



BANK OF CANADA
BANQUE DU CANADA

Working Paper/Document de travail
2012-2

Time-Varying Effects of Oil Supply Shocks on the U.S. Economy

by Christiane Baumeister and Gert Peersman

Bank of Canada Working Paper 2012-2

January 2012

Time-Varying Effects of Oil Supply Shocks on the U.S. Economy

by

Christiane Baumeister¹ and Gert Peersman²

¹International Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
cbaumeister@bankofcanada.ca

²Department of Financial Economics
Ghent University
9000 Gent, Belgium
gert.peersman@ugent.be

Bank of Canada working papers are theoretical or empirical works-in-progress on subjects in economics and finance. The views expressed in this paper are those of the authors.

No responsibility for them should be attributed to the Bank of Canada.

Acknowledgements

We thank the editor, three anonymous referees, Luca Benati, Jean Boivin, Luca Gambetti, Domenico Giannone, James Hamilton, Morten Ravn, Tatevik Sekhposyan, Frank Smets, Herman van Dijk and Timo Wollmershäuser as well as numerous conference and seminar participants for useful comments and suggestions. Special thanks to Lutz Kilian for many insightful discussions and feedback on earlier versions of the paper. We acknowledge financial support from the IUAP Programme – Belgian Science Policy [Contract No. P6/07] and the Belgian National Science Foundation. All remaining errors are ours.

Abstract

We use vector autoregressions with drifting coefficients and stochastic volatility to investigate how the dynamic effects of oil supply shocks on the U.S. economy have changed over time. We find a substantial decline in the short-run price elasticity of oil demand since the mid-eighties. This finding helps explain why an oil production shortfall of the same magnitude is associated with a stronger response of oil prices and more severe macroeconomic consequences over time, while an oil price increase of the same magnitude is associated with a smaller decline in oil production and smaller losses in U.S. output in more recent years. We also show that oil supply shocks more recently account for a smaller fraction of the variability of the real price of oil, implying a greater role for oil demand shocks. Notwithstanding this time variation, the overall cumulative effect of oil supply disruptions on the U.S. economy has been modest. Oil supply shocks contributed to some extent to the 1991 recession and slowed the economic boom of 1999-2000, but they do not explain other U.S. recessions nor do they help explain the “Great Inflation” of the 1970s and early 1980s.

JEL classification: E31, E32, Q43

Bank classification: Econometric and statistical methods; International topics

Résumé

Les auteurs se servent de modèles vectoriels autorégressifs dotés de coefficients variables et d'une volatilité stochastique pour étudier comment ont évolué au fil du temps les effets dynamiques des chocs d'offre pétroliers sur l'économie américaine. Les auteurs notent une baisse prononcée de l'élasticité-prix de la demande de pétrole à court terme depuis le milieu des années 1980. Cette observation aide à expliquer pourquoi, à ampleur égale, un déficit de production s'accompagne d'une réaction du prix de l'or noir et de dégâts à l'économie de plus en plus importants au fil de la période étudiée, alors que des hausses des prix du pétrole d'ampleur inchangée sont associées à un recul moins marqué de la production pétrolière et à un repli plus modeste de l'activité économique aux États-Unis ces dernières années. Les auteurs montrent aussi que la part de la variabilité du prix réel de l'or noir imputable aux chocs d'offre a diminué récemment, avec pour corollaire un rôle plus important pour les chocs de demande. Malgré la variation dans le temps, l'incidence cumulative globale des ruptures d'approvisionnement a été modeste sur l'économie américaine. Les chocs d'offre pétroliers ont concouru dans une certaine mesure à la récession de 1991 et tempéré le boom de 1999-2000, mais ils n'expliquent pas d'autres récessions aux États-Unis ni la « Grande Inflation » des années 1970 et du début des années 1980.

Classification JEL : E31, E32, Q43

Classification de la Banque : Méthodes économétriques et statistiques; Questions internationales

1 Introduction

This paper studies the relationship between oil prices and U.S. macroeconomic performance over the period 1974 up to now. There is considerable evidence that this relationship has been unstable over time (see, e.g., Edelstein and Kilian 2009; Herrera and Pesavento 2009; Blanchard and Galí 2010; Ramey and Vine 2012). In particular, several researchers have noted a substantial decline in the macroeconomic consequences of oil price shocks. One of the reasons for this temporal instability is the fact that oil price shocks are merely symptoms of underlying oil supply and demand shocks. As the composition of these structural oil supply and demand shocks evolves, so does the dynamic correlation between the U.S. economy and the real price of oil (see, e.g., Kilian 2009a).

Even if we distinguish between oil demand and oil supply shocks, however, there are additional reasons why the response of the U.S. economy to these shocks may have changed over time. Such time-varying effects may come about, for example, because of variation in the oil intensity of economic activity, because of changes in the regulation of energy markets, or because of changes in the composition of automobile production and the overall importance of the U.S. automobile sector (see, e.g., Edelstein and Kilian 2009; Ramey and Vine 2012). Likewise, changes in capacity utilization rates in crude oil production and the transition toward a market-based system of oil trading in the 1980s may account for additional instabilities in the response to oil supply shocks (see, e.g., Hubbard 1986). Indeed, there is a consensus in the literature that the short-run price elasticity of oil demand has declined since the mid-1980s, although the extent of these changes has proved hard to pin down (see, e.g., Hughes, Knittel, and Sperling 2008; Dargay and Gately 1994, 2010).

Time-varying responses are not allowed for in the existing literature on the effects of oil supply shocks on U.S. macroeconomic aggregates. It is by now well understood that oil price shocks are not the same as oil supply shocks (see Kilian 2008a). As a result, in recent years, considerable work has been done to explicitly identify oil supply shocks and to understand their effect on the real price of oil and on the U.S. economy. Hamilton (1983) first remarked on the coincidence of oil supply disruptions and subsequent oil price surges during 1948-1972. Building on this insight, Hamilton (2003) developed a quantitative dummy measure based

on physical oil supply disruptions associated with major political crises in oil-producing countries and investigated its predictive power for changes in the price of oil and for U.S. real GDP. Kilian (2008b,c) proposed an alternative measure of exogenous oil supply shocks obtained by constructing explicit counterfactuals for all major oil producers and studied the responses of inflation and real output in major industrial economies to this measure of oil supply shocks. Yet another approach developed by Kilian (2009a) has relied on exclusion restrictions in structural vector autoregressive (VAR) models of the global oil market to identify oil supply shocks.

A common feature of all these empirical studies is that they rely on time-invariant regressions. Implicitly, it is assumed that the effect of oil supply shocks on the real price of oil and on macroeconomic aggregates has not changed over time. This observation raises the question of how reliable existing estimates of the effects of oil supply shocks are.

