

CAMA

Centre for Applied Macroeconomic Analysis

Inflation Dynamics and The Role of Oil Shocks: How Different Were the 1970s?

CAMA Working Paper 59/2013
August 2013

Benjamin Wong

Centre for Applied Macroeconomic Analysis (CAMA), ANU

Abstract

This paper presents evidence on why inflation pass-through from oil shocks in the 21st century relative to the 1970s has dampened. First, results suggest global business cycle demand driven oil shocks are not inflationary. Second, there has been a reduction in inflation pass-through from oil supply and speculative oil demand shocks. Movements in oil inventories and production suggest oil supply and speculative oil demand shocks in the 1970s were different. Oil market participants expect higher oil prices to persist into the future. The analysis highlights the importance of modelling inventories as a means of capturing expectations in the oil market.

Keywords

Oil Shocks, Time-Varying Parameters, Inflation Pass-through

JEL Classification

C11, C32, E31, Q43

Address for correspondence:

(E) cama.admin@anu.edu.au

[The Centre for Applied Macroeconomic Analysis](#) in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

Inflation Dynamics and The Role of Oil Shocks: How Different Were the 1970s?

Benjamin Wong ^{*†}

Abstract

This paper presents evidence on why inflation pass-through from oil shocks in the 21st century relative to the 1970s has dampened. First, results suggest global business cycle demand driven oil shocks are not inflationary. Second, there has been a reduction in inflation pass-through from oil supply and speculative oil demand shocks. Movements in oil inventories and production suggest oil supply and speculative oil demand shocks in the 1970s were different. Oil market participants expect higher oil prices to persist into the future. The analysis highlights the importance of modelling inventories as a means of capturing expectations in the oil market.

Keywords: Oil Shocks, Time-Varying Parameters, Inflation Pass-through

JEL Classification: C11, C32, E31, Q43

^{*}Centre for Applied Macroeconomic Analysis, The Australian National University. Canberra, ACT 0200, Australia. Email: ben.wong@anu.edu.au. Tel: (+61) 2 612 51833.

[†]I thank Renée Fry-McKibbin, Warwick McKibbin and Shaun Vahey for their guidance and advice, Christiane Baumeister, Hilde Bjørnland and Lutz Kilian for extensive suggestions and discussions as well as Jean Pierre Allégret, Andrew Binning, Joshua Chan, Soojin Jo, Timothy Kam, Francesco Ravazzolo, Tovoony Razafindrabe, Tara Sinclair, Rodney Strachan, Leif Anders Thorsrud, Varang Wiriyawit, Sen Xue, seminar participants at The Australian National University, BI Norwegian Business School, Reserve Bank of New Zealand, the 21st Symposium for the Society of Nonlinear Dynamics and Econometrics and the 2nd PhD conference in International Macroeconomics and Financial Econometrics for helpful discussions and comments. I gratefully acknowledge Gary Koop, Dimitris Korobilis and Andreas Tudyka for code which I adapted. This research is supported by the Australian Government's International Postgraduate Research Scholarship and ARC Discovery Project DP120103443. Errors and omissions remain my own.

1 Introduction

How different were inflation outcomes from oil shocks in the 1970s and thereafter? Figure 1 juxtaposes quarter on quarter changes in the real oil price against U.S. core Personal Consumption Expenditure (PCE) inflation since 1970.¹ An increase in the change in the real oil price appears to accompany a similar large rise in inflation in the 1970s, though this observation appears to break down thereafter. Structural change stands out as an obvious candidate explaining this dichotomy in the U.S. oil price-inflation relationship since the 1970s. While monetary policy is an evident possible source of this structural change, there has also been a large structural shift in the oil market since the collapse of OPEC in the mid-1980s. However, it is also true that the types of oil shocks emanating from the oil market are very different today compared to the 1970s. While oil supply disruptions were certainly a feature of the oil shocks in the 1970s (e.g. OPEC, Yom Kippur War, Iranian Revolution etc.), they are almost conspicuous by their absence with oil price increases thereafter.

The main contribution of this paper is to model different oil shocks in a Time-Varying Parameter environment. This allows, in a single model framework, both the consideration of structural change and the identification of different types of oil shocks in assessing the source of this breakdown. In doing so, the analysis seeks to answer whether the breakdown in the oil-inflation relationship is due to a change in the source of oil shocks (i.e. oil supply disruptions) or whether inflation pass-through from oil shocks has reduced in general?

