International Journal of Energy Economics and Policy

ISSN: 2146-4553

available at http: www.econjournals.com

International Journal of Energy Economics and Policy, 2016, 6(2), 167-182.

The Asymmetric Effects of Oil Price Shocks on the Canadian Economy

Luiggi Donayre[1]*, Neil A. Wilmot[2]

[1]Department of Economics, University of Minnesota - Duluth, 1318 Kirby Drive, Duluth, Minnesota, United States, [2]Department of Economics, University of Minnesota - Duluth, 1318 Kirby Drive, Duluth, Minnesota, United States. *Email: adonayre@d.umn.edu

ABSTRACT

A threshold vector autoregression (TVAR) is estimated to study the effects of oil price shocks on Canadian output and price level. While much of the literature has investigated potential asymmetric effects of positive and negative oil price shocks within a linear vector autoregression (VAR), we do so within a nonlinear VAR. Further, we extend the analysis to consider the correlation between asymmetries associated with the business cycle phase and size/sign asymmetries. Positive oil price shocks are found to have a stronger effect on output than negative oil price shocks. This asymmetry is significant in recessions, but lessened during expansions. The results also suggest that the reduction in inflation due to a negative oil price shock is larger than the increase in inflation following a positive oil price shock, especially during periods of low output growth. Yet, neither inflation nor output growth seems to vary disproportionately with the size of the oil price shock. In general, the results are robust to the ordering of the variables in the VAR process and to the time window over which the net oil price change is computed.

Keywords: Net Oil Price Increase, Asymmetry, Vector Threshold Autoregressions, Nonlinear Impulse Response Functions
JEL Classifications: C32, E32, Q43

1. INTRODUCTION

A number of studies have suggested that oil price shocks are one of the main sources of fluctuations in aggregate economic activity (Hamilton, 1996, 2003). However, linear, symmetric models of the transmission mechanism of oil price shocks have not been able to account for large business cycle fluctuations. As a consequence, several authors have made use of nonlinear transformations of oil prices and have found evidence of asymmetries in the response of output to oil price shocks (Jiménez-Rodríguez and Sánchez 2005; Hamilton, 2008, 2011; Herrera et al., 2011; Engemann et al., 2011; Köse and Baimaganbetov, 2015).

Much of this literature has been motivated by Hamilton (1983, 1996, 2008), who has found that ten out of the last eleven recessions were preceded by oil price increases. Yet, periods of high economic growth did not necessarily follow large declines in oil prices[1]. Following this evidence, the nonlinear relationship between oil price shocks and aggregate economic activity was empirically supported in subsequent work[2]. While this fact suggests that output may respond differently to positive and negative oil price shocks, all previous studies have omitted any potential correlation between this sign asymmetry and the business cycle phase. This potential correlation motivates our analysis.

We expand the scope of the literature in the following directions. First, we study the effects of oil price shocks during contractionary and expansionary times. Unlike previous studies, we do so within a threshold vector autoregression (TVAR). This potential asymmetry is motivated by the relationship between the effects of oil price shocks and the business cycle phase (Hamilton,

1 Hamilton (1983) originally noted that 7 out of the 8 postwar U.S. recessions prior to 1983 followed sharp increases in the price of oil. This pattern has continued until the 2007-2009 recession.

2 While recent work has called into question whether the response of output to the sign of oil price shocks is asymmetric (Kilian and Vigfusson, 2011a and b; Herrera et al., 2015), Hamilton (2011) argues that the work of these authors does not challenge the strong empirical evidence in favor of asymmetry. Indeed, An et al. (2014) find strong evidence that positive oil price shocks have larger effects than negative oil price shocks, even after taking into account the criticism of Kilian and Vigfusson (2011a and b).

2008, 2011)[3]. Second, we extend the analysis in previous studies by evaluating potential correlations between sign and business cycles asymmetries. For example, we consider the possibility that positive oil price innovations may generate a large response in output during recessions, but not during expansions. Third, we also examine potential correlations between size and business cycle asymmetries. Fourth, while much of the empirical literature focuses on U.S. data, the attention is shifted to Canada, motivated by its status as a net oil exporter country that is becoming more energy-independent and which exhibits a large energy sector share of gross domestic product (GDP)[4].

Understanding whether oil price shocks generate an asymmetric response of output is important for several reasons. Finding evidence that measures of economic activity respond differently to positive and negative oil price innovations could provide insight into why large recessions have been followed by smaller expansions. Furthermore, it can have important policy implications, especially in assessing the effects of large unexpected oil price declines (for example, in 1986, 1998, 2008 or the ongoing fall in the price of crude oil which started in mid-2014) and the associated response of monetary authorities. Furthermore, evaluating these potential asymmetric effects is crucial to correctly model oil prices and selecting among alternative theories of the transmission mechanism behind oil price shocks.

Our paper is similar in spirit to those in the vast literature that makes use of VAR analysis to study the asymmetric effects of oil price shocks on output. See, for example, Lee and Ni (2002); Jiménez-Rodríguez and Sánchez (2005); Gronwald (2012); An et al. (2014); Köse and Baimaganbetov (2015). However, with the exception of An et al. (2014), the aforementioned studies make use of a nonlinear transformation of oil prices to study sign asymmetries, but do so within a linear VAR. Unlike them, we formally estimate and test a TVAR, which allows for potentially asymmetric effects of shocks within a multivariate environment, modeled in a parsimonious way. The econometric specification allows the regime-switching (due to different business cycle phases, for example) to be endogenously estimated. Meanwhile, unlike An et al. (2014) and all previous studies, we explicitly allow for the possibility that positive and negative oil price shocks have different effects during periods of low growth or high growth. Further, we then capture the dynamic propagation of oil price innovations by means of nonlinear impulse response functions (IRF), unlike previous studies such as Jim´enez-Rodr´ıguez and Sánchez (2005) and Köse and Baimaganbetov (2015). This is important because nonlinear dynamic responses are contingent on the nature of the shock (i.e., its size and sign) and the history of the system (e.g., expansionary or recessionary initial conditions).

Using Canadian data, the results strongly support a nonlinear relationship between oil price and output. In particular, output responds asymmetrically to the direction of oil price innovations and this asymmetry is correlated with the business cycle phase. That is, positive oil price shocks have a larger effect on output than negative oil price shocks during recessions, consistent with most findings in the literature, although at odds with the findings in Köse and Baimaganbetov (2015). However, this asymmetry is lessened during expansionary times. Similarly, the results suggest that a reduction in inflation due to a negative oil price shock is larger than the increase in inflation due to a positive oil price shock, and the different inflation dynamics are exacerbated during periods of low output growth. Meanwhile, neither the response of output growth nor that of inflation varies disproportionately with the size of the oil price shock. Further, the results are robust to the ordering of the variables in the VAR process and to the time window over which the net oil price change is computed.

The remainder of the paper is organized as follows. The second section reviews the literature and motivates potential asymmetries in the oil price-output relationship with respect to the sign and the size of oil price shocks, as well as the business cycle phase. In the third section, the empirical procedure is explained. The fourth section describes the data and discusses the results. Section five conducts two robustness checks. Some concluding remarks are provided in the fifth section.

2. MOTIVATION

Oil price shocks have become, in many economic circles, the leading alternative to monetary policy as an explanation to macroeconomic fluctuations. As a consequence of the stylized facts pointed out by Hamilton (1983, 1996, 2008), discussed in the previous section, a subsequent body of empirical work has found evidence of asymmetries in the response of output to positive and negative oil price shocks.

To motivate this asymmetry, it is important to note that the effects of oil prices are different for oil importers and oil exporters. In the aftermath of a positive oil price innovation, net importers face a negative demand effect, whereas net exporters benefit from positive supply effect. In this environment, the asymmetric response of output to oil price shocks of different sign has been motivated by different theoretical models. Davies (1987a and b), Bresnahan and Ramey (1993) and Davies and Haltiwanger (2001), for example, suggest that oil price innovations cause labor and capital to reallocate from contracting to expanding sectors. In importing countries, costly sectoral reallocation thus amplifies the recessionary effects of an oil price increase and mitigates the expansionary effects of a negative oil price shock. Similarly, Hamilton (1988) proposes a model where the asymmetry arises because workers choose not to relocate to other sectors, given a positive probability that their sectors will improve after a positive oil price shock. Hence, in this model, the negative effect of an oil price increase is also amplified for oil importers.

3 The effects of oil price innovations might be different in recessions and expansions, much in the same way that monetary policy shocks have asymmetric effects in different business cycle phases. (García and Schaller, 2002; Peersman and Smets, 2002; Lo and Piger, 2005).

4 For example, energy products represent, roughly, 25% of total merchandise exports in Canada. It is also important to analyze other economies to better understand potential asymmetric effects of oil prices. For example, if asymmetries are only present for U.S. data, then there is no need to revise the way in which we model asymmetries. If not, we should re-evaluate the econometric methodology used and the theoretical models that motivate the nonlinear relationship.

Figure 1: Canada: Oil production, output growth and oil shocks, (a) Oil production (% of GDP): 1980-2013, (b) Output growth and oil price shocks: 1989-2013

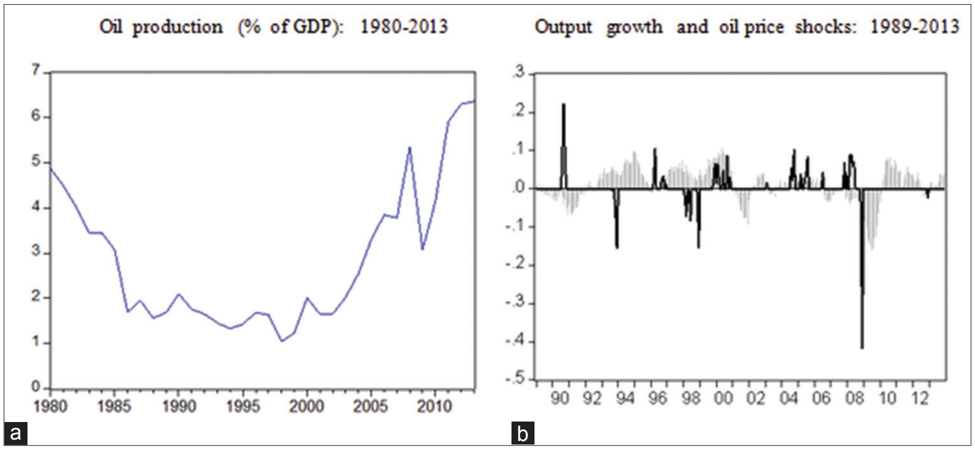

Oil production as a share of gross domestic product (GDP) is calculated as the production of oil, times the import price of oil, as a fraction of nominal GDP. Output growth (gray bars) is measured as the growth rate of industrial production. The oil price shock (solid line) is defined as Hamilton's net oil price change (increases and decreases).