In this paper, we use a time-varying parameter structural VAR model to investigate how the effects of oil supply shocks on the U.S. economy have changed over time. Our analysis incorporates several innovations relative to the previous literature. First, rather than imposing an arbitrary sample split as some previous studies have done, we let all model coefficients evolve continuously, allowing the data to tell us when and how any changes may have occurred. Second, we identify oil supply shocks not based on contemporaneous exclusion restrictions but instead based on the sign restrictions that an oil supply shock moves oil prices and oil production in opposite directions. Ours is the first application of this identification approach in the context of the global oil market.¹ Third, we allow for time-varying heteroskedasticity in the VAR innovations that accounts for changes in the magnitude of structural shocks and their immediate impact. This feature is particularly important in the present setting given the observed increase in oil price volatility and the reduction in macroeconomic volatility during the Great Moderation. Fourth, in addition to shedding light on time variation in the responses of U.S. real GDP and consumer prices to oil supply shocks, our analysis also allows us to obtain estimates of changes over time

¹Sign restrictions for modeling oil market dynamics have subsequently been adopted by Peersman and Van Robays (2009, 2011), Baumeister and Peersman (2011), Kilian and Murphy (2010, 2011), and Lippi and Nobili (2011), among others.

in the short-run price elasticity of oil demand, complementing the existing literature. By construction, this demand elasticity corresponds to the ratio of the impact responses of world oil production and of the real price of oil in response to an exogenous shift of the world oil supply curve along the oil demand curve.

Our analysis yields several intriguing results. First, we find that even though the effect on economic aggregates of a 1% oil supply disruption has increased over time, the effect on U.S. real GDP of an oil supply disruption associated with a 10% increase in the real price of oil has declined over time. We show that this evidence cannot be explained merely by changes in the variance of oil supply shocks but requires a decline in the price elasticity of oil demand over time, such that a given decrease in the quantity of oil supplied is associated with a larger increase in the real price of oil. Second, the contribution of oil supply shocks to the variability of the real price of oil has moderately declined over time, indicating a larger role for oil demand shocks. It is reassuring that accounting for time variation does not overturn this important insight from the recent literature. Third, although the contribution of oil supply shocks to the variance of macroeconomic aggregates is non-negligible, they explain only a small part of the recessions since the 1970s and of the "Great Inflation".

The remainder of the paper is organized as follows. Section 2 presents the econometric methodology and describes the identification strategy in more detail. Section 3 discusses the main empirical results and evaluates the robustness of our findings. In section 4, we discuss some potential explanations for a less price-elastic oil demand curve since the second half of the 1980s. Section 5 contains the concluding remarks.

2 Empirical methodology

The apparent instabilities in the oil-macro relationship suggest that modeling the transmission of oil supply shocks adequately requires an empirical framework that can account for changes over time. The anecdotal evidence presented in the introduction points towards a gradual evolution of the interaction between the oil market and the U.S. economy. The idea of a slow-moving but continuous adjustment is also in line with the adaptive behavior

of economic agents that results from an ongoing learning process. For example, Primiceri (2005) makes the case that the aggregation among agents' assessments tends to be reflected in smooth changes since agents would not be expected to update their beliefs all at once. This line of reasoning suggests that the appropriate modeling approach is a time-varying parameter model (TVP-VAR) featuring smoothly evolving coefficients and heteroskedasticity in the innovations.

Although the possibility of abrupt breaks cannot be excluded a priori, it can be shown that a time-varying parameter model is capable of capturing such discrete shifts should they occur. As the Monte Carlo study in the online appendix illustrates, drifting coefficient models in practice can successfully track processes subject to structural breaks or regime shifts. In contrast, models of discrete shifts cannot accommodate smooth structural change.

2.1 A VAR with time-varying parameters and stochastic volatility

We model the joint behavior of global oil production, the real U.S. refiners' acquisition cost of imported crude oil (obtained by deflating the nominal price by U.S. CPI), U.S. real GDP, and U.S. consumer prices as a VAR(p) with time-varying parameters and stochastic volatility:

$$y_t = c_t + B_{1,t}y_{t-1} + \dots + B_{p,t}y_{t-p} + u_t \equiv X_t'\theta_t + u_t \quad (1)$$

where $y_t \equiv [\Delta q_t^{oil}, \Delta p_t^{oil}, \Delta gdp_t, \Delta cpi_t]'$ and Δ denotes the first difference operator.² All variables were transformed to non-annualized quarter-on-quarter rates of growth by taking the first difference of their natural logarithm. The time-varying intercepts c_t and the matrices of time-varying coefficients $B_{1,t \dots p,t}$ are stacked in θ_t , and X_t is a matrix including lags of y_t and a constant to obtain the state-space representation of the model. The data frequency is quarterly. The overall sample covers the period 1947Q1 – 2011Q1, but the first twenty-five years of data are used as a training sample to generate the priors for estimation over the actual sample period which starts in 1974Q1. This is the earliest possible starting date given that before 1974 the oil price was regulated which impairs the use of standard time-series

²A detailed description of the data used in this paper can be found in the online appendix.

models of the oil market even when time variation is allowed for (see, e.g., Alquist, Kilian, and Vigfusson 2011). The lag length is set to $p = 4$ to allow for sufficient dynamics in the system and capture lags in the transmission of oil shocks (see, e.g., Hamilton and Herrera 2004). The u_t in the observation equation is an unconditionally heteroskedastic disturbance term that is normally distributed with a zero mean and a time-varying covariance matrix Ω_t which can be decomposed in the following way:

$$\Omega_t = A_t^{-1} H_t (A_t^{-1})' \quad (2)$$

A_t is a lower triangular matrix that models the contemporaneous interactions among the endogenous variables and H_t is a diagonal matrix that contains the stochastic volatilities:

$$A_t = \begin{bmatrix} 1 & 0 & 0 & 0 \\ \alpha_{21,t} & 1 & 0 & 0 \\ \alpha_{31,t} & \alpha_{32,t} & 1 & 0 \\ \alpha_{41,t} & \alpha_{42,t} & \alpha_{43,t} & 1 \end{bmatrix} \quad H_t = \begin{bmatrix} h_{1,t} & 0 & 0 & 0 \\ 0 & h_{2,t} & 0 & 0 \\ 0 & 0 & h_{3,t} & 0 \\ 0 & 0 & 0 & h_{4,t} \end{bmatrix} \quad (3)$$

The drifting coefficients are meant to capture possible nonlinearities or time variation in the lag structure of the model. The multivariate time-varying variance-covariance matrix allows for heteroskedasticity of the shocks and time variation in the simultaneous relationships between the variables in the system. Allowing for time variation in both the coefficients and the variance-covariance matrix, leaves it up to the data to determine whether the time variation of the linear structure comes from changes in the size of the shock and its contemporaneous impact or from changes in the propagation mechanism. Let $\alpha_t = [\alpha_{21,t}, \alpha_{31,t}, \alpha_{32,t}, \alpha_{41,t}, \alpha_{42,t}, \alpha_{43,t}]'$ be the vector of the non-zero and non-one elements of the matrix A_t , and $h_t = [h_{1,t}, h_{2,t}, h_{3,t}, h_{4,t}]'$ be the vector of the diagonal elements of H_t . Following Primiceri (2005), the three driving processes of the system are postulated to evolve as follows:

$$\theta_t = \theta_{t-1} + \nu_t \quad \nu_t \sim N(0, Q) \quad (4)$$

$$\alpha_t = \alpha_{t-1} + \zeta_t \quad \zeta_t \sim N(0, S) \quad (5)$$

$$\ln h_{i,t} = \ln h_{i,t-1} + \sigma_i \eta_{i,t} \quad \eta_{i,t} \sim N(0, 1) \quad (6)$$

