sup, dloil_sup
order1, dloil_sup, dlpoil
order1, dloil_sup, dlusoil_inv
order1, dloil_sup, dlwip
order1, dlpoil, dloil_sup
order1, dlpoil, dlpoil
order1, dlpoil, dlusoil_inv
order1, dlpoil, dlwip
order1, dlusoil_inv, dloil_sup
order1, dlusoil_inv, dlpoil
order1, dlusoil_inv, dlusoil_inv
order1, dlusoil_inv, dlwip
order1, dlwip, dloil_sup
order1, dlwip, dlpoil
order1, dlwip, dlusoil_inv
order1, dlwip, dlwip
step
95% CI
impulse response function (irf)
Graphs by irfname, impulse variable, and response variable

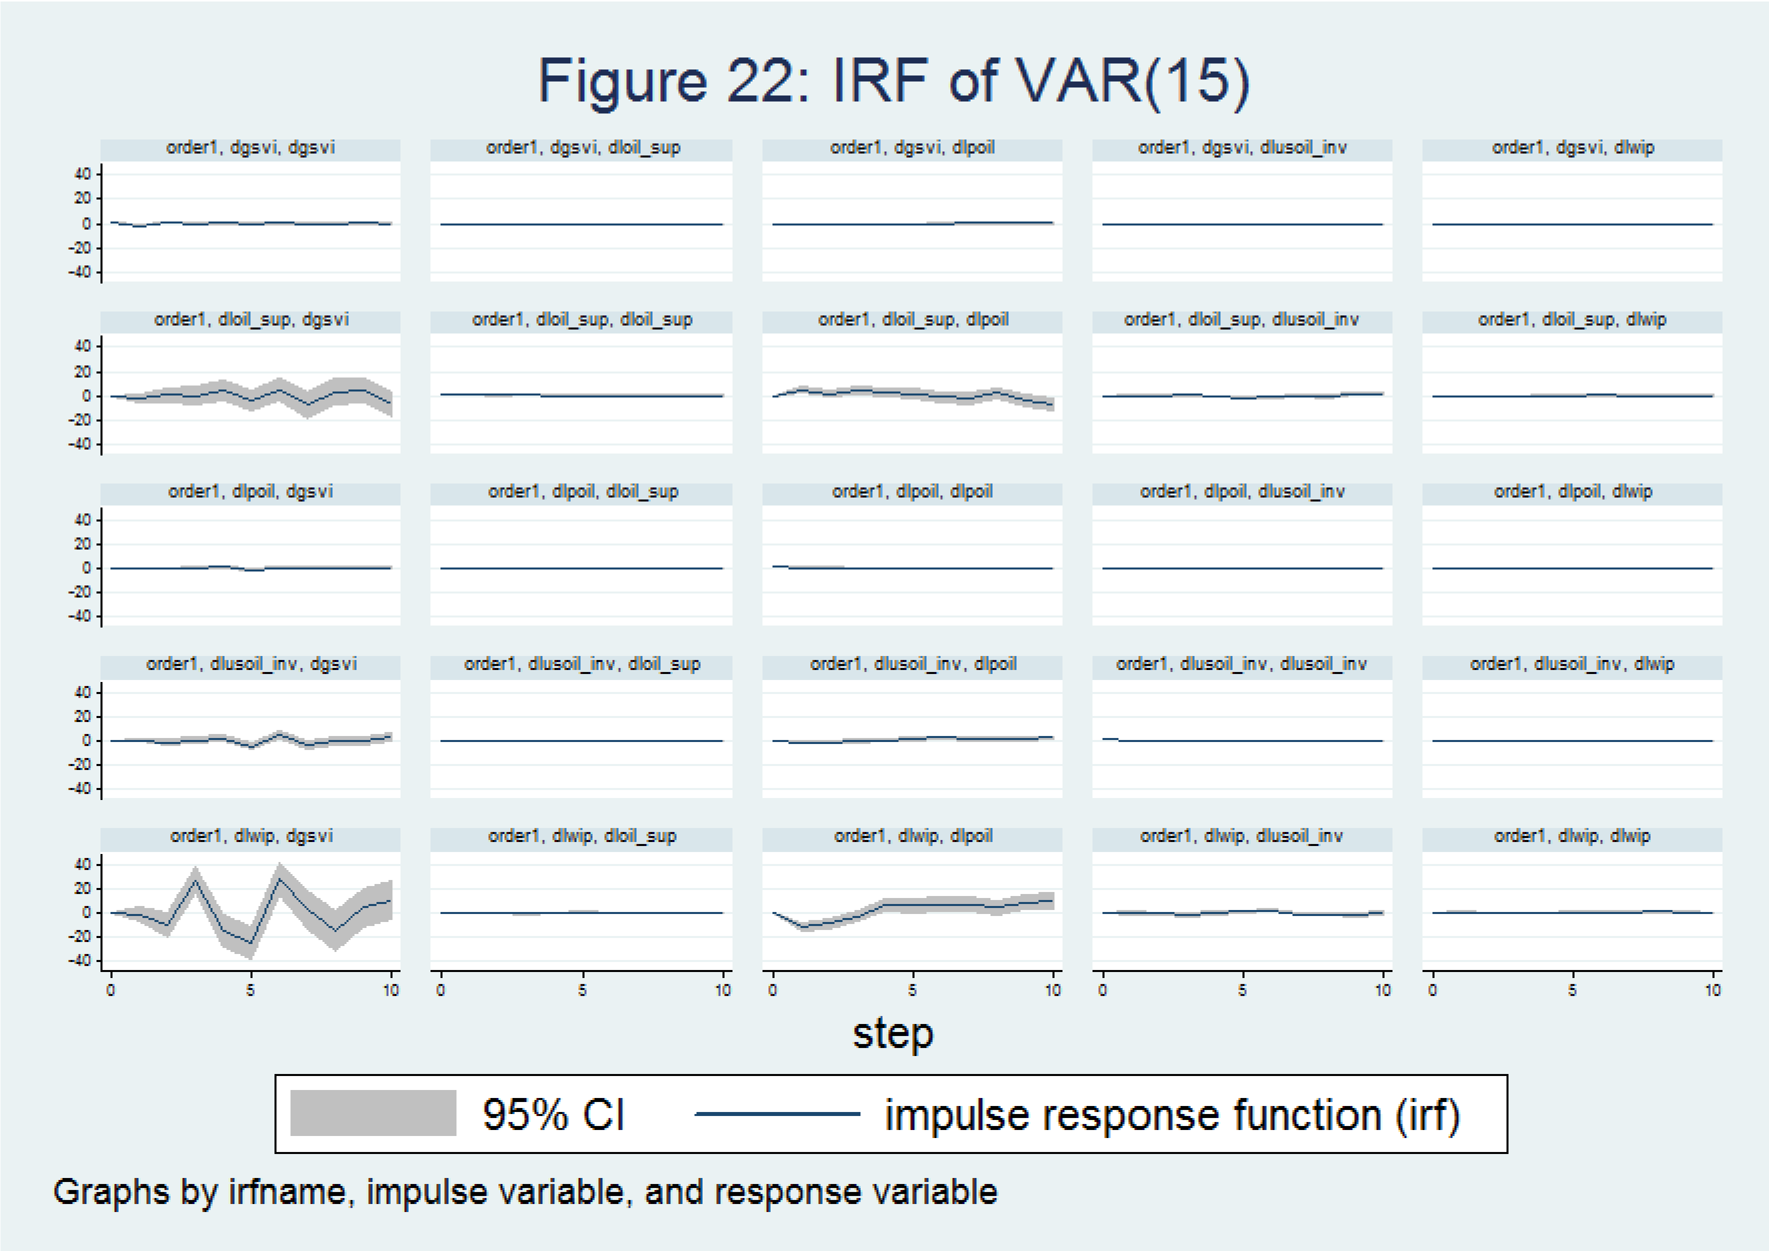

Figure 22: IRF of VAR(15)
order1, dgsvi, dgsvi
order1, dgsvi, dloil_sup
order1, dgsvi, dlpoil
order1, dgsvi, dlusoil_inv
order1, dgsvi, dlwip
order1, dloil_sup, dgsvi
order1, dloil_sup, dloil_sup
order1, dloil_sup, dlpoil