Edelstein and Kilian (2007, 2009) propose a model that explains this type of asymmetry by means of precautionary savings motives. For an oil importing country, a positive oil shock may cause concern about future income and employment, leading to an increase in precautionary savings. To the extent than declines in oil prices are not associated with higher uncertainty, output may respond asymmetrically to positive and negative oil price innovations.

Asymmetry may also stem from the irreversibility of investment. Bernanke (1983) and Pindyck (1991) develop a model where oil price increases generate uncertainty in individuals and firms, forcing them to postpone their purchases of durable goods and capital goods, respectively. Therefore, for an oil importing economy, increased uncertainty amplifies the recessionary effects of an oil price increase, in the same way as explained above.

Much of the empirical literature has found support for this type of asymmetry. Jiménez-Rodríguez and Sánchez (2005), for example, use different linear and nonlinear specifications to empirically assess the effects of oil price shocks on the real economic activity of G-7 countries, Norway and the Euro zone. They find strong evidence that oil price increases generate a larger effect on real GDP than oil price decreases, except in Japan. Engemann et al. (2011), similarly, estimate business cycles turning points, together with oil effects, in a time-varying Markov-switching model and find that oil price shocks increase the probability of entering a recession. At a more disaggregated level, Herrera et al. (2011) test different asymmetric specifications between the real price of oil to U.S. industrial production and its sectoral components. They find that there is strong evidence of an asymmetric effect at the disaggregated level, especially for energy-intensive sectors, although the evidence is weaker at the aggregate level.

From the perspective of an oil exporter, the effect of a positive oil price shock is less clear. In this case, the negative demand effects explained in the preceding paragraphs of this section are offset by the positive supply effects of a higher-priced commodity. Meanwhile, because extraction and exploration necessary for oil production involve long-term investments, these supply-side positive effects may only appear in the long-run.

It is also important to note that crude oil production represents a large fraction of GDP for Canada, as depicted in the left panel of Figure 1. The majority of Canada's increase in crude oil production derives from Alberta's oil sands region. This non-conventional crude oil production requires a significantly high level of up-front investment, generally stretching over several years. Such expensive capital projects have the potential to drive asymmetric responses between oil prices and output. For example, a negative oil price shock during the construction phase may not delay spending/construction, particularly if the shock is expected to be transitory[5]. To the extent that crude oil production affects a large fraction of GDP in Canada, potential asymmetric effects of oil price innovations could have important implications for living standards[6]. These facts drive our focus on a net oil exporter, like Canada.

To motivate the different potential asymmetries considered in this paper, the right panel of Figure 1 exhibits the growth rate of output in Canada (in gray bars) and a Hamilton-type oil price shock (in the dark line) for the 1989-2013 period. The latter is defined as the net amount by which oil prices at time t exceed the maximum

5 Additionally, the vast increase in non-conventional crude oil production has brought about an increased need for transportation facilities. There are currently several proposed pipeline projects being examined in Canada, which would transport the oil from land-locked Alberta to the coastal regions of North America. One such project, which would transport crude oil to the east coast of Canada (New Brunswick), is the Energy East pipeline, expected to require an investment of $12 Billion dollars (Canadian).

6 Canada has, in recent years, received increasing attention as an important producer of oil. As reported in the left panel of Figure 1, oil production in Canada has increased by, roughly, 50% in the last 15 years. Further, it is forecast that oil production will reach 4.3 million bbl/d by 2035 (International Energy Agency, 2012).

value over the previous 3 years[7]. As can be observed, negative growth in Canadian output is always preceded by a positive oil price shock. However, output growth seems to be unaffected by the oil price drops of 1993, 1998 and 2012, consistent with the evidence in favor of the sign asymmetry documented in the literature. The right panel of Figure 1 also exhibits positive oil price shocks during times of economic expansion that seem to have little effect on output growth (for example, in 2006). Meanwhile, some negative oil price shocks occurred during recessionary times, like that of 2009, with large, negative effects on output growth. This suggests that the effects of oil price shocks might be different during recessions and expansions. Finally, Hamilton (1996, 2003) argues that economic agents do not change their behavior in response to small fluctuations in oil price changes. Therefore, output may respond asymmetrically to the size of the oil price shock. Graphically, the right panel of Figure 1 suggests that the effects of oil price shocks might, indeed, depend on the size of the shock. This may explain, for example, why the oil price drops of 1993, 1998 and 2012, all smaller in magnitude than the oil price drop of 2009, appear not to have affected output growth.

3. EMPIRICAL APPROACH

In the following subsections, we define oil price shocks, describe the empirical models, the testing and estimation procedures and the computation of nonlinear IRF.

3.1. Oil Price Shocks and Empirical Models

We follow Hamilton (1996) and consider net oil price increases (NOPI) as a measure of oil price shock. Hamilton (1996, 2003) argues that economic agents do not change their behavior in response to small fluctuations in oil price changes. Consequently, an oil price shock is defined as the amount by which the change in the oil price in month t exceeds the maximum value over the previous 3 years[8]. Formally, the oil price shock, s_t is defined as

$$s_t = max\left\{0, ln\frac{p_t}{max\{p_{t-1},\ldots,p_{t-36}\}}\right\} \quad (1)$$

where p_t is a measure of oil price. Note that the oil price shock, as defined in equation (1), implicitly assumes that only increases in oil prices that go beyond simple corrections to earlier declines are relevant to economic agents.

To account for the interdependence of macroeconomic aggregates, we consider a standard VAR as the linear, benchmark model:

$$X_t = \Psi_0 + \Psi_1(L)X_{t-1} + \varepsilon_t \quad (2)$$

where X_t is a k × 1 vector of macroeconomic aggregates; the autoregressive polynomial matrix Ψ_1 (L) is assumed to have roots that satisfy stationarity conditions; Ψ_0 is a k × 1 vector of intercept coefficients; and ε_t is a Gaussian shock with mean zero and covariance matrix Σ. All estimations below allow for possibly heteroskedastic errors.

A linear VAR provides a suitable starting point to analyze multivariate relationships between economic variables. Because they are parsimonious and the IRF are easy to construct and interpret, they have become a standard tool to evaluate the effects of shocks, like oil price shocks, on output and other macroeconomic aggregates. See, for example, Lee and Ni (2002), Jiménez-Rodríguez and Sánchez and Sanchez (2005), Gronwald (2012), Baumeister and Peersman (2013), among others.

In environments where responses to shocks are, potentially, asymmetric or nonlinear, the IRF associated with a linear VAR fail to capture such asymmetry. To cope with these issues, we estimate a threshold VAR (TVAR), a natural extension of the linear VAR framework, which can capture potentially asymmetric effects of oil price shocks by splitting the sample endogenously into different regimes. Weise (1999) shows that a broad class of structural models can be represented in a reduced form with changing coefficients and, hence, we adopt this more general specification and allow coefficients in the lag polynomials to change across regimes. In particular, the specification of the TVAR is given by

$$X_t = \left(\Psi_0^1 + \Psi_1^1(L)X_{t-1}\right) + \left(\Psi_0^2 + \Psi_1^2(L)X_{t-1}\right)I\left[q_{t-d} \le \gamma\right] + \varepsilon_t \quad (3)$$

where the autoregressive polynomial matrices under each regime, $\Psi_1^1(L)$ and $\Psi_1^2(L)$, are assumed to have roots that satisfy stationarity conditions; Ψ_0^1 and Ψ_0^2 are k × 1 vectors of intercept coefficients under regimes 1 and 2, respectively; q_{t-d} is the threshold variable that determines the prevailing regime of the system, lagged d periods; γ is the threshold parameter around which the dynamics of the system changes; and I[.] is an indicator function that equals 1 when $q_{t-d} \le \gamma$ and zero otherwise[9].

3.2. Estimation and Testing Procedures

To estimate the model, we follow Hansen (1996, 1997) and Teräsvirta and Yang (2014, 2016). In general, the models are estimated by least squares (LS). The LS estimators $(\hat{\Psi}, \hat{\Sigma}, \hat{\gamma})$ jointly minimize the sum of squared residuals, S_n, through concentration. Conditional on γ, S_n is linear in Ψ and Σ. Therefore, the model can be estimated sequentially for each possible value of γ, yielding a γ-dependent concentrated sum of squared errors, $S_n(\gamma)$. The LS

7 This variable is described in more detail in section 3.

8 Hamilton (1996) originally uses changes in oil prices that are larger than the maximum value over the previous four quarters. Hamilton (2008) uses a 3-year window, which provides a better fit. Because oil price increases in net oil exporters could have long-term effects associated with exploration and production of oil, we follow Hamilton (2008) and consider a 3-year window. As a robustness check, we also considered changes in oil that are larger than the maximum value over the previous 2 years. The results did not change much and are shown in section 5.

9 In a threshold model, the threshold variable q_{t-d} is observable, in contrast to other switching models such as Markov-switching models, where the transition variable is probabilistic. Therefore, threshold models allow us to obtain values for q_{t-d} that determine the change in dynamics in the relationship between oil prices and output. This is important as it provides policymakers with observable and measurable ways to address possible policy changes.

estimate of γ is thus the value of this parameter that minimizes $S_n(\gamma)$. Formally,

$$\hat{\gamma} = \text{argmin}_{\gamma \in \Gamma_n} \{S_n(\gamma)\} \quad (4)$$

where $\Gamma_n = \Gamma \cap \{q_1, q_2, \ldots, q_n\}$ and the bounded set $\Gamma = [\underline{\gamma}, \overline{\gamma}]$ is defined a priori to contain the middle 70% of all possible threshold values to ensure that the model is well identified[10]. The LS estimates of Ψ and Σ are then given by the parameters associated with $S_n(\hat{\gamma})$.

To test whether the TVAR model in equation (3) is statistically significant relative to the linear benchmark VAR in equation (2), we modify the procedure developed by Hansen (1996). Considering these modifications, the null hypothesis is given by $H_0 : \Psi_0^2 = \Psi_1^2(L) = 0$. Because this problem is tainted by the existence of nuisance parameters (specifically, the threshold γ is not identified under the null hypothesis), a test with near-optimal power against a wide range of alternative hypotheses is given by the following F_n sup-statistic:

$$F_n = \sup_{\gamma \in \Gamma_n} \{F_n(\gamma)\} \quad (5)$$

where $F_n(\gamma)$ is a heteroskedasticity-consistent Wald statistic against the alternative when γ is known. To the extent that γ is not identified under the null hypothesis, the distribution of (5) is nonstandard. However, its asymptotic distribution can be approximated by a bootstrap procedure (Hansen, 1996).