The time-varying parameters θ_t and α_t are modeled as driftless random walks. We impose a stability constraint on the evolution of the time-varying parameters to enforce stationarity of the VAR system. Specifically, we include an indicator function that selects only stable draws i.e. the indicator function $I(\theta_t) = 0$ if the roots of the associated VAR polynomial are inside the unit circle as in e.g. Cogley and Sargent (2005). The elements of the vector of volatilities h_t are assumed to evolve as geometric random walks independent of each other. The error terms of the three transition equations are independent of each other and of the innovations of the observation equation. In addition, we impose a block-diagonal structure for S of the following form:

$$S \equiv Var(\zeta_t) = \begin{bmatrix} S_1 & 0_{1 \times 2} & 0_{1 \times 3} \\ 0_{2 \times 1} & S_2 & 0_{2 \times 3} \\ 0_{3 \times 1} & 0_{3 \times 2} & S_3 \end{bmatrix} \quad (7)$$

which implies independence also across the blocks of S with $S_1 \equiv Var(\zeta_{21,t})$, $S_2 \equiv Var([\zeta_{31,t}, \zeta_{32,t}]')$, and $S_3 \equiv Var([\zeta_{41,t}, \zeta_{42,t}, \zeta_{43,t}]')$ so that the covariance states can be estimated equation by equation. We estimate this model using Bayesian methods described in Kim and Nelson (1999). An overview of the prior specifications and the estimation strategy (Markov Chain Monte Carlo algorithm) is provided in the online appendix.

2.2 Identification of oil supply shocks

It is now widely accepted that oil prices are not only determined by supply-side factors but also driven by demand conditions (see Barsky and Kilian 2002, 2004; Kilian 2008a; Hamilton 2009a,b; Kilian 2009a,b). Innovations in the oil price equation of a VAR model are not an adequate measure of exogenous variation in oil supply because they also capture shifts in the demand for crude oil. The resulting estimates of macroeconomic effects only represent the consequences of an average oil price shock determined by a combination of supply and demand factors. Blanchard and Galí (2010), among others, make the case that this distinction does not matter because an oil price shock triggered by increased demand

for oil in one country can still be experienced as a supply shock by the remaining countries. This presumption is, however, very stringent in light of the results of Kilian (2009a) who shows that there exist important differences in the responses of macroeconomic aggregates depending on the underlying source of the oil price shock. Intuitively it is clear that an increase in the real price of oil induced by favorable global economic conditions exerts a different influence on the macroeconomic performance than one due to oil supply disruptions resulting from a war (see Kilian 2009a; Rotemberg 2010).

Kilian (2009a) disentangles oil supply from oil demand shocks based on contemporaneous exclusion restrictions in a monthly vector autoregression that includes world oil production and the real price of crude oil. An oil supply shock is identified as the sole disturbance that has an immediate influence on the level of oil production. Accordingly, global oil production does not respond instantaneously to oil demand shocks which implies that the short-run oil supply curve is vertical. This assumption, while arguably tenable at the monthly frequency as discussed in Kilian and Vega (2011), is however, less appropriate when quarterly data are used as in our study.³

Therefore, we propose a new approach to identifying structural oil supply shocks that involves sign restrictions on the estimated time-varying impulse responses. This approach builds on other applications of sign restrictions including the work of Faust (1998), Davis and Haltiwanger (1999), Canova and De Nicoló (2002), Uhlig (2005), and Peersman (2005), among others. Related work on time-varying parameter VAR models with sign restrictions includes Benati and Mumtaz (2007) and Canova and Gambetti (2009). Our identifying restrictions are based on insights derived from a basic supply and demand model for the global oil market which is represented by world oil production and the real price of crude oil. An oil supply shock is identified as a disturbance that shifts the upward-sloping oil supply curve along the downward-sloping oil demand curve and hence, results in an opposite

³An alternative approach in the literature has been to quantify oil supply shocks directly through suitable counterfactual thought experiments. Hamilton (2003) and Kilian (2008b,c) have developed measures of oil supply shocks based on physical oil supply disruptions in the wake of major exogenous political events in oil-producing countries. The advantage of defining oil supply shocks within a structural VAR model is that it also captures endogenous production responses among oil producers to exogenous oil supply disruptions.

movement in oil production and in the real price of oil. The identifying assumptions are that after a negative oil supply shock world oil production declines and the real price of oil increases. No constraints are imposed on the responses of U.S. real output and of consumer prices. The reactions of these variables will be determined by the data. The sign restrictions are imposed to hold for four quarters following the shock, consistent with widely held beliefs about oil price dynamics in the literature (Hamilton 2003, Kilian 2008c). As a consequence of these inequality constraints, our identification scheme does not deliver exact identification (see, e.g., Fry and Pagan 2011).

Also note that the model is only partially identified in the sense that only the oil supply shock is explicitly identified and the conglomerate of residual oil demand shocks has no structural interpretation which is in line with the focus of this paper, as outlined in the introduction. In Baumeister and Peersman (2011), we extend the present identification strategy to explore in more detail changes in global oil price dynamics. Notably, we include a measure of global real economic activity which enables us to differentiate between oil demand shocks driven by the global business cycle and oil demand shocks related to an expectations-driven component. While Baumeister and Peersman (2011) study the time-varying dynamics of the real price of oil as a function of all three types of oil market shocks, here we are concerned with understanding the transmission of oil supply shocks to the domestic economy.

3 Results

We begin our analysis with a discussion of the impulse response estimates in section 3.1. The computation of the time-varying impulse responses and the implementation of the sign restrictions are described in the online appendix. In section 3.2, we assess the quantitative importance of oil supply shocks and in section 3.3, we conduct several sensitivity analyses.

3.1 Responses to an oil supply shock

In structural VAR models, it is conventional to report the responses of the endogenous variables to one-standard-deviation shocks. The problem in time-varying structural VAR

models is that a one-standard-deviation shock corresponds to a different-sized shock at each point in time which in turn may affect the scale and shape of the impulse response functions. This fact complicates the analysis of the dynamic effects of an oil supply shock on the macroeconomy as we move across time. In order to compare the economic consequences across episodes, it is necessary to establish a benchmark scenario against which the time-varying responses can be assessed. Since oil supply shocks are characterized by the joint response of oil production and the real price of oil, both variables are suitable candidates for such a benchmark scenario. Granting that in a time-varying parameter model one cannot make the nature of an oil supply shock identical over time, one can consider the less ambitious thought experiment of making these shocks comparable along one of those dimensions at a time.