The paper extends on several important contributions by residing at the intersection of two main strands of the literature. On the one hand, the paper seeks to understand how the oil-inflation relationship has broken down. This is a large literature (see, *inter alia* Hooker, 2002; Blanchard and Galí, 2009; Herrera and Pesavento, 2009; Bachmeier and Cha, 2011). An obvious antecedent of this paper in this area of the literature is Clark and Terry (2010). They use Time-Varying Parameter Vector Autoregression (TVP-VAR) techniques to model this breakdown, a similar approach to that adopted by this paper. Even though Clark and Terry (2010) study energy, and not specifically oil, shocks, their motivation for doing so is similar to this paper.

On the other hand, the main point of departure from the approach of Clark and Terry (2010) is the explicit regard for the source of oil shocks. In this tradition, this paper also resides in an area of the literature which is concerned with identifying different demand and supply shocks in the oil market and viewing them as distinct. The desire for such an approach is apparent. Without a taxonomy of demand and supply shocks in the oil market, one cannot answer whether it was the different types of oil shocks in the 1970s which drove the breakdown of the oil-inflation pass-through. The paper builds on several important contributions in this region of the literature. First, the paper

¹All other inflation series like the Consumption Price Index produce similar patterns. The choice of core PCE is only because it is used for the main empirical analysis later.

extends the analysis of Baumeister and Peersman (2013b), who study the time-varying role of oil supply shocks with regard to U.S. inflation, by also considering oil demand. Like Baumeister and Peersman (2013a,b), the paper is also in line with oil market models which study time variation in the oil market. Second, the paper builds on the identification strategy of Kilian and Murphy (2013). In doing so, the paper also demonstrates how to extend the sign restriction identification strategy of Kilian and Murphy (2013). The Kilian and Murphy (2013) identification strategy is combined with another common and plausible identification strategy which regards the oil price as being pre-determined (Kilian and Vega, 2011). The Kilian and Murphy (2013) model is the latest incarnation of oil market models which consider both oil supply and demand, with a key feature being the explicit identification of speculative oil demand shocks. Speculative demand shocks in the oil market have been hypothesised to have an important role in the oil crisis of 1979/80 (e.g Barsky and Kilian, 2002; Kilian, 2009) and should not be ignored. The Kilian and Murphy (2013) model is thus a useful framework to build on.

Briefly, the results are as follows. Using data from 1974Q1 to 2012Q2, oil supply shocks and oil speculative demand shocks in the oil market have had a more muted impact on U.S. core inflation since the mid-1980s. Therefore, these two types of oil shocks stand out as possible candidates for why oil shocks are no longer inflationary. Movements in oil production and oil inventories hint at participants in the oil market during the 1970s reacting to oil supply and speculative oil demand shocks in a manner which indicates their expectations that a high oil price would persist through a future shortfall in oil supply relative to demand. This demonstrates the key insight of modelling inventories in the Kilian and Murphy (2013) framework in order to capture and understand market expectations. Such behaviour has not been apparent since the OPEC collapse in the mid-1980s, coinciding with a time where oil shocks are no longer inflationary.

The paper proceeds with a description of the empirical specification in section 2. In particular, it offers a treatment of the econometric and economic model underpinning the analysis. Section 3 discusses the results and interpretations of these findings. There will be some concluding remarks thereafter.

2 Empirical Specification

The empirical strategy consists of two main steps. The first employs the use of a TVP-VAR in the spirit of Cogley and Sargent (2001, 2005) and Primiceri (2005) to estimate the reduced form Vector Autoregression (VAR) coefficients. This is a flexible framework which allows for structural changes both in the oil market and the oil-inflation relationship. The second step involves the identification of the different oil shocks. This occurs through a mixture of exclusion and sign restrictions, largely guided by Kilian and Murphy (2013).

2.1 Econometric Model

Consider the following VAR model,

$$y_t = c_t + B_{1,t}y_{t-1} + \dots + B_{p,t}y_{t-p} + \Gamma_t\epsilon_t \quad (1)$$

where y is a 5×1 vector of endogenous variables and c a 5×1 vector of time-varying constants. The B matrices are each a 5×5 matrix of time-varying coefficients with p the maximum lag order of the estimated VAR. Γ_t is of dimension 5×5 , sometimes referred to as the impact matrix in structural VAR models. ϵ_t are serially uncorrelated structural innovations with distribution $N(0, I_5)$.