In terms of the implementation of the test, we generated 1,000 realizations of the Wald statistics for each grid point under the null hypothesis. Then, we construct the empirical distribution for F_n over the grid space. The approximation to the asymptotic p-value of the test is obtained by counting the number of instances in which the bootstrap sample for F_n exceeds the observed statistic.

3.3. Nonlinear IRF

Given the dynamic nature of the system, which evolves between different regimes over time, it is important to evaluate the responses of variables in the system to given shocks over time. These dynamic responses can be captured by an IRF over h periods, which is defined as the change in the conditional expectation of Y_{t+h} resulting from an exogenous shock ε_t:

$$E[Y_{t+h}|\Omega_{t-1}, \varepsilon_t] - E[Y_{t+h}|\Omega_{t-1}] \quad (6)$$

where Ω_{t-1} is the information set at time t–1. Linear IRFs, however, are not able to capture potentially asymmetric responses. In particular, the responses of the variables to shocks in a linear VAR do not depend on the direction of the shocks, their size or the history of the system. For example, the response of output to a positive oil price shock is the same as the response to a negative oil price shock, multiplied by negative one, by construction. Similarly, the response of output to an oil price shock is, necessarily, the same in expansions and recessions.

To overcome these issues, we construct nonlinear IRF in the spirit of Koop et al. (1996) and Teräsvirta and Yang (2014, 2016). In a nonlinear environment, the conditional expectations in (6) must be simulated since the impact of the shocks depend on the sign and size of the shock, as well as the history of the system. The model is assumed to be known, so sample variability is not taken into account. To compute the nonlinear IRF, the following procedure is implemented (a detailed description of the simulation method and the way to compute nonlinear IRF is provided in Appendix A): First, shocks for periods 1 to h are drawn, with replacement, from the residuals of the estimated TVAR given in equation (3) and, for a given history (or initial values) of the system, fed through the estimated model to produce a simulated data series[11]. This produces a forecast of the variables conditional on a particular history and sequence of shocks for h periods ahead. Second, the same procedure is carried out, given the same particular history and sequence of shocks, with the exception that the shock in period 0 is fixed at a particular value. The shocks are then fed through equation (3) and a forecast is produced as explained above. The difference between this forecast and the baseline model is the IRF for a given sequence of shocks and a given history. Third, these steps are computed in this way for 1,000 draws of the residuals and averaged out to produce an IRF conditional on initial values only. Fourth, given the arbitrary shock and particular history, the difference between the averaged forecasts is averaged out over initial values taken from the sub-sample of the data.

4. RESULTS

4.1. Data

The sample period considered spans between January 1986 and December 2013, implying a total of 336 observations. The starting date of 1986 is determined by the availability of the oil price variable. All data, other than the oil price, are taken from the Canadian Socioeconomic Database (CANSIM) and are seasonally adjusted at the source. For the oil price, we consider the price of the benchmark crude oil for the Canadian market, the Edmonton Par spot price, taken from Natural Resources Canada (NRCAN). We follow Lee and Ni (2002) and Gronwald (2012) and set up the linear VAR that serves as a benchmark to include the following variables: The 12-month natural logarithm difference in industrial production (IP), the 12-month natural logarithm difference in the consumer price index (CPI), the 12-month natural logarithm difference in M2, the 12-month difference in the 3-month Treasury bill rate, and the oil price shock, s_t, as defined in equation (1).

10 It is standard practice for threshold models to exclude the 15% of each end of the vector of ordered threshold values to avoid distortions in inference. If possible thresholds that are too close to the beginning or the end of the ordered data were considered, there would not be enough observations to strongly identify the sub-sample parameters.

11 The nonlinear IRFs are averaged over different histories taken from sub-samples of the data. For instance, the nonlinear IRF for regime 1 are computed averaging out over histories corresponding to all dates in which the $q_{t-d} \leq \hat{\gamma}$

Output growth, CPI inflation and M2 growth are all detrended to avoid detecting non-linearities arising, possibly, due to a time trend[12]. To estimate the VAR, we assume the following ordering for the Cholesky identification strategy: Change in the 3-month Treasury bill rate, oil price shock, M2 growth, IP growth, CPI inflation. This ordering implies that short-term rates affect all other variables in the system. Oil prices are relatively exogenous, but have an effect on IP growth and the rate of inflation, given that oil is an important input in the production process[13].

To find the suitable lag length for the TVAR, standard specification tests are used. Specifically, the Akaike information criterion and Schwartz information criterion both suggest an optimal lag length of 2. The same order specification is used in the estimation and testing of the nonlinear model for comparison. Finally, the results from augmented Dickey-Fuller tests suggest that all variables are stationary.

4.2. Discussion of Results

On the basis of the stylized facts documented by Hamilton (1983, 2008, 2011), the switching variable q_{t-d} is defined as the growth rate of industrial production lagged d periods[14]. Given a lag length of 2 for the VAR, we estimate models with d = 1 and d = 2. The optimal value for d is ultimately chosen based on the minimum p-value of the sup-F test of linearity.

Table 1 reports the bootstrapped p-values of the sup-F test for the threshold VAR against the benchmark linear model, the value of the sup-F statistic and the estimated coefficient for the threshold parameter. While there is evidence of non-linearities in the models with either d = 1 or d = 2, the lowest p-value occurs when d = 2. Hence, all results reported below are based on the estimates of the TVAR model when the switching variable is q_{t-2}.

Figure 2 plots the threshold variable and the threshold parameter, estimated as in equation (4). The estimated threshold value, $\hat{\gamma} = -1.396$, splits the sample in two regimes, which correspond closely to the recession dates established by the Organization for Economic Co-operation and Development (OECD), shown in shaded bars. In particular, when $q_{t-2} \leq -1.396$, the Canadian economy is in the low-growth regime. Conversely, when $q_{t-2} > -1.396$, the economy is in the high-growth regime. According to the estimated threshold, the Canadian economy remains in the high-growth regime the majority of the sample period, switching to the low-growth regime only during periods of very slow growth.

Even though the TVAR fits the data better than the benchmark linear VAR, it is important to evaluate the responses of output to oil price shocks over time, given the nonlinear nature of the model. In this kind of setting, there are two potential sources of variability in the impulse responses. First, parameter and estimation uncertainty are present, as in any econometric model. Second, IRF exhibit variability due to history dependence for the evolving system. Even if the VAR parameters and the threshold parameter are fixed, the impulse responses will depend on the initial state. If the economy starts below the threshold, but close to it, the response may be different than the case when the economy starts far from the threshold, because the economy may switch to the high state sooner. For this reason, we construct nonlinear IRFs following Koop et al. (1996) and Teräsvirta and Yang (2014, 2016), as detailed in the previous section. It is important to note that, in all cases below, the IRFs are computed allowing the economy to endogenously switch between regimes after the initial oil price shock.

Table 1: Linearity test and delay lag selection

Switching variable	sup-F statistic	Bootstrapped p-value	Estimated threshold value
q_{t-1}	75.071	0.056	–2.696
q_{t-2}	80.096	0.016	–1.396

This table reports the values of the sup-F statistics defined in equation (5), as well as the bootstrapped p-values of the tests when the switching variable is the growth rate in IP, delayed d periods. The threshold variable was set to contain the 70% middle part of the observations to avoid over-fitting. The effective sample period ranges from January 1989 through December 2013.

Figure 2: Threshold variable and estimated threshold parameter

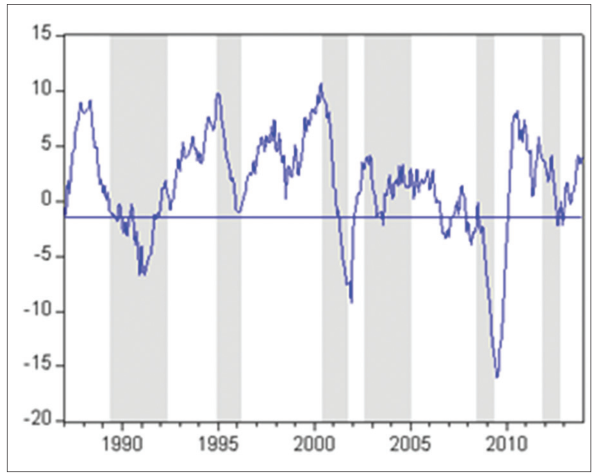

Threshold variable q_{t-2} and estimated threshold parameter, $\hat{\gamma} = -1.396$. OECD-dated recession dates are shown in the gray, shaded bars. The sample shown spans January 1987 through December 2013.

Figure 3 through 5 exhibit the accumulated median forecasts for industrial production growth and inflation to different oil price shocks for a 12-month horizon. Tables 2 and 3 summarize some key features of the accumulated responses of output growth and inflation over 12 months.

The left panel of Figure 3 plots the estimated cumulative responses of industrial production growth to a positive, one standard deviation shock in oil prices considering different initial conditions. When the economy begins in the low-growth state, a positive oil price shock permanently reduces output growth.

12 The results of the test, however, are not sensitive to data detrending. Linearity is rejected in both cases.

13 We also considered different orderings of the variables in the VAR as a robustness check. The overall results were robust to different orderings, as shown in section 5.

14 The choice of the threshold variable is motivated by our interest in studying the effects of positive and negative oil price shocks in different states of the economy. The results for other measures of economic conditions, such as the joblessness rate, are similar.

Table 2: Cumulative change in output growth

Size of shock	Model and initial state	Positive shock			Negative shock		
		6 months	12 months	Largest change (month)	6 months	12 months	Largest change (month)
1. S.D. oil price shock	TVAR (low–growth regime)	–0.585	–1.504	–1.504 (month 12)	–0.277	–0.874	–0.874 (month 12)
	TVAR (high–growth regime)	–0.309	–0.793	–0.793 (month 12)	–0.122	–0.396	–0.396 (month 12)
2. S.D. oil price shock	TVAR (low–growth regime)	–0.508	–1.347	–1.347 (month 12)	–0.354	–1.032	–1.032 (month 12)
	TVAR (high–growth regime)	–0.262	–0.694	–0.694 (month 12)	–0.169	–0.495	–0.495 (month 12)

This table reports selected cumulative responses of output growth to different oil-price shocks, derived from the nonlinear impulse response functions described in Appendix A. Negative shocks are multiplied by −1 and two-standard deviation shocks are divided by 2 to allow for a direct comparison to the responses of output to positive and one-standard deviation oil price shocks, TVAR: Threshold vector autoregression.