Normalization on oil quantity. For example, the oil supply shock can be normalized on the quantity of oil supplied which helps relate our analysis to the previous literature on oil supply shocks which has focused on physical disruptions in the production of crude oil (see Hamilton 2003; Kilian 2008b,c). The dynamic effects of exogenous oil supply shocks normalized such that they correspond to a 1 percent decrease in global oil production on impact at each point in time are shown in panel A of Figure 1 for the median responses together with the 16th and 84th percentiles of the posterior distribution. We plot the reaction of the real price of oil for the quarter in which the shock occurs. For the macroeconomic variables instead, we depict the responses four quarters after the shock given that the greatest effect on real GDP is expected to occur with a delay of about one year (see Hamilton 2008). This convention is adopted throughout the paper. The estimated responses have been accumulated and are shown in levels. The response of the real price of crude oil with respect to a 1 percent shortfall in world oil production increases substantially over time, from an average value of 3 percent in the 1970s and 1980s to 8 percent in the 1990s up to 15 percent in the 2000s with a spike of 28 percent in 2008. The oil price increases triggered by a given reduction in oil production are in turn more disruptive to the economy in the second part of the sample. The accumulated loss in real GDP growth is about twice as big in the 1990s and almost three times as big in the 2000s as in the 1970s. The responses of

consumer prices get more pronounced from the mid-1990s onwards and climb considerably in the 2000s. This evidence underscores the importance of allowing for time variation in studying the effects of physical oil production shortfalls.

Normalization on oil price. Given that the focus of much previous research has been on the effects of an unanticipated increase in the price of oil, we now consider alternatively the effect of an oil supply shock normalized such that it raises the real price of oil by 10 percent on impact at each point in time. The latter thought experiment is used by Blanchard and Galí (2010) as a benchmark for their intertemporal comparison. The normalized time-varying impulse responses are shown in panel B of Figure 1. For this scenario, we find a more muted reaction of economic activity in the latter part of the sample. This finding complies well with existing empirical evidence on the time-varying effects of oil price shocks (e.g. Edelstein and Kilian 2007; Herrera and Pesavento 2009; Blanchard and Galí 2010; Ramey and Vine 2012). This experiment shows that a 10 percent rise in the real price of oil is currently accompanied by an oil production shortfall of 0.5 to 1.5 percent. To elicit the same oil price move in the 1970s, a decline in the physical supply of crude oil of up to 12 percent is required. Thus, despite the assertion by Blanchard and Galí (2010) that "what matters [...] to any given country is not the level of global oil production, but the price at which firms and households can purchase oil" (p. 17), it appears that a larger reduction in oil production causes more severe macroeconomic consequences even if the oil price increases by the same amount.⁴

How can we interpret the decrease in the response of global oil production to a given rise in the real price of oil over time? The simple supply-and-demand diagram of the oil market displayed in Figure 2 illustrates the implications of this finding. A shift of the oil supply curve along a given demand curve implies that the ratio of the quantity response over the price response is invariant to the extent to which the oil supply curve is shifted

⁴To gain a better idea about the extent of time variation over the sample period, we also examined the joint posterior distribution of impulse responses for all models across selected pairs of representative dates following the approach proposed by Cogley, Primiceri, and Sargent (2010). The results of these bilateral diagnostics, which are relegated to the online appendix, largely confirm the pattern of time variation described for both normalizations.

exogenously, i.e., the ratio $\Delta Q/\Delta P$ is the same for moving from S_1 to S_2 and from S_1 to S'_2 as shown in panel A. This means that a change in the size of oil supply shocks *alone* cannot explain the fact that the impact responses of world oil production decrease from Q_2 to Q'_2 over time for a given increase in the real price of crude oil. This result can only be explained by a steepening of the oil demand curve as illustrated in panel B. To obtain an equilibrium outcome characterized by the intersection of the oil supply and demand curves at P_2 and Q'_2 , it is not enough for the supply curve to shift by less, but we need the slope of the demand curve to steepen. In particular, the intersection of S'_2 with D is not consistent with the equilibrium solution.

While we cannot exclude the possibility that the variability of the underlying oil supply shocks has changed over time, only a steepening of the oil demand curve which is equivalent to a decline in the short-run price elasticity of oil demand over time can reconcile the fact that the quantity response is smaller for a given price increase. This is indeed what we find in the data. While the average value of the price elasticity is around -0.6 in the early part of the sample with the exception of the 1979/80 episode, that elasticity declines considerably starting in the mid-eighties and reaches a low of -0.1 towards the end of the sample. In other words, it now takes a smaller exogenous reduction in world oil production to push up the price of oil by the same amount.

In sum, we can reproduce the findings of other researchers that a given supply-driven oil price increase is associated with a decreasing effect on real economic activity later in the sample, and we show that this finding is fully consistent with the result that a given reduction in oil production leads to a greater decline in real GDP over time, once we take the steepening of the oil demand curve into account.

3.2 Quantitative importance of oil supply shocks

Impulse responses are only informative of the transmission of a one-time oil supply shock but do not tell us how important oil supply shocks have been on average nor do they tell us how much of the historical variation in oil market and macroeconomic variables is due to oil supply shocks. To shed some light on these questions, we now examine the forecast error

variance decompositions and historical decompositions based on our TVP-VAR model.

Variance decomposition. Figure 4 displays the evolution of the median of the contribution of oil supply shocks to the forecast error variance after 20 quarters, along with the 16th and 84th percentiles of the posterior distribution. The contribution of oil supply disturbances to the variance of U.S. real GDP growth and CPI inflation consistently ranges between 15 and 20 percent. The share of output volatility attributable to oil supply shocks oscillates moderately over time, whereas the fraction of movements in consumer price inflation induced by oil supply shocks gradually increased since the early 2000s. The latter finding is not surprising given that the volatility of consumer price inflation has decreased over time, while the impact of oil supply shocks on inflation has increased slightly. We conclude that oil supply shocks are still relevant for macroeconomic fluctuations.

The fraction of the variance of global crude oil production growth explained by oil supply shocks fluctuates between 25 and 35 percent in the early part of the sample, but has stabilized since the early 1990s at around 30 percent with the exception of a sharp drop in 2008. Figure 4 also shows that the contribution of oil supply shocks to the variability of changes in the real price of oil declines from around 30 to 35 percent in the first half of the sample to around 20 to 25 percent in the second half.

These estimates indicate that oil supply shocks have become a less important source of oil price movements in recent years. Although we do not identify specific oil demand shocks, we can view all shocks other than the oil supply shock collectively as an oil demand shock. While the contribution of these oil demand shocks to the variability in world oil production has remained relatively stable over time, they are responsible for a substantial share of the volatility in the real oil price, which has increased notably since the early 1990s. This evidence is consistent with the empirical results of Kilian (2008a, 2009a), reflected in the conventional wisdom that "demand increases rather than supply reductions have been the primary factor driving oil prices over the last several years" (Hamilton 2008, p. 175).

Historical decomposition. The historical contribution of oil supply disturbances to the four endogenous variables is presented in Figure 4. The dashed line shows the actual time

series relative to its average growth rate and the solid line shows the cumulative effect of oil supply shocks on the evolution of each variable, while turning off all other shocks. In other words, the historical decomposition shows how the series in question would have evolved had only oil supply shocks occurred. This implies that the difference between the actual data and the contribution of oil supply shocks is driven by the composite of oil demand shocks.

With regard to inflation, Figure 4 reveals that oil supply shocks explain little of the Great Inflation. Despite the fact that there is some contribution, the bulk of the inflation in the 1970s and early 1980s is explained by other shocks. While this insight is apparently in contrast with popular perception, it is consistent with related evidence in the literature. In particular, Barsky and Kilian (2002, 2004) showed that shifts in monetary policy regimes associated with demand shifts in the oil market were the source of the stagflationary experience of the 1970s.