In state space form, equation (1) is the observation equation. A reduced form time-varying covariance matrix, Ω_t , satisfying $\Omega_t = \Gamma_t'\Gamma_t$ can be constructed using

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

where

$$A_t = \begin{pmatrix} 1 & 0 & \dots & 0 \\ \alpha_{21,t} & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ \alpha_{51,t} & \dots & \alpha_{54,t} & 1 \end{pmatrix}, \Sigma_t = \begin{pmatrix} \sigma_{1,t} & 0 & \dots & 0 \\ 0 & \sigma_{2,t} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \sigma_{5,t} \end{pmatrix} \quad (3)$$

The reduced form covariance matrix Ω has a t subscript, indicating the covariance matrix is time-varying. The impact matrix Γ thus also has a t subscript as it is also time-varying by construction. Identification will recover Γ_t from Ω_t .

Stacking the matrices of time-varying parameters in a vectorised fashion where $B_t = \text{vec}(B_{i,t})$, $i \in \{1, 2, \dots, p\}$, $\alpha_t = \text{vec}(A_t)$, $\forall \alpha_{ij,t} \notin \{0, 1\}$ and $\sigma_t = \text{diag}(\Sigma_t)$, the state equations can be written as

$$B_t = B_{t-1} + v_t \quad (4)$$

$$\alpha_t = \alpha_{t-1} + \zeta_t \quad (5)$$

$$\log\sigma_t = \log\sigma_{t-1} + \eta_t \quad (6)$$

All the time-varying parameters have a random walks structure, similar to Cogley and Sargent (2005) and Primiceri (2005). v , ζ and η are jointly normally distributed. A block diagonal matrix expresses the covariance matrix for the innovations

$$V \begin{pmatrix} \epsilon_t \\ v_t \\ \zeta_t \\ \eta_t \end{pmatrix} = \begin{pmatrix} I_n & 0 & 0 & 0 \\ 0 & R & 0 & 0 \\ 0 & 0 & S & 0 \\ 0 & 0 & 0 & W \end{pmatrix} \quad (7)$$

where R , S and W are all positive definite with the off diagonals set to zero as commonly done. S is also block diagonal where

$$S = V(\zeta_t) = \begin{pmatrix} S_1 & 0 & \cdots & 0 \\ 0 & S_2 & \cdots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & S_5 \end{pmatrix} \quad (8)$$

and $S_i = \mathbb{E}(\zeta_1 \cdots \zeta_i)'(\zeta_1 \cdots \zeta_i)$.

The y vector contains world oil production, the real oil price, OECD oil inventories stock, an index for global economic activity and the U.S. core Personal Consumption Expenditure (PCE) index. All variables are seasonally adjusted. Raw data which are not seasonally adjusted are filtered using quarterly dummies. Both world oil production and the core PCE index enter the VAR as the first difference of the log-level. The inventories stock enters as a level difference specification because it is needed to explicitly calculate the oil demand elasticity in use for identification. The variable set is largely an augmentation of Kilian and Murphy's (2013) four variable VAR with the addition of a price index.

All the data except the global economic activity index are sourced or constructed from the Energy Information Agency (EIA) or the FRED database. The global economic activity index is from Lutz Kilian's website.² Similar to Kilian and Murphy (2013), OECD inventories are taken to be a proxy for global inventories.³ The appendix contains details on the data constructions and sources. A lag order of four is chosen to capture dynamics at longer horizons in oil market VAR, which Hamilton and Herrera (2004) argue as being important.

The use of the core PCE index in measuring inflation is deliberate. First, the Federal Reserve uses PCE inflation as its main guide for gauging inflationary pressure. This indicates such a measure for inflation is of greatest concern to policymakers. Second, the use of a core index strips out direct energy price increases by definition. To some extent, the concept of core inflation indicates the belief policymakers can do very little about energy price inflation. Therefore, it is only when such energy price inflation spills over to core inflation that energy prices become a policy concern. Moreover, not stripping out energy prices will increase the price index by construction. Last, but not least, calculation of the Consumer Price Index (CPI) has the drawback of treating shelter differently before 1981. The use of PCE at least makes any comparison as consistent as possible. It is also for this reason that Hooker (2002) chooses to use core PCE rather than core CPI. Using the same measure will allow for some direct comparison with Hooker's conclusions. A natural question is why

²<http://www-personal.umich.edu/~lkilian/reaupdate.txt>

³Kilian and Lee (2013) use proprietary data from the Energy Intelligence Group to measure the global inventory stock and do not document fundamental differences compared to using OECD inventories as a proxy.

to strip out food price inflation if the concern is just inflation pass-through from oil shocks. Beside the fact that food prices, like oil shocks, are also largely beyond the control of the policymaker, a feature of Kilian's index in the 1970s is that they appear to be largely based on shipping costs for grains. Given there were large food price shocks in this period, there is a strong possibility that the response of an inflation measure which includes food may merely capture inflation dynamics stemming from food price shocks, and have nothing to do with oil price shocks (see Blinder and Rudd, 2012).