Table 3: Cumulative change in inflation

Size of shock	Model and initial state	Positive shock			Negative shock		
		6 months	12 months	Largest change (month)	6 months	12 months	Largest change (month)
1. S.D. oil price shock	TVAR (low–growth regime)	0.304	0.521	0.521 (month 12)	0.477	0.745	–0.874 (month 12)
	TVAR (high–growth regime)	0.175	0.292	0.292 (month 12)	0.217	0.341	–0.396 (month 12)
2. S.D. oil price shock	TVAR (low–growth regime)	0.347	0.577	–1.347 (month 12)	6	12	–1.032 (month 12)
	TVAR (high–growth regime)	0.185	0.304	–0.694 (month 12)	0.433	0.689	–0.495 (month 12)

This table reports selected cumulative responses of CPI inflation to different oil-price shocks, derived from the nonlinear impulse response functions described in Appendix A. Negative shocks are multiplied by −1 and two-standard deviation shocks are divided by 2 to allow for a direct comparison to the responses of CPI inflation to positive and one-standard deviation oil price shocks, TVAR: Threshold vector autoregression.

Figure 3: Nonlinear impulse response functions (IRF): Effects of an oil price increase by initial state, (a) Cumulative response of output growth, (b) Cumulative response of inflation

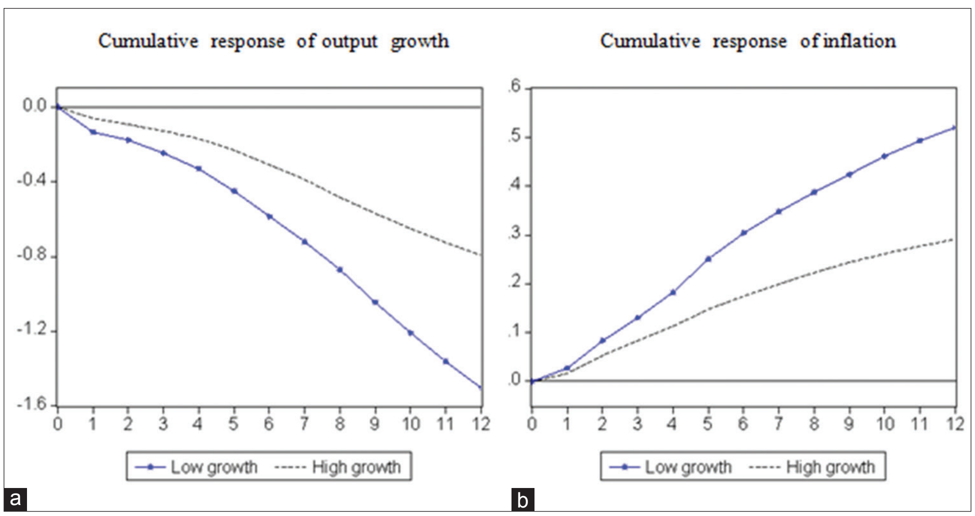

Accumulated responses of output growth and inflation to a positive oil price shock, as defined in section 3. The nonlinear IRF are computed according to the details provided in Appendix A.

As seen in Table 2, the estimated cumulative response in output growth is -1.504 percentage points after 1 year. When the economy begins in the high-growth state, a positive oil price shock reduces output permanently as well. However, the fall in output growth is smaller. As seen in Table 2, the estimated cumulative response in output growth, in this case, peaks at -0.793 percentage points after 12 months.

When evaluating the dynamic behavior of the system, the results show that positive oil price shocks generate an asymmetric response in output growth: The reduction in output growth is larger in recessions than it is in expansions. To the extent that the Canadian economy is a net exporter of oil, this result suggests that, at least in periods of low growth, the negative demand effects associated with a positive oil price shock more than offset the positive supply effects, as discussed in section 2. Intuitively, if the reduction in real income of final consumers as a consequence of a higher oil price prevails over the increase in profitability in the oil sector, then relevant implications for policy could be elicited. Namely, the policy function behind the decision-making process of monetary authorities should respond to oil price increases more strongly than oil price reductions. While this is consistent with the behavior of most central banks (Yellen, 2012), our results suggest that this behavior should be emphasized in periods of low economic growth.

The right panel of Figure 3 reports the estimated cumulative responses of inflation to a positive, one standard deviation shock in oil prices by initial state. Inflation increases in both cases, but the increase is larger in the low-growth state. Specifically, the estimated cumulative responses of inflation in the low-growth and high-growth regimes peak at 0.521 and 0.292 percentage points after 12 months, respectively (Table 3). When the economy starts in a low-growth regime, prices tend to be further below long-run inflation targets due to insufficient aggregate spending, thus leaving more room for them to increase. While resources are more fully employed during periods of high growth (thus potentially bidding up prices more rapidly), the dynamics of inflation shown in the right panel of Figure 3 are consistent with the notion that monetary authorities seem to have a bias towards missing long-run inflation target from above. In this sense, our results are in line with de facto practices of most central banks (Yellen, 2012).

To evaluate asymmetries associated with the direction of the oil price shock, Figure 4 depicts the cumulative responses of output growth and inflation to a positive one standard deviation and to a negative one standard deviation oil price shocks, by initial state. To allow for a direct comparison, the responses of output and inflation to a negative oil price shock have been multiplied by –1.

In the upper left panel of Figure 4, when the economy begins in the low-growth state, the response of output growth to a positive oil price innovation is larger than its response to a negative oil price innovation in every month after the initial shock. As shown in Table 2, the estimated 12-month cumulative responses of output growth to positive and negative oil price shocks are –1.504 and –0.874, respectively. Therefore, in recessionary times, output responds asymmetrically to positive and negative oil shocks, consistent with the findings in the empirical literature.

During expansionary times, the response of output growth to positive and negative oil price innovations is different, as shown in the upper right panel of Figure 4. However, those differences are smaller relative to those prevailing in the low-growth regime. For example, the drop in output growth in response to a positive and to a negative oil price innovation peaks at –0.793 and –0.396 after a year, respectively (Table 2). While formally testing for possible reasons that explain why this sign asymmetry is lessened during periods of high economic growth goes beyond the scope of our TVAR model, certain characteristics of the behavior of economic agents that change across business cycle phases may be at the heart of these differences. The reallocation of resources hypotheses to explain why output may respond differently to positive and negative oil price shocks (Hamilton, 1988; Bresnahan and Ramey, 1993; Davies and Haltiwanger, 2001), for example, is more likely

Figure 4: Cumulative effects of positive and negative oil price shocks (a) Output growth, low-growth state, (b) Output growth, high-growth state, (c) Inflation, low-growth state, (d) Inflation, high-growth state

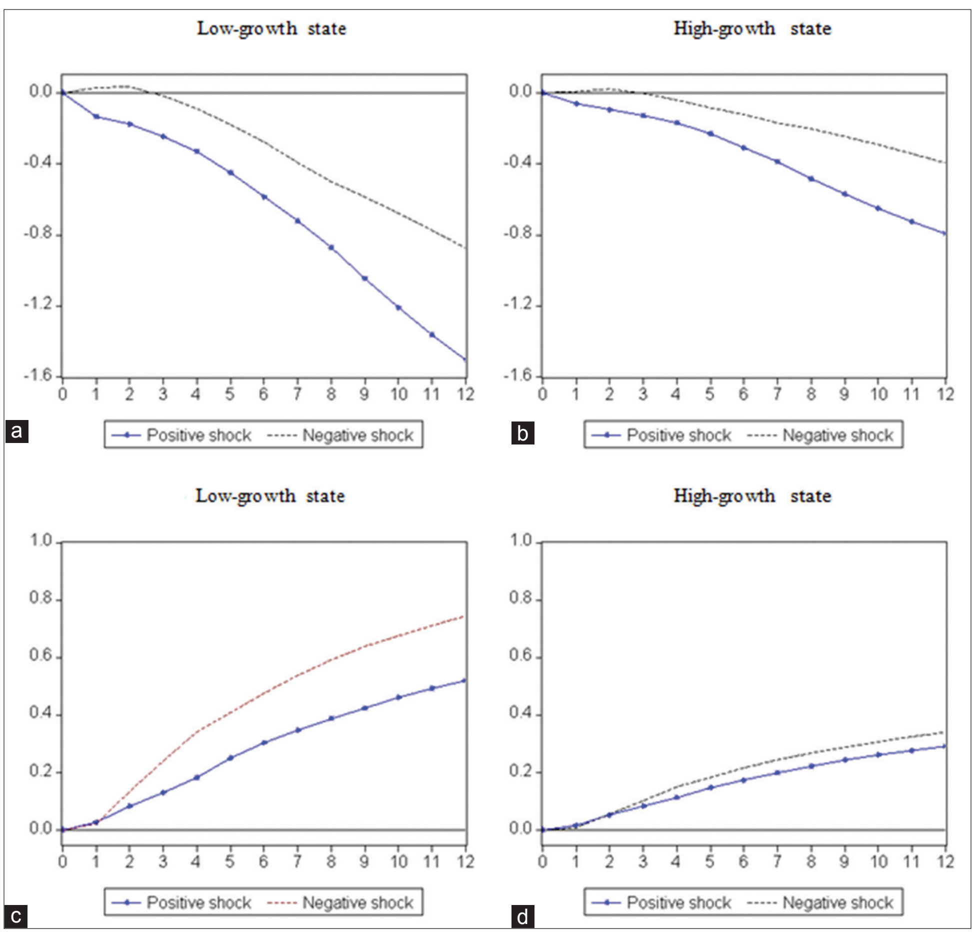

Accumulated response of output growth and inflation to positive and negative one standard deviation oil price shocks. The nonlinear IRFs are computed according to the details provided in Appendix A.

to occur in periods of underutilization of resources. Similarly, while higher oil prices may lead to putting off consumption of durable goods and investment in capital goods as in Pindyck (1991), the reasons for delaying spending in these goods are minimized if the economy is growing steadily. Meanwhile, the heightened uncertainty associated with positive oil price shocks that leads to precautionary savings (Edelstein and Kilian, 2007, 2009) or amplication of financial channels (Baskaya et al., 2013; Plante and Traum, 2014) may be lessened (or its costs downplayed) in periods of high economic growth. Finally, to the extent that inflation increases more rapidly when resources are fully utilized, monetary authorities that respond more strongly to deviations of inflation above long-run inflation targets (Herrera and Pesavento, 2009; Kilian and Vigfusson, 2011a; Yellen, 2012) could attenuate the negative effects of higher oil prices in periods of high economic growth.