The contribution of oil supply shocks to the evolution of the real price of crude oil and of real economic activity changes from episode to episode. We find that the contribution of oil supply shocks to the course of the real price of oil during the events of 1978-80 is relatively limited which indicates that this was primarily a demand-driven oil price surge reflecting rising oil demand at a time of low spare capacity. This finding confirms the conjecture in Barsky and Kilian (2002) who argue that oil supply shocks were never the sole driving force behind the fluctuations in the real price of crude oil in 1979. It also matches the empirical evidence in Kilian (2009a). Similarly, oil supply shocks only made a small contribution during the oil price collapse in 1986. They played a somewhat bigger role in 1990 and 1999. While a substantial share of the oil price spike after Iraq's invasion of Kuwait in August 1990 can be attributed to oil supply disruptions, the subsequent downward spike in the price of oil cannot be explained by oil supply shocks. Only a fraction of the oil price movements since 2002 can be explained by oil supply shocks. This implies that the sustained oil price surge was to a large extent driven by oil demand shocks which is consistent with Kilian (2009a) and helps explain why these shocks were not accompanied by a major recession in the U.S. economy.

Even during earlier episodes the recessionary effects of oil supply shocks have been weak.

For example, the contribution of oil supply shocks to the economic downturn of 1980 was quite modest and there was virtually no contribution of oil supply shocks to the recessions of 1982 and 2001. Likewise, the oil supply shock of 1990 only accounted for about one-fifth of the decline in real GDP growth during the ensuing recession. In contrast, the negative oil supply shocks in 1999, when OPEC and non-OPEC countries jointly decided to cut oil production, slowed the ongoing economic boom as a result of the oil price increase. Again, it is noteworthy that notwithstanding the striking differences in impulse response dynamics over time, the historical decomposition for real GDP growth is broadly similar to results from a time-invariant model in Kilian (2009b).

3.3 Sensitivity analysis

Alternative models of time variation. In the benchmark model we postulated a drifting coefficient model. We now illustrate the importance of allowing for smooth time variation as opposed to a one-time break in 1986. Figure 5, panel A displays the median responses of the four endogenous variables after an oil supply shock identified with sign restrictions and normalized to a 1% decrease in oil production in a fixed-coefficient VAR model estimated over the two subperiods 1974Q1 – 1985Q4 and 1986Q1 – 2011Q1, together with the 16th and 84th percentiles. It turns out that the response of the real price of oil is much larger in the second subsample, reaffirming our substantive finding of a less price-elastic oil demand curve. In contrast to our baseline results, however, there is no compelling evidence for time variation in the responses of real GDP and consumer prices across subsamples given that the posterior intervals overlap. This analysis highlights that one would not have been able to uncover the same changes in the macroeconomic consequences of oil supply shocks by imposing a one-time break in the oil-macro relationship.

Alternative identifying assumptions. Our second methodological contribution was the identification of oil supply shocks based on sign restrictions. An alternative assumption would have been to impose a vertical short-run oil supply curve which implies three exclusion restrictions on the first row of the structural impact multiplier matrix. This corresponds

to a simplified version of the monthly oil market VAR model proposed by Kilian (2009a). Although these identifying restrictions were not intended for quarterly data, it is useful to compare the results to our baseline model. Unlike Kilian (2009a), we implement this procedure allowing for time variation in the parameters. Figure 5, panel B presents the time-varying median responses of the real price of oil and of the macroeconomic aggregates to an oil supply shock identified by exclusion restrictions to a 1% oil production shortfall, together with the 16th and 84th percentiles. Although there is some evidence of time variation in the response of the real price of oil, the posterior intervals are so wide to leave open the possibility that the responses remained unchanged. In contrast, the time-varying estimates in our baseline model cannot be explained by estimation uncertainty only. Moreover, a puzzling finding in Figure 5, panel B is that the median real GDP response to a negative oil supply shock is positive, especially toward the end of the sample. This counterintuitive result suggests that the recursive identification scheme is inappropriate at the quarterly frequency and highlights the potential benefits of using sign restrictions.

Other model specifications. We also investigated the sensitivity of our results to changes in the variables included in the benchmark model. When U.S. unemployment instead of real GDP is used as the indicator of real economic activity, we find a substantial rise in unemployment following a negative oil supply shock again normalized on oil production. The strength of this response increases notably in the most recent past. When we replace the consumer price index with the implicit GDP deflator as a measure of inflation, we observe a somewhat more subdued increase in the price level after a negative oil supply shock that corresponds to a 1% oil production shortfall, but a similar pattern of time variation emerges. Likewise, using different oil price measures such as the real refiners' acquisition cost of composite crude oil or the West Texas Intermediate spot oil price does not affect our conclusions. Finally, augmenting our model by the federal funds rate, as is common in the literature on monetary policy, does not change our findings about the dynamic response of the U.S. economy to oil supply shocks and the structural change in the crude oil market.

Timing of the sign restrictions. Since in the early part of our sample the nominal oil price was constrained by long-term bilateral agreements which were subject to revision only periodically, an obvious concern arises with regard to the timing of the sign restrictions imposed in the baseline model. When nominal oil prices are sluggish, a positive demand shock in the U.S. economy, which raises world oil production and the consumer price level, causes a fall in the real price of oil, unless nominal contracts are renegotiated timely to reflect the new macroeconomic conditions. The resulting opposite movement in world oil production and in the real price of oil would imply that this demand shock is erroneously identified as a positive oil supply shock. One way of addressing this problem is to impose that the sign restrictions are only binding from the fourth quarter after the shock onwards such that the immediate reaction of oil production and of the real price of oil is unconstrained. It can be shown that our findings are not sensitive to this change in the timing of the identification restrictions. For further details the reader is referred to the online appendix.

4 Reasons for the steepening of the oil demand curve

In this section, we consider developments in the economy and in the crude oil market that help explain the substantial reduction of the short-run price elasticity of oil demand we documented earlier. These developments include a decline in the oil intensity of U.S. real output, fuel switching, shifts in the composition of total oil demand, and variation in the global spare capacity of oil production. We do not preclude that other factors contributed to the steepening of the oil demand curve.

Energy efficiency and sectoral shifts. Following the oil price surges of the 1970s, the extent to which the U.S. economy and other industrialized countries rely on oil has changed substantially since the mid-1980s. Industries gradually switched away from oil to alternative sources of energy, developed more energy-efficient technologies and improved energy conservation. These efforts were supported by government policies aimed at reducing oil usage and increasing energy awareness. This transition together with service-biased growth in the U.S. caused oil intensity in aggregate economic activity to fall steadily, reflected in a reduction

in the use of oil input per unit of output as shown in Figure 6, panel A. This development helps account for the decline of the short-run price elasticity of oil demand (see, e.g., Ryan and Plourde 2002; Hughes, Knittel, and Sperling 2008; Dargay and Gately 2010).⁵

In addition, the composition of total oil demand has changed over time with oil consumption now being concentrated in sectors such as transportation, which traditionally were characterized by a low own-price elasticity of demand due to the lack of substitutes for transportation fuels. The increased share of transportation in total oil demand has also contributed to a steepening of the oil demand curve (see Dargay and Gately 2010; Ramey and Vine 2012).