order1, dloil_sup, dlusoil_inv
order1, dloil_sup, dlwip
order1, dlpoil, dgsvi
order1, dlpoil, dloil_sup
order1, dlpoil, dlpoil
order1, dlpoil, dlusoil_inv
order1, dlpoil, dlwip
order1, dlusoil_inv, dgsvi
order1, dlusoil_inv, dloil_sup
order1, dlusoil_inv, dlpoil
order1, dlusoil_inv, dlusoil_inv
order1, dlusoil_inv, dlwip
order1, dlwip, dgsvi
order1, dlwip, dloil_sup
order1, dlwip, dlpoil
order1, dlwip, dlusoil_inv
order1, dlwip, dlwip
40
20
0
-20
-40
0
5
10
step
95% CI
impulse response function (irf)
Graphs by irfname, impulse variable, and response variable

Figure 23: Historical Shock Decompositions: VAR(2)

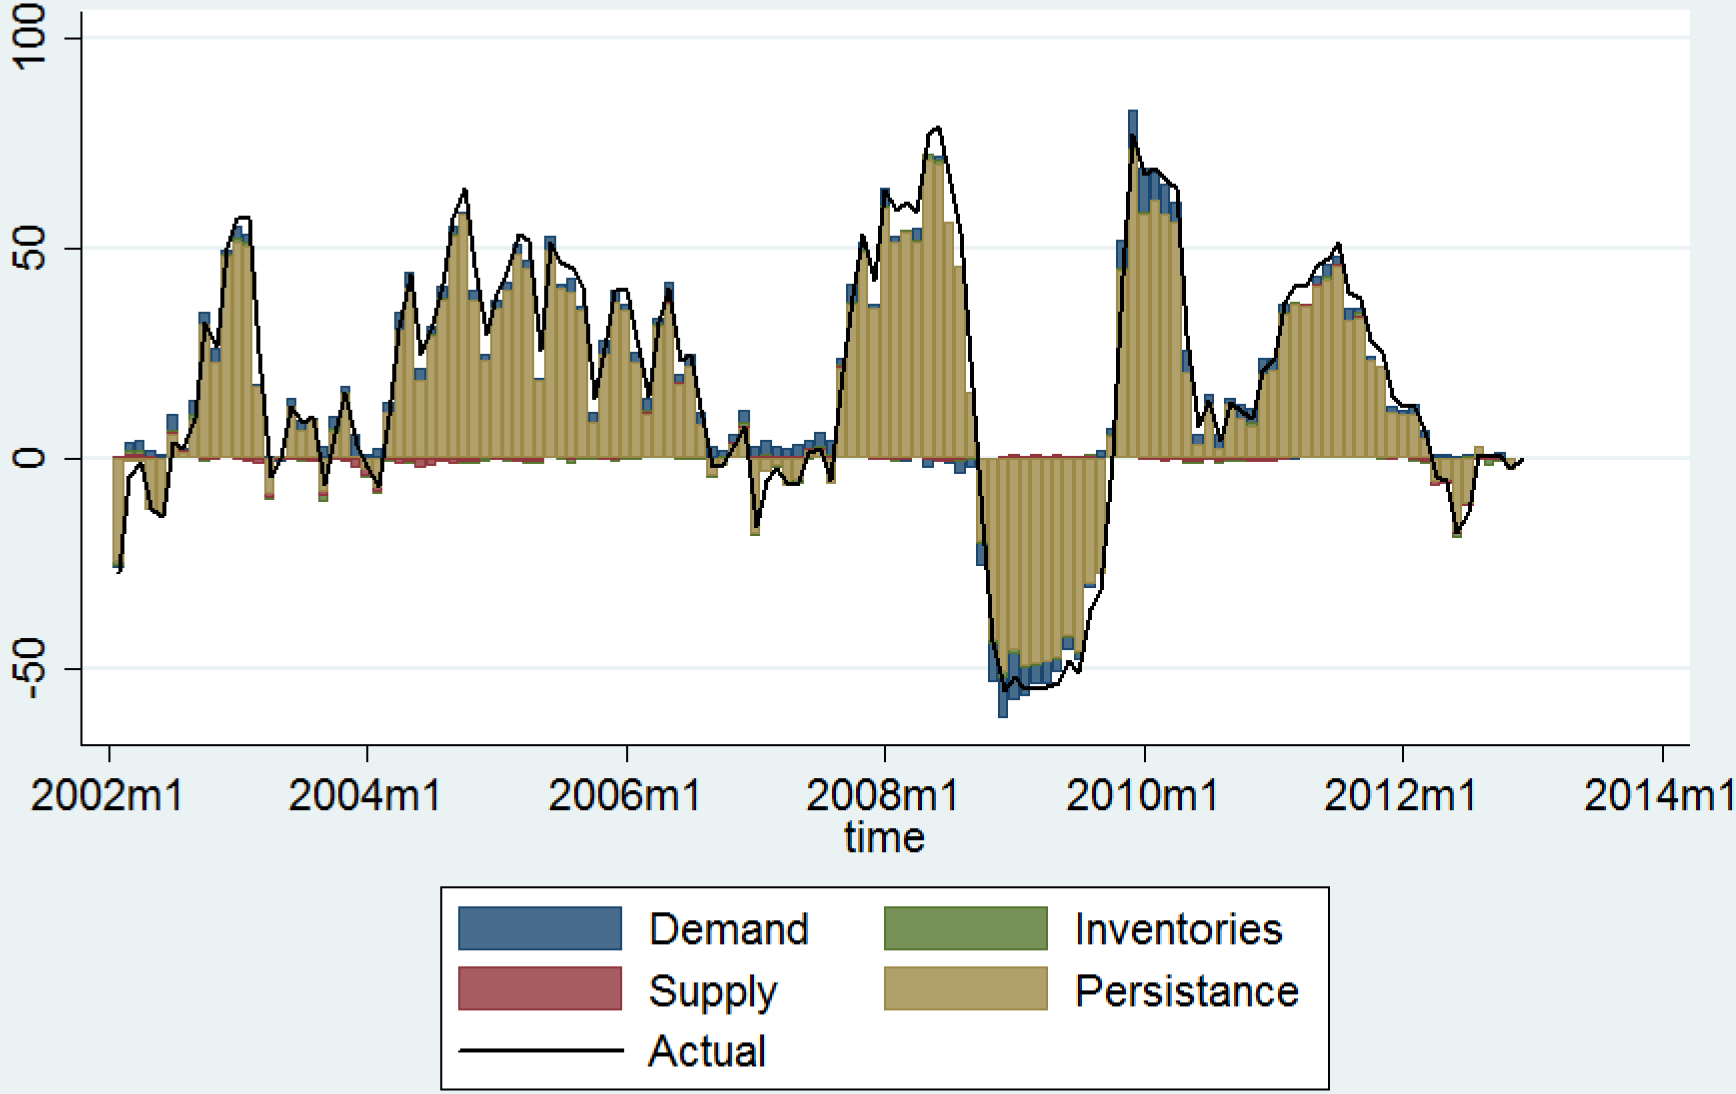

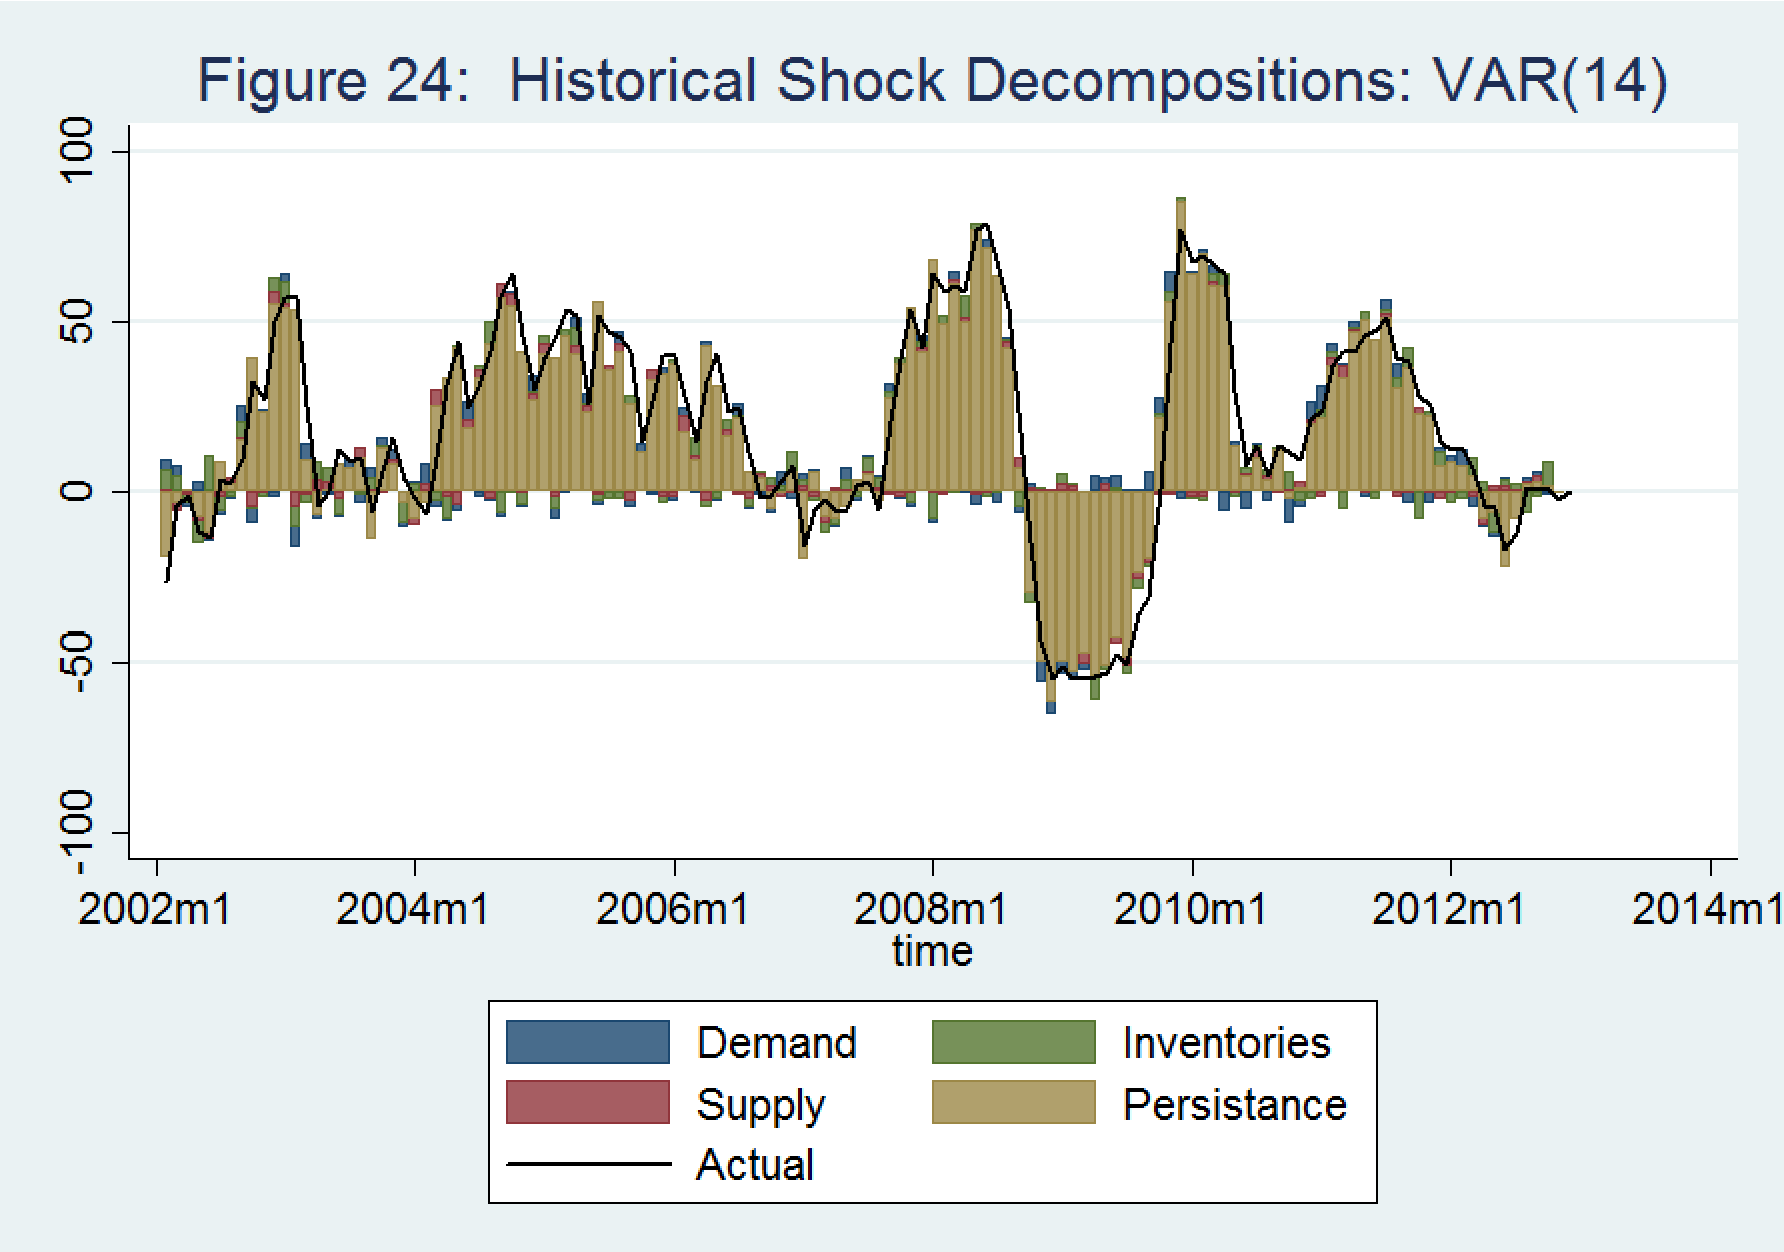
Figure 24: Historical Shock Decompositions: VAR(14)
100
50
0
-50
-100
2002m1