The lower left and the lower right panels of Figure 4 exhibit the cumulative responses of inflation to positive and negative one standard deviation oil price shocks, by initial state. Consistent with the results shown in Figure 3, the responses of inflation to oil price shocks of different sign are larger in the low-growth state, relative to the responses in the high-growth state. In the case of the economy starting in the low-growth regime, the cumulative responses of inflation to a positive and to a negative oil price shock are 0.521 and 0.745 percentage points after a year, respectively. Meanwhile, similar responses when the economy starts in the high-growth regime peak at 0.292 and 0.341 after 12 months (Table 3). This evidence suggests that the fall in inflation due to a negative oil price shock is larger in magnitude than the increase in inflation due to a positive oil price shock, and this difference is exacerbated in the low-growth regime. Intuitively, a potential bias in the response of monetary authorities to inflation above or below the long-run target may explain these differences (Herrera and Pesavento, 2009; Kilian and Vigfusson, 2011a; Yellen, 2012). Inflation is more likely to be below its long-run target during recessionary times. Therefore, a negative oil price shock will be followed by a highly accommodative response from monetary authorities. However, a positive oil price shock will bring inflation back to target and, consequently, monetary authorities will not need to respond.

Motivated by the fact that some small oil price drops (for example, in 1993, 1998 and 2012) appear to have had no effect on output growth, relative to the significant drop in output in the aftermath of the large oil price drop of 2009, we analyze potential asymmetries with respect to the size of oil price shocks, as argued in Hamilton (1996, 2003). Figure 5 reports the cumulative responses of output growth and inflation to one standard deviation and two standard deviation oil price shocks, by initial state. To allow for a direct comparison, the responses of output growth and inflation to a two standard deviation shock are rescaled (divided by 2).

Figure 5: Cumulative effects of small and large oil price shocks (a) Output growth, low-growth state, (b) Output growth, high-growth state, (c) Inflation, low-growth state, (d) Inflation, high-growth state

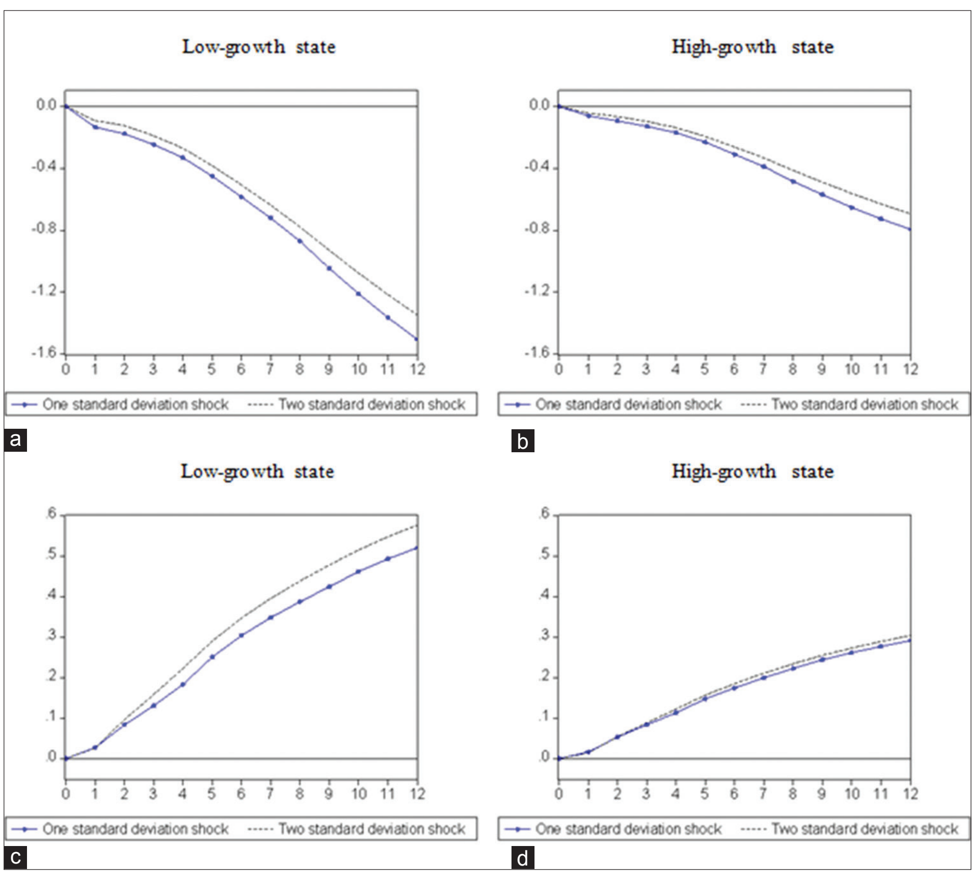

Accumulated response of output growth and inflation to small (one standard deviation) and large (two standard deviations, rescaled by 0.5) oil price shocks. The nonlinear impulse response functions are computed according to the details provided in Appendix A.

Figure 6: Effects of an oil price increase by initial state (alternative vector autoregression ordering), (a) Cumulative response of output growth, (b) Cumulative response of inflation

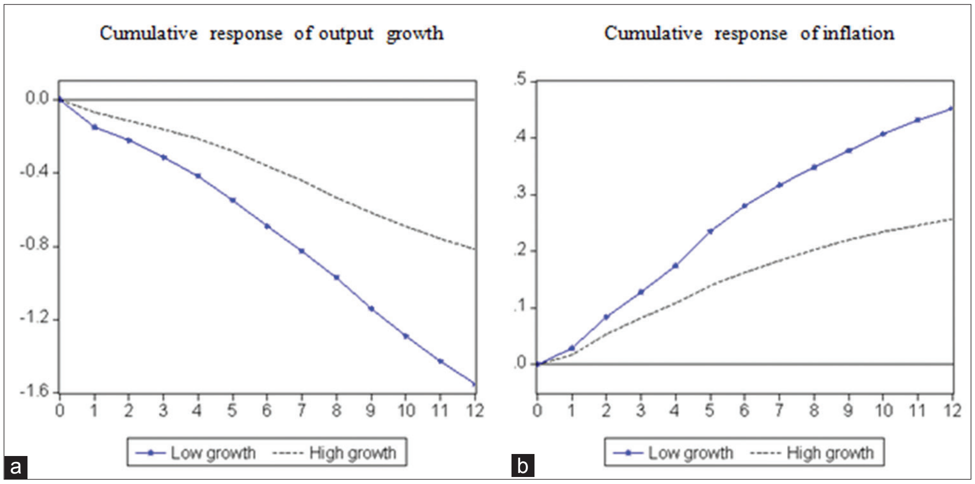

Accumulated response of output growth and inflation to a positive oil price shock, as defined in section 3. The nonlinear impulse response functions are computed according to the details provided in Appendix A.

The upper left and upper right panels of Figure 5 plot the responses of industrial production growth to a one standard deviation and to a two standard deviation oil price shocks, in the low-growth and high-growth regimes, respectively. Similarly, the lower left and lower right panels of Figure 5 show the responses of inflation to a one standard deviation and to a two standard deviation oil price innovations, by initial state. In general, there is no evidence that the response of output or inflation varies disproportionately with the size of the oil price shock. In all cases, the responses to small and large shocks, in low-growth and high-growth regimes, are very similar. It is important to note, however, that the responses are larger in the low-growth state relative to those in the high-growth state, consistent with the results in Figure 3.

In synthesis, the results suggest that oil price shocks feed back into economic growth conditions and, thus, play an important role in the evolution of low-growth and high-growth regimes. More importantly, there is evidence that the response of output is asymmetric to the sign of the oil price shock and that this asymmetry is correlated with the business cycle phase, even as the economy is allowed to evolve across regimes after the initial shock and initial state. Specifically, positive oil price shocks have a larger effect on output in recessions, but this asymmetry is lessened during expansionary times. This result is consistent with previous findings in the literature (Jiménez-Rodríguez and Sánchez, 2005; Hamilton, 2008, 2011; An et al., 2014), but at odds with those in Köse and Baimaganbetov (2015). This discrepancy is explained by the fact that Köse and Baimaganbetov (2015) use linear IRFs to analyzie nonlinear dynamic responses. Meanwhile, the response of inflation to a negative oil price shock is larger than its response to a positive oil price shock, and this difference is exacerbated during recessionary times.

5. ROBUSTNESS

In this section, we conduct robustness checks with respect to the ordering of the variables in the VAR and to the time window over which the net oil price change is computed. Results are shown in the next two subsections.

5.1. VAR Ordering

It is well-known that the estimation of coefficients and IRF from VAR processes may depend on the Cholesky ordering imposed to identify the shocks. In this subsection, we consider an alternative ordering representing a more exogenous ranking of the oil price shock: Oil price shock, output growth, M2 growth, interest rate, and CPI inflation. In this case, it is assumed that the oil price variable does not react contemporaneously to the rest of variables in the system. However, it has an immediate impact on output growth. The latter is then allowed to feed into changes on interest rates which, in turn, generates changes in the rate of inflation.

Under this alternative ordering, the linearity test suggests that the optimal delay lag, d, is also 2. Further, the p-value associated with the null hypothesis of linearity is 0.019, consistent with the results shown in Table 1. Figure 6 shows the median estimated cumulative responses of industrial production growth and inflation to a positive, one standard deviation shock in oil prices considering different initial conditions. Consistent with the results of Figure 3, the responses of output growth and inflation are larger in the low-growth regime, relative to the high-growth regime.