Finally, industrialized countries have increasingly outsourced their industrial production to emerging economies. Emerging economies typically rely heavily on oil as an input factor. Their oil demand tends to be less sensitive to changes in oil prices than advanced economies. In addition, it has been suggested that governments in developing economies have used fuel subsidies to shield consumers from the impact of rising global oil prices thereby stimulating oil consumption and making consumer demand unresponsive to international price signals.⁶ Given the larger share of emerging economies in global oil demand over time, this observation is consistent with the view that the elasticity of oil demand in global markets has declined.

Capacity constraints in crude oil production. Yet another likely explanation is that in the presence of capacity constraints in crude oil production, the composition of oil demand is likely to change in favor of less elastic speculative or precautionary buying. At times of low spare capacity, even small supply disruptions can lead to large price increases because market participants anticipate that an unexpected loss in oil production cannot be replaced by other oil producers. In that sense, increasing capacity utilization signals some tightness in the market which increases the willingness of oil consumers to pay a higher price for a

⁵Dargay and Gately (1994) attribute this phenomenon to the irreversibility of technical knowledge, the durability of efficiency-improving investments and the non-abrogation of laws regarding energy-cost labeling and energy-efficiency standards.

⁶"Both wholesale and retail prices of oil products in the domestic market are lower than they are in the global market" as exemplified for China by Hang and Tu (2007, p. 2978). An estimate by Morgan Stanley shows that almost a quarter of the world's petrol is sold at less than the market price (The Economist, 2008).

barrel of oil at the margin that provides insurance against potential scarcity, i.e., they pay an insurance premium (see Alquist and Kilian 2010).⁷ Kilian (2008b) has documented that annual global capacity utilization rates in crude oil production have been steadily increasing from the mid-eighties to the early 1990s and stayed at levels close to full capacity ever since (see Figure 6, panel B). This observation is in line with the gradual decline in the short-run price elasticity of oil demand uncovered in our empirical analysis.

5 Conclusions

In this paper, we have analyzed the time-varying effects of oil supply shocks on the U.S. economy and the oil market from 1974 onwards. There are several a priori reasons to expect this relationship to have evolved over time. For example, the transition from a regime of administered oil prices to a market-based system, changing capacity utilization rates in crude oil production, and changes in the energy dependence of the U.S. economy all have implications for the effects of oil supply shocks on oil market variables and U.S. macroeconomic aggregates.

Our analysis combines a novel identification strategy for oil supply shocks based on inequality constraints with the estimation of a time-varying parameter VAR model. The first generation of structural oil market models has relied on exclusion restrictions on the impact multiplier matrix that can be interpreted as a vertical short-run oil supply curve. Instead, we propose to identify oil supply shocks based on sign restrictions derived from a simple supply and demand model of the crude oil market. Specifically, we identify an oil supply shock as a disturbance in the global oil market which moves oil production and

⁷As has been noted by Gately (1984, p. 1103), "aggravating the market tightness was an extended period of aggressive stockbuilding by the importing countries for much of 1979 and 1980. Such a stockbuilding "scramble" during a disruption was certainly perverse. It undoubtedly drove the price higher than it would have gone otherwise." This aggressive hoarding behavior could hint at the increased importance of less elastic precautionary demand in total oil demand in a tightening market. In fact, Adelman (2002, p. 179) states that "when buyers fear damage from sudden dearth, there is also a precautionary motive; which may be joined to a speculative motive, to profit by buying sooner."

the real price of oil in opposite directions. This approach is particularly appealing when quarterly data are used, since traditional delay restrictions are credible only at the monthly frequency.

Until now, time variation in oil markets has been analyzed by splitting the sample or by estimating bivariate VARs on rolling windows. The first approach is not appealing when dealing with smooth structural change. The second approach does not allow the identification of structural oil supply shocks because of the degrees-of-freedom problem. We dealt with these challenges by estimating a multivariate structural VAR with time-varying parameters and stochastic volatility in the spirit of Cogley and Sargent (2005), Primiceri (2005), and Benati and Mumtaz (2007), among others.

We showed that when an exogenous oil supply shock in the time-varying parameter model is normalized to hold constant across time the implied change in the real price of oil, the response of real GDP declines over time, which is consistent with other recent evidence. When normalizing the shock to hold constant across time the implied change in global oil production, however, the response of U.S. output and inflation have become larger in magnitude in more recent years. We showed that these two findings can be reconciled by a decline in the short-run price elasticity of oil demand such that a given shortfall in oil production is associated with a greater price response. Indeed, we found that the oil demand curve is currently much steeper relative to the 1970s and early 1980s.

We further showed that the share of the volatility of the real price of oil explained by oil supply shocks has moderately decreased over time, indicating that oil supply shocks are not the primary driver of oil price movements in more recent periods. The contribution of oil supply shocks to the variability of real activity and inflation is economically relevant, ranging fairly steadily between 15 and 20 percent with the exception of 2008. Our analysis also reaffirmed that even allowing for time variation, the Great Inflation of the 1970s and early 1980s cannot be accounted for by negative oil supply shocks. Nor have the recessionary effects of oil supply shocks been very pronounced. We found that oil supply disruptions mattered for real economic activity mainly during two episodes. They contributed to the 1991 recession and they slowed the ongoing boom at the end of the millennium. There is

little or no evidence that oil supply shocks mattered much for the recessions of the early 1980s and for the downturns in 2001 and in 2008, in contrast.

Our analysis adds to a growing literature on endogenous oil prices and their implications for the macroeconomy, including Baumeister and Peersman (2011). One key difference between our analysis in this paper and that in the other paper is that here we focused on studying and understanding the responses of U.S. macroeconomic aggregates, whereas Baumeister and Peersman (2011) focus on explaining oil price dynamics. A second key difference is that in this paper we were solely concerned with the effects of oil supply shocks without taking a stand on the difficult problem of disentangling different oil demand shocks. Indeed, trying to address both of these problems at the same time in a TVP-VAR framework would cause a degrees-of-freedom problem. Thus, the analysis in these two papers is complementary.

References

- [1] **Adelman, M. A.** 2002. "World Oil Production & Prices 1947-2000." *The Quarterly Review of Economics and Finance*, 42(2): 169-91.
- [2] **Alquist, Ron, and Lutz Kilian.** 2010. "What Do We Learn from the Price of Crude Oil Futures?" *Journal of Applied Econometrics*, 25(4): 539-73.
- [3] **Alquist Ron, Lutz Kilian, and Robert Vigfusson.** 2011. "Forecasting the Price of Oil." In *Handbook of Economic Forecasting*, 2nd ed., ed. Graham Elliott G and Allan Timmermann. Amsterdam, North-Holland.
- [4] **Barsky, Robert B., and Lutz Kilian.** 2002. "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative." In *NBER Macroeconomics Annual 2001*, ed. Ben S. Bernanke and Kenneth Rogoff, 137-83. Cambridge, MA: MIT Press.
- [5] **Barsky, Robert B., and Lutz Kilian.** 2004. "Oil and the Macroeconomy since the 1970s." *Journal of Economic Perspectives*, 18(4): 115-34.
- [6] **Baumeister, Christiane, and Gert Peersman.** 2011. "The Role of Time-Varying Price Elasticities in Accounting for Volatility Changes in the Crude Oil Market." Bank of Canada Working Paper 2011-28.
- [7] **Benati, Luca, and Haroon Mumtaz.** 2007. "U.S. Evolving Macroeconomic Dynamics: A Structural Investigation." ECB Working Paper 746 .
- [8] **Blanchard, Olivier J., and Jordi Galí.** 2010. "The Macroeconomic Effects of Oil Price Shocks: Why are the 2000s so different from the 1970s?" In *International Dimensions of Monetary Policy*, ed. Jordi Galí and Mark J. Gertler, 373-421.
- [9] **Canova, Fabio, and Gianni De Nicoló.** 2002. "Monetary Disturbances Matter for Business Fluctuations in the G-7." *Journal of Monetary Economics*, 49(6): 1131-59.
- [10] **Canova, Fabio, and Luca Gambetti.** 2009. "Structural Changes in the U.S. Economy: Is There a Role for Monetary Policy?" *Journal of Economic Dynamics and Control*, 33(2): 477-90.