2004m1
2006m1
2008m1
2010m1
2012m1
2014m1
time
Demand
Inventories
Supply
Persistance
Actual

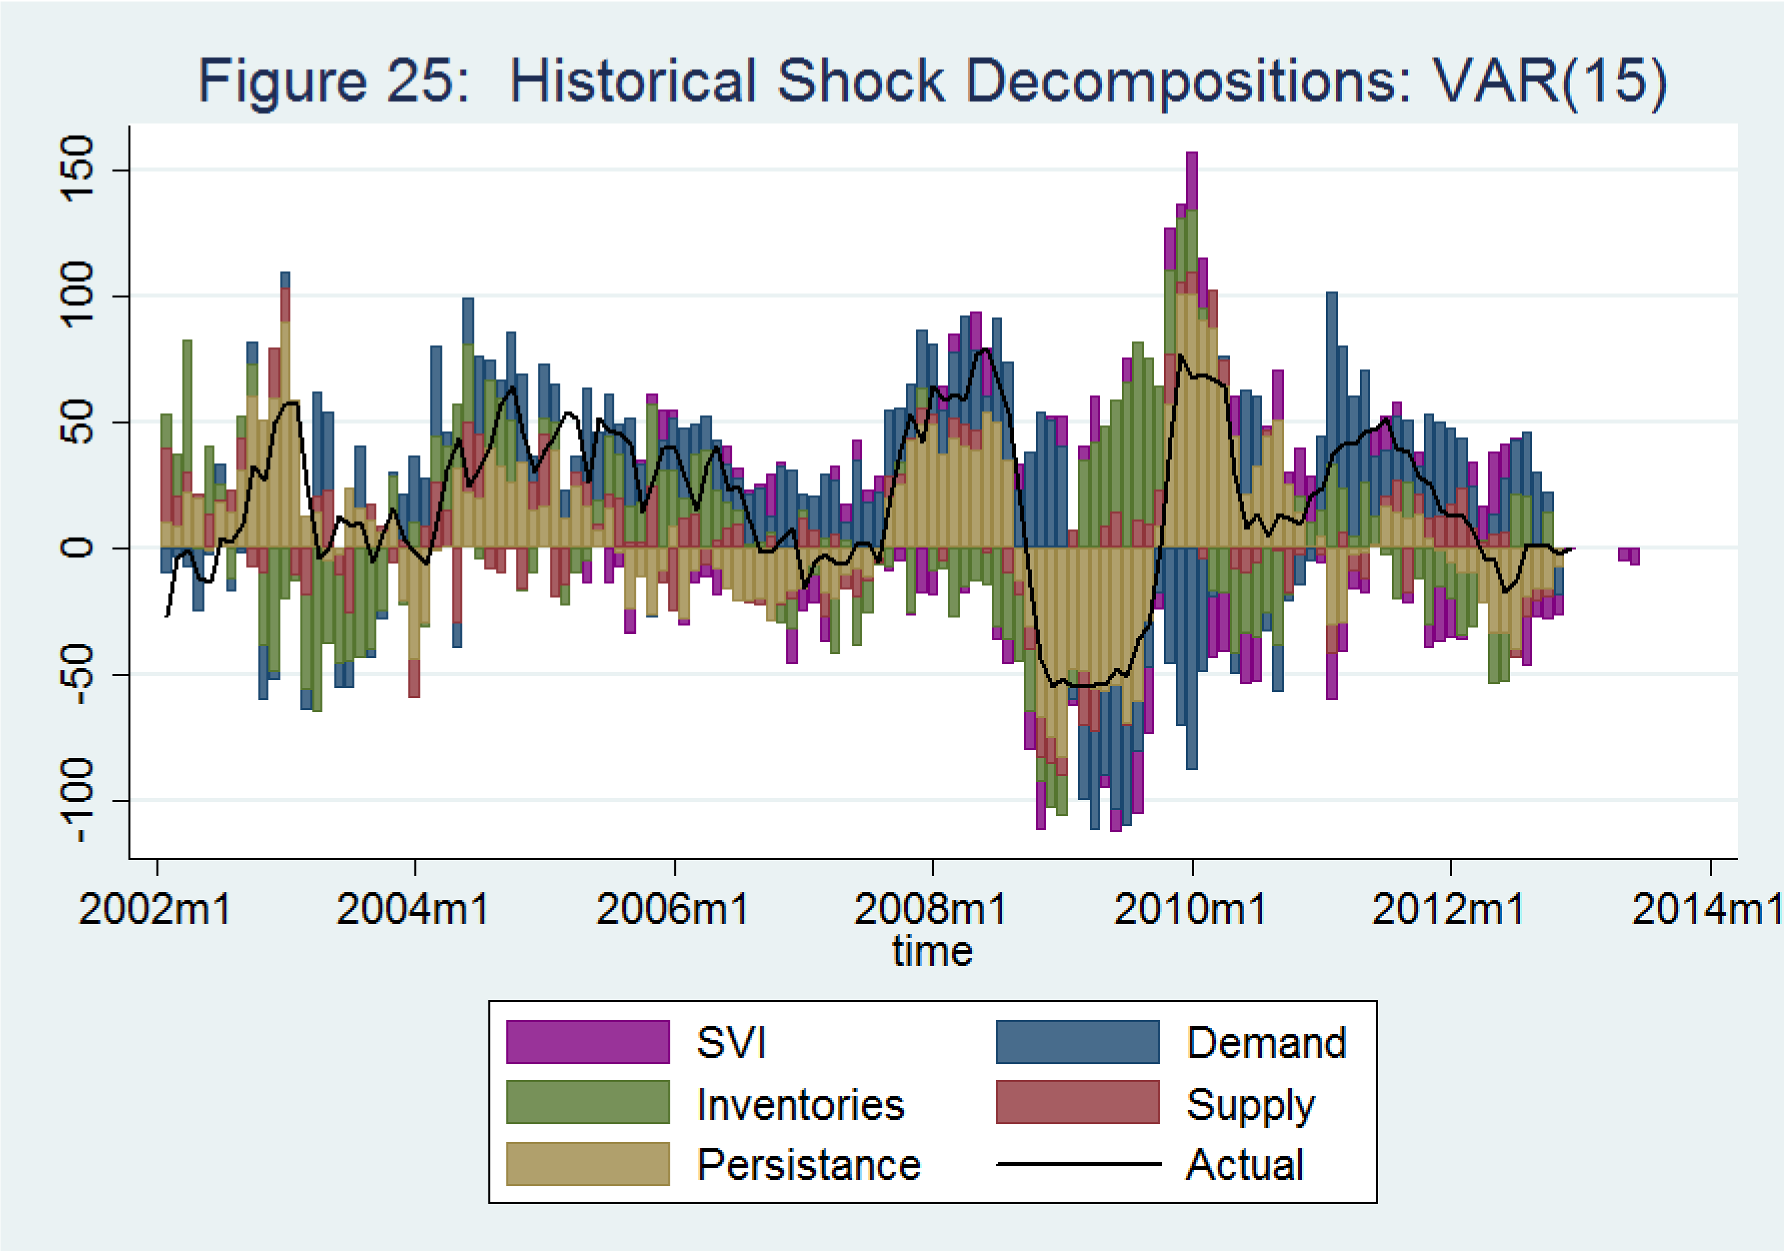
Figure 25: Historical Shock Decompositions: VAR(15)
150
100
50
0
-50
-100
2002m1
2004m1
2006m1
2008m1
2010m1
2012m1
2014m1
time