Figure 7 depicts the cumulative responses of output growth and inflation to a positive one standard deviation and to a negative one standard deviation oil price shocks, by initial state. To allow for a direct comparison, as before, the responses of output and inflation to a negative oil price shock have been multiplied by –1. Consistent with the findings from the upper row of Figure 4, the upper row of Figure 7 shows that the response of output to a positive oil price shock is larger than the response to a negative oil price shock, and that the difference is exacerbated when the economy starts in the low-growth regime. In particular, the responses of output growth to positive and negative oil price shocks when the economy starts in the low-growth regime peak at -1.550 and -0.921 percentage points after a year, respectively. However, similar responses when

Figure 7: Cumulative effects of positive and negative oil price shocks (alternative vector autoregression ordering), (a) Output growth, low-growth state, (b) Output growth, high-growth state, (c) Inflation, low-growth state, (d) Inflation, high-growth state

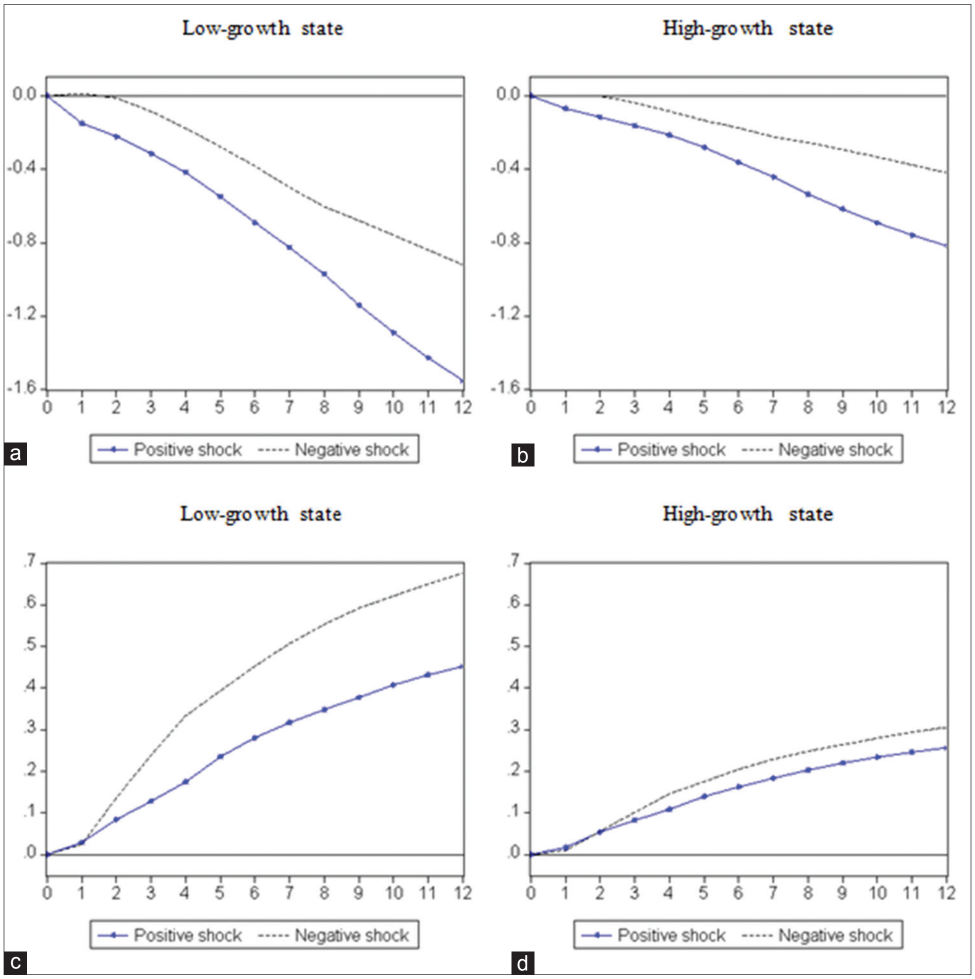

Accumulated response of output growth and inflation to positive and negative one standard deviation oil price shocks. The nonlinear IRFs are computed according to the details provided in Appendix A.

the economy starts in the high-growth regime peak at -0.871 and -0.456 percentage points. For the case of inflation, the results are also robust to the alternative VAR ordering. When the economy starts in the low-growth regime, the response of inflation to positive and negative oil price shocks reach a maximum of 0.453 and 0.677 percentage points after 1 year. When the economy starts in the high-growth regime, such responses reach a maximum of 0.257 and 0.306 percentage points. Hence, the reduction in inflation due to a negative oil price shock is larger in magnitude than the increase in inflation due to a positive oil price shock, and this difference is exacerbated in the low-growth regime.

Figure 8 reports the cumulative responses of output growth and inflation to one standard deviation and two standard deviation oil price shocks, by initial state. To allow for a direct comparison, as before, the responses of output growth and inflation to a two standard deviation shock are rescaled (divided by 2). The results are consistent with those documented in Figure 5. In general, there is little evidence of an asymmetric response of neither output growth nor inflation to the size of the oil price shock. However, the responses are larger in magnitude when the economy starts in the low-growth regime, relative to such responses when the economy starts in the high-growth regime, consistent with the evidence in Figures 3 and 6.

The robustness of the results to a more exogenous ranking of the oil price shock is in line with the findings in Kilian and Vega (2011), who show that energy prices are predetermined and that there is no evidence of any significant feedback at monthly horizons. Their findings thus support our assumption that oil prices are exogenous and imply that a less exogenous ordering would be inadequate[15]. Overall, the results are robust to the VAR ordering.

5.2. Oil Price Correction Window

We also consider NOPI using a 2-year window, rather than the 3-year window considered in equation (1). In this way, the oil

15 For the sake of completeness, we also consider a different VAR ordering representing a less exogenous ranking of the oil price shock: Output growth, CPI inflation, interest rate, M2 growth, and oil price variable. In this ordering, the oil price variable is allowed to react to all monetary variables in the system, given that it is an asset price. As in the alternative ordering considered previously in this subsection, the general pattern of the results remains largely unchanged. The results are available from the authors upon request.

Figure 8: Cumulative effects of small and large oil price shocks (alternative VAR ordering), (a) Output growth, low-growth state, (b) Output growth, high-growth state, (c) Inflation, low-growth state, (d) Inflation, high-growth state

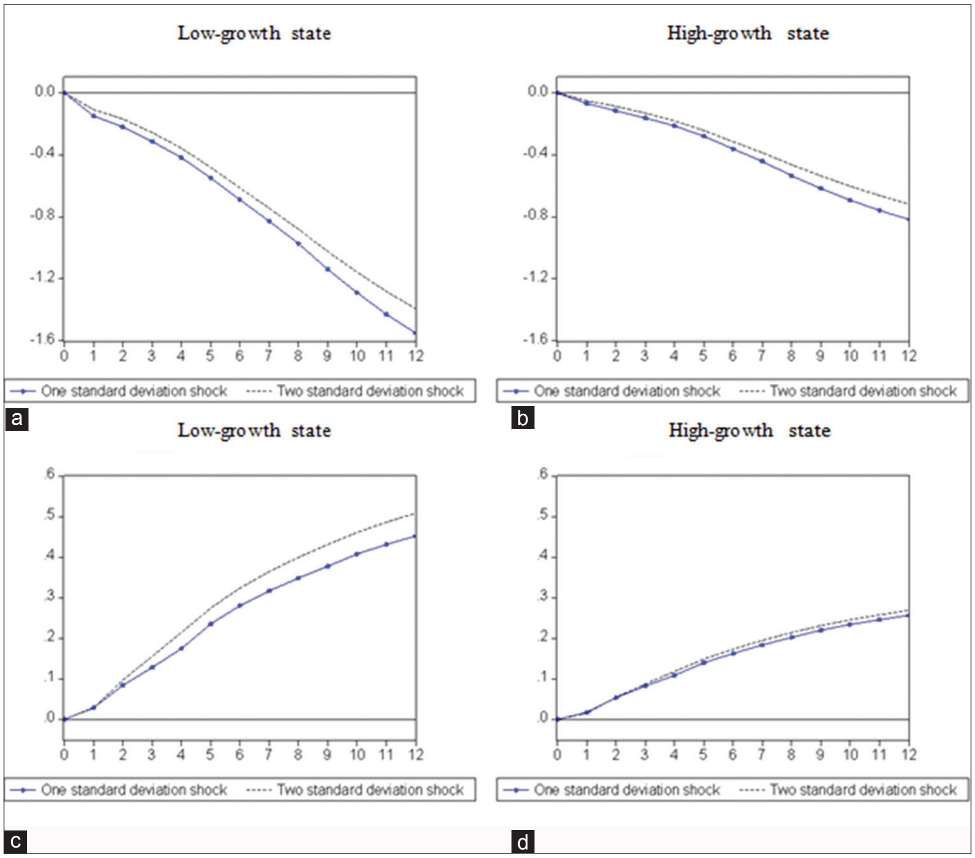

Accumulated response of output growth and inflation to small (one standard deviation) and large (two standard deviations, rescaled by 0.5) oil price shocks. The nonlinear IRFs are computed according to the details provided in Appendix A.

price shock is defined as the amount by which the change in oil price in month t exceeds the maximum value over the previous 2 years. Formally:

$$s_t = max\left\{0, ln\frac{p_t}{max\left\{p_{t-1},\ldots,p_{t-24}\right\}}\right\} \quad (7)$$

In this case, the optimal delay lag, d, is 1 and the p-value associated with the null hypothesis of linearity is 0.046. There is clear evidence of nonlinearities, although the results of the linearity test suggest that the 36 years window provides a better fit, as argued in Hamilton (2008).

Figure 9 plots the median estimated cumulative responses of industrial production growth and inflation to a positive, one standard deviation shock in oil prices, considering different initial conditions. The responses of output growth and inflation are larger in the low-growth regime, relative to the high-growth regime, consistent with the results of Figures 3 and 6.

To evaluate asymmetries with respect to the sign of the oil price shock under different starting conditions, cumulative responses of output growth and inflation to a positive one standard deviation and to a negative one standard deviation oil price shocks, by initial state, are obtained. The results are reported in Figure 10. As before, the responses to a negative oil price shock have been multiplied by –1. In general, the responses of output growth and inflation are asymmetric with respect to the sign of the oil price shock. The fall in output resulting from a positive oil price innovation is larger than the increase in output resulting from a negative oil price shock, and this difference is larger when the economy starts in the low-growth regime. Similarly, the reduction in inflation due a negative oil price shock is larger than the increase in inflation resulting from a positive oil price innovation, and this difference is exacerbated in the low-growth regime. Therefore, these results are consistent with those shown in Figures 4 and 7.