- [11] **Cogley, Timothy, Giorgio E. Primiceri, and Thomas J. Sargent.** 2010. "Inflation-Gap Persistence in the US." *American Economic Journal: Macroeconomics*, 2(1): 43-69.
- [12] **Cogley, Timothy, and Thomas J. Sargent.** 2005. "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII US." *Review of Economic Dynamics*, 8(2): 262-302.
- [13] **Dargay, Joyce M., and Dermot Gately.** 1994. "Oil Demand in the Industrialized Countries." *Energy Journal*, 15 (Special Issue: The Changing World Petroleum Market): 39-68.
- [14] **Dargay, Joyce M., and Dermot Gately.** 2010. "World Oil Demand's Shift toward Faster Growing and Less Price-Responsive Products and Regions." *Energy Policy*, 38: 6261-77.
- [15] **Davis, Steven J., and John Haltiwanger.** 1999. "On the Driving Forces behind Cyclical Movements in Employment and Job Reallocation." *American Economic Review*, 89(5): 1234-58.
- [16] **Edelstein, Paul, and Lutz Kilian.** 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) (Contributions): Article 35.
- [17] **Edelstein, Paul, and Lutz Kilian.** 2009. "How Sensitive are Consumer Expenditures to Retail Energy Prices?" *Journal of Monetary Economics*, 56(6): 766-79.
- [18] **Faust, Jon.** 1998. "The Robustness of Identified VAR Conclusions about Money." *Carnegie-Rochester Conference Series on Public Policy*, 49(1): 207-44.
- [19] **Fry, Renée, and Adrian Pagan.** 2011. "Sign Restrictions in Structural Vector Autoregressions: A Critical Review." *Journal of Economic Literature*, forthcoming.

- [20] **Gately, Dermot.** 1984. "A Ten-Year Retrospective: OPEC and the World Oil Market." *Journal of Economic Literature*, 22(3): 1100-14.
- [21] **Hamilton, James D.** 1983. "Oil and the Macroeconomy Since World War II." *Journal of Political Economy*, 91(2): 228-48.
- [22] **Hamilton, James D.** 2003. "What is an Oil Shock?" *Journal of Econometrics*, 113(2): 363-98.
- [23] **Hamilton, James D.** 2008. "Oil and the Macroeconomy." In *The New Palgrave Dictionary of Economics*, 2nd ed., ed. Steven N. Durlauf and Lawrence Blume, Palgrave Macmillan.
- [24] **Hamilton, James D.** 2009a. "Understanding Crude Oil Prices." *Energy Journal*, 30(2): 179-206.
- [25] **Hamilton, James D.** 2009b. "Causes and Consequences of the Oil Shock of 2007-08." *Brookings Papers on Economic Activity*, 40(1): 215-261.
- [26] **Hamilton, James D. and Ana María Herrera.** 2004. "Comment: Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy." *Journal of Money, Credit, and Banking*, 36(2): 265-86.
- [27] **Hang, Leiming, and Meizeng Tu.** 2007. "The Impacts of Energy Prices on Energy Intensity: Evidence from China." *Energy Policy*, 35(5): 2978-88.
- [28] **Herrera, Ana María, and Elena Pesavento.** 2009. "Oil Price Shocks, Systematic Monetary Policy, and the 'Great Moderation'." *Macroeconomic Dynamics*, 13(1): 107-37.
- [29] **Hubbard, R. Glenn.** 1986. "Supply Shocks and Price Adjustment in the World Oil Market." *Quarterly Journal of Economics*, 101(1): 85-102.
- [30] **Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling.** 2008. "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand." *Energy Journal*, 29: 93-114.

- [31] **Kilian, Lutz.** 2008a. "The Economic Effects of Energy Price Shocks." *Journal of Economic Literature*, 46(4): 871-909.
- [32] **Kilian, Lutz.** 2008b. "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?" *Review of Economics and Statistics*, 90(2): 216-40.
- [33] **Kilian, Lutz.** 2008c. "A Comparison of the Effects of Exogenous Oil Supply Shocks on Output and Inflation in the G7 Countries." *Journal of the European Economic Association*, 6(1): 78-121.
- [34] **Kilian, Lutz.** 2009a. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review*, 99(3): 1053-69.
- [35] **Kilian, Lutz.** 2009b. "Comments on 'Causes and Consequences of the Oil Shock of 2007-08' by James D. Hamilton." *Brookings Papers on Economic Activity*, 40(1): 267-78.
- [36] **Kilian, Lutz, and Daniel P. Murphy.** 2011. "Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models." *Journal of the European Economic Association*, forthcoming.
- [37] **Kilian, Lutz, and Daniel P. Murphy.** 2010. "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil." Mimeo, University of Michigan.
- [38] **Kilian, Lutz, and Clara Vega.** 2011. "Do Energy Prices Respond to U.S. Macroeconomic News? A Test of the Hypothesis of Predetermined Energy Prices." *Review of Economics and Statistics*, 93(2): 660-71.
- [39] **Kim, Chang-Jin, and Charles R. Nelson.** 1999. *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*. Cambridge, MA: MIT Press.
- [40] **Lippi, Francesco, and Andrea Nobili.** 2011. "Oil and the Macroeconomy: A Quantitative Structural Analysis." *Journal of the European Economic Association*, forthcoming.

- [41] **Peersman, Gert.** 2005. "What Caused the Early Millennium Slowdown? Evidence Based on Vector Autoregressions." *Journal of Applied Econometrics*, 20(2): 185-207.
- [42] **Peersman, Gert, and Ine Van Robays.** 2009. "Oil and the Euro Area Economy." *Economic Policy*, 24(60): 603-51.
- [43] **Peersman, Gert, and Ine Van Robays.** 2011. "Cross-Country Differences in the Effects of Oil Shocks." *Energy Economics*, forthcoming.
- [44] **Primiceri, Giorgio E.** 2005. "Time Varying Structural Vector Autoregressions and Monetary Policy." *Review of Economic Studies*, 72(3): 821-52.
- [45] **Ramey, Valerie A., and Daniel J. Vine.** 2012. "Oil, Automobiles, and the U.S. Economy: How Much Have Things Really Changed?" In *NBER Macroeconomics Annual 2011*, 333-68. Cambridge, MA: MIT Press.
- [46] **Rotemberg, Julio J.** 2010. "Comment on Blanchard-Galí: The Macroeconomic Effects of Oil Price Shocks: Why Are the 2000s so Different from the 1970s?" In *International Dimensions of Monetary Policy*, ed. Jordi Galí and Mark J. Gertler, 422-428.
- [47] **Ryan, David L., and André Plourde.** 2002. "Smaller and Smaller? The Price Responsiveness of Nontransport Oil Demand." *Quarterly Review of Economics and Finance*, 42(2): 285-317.
- [48] "Crude Measures." *The Economist*, 2008, May 31, 75.
- [49] **Uhlig, Harald.** 2005. "What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure." *Journal of Monetary Economics*, 52(2): 381-419.