SVI
Demand
Inventories
Supply
Persistance
Actual

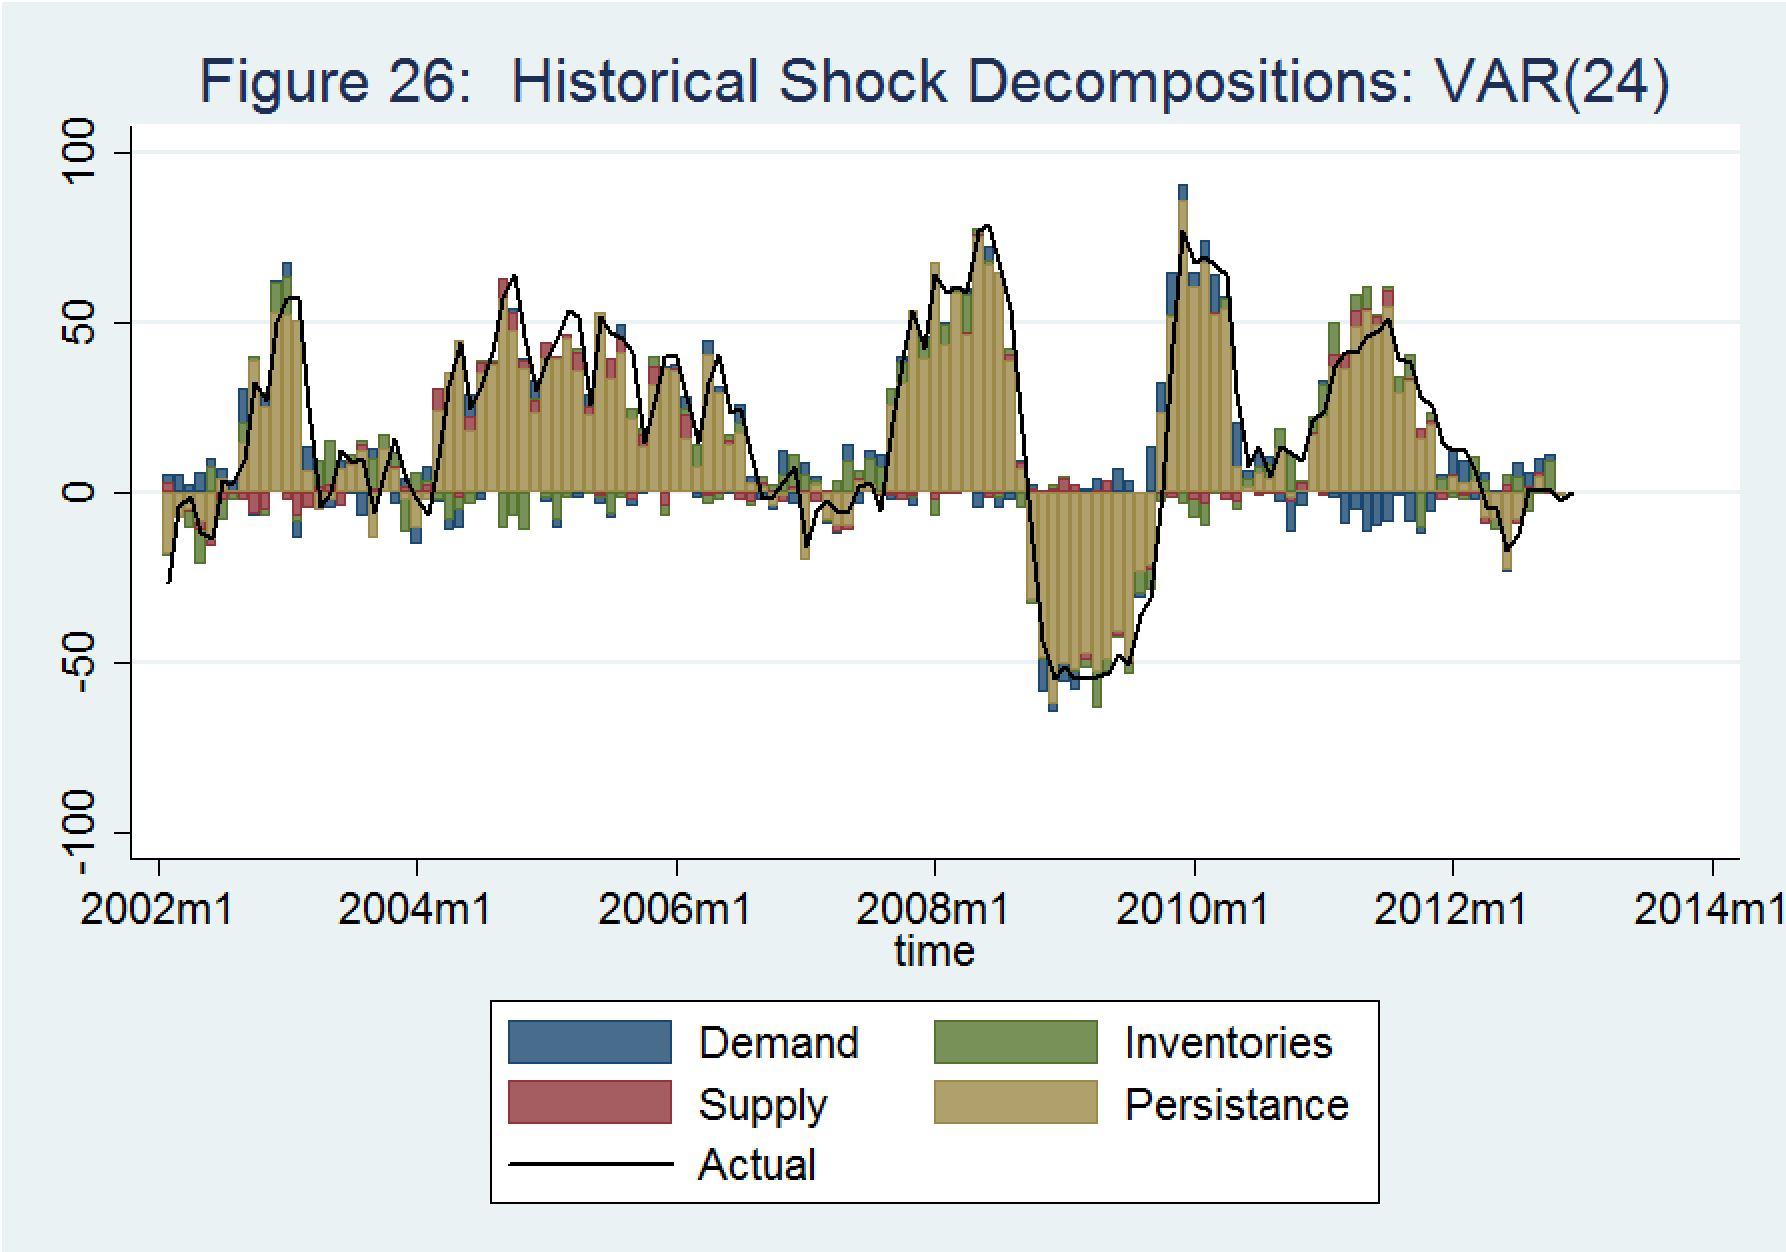
Figure 26: Historical Shock Decompositions: VAR(24)
100
50
0
-50
-100
2002m1
2004m1
2006m1
2008m1
2010m1
2012m1
2014m1
time
Demand
Supply
Actual
Inventories
Persistance