To evaluate possible asymmetric responses with respect to the size of the oil price shock, cumulative responses of output growth and inflation to one standard deviation and two standard deviation oil price shocks, by initial state, are computed and the results reported in Figure 11. As before, the responses of output growth and inflation to a two standard deviation shock are rescaled (divided by 2). Figure 11 shows little evidence of an asymmetric response for neither output growth nor inflation to the size of the oil price shock, consistent with the results reported before. At the same time, the responses are larger in magnitude when

Figure 9: Effects of an oil price increase by initial state (2-year window NOPI), (a) Cumulative response of output growth, (b) Cumulative response of inflation

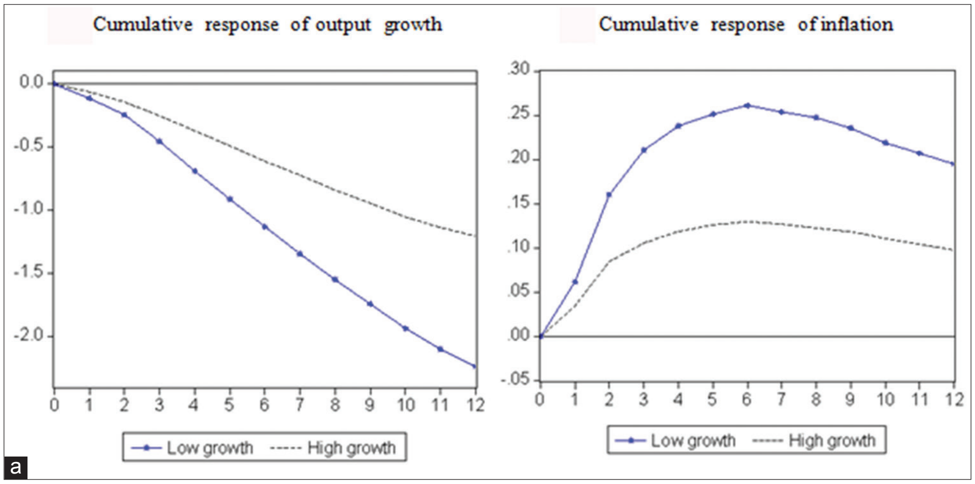

Accumulated response of output growth and inflation to a positive oil price shock, as defined in equation (1). The nonlinear IRF are computed according to the details provided in Appendix A.

Figure 10: Cumulative effects of positive and negative oil price shocks (2-year window NOPI), (a) Output growth, low-growth state, (b) Output growth, high-growth state, (c) Inflation, low-growth state, (d) Inflation, high-growth state

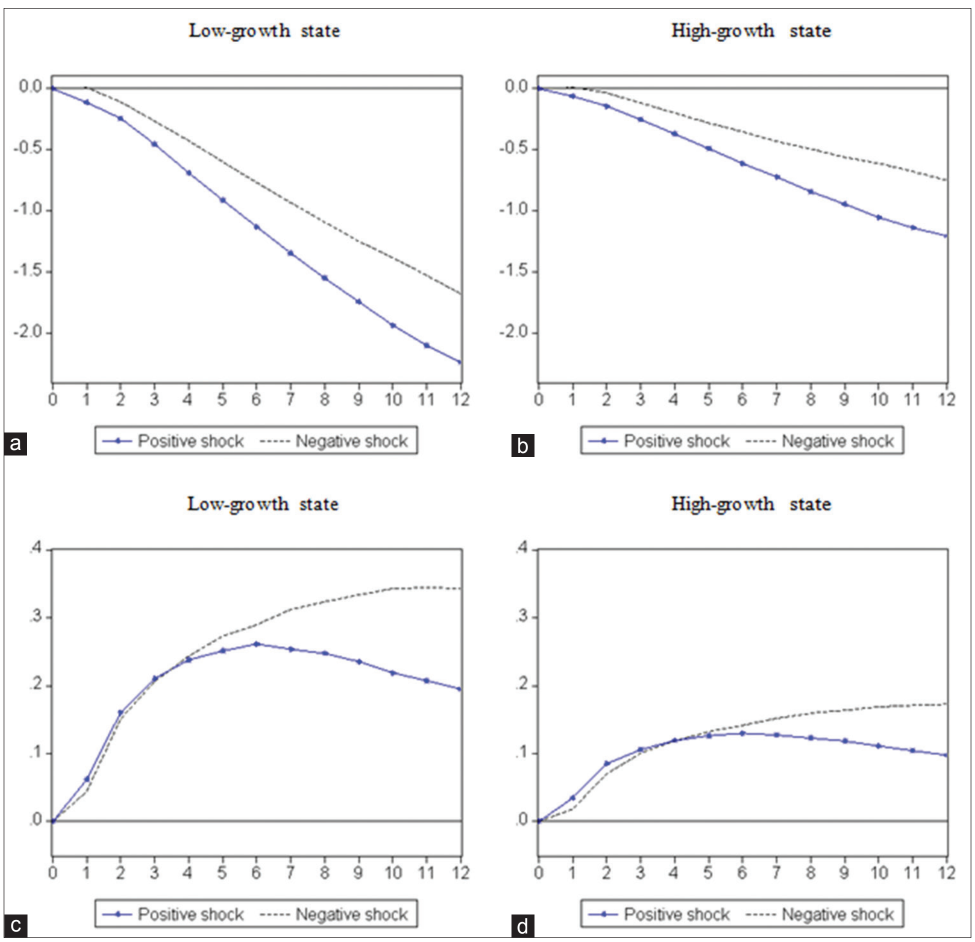

Accumulated response of output growth and inflation to positive and negative one standard deviation oil price shocks. The nonlinear IRF are computed according to the details provided in Appendix A.

the economy starts in the low-growth regime, relative to the responses when the economy starts in the high-growth regime, as in Figures 3, 6 and 9.

Overall, the analysis in this section suggests that the results are robust to the ordering of the variables in the VAR process and to the time window over which the net oil price change is computed.

Figure 11: Cumulative effects of small and large oil price shocks (2-year window NOPI), (a) Output growth, low-growth state, (b) Output growth, high-growth state, (c) Inflation, low-growth state, (d) Inflation, high-growth state

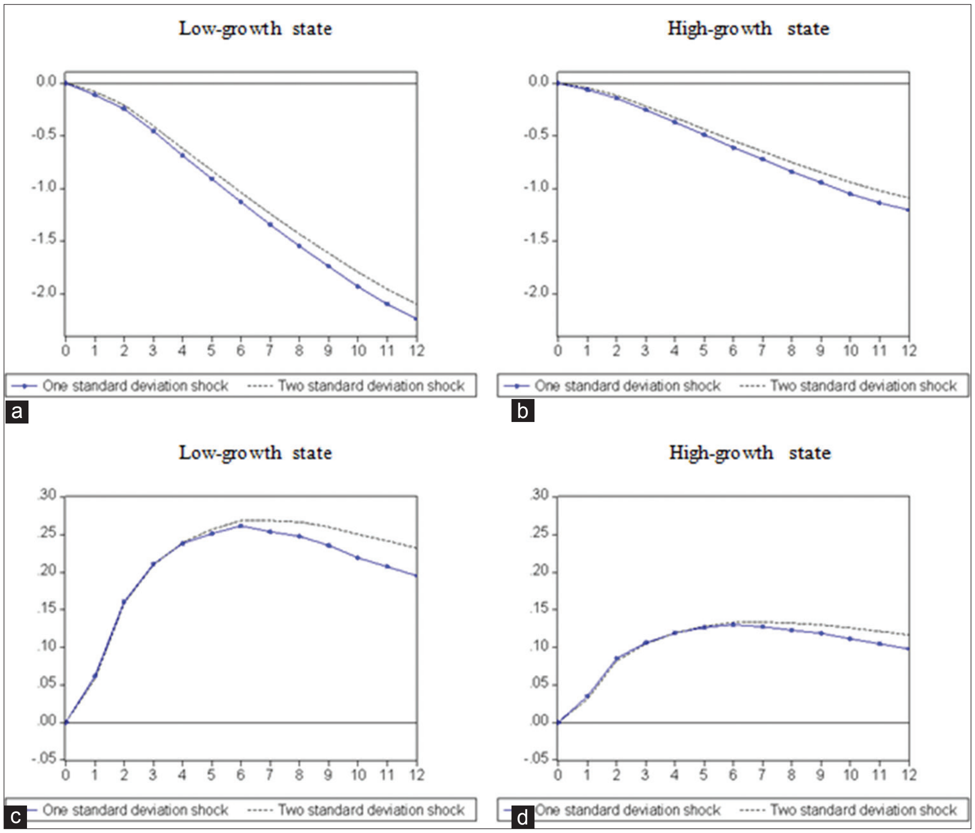

Accumulated response of output growth and inflation to small (one standard deviation) and large (two standard deviations, rescaled by 0.5) oil price shocks. (1). The nonlinear IRFs are computed according to the details provided in Appendix A.

6. CONCLUSION AND POLICY IMPLICATIONS

We have extended the literature that studies the effects of positive and negative oil price shocks on output and inflation by considering, unlike previous studies, a nonlinear VAR framework that allows us to analyze potential correlations between asymmetries associated with the business cycle phase and the size/sign of oil price shocks.

Our results are at odds with the conclusions from Kilian and Vigfusson (2011a and b) and Herrera et al. (2015) who, unlike much of the empirical literature, find little evidence of an asymmetric response of output to the sign of oil price shocks. However, the approach we propose refines previous findings and adds important new insights associated with the interconnections between the different asymmetries considered that might explain this mixed evidence. Specifically, we find evidence that the response of output growth to positive oil price shocks is larger than its response to negative oil price shocks, consistent with the results from An et al. (2014), even after taking into account the criticism of Kilian and Vigfusson (2011a and b). Meanwhile, this asymmetric response is exacerbated during low-growth periods and lessened during high-growth periods. In general, our findings are consistent with those in previous studies (Jiménez-Rodríguez and Sánchez, 2005; Hamilton, 2008, 2011; An et al., 2014), although contrary to the results of Köse and Baimaganbetov (2015), who find a larger effect of negative oil price shocks. The discrepancy is explained, notwithstanding, by the fact that Köse and Baimaganbetov (2015) use linear IRF to evaluate nonlinear dynamic responses. The analysis of the nonlinear IRF also suggests that the negative demand effects of an oil price increase offset the positive supply side effects for the case of a net oil exporter, like Canada. Intuitively, this could occur because the positive supply side effects of higher-priced commodities are delayed, since extraction and production of crude oil involve long-term investments.

An analysis of the results also shows that inflation evolves differently after positive and negative oil price shocks. Specifically, the reduction in inflation due to a negative oil price shock is larger than the increase in inflation after a positive oil price innovation. This difference is particularly large when the economy starts in the low-growth regime. Meanwhile, we did not find much evidence that the response of output growth, nor that of inflation, varies disproportionately with the size of the shock. In general, the results are robust to the ordering of the variables in the VAR process and to the time window over which the net oil price change is computed.