Figure 1: Median impact responses of the two oil market variables and median responses of macroeconomic variables four quarters after the shock to a negative oil supply shock where shaded areas indicate 68% posterior credible sets.
 Panel A: Oil supply shock normalized to a 1% decrease in world oil production.
 Panel B: Oil supply shock normalized to a 10% increase in the real price of crude oil.




Figure 2: Structural changes in the global crude oil market. Panel A: Change in the volatility of oil supply shocks. Panel B: Steeper oil demand curve.




Figure 3: Median of the contribution of oil supply shocks to the forecast error variance of all four endogenous variables after 20 quarters with 16th and 84th percentiles of the posterior distribution.




Figure 4: Historical decomposition of world oil production growth, changes in the real price of oil, real GDP growth and CPI inflation.

Note: The vertical lines indicate major events in the crude oil market, in particular the outbreak of the Iranian revolution in 1978Q3 and of the Iran-Iraq war in 1980Q3, the collapse of OPEC in 1985Q4, the invasion of Kuwait in 1990Q3, the coordinated supply cut of OPEC and non-OPEC countries in 1999Q1, and the Iraq war in 2003Q1. The grey bars indicate NBER recessions.

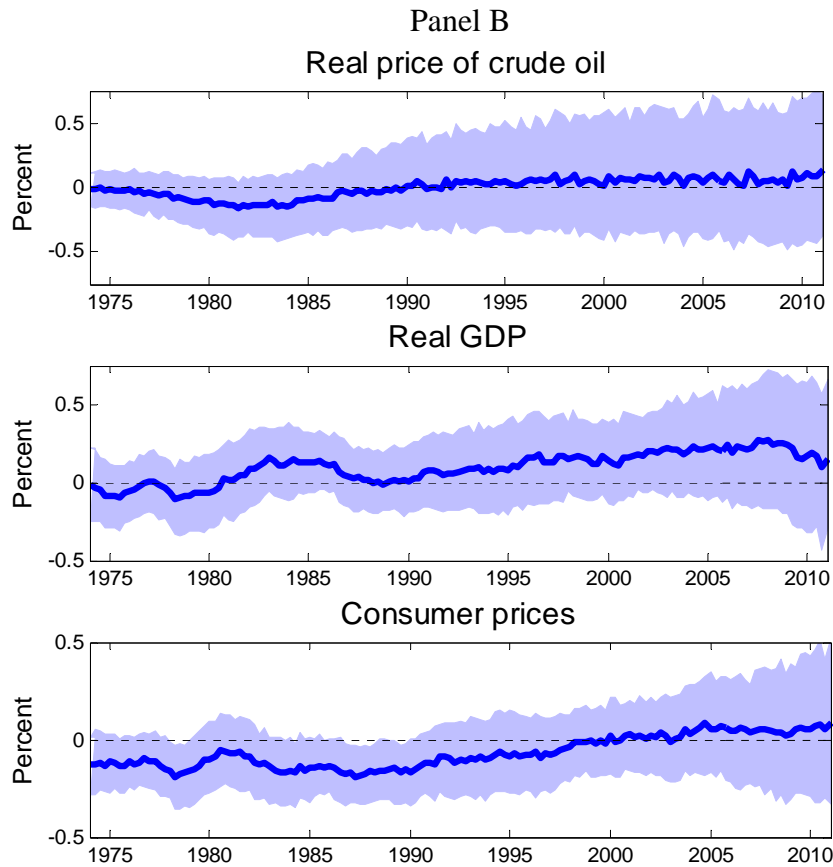
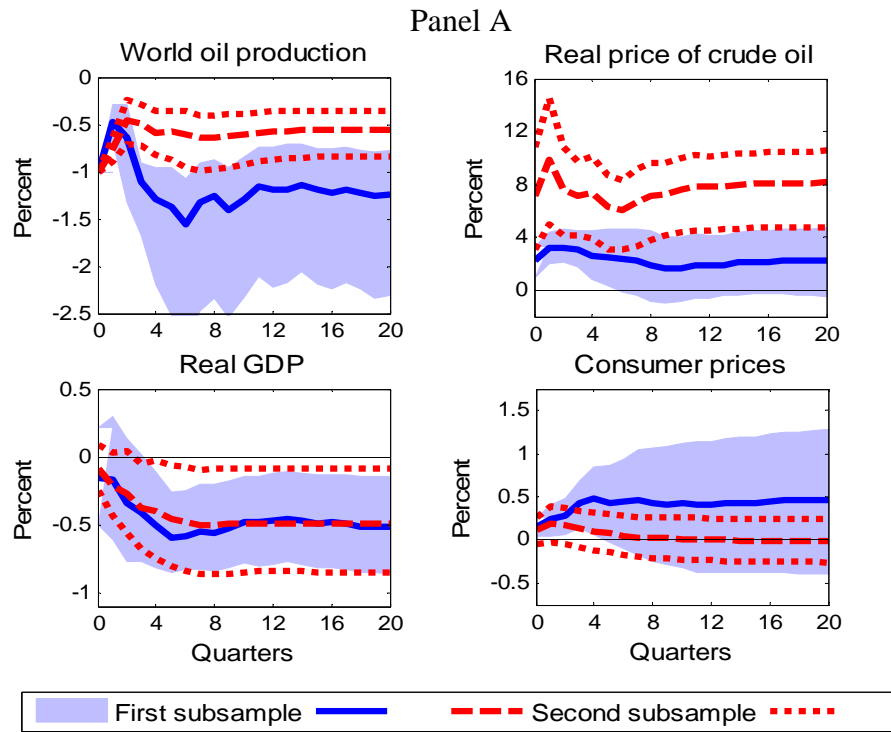



Figure 5: Posterior median responses after an oil supply shock normalized to a 1% decrease in oil production where shaded areas indicate 68% posterior credible sets.
 Panel A: Identification based on sign restrictions in a fixed-coefficient VAR estimated over two subsamples, 1974Q1-1985Q4 and 1986Q1-2011Q1.
 Panel B: Identification based on exclusion restrictions in a time-varying VAR.

Panel A



Panel B




Figure 6: Panel A: U.S. oil intensity of production by year.

Panel B: Global capacity utilization rates in crude oil production by year.

Note: U.S. oil intensity of production has been computed as the annual petroleum consumption in British Thermal Units (BTU) by the industrial, commercial, and part of the transportation sector per dollar of U.S. real GDP. The data sources are discussed in the online appendix.