The sample is from 1974Q1 to 2012Q4. Apart from the fact that most of the time series are only available from the EIA post 1973, specific features of the oil market also make 1974Q1 a suitable choice. A feature of oil data predating the 1970s is the oil price staying at some constant level for an extended period of time followed by one off increases. This is a reflection of regulation in the domestic U.S. oil industry before 1972. The one off increases are the periodic price adjustments by regulators (see, e.g. Hamilton, 1983; Alquist et al., 2011). This brings forth two conceptual issues if the sample predates the 1970s. First, the U.S. did not become a net importer of crude oil until the 1970s and so domestic regulation effectively set the price for crude oil consumption in the U.S. Oil price shocks in the data are thus more akin to regulatory shocks. Second, and more importantly, it is not sensible to fit pre-1970s oil data in a VAR process when the data generating process is a step function.

2.2 Economic Model

It is useful to think about the economic model as one of the oil market augmented by inflation. Oil shocks are identified from a structural model of the oil market with an analysis of its joint dynamics with U.S. inflation. The relevant model of interest for the oil market is the one which Kilian and Murphy (2013) introduce. Unlike work which identifies oil shocks from a generic oil price innovation, the model of the oil market does not *a priori* view all oil shocks being alike. Instead the oil price is regarded as a mere outcome of demand and supply pressures within the oil market. The three oil shocks of interest arising from the oil market are a flow supply shock, a flow demand shock and a speculative demand shock. A key feature is that the underpinnings and fundamentals of the various oil market shocks differ. Oil shocks are thus viewed as textbook shifts of the demand and supply curve in the oil market.

Alternative models of the oil market have previously been specified within the literature with the use of labels such as oil supply shocks, aggregate demand shock, other oil demand shocks and oil-market specific demand shocks.⁴ The key difference between the economic model which Kilian and Murphy (2013) introduce, and which this analysis will build on, lies in the explicit specification

⁴Kilian (2009) is perhaps the intellectual foundation which much of this work rests. It has become more natural to think of the oil market within such a framework with a plethora of work adopting such an approach (e.g. Chen, 2009; Kilian and Park, 2009; Kilian, 2010; Peersman and Van Robays, 2012; Baumeister and Peersman, 2013a,b)

of speculative demand. Terms like other demand shocks and oil-market specific demand shocks have variously been conjectured to contain a large element of demand stemming from a precautionary or speculative nature. However, it is only within the Kilian and Murphy (2013) framework with inventories that there the ability to explicitly label it as a speculative demand.

Inventories can be thought of as serving three distinct economic motives. First, inventories are used to smooth consumption. If oil production falls short, inventories are a means of avoiding large adjustments in oil consumption. Second, inventories serve a convenience yield motive. Firms will always hold some positive level of inventories because this will help cover any temporary shortfall and avoid disruption in the production process. This is a hedge against uncertainty. Alquist and Kilian's (2010) analysis suggests uncertainty in the oil market increases the motivation to hold more inventory. This is a form of precautionary demand which is related to the firm's desire to minimise disruption to their production. Third, given oil is storable, inventories can be a signal of arbitrage motives. Suppose one expects the price of oil to increase in the future, an arbitrage opportunity exists between the spot and futures markets. Now, suppose that oil refiners expect the price of oil to rise in the future, perhaps through expected shortfalls in future production relative to future demand. Hamilton (2009) illustrates this in a simple model where implying the marginal benefit of holding inventories will outweigh the cost of more costly oil in the future. An optimal course of action is to increase their inventory holdings. This puts upward pressure on the spot price until a no arbitrage condition between the spot and futures market is satisfied. This link between the spot and future markets is useful as in the model by Kilian and Murphy (2013), this obviates the need to model the futures market. One can infer speculation through spot price and inventory movements.

A flow supply shock is an exogenous shift in the supply curve in the oil market. This is sometimes referred to as an oil supply shock. An example is exogenous oil supply disruptions when oil producers decide to cut/increase oil production for political motives. A shift in the supply curve will see production increase and prices increase with a simultaneous slowing down of global economic activity. The response of inventories is left unconstrained because it is not clear whether the endogenous motive to smooth consumption, or the motive to mitigate against uncertainty and to arbitrage, dominate. The former will see inventories being drawn down, while the latter two will see the opposite.

Flow demand shocks are shifts in the oil demand curve which correspond to the global business cycle. Therefore, accompanying a rise in global economic activity which characterises flow demand shocks, one can expect oil production and prices to increase due to a shift in the demand curve. Once again, like flow supply, the response of inventories is uncertain as it is again unclear which motive of inventories accumulation dominate.