Better understanding the asymmetric effects of positive and negative oil price shocks, as well as their dependence on the state of the business cycle, can have relevant policy implications. First, assessing the effects of unexpected oil declines is important to determine the appropriate response of monetary authorities. For example, our results suggest that the recent decline in the price of crude oil which started in mid-2014 should not have a large effect in aggregate economic activity if the economy is in a high-growth state and, therefore, should not affect the stance of monetary policy. Second, policymakers could engage in expansionary economic policy to counteract the detrimental effects of positive oil shocks during low-growth regimes, while understanding that negative oil shocks may require less accommodation. For instance, the Canadian economy experienced numerous positive oil price shocks during the oil boom of the mid-to-late 2000s. In the later stages of this boom (late 2007 and early 2008), policymakers at the Bank of Canada reduced the target for the overnight rate. This accommodative policy was set against the backdrop of a record high for world oil prices, a Canadian economy that was experiencing moderate growth and inflation levels that remained inside the Banks stated target bound. Third, policies that promote diversification of the economy, reducing the reliance on oil-related production, can be beneficial for output growth given the effect of positive oil shocks during low-growth states.

Finally, the results from our model leave open the important question of why oil price shocks would have larger effects on output during recessions and why inflation seems to respond more strongly to negative oil price shocks. Theoretical explanations of the correlation between the different asymmetries explored here and tests of different asymmetry-generating models are left for future research.

7. ACKNOWLEDGMENTS

The authors would like to thank Sebastian Fossati, Junsoo Lee, Tatsuma Wada and the participants at the 2015 Symposium of the Society for Nonlinear Dynamics and Econometrics, the 2015 Meetings of the Midwest Economics Association, the 2015 International Conference of the Society for Computational Economics and the University of Memphis for helpful comments. All remaining errors are our own.

REFERENCES

An, L., Jin, X., Ren, X. (2014), Are the macroeconomic effects of oil price shock symmetric? A factor-augmented vector autoregressive approach. Energy Economics, 45, 217-228.

Baskaya, Y.S., Hulagu, T., Kucuk, H. (2013), Oil price uncertainty in a small open economy. IMF Economic Review, 61(1), 168-198.

Baumeister, C., Peersman, G. (2013), Time-varying effects of oil supply shocks on the U.S. economy. American Economic Journal: Macroeconomics, 5(4), 1-28.

Bernanke, B.S. (1983), Irreversibility, uncertainty, and cyclical investment. Quarterly Journal of Economics, 98, 85-106.

Bresnahan, T.F., Ramey, V.A. (1993), Segment shifts and capacity utilization in the U.S. automobile industry. American Economic Review Papers and Proceedings, 83(2), 3213-218.

Davies, S.J. (1987a), Allocative disturbances and specific capital in real business cycle theories. American Economic Review, 77(2), 326-332.

Davies, S.J. (1987b), Fluctuations in the pace of labor reallocation. In: Brunner, K., Meltzer, A.H., editors. Empirical Studies of Velocity, Real Exchange Rates, Unemployment and Productivity. Vol. 27. Carnegie-Rochester Conference Series on Public Policy. p335-402.

Davies, S.J., Haltiwanger, J. (2001), Sectoral job creation and destruction responses to oil price changes. Journal of Monetary Economics, 48(3), 465-512.

Edelstein, P., Kilian, L. (2007), The response of business fixed investment to changes in energy prices. A test of some hypotheses about the transmission of energy price shocks. The B.E. Journal of Macroeconomics, 7(1), 1-41.

Edelstein, P., Kilian, L. (2009), How sensitive are consumer expenditures to retail energy prices? Journal of Monetary Economics, 56(6), 766-779.

Engemann, K.M., Kliesen, K.L., Owyang, M.T. (2011), Do oil shocks drive business cycles? Some U.S. and international evidence. Macroeconomic Dynamics, 15(3), 498-517.

García, R., Schaller, H. (2002), Are the effects of monetary policy asymmetric? Economic Inquiry, 40(1), 102-119.

Gronwald, M. (2012), Oil and the U.S. macroeconomy: A reinvestigation using rolling impulse responses. Energy Journal, 33(4), 143-159.

Hamilton, J.D. (1983), Oil and the macroeconomy since world war II. Journal of Political Economy, 91(2), 228-248.

Hamilton, J.D. (1988), A neoclassical model of unemployment and the business cycle. Journal of Political Economy, 96(3), 593-617.

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

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

Hamilton, J.D. (2008), Oil and the macroeconomy. In: Durlauf, S., Blume, L., editors. The New Palgrave Dictionary of Economics. Vol. 1. New York: Palgrave McMillan Ltd.

Hamilton, J.D. (2011), Nonlinearities and the macroeconomic effects of oil prices. Macroeconomic Dynamics, 15(3), 364-378.

Hansen, B.E. (1996), Inference when a nuisance parameter is not identified under the null hypothesis. Econometrica, 64, 413-430.

Hansen, B.E. (1997), Inference in TAR models. Studies in Nonlinear Dynamics and Econometrics, 2(1), 1-14.

Herrera, A.M., Lagalo, L.G., Wada, T. (2011), Oil price shocks and industrial production: Is the relationship linear? Macroeconomic Dynamics, 15(3), 472-497.

Herrera, A.M., Lagalo, L.G., Wada, T. (2015), Asymmetries in the response of economic activity to oil price increases and decreases? Journal of International Money and Finance, 50, 108-133.

Herrera, A.M., Pesavento, E. (2009), Oil price shocks, systematic monetary policy, and the great moderation. Macroeconomic Dynamics, 13, 107-137.

International Energy Agency. (2012), World Energy Outlook. Available from: http://www.worldenergyoutlook.org.

Jiménez-Rodríguez, R., Sánchez, M. (2005), Oil price shocks and real GDP growth: Empirical evidence for some OECD countries. Applied Economics, 37, 201-228.

Kilian, L., Vega, C. (2011), Do Energy prices respond to U.S. macroeconomic news? A test of the hypothesis of predetermined energy prices. The Review of Economics and Statistics, 93(2), 660-671.

Kilian, L., Vigfusson, R.J. (2011a), Are the Responses of the U.S. Economy asymmetric in energy price increases and decreases? Quantitative Economics, 2(3), 419-453.

Kilian, L., Vigfusson, R.J. (2011b), Nonlinearities in the oil price-output relationship. Macroeconomic Dynamics, 15(3), 337-363.

Koop, G., Pesaran, M.H., Potter, S. (1996), Impulse-response analysis in nonlinear multivariate models. Journal of Econometrics, 74, 119-147.

Köse, N., Baimaganbetov, S. (2015), The asymmetric impact of oil price shocks on Kazakhstan macroeconomic dynamics: A structural vector autoregression approach. International Journal of Energy Economics and Policy, 5(4), 1058-1064.
Lee, K., Ni, S. (2002), On the dynamic effects of oil price shocks: A study using industry level data. Journal of Monetary Economics, 49, 823-852.
Lo, M.C., Piger, J. (2005), Is the response of output to monetary policy asymmetric? Evidence from a regime-switching coefficients model. Journal of Money, Credit and Banking, 37, 865-887.
Peersman, G., Smets, F. (2002), Are the effects of monetary policy in the Euro Area greater in recessions than in booms? In: Mahadeva, L., Sinclair, P., editors. Monetary Transmission in Diverse Economies. Cambridge, UK: Cambridge University Press. p28-48.
Pindyck, R.S. (1991), Irreversibility, uncertainty, investment. Journal of Economic Literature, 29(3), 1110-1148.
Plante, M., Traum, N. (2014), Time-Varying Oil Price Volatility and Macroeconomic Aggregates, Federal Reserve Bank of Dallas Working Paper, 1201.
Teräsvirta, T., Yang, Y. (2014), Specification, Estimation and Evaluation of Vector Smooth Transition Autoregressive Models with Applications. Creates Research Paper 2014-09, University of Aarhus.
Teräsvirta, T., Yang, Y. (2016), Linearity and Misspecification Tests for Vector Smooth Transition Regression Models forthcoming at Macroeconomic Dynamics.
Weise, C. (1999), The asymmetric effects of monetary policy: A nonlinear vector autoregression approach. Journal of Money, Credit and Banking, 31, 85-108.
Yellen, J. (2012), Revolution and Evolution in Central Bank Communications. Speech at at the Haas School of Business, University of California - Berkeley.

Appendix A

COMPUTATION OF NONLINEAR IMPULSE-RESPONSE FUNCTIONS

The procedure to compute the nonlinear impulse-response functions (IRF) follows Koop et al. (1996) and Teräsvirta and Yang (2014, 2016). The reader is referred to the original documents for further details.

A nonlinear IRF can be defined as the effect of a one-time shock on the forecast of variables in a particular model, given a specific history. The response constructed must then be compared to a benchmark 'no shock' scenario. In this way, the nonlinear IRF can be expressed as follows:

$$GI_Y(h, \nu_t, \omega_t) = E[Y_{t+h}/\nu_t, \omega_{t-1}] - E[Y_{t+h}/\omega_{t-1}]$$

where GI_Y is the nonlinear impulse-response function of a variable Y for period h, given the specific history ω_{t-1} and initial shock ν_t, and E[.] is the expectations operator.

To compute the nonlinear IRF, the conditional expectations in the equation above are simulated. The nonlinear model is assumed to be known (i.e., sample variability is ignored). The shock to Y, ν_0, occurs in period 0, and responses are computed for h periods ahead. Thus, the GI_Y function is generated according to the following steps:

Step 1: Pick a history ω_{t-1}. The history is the actual value of the lagged endogenous variables at a particular date, or for a particular episode (e.g., those values of the endogenous variables that fall under the low-growth regime).

Step 2: Pick a sequence of shocks $\nu_{j,t+h}$, h = 0, 1, ..., n for j = 0, 1, ..., k. They are drawn with replacement from the estimated residuals of the VAR model.

Step 3: Using $\omega_{i,t-1}$ and $\nu_{j,t+h}$, simulate the path for y_{t+h}, over n periods according to equation (3). This benchmark path is denoted as $Y_{t+h}(\omega_{i,t-1}, \nu_{j,t+h})$ for h = 1, ... , n.

Step 4: Using the Cholesky decomposition of Ω_t to orthogonalize the shocks, solve for the oil price shock and reconstruct the implied vector of errors. Using the same $\omega_{i,t-1}$ and $\nu_{j,t+h}$, plus the additional initial oil price shock ν_0, simulate the path for y_{t+h} over n + 1 periods according to equation (3). This profile path is denoted $Y_{t+h}(\nu_0, \omega_{i,t-1}, \nu_{j,t+h})$ for h = 0, 1, ..., n.

Step 5: Repeat steps 2 to 4 B times.

Step 6: Repeat steps 1 to 5 R times and compute the quantiles of the difference between the profile and benchmark paths $Y_{t+h}(\nu_0, \omega_{i,t-1}, \nu_{j,t+h}) - Y_{t+h}(\omega_{i,t-1}, \nu_{j,t+h})$.