Speculative demand shocks are also shifts in the demand curve arising from expectations of future shortfalls in production relative to demand. Therefore, they are characterised by increases

in the oil price and oil production, due to it being a demand shock, and stocking of inventories as an optimising response to an expected higher future price.⁵ Given the stocking of inventories, oil consumption goes down, and therefore real activity falls. Speculation within the Kilian and Murphy (2013) framework includes both arbitrage and precautionary demand from uncertainty. These represent changes in the first and second moment expectations of the oil price respectively. This notion of speculation takes a much larger definition than is traditionally taken by the public and mass media.

The economic model does not take a view of how oil shocks feed into inflation. While oil is used in the production process, the possibility for firms to substitute and the imperfect ability of firms to fully pass on this increase thus does not necessarily constitute higher prices. Oil shocks may impact inflation through first round or second round effects. Given that oil is a direct input, an oil shock raises prices as long as producers pass on increases in production costs. The described mechanism is only what one might describe as a direct, or first round, effect of oil shocks. Second round effects due to wage-price spirals triggered by oil shocks, perhaps as a result of monetary policy or inflexibility of labour markets to adjust, are always possible. In this respect, the question under study does not distinguish whether these effects, if any, are direct or secondary. It is also probably true in the absence of runaway inflation expectations, one can expect only a one off increase in inflation from oil shocks (see Blinder and Rudd, 2012).

2.3 Identification Strategy

Guided by the intuition of the economic model, the identification strategy uses a mixture of sign restriction and contemporaneous restrictions in an attempt to identify the three aforementioned oil shocks before quantifying their impact on inflation.

2.3.1 Restrictions Upon Impact

To recap, the three oil market shocks identified are flow supply, flow demand and speculative demand shocks. The first three columns of Table 1 present the impact sign restrictions which identify the three oil market shocks. These are motivated by the economic model in the previous subsection.

2.3.2 Elasticity Restrictions

Kilian and Murphy (2012) argue the mere use of sign restrictions is too weak to identify shocks in oil market VARs. They propose strengthening identification through augmenting the sign re-

⁵Note in Hamilton (2009), oil production does not necessarily need to increase since the oil supply curve may be perfectly inelastic. As the identification will make clear later on, the analysis retains the feature where oil supply is extremely, but not perfectly, inelastic. Whichever the case, oil production will weakly increase.

restrictions with a series of restrictions on the demand and supply elasticities to reduce the set of admissible structural models which are empirically sensible. The key underpinning behind this idea is that impulse response functions contain price and quantity information. Given knowledge of the inelasticity of the supply and demand curve in the oil market, such restrictions can be built into aiding the identification of oil shocks.

The largest upper bound for the oil supply elasticity considered by Kilian and Murphy (2012) is 0.1. A bound three times larger given the quarterly data, namely 0.3, is adopted. This is smaller than the largest bound which Baumeister and Peersman (2013a) consider, though the results are largely robust to that choice.

Unlike the consensus regarding the price elasticity of supply for crude oil, the literature does not provide any clear guide regarding the short-run demand elasticity. Kilian and Murphy (2013) suggest bounding the short-run elasticity to be more inelastic than some long-run elasticity. They suggest Hausman and Newey's (1995) empirical estimate of -0.8 as a reasonable lower bound for the oil demand elasticity in use. A similar bound is adopted here. Note the oil demand elasticity in use is different from the oil demand elasticity in production as the former considers only crude oil which is consumed. This is because with inventories, not all oil produced is consumed by construction. Details on constructing the oil demand elasticity in use are in the appendix. The approach is similar to Kilian and Murphy (2013).

2.3.3 Contemporaneous Restrictions

In addition to the three oil market shocks, there are two residual shocks identified in the model. These two shocks have no economic interpretation and the only requirement is that they are orthogonal to all the other structural shocks (as in all structural VAR exercises) and have no contemporaneous impact on the oil market. This has strong parallels to empirical work which identifies the oil or energy price as pre-determined in structural VARs (e.g. Blanchard and Galí, 2009; Clark and Terry, 2010; Bachmeier and Cha, 2011), usually through a recursive identification scheme. To the extent where the oil market is regarded as pre-determined, some work with first generation oil VAR models based on the framework by Kilian (2009) has found it helpful to order variables in a recursive ordering after oil market variables (e.g. Chen, 2009; Kilian and Park, 2009; Kilian, 2010). The implicit assumption thus is treating the oil block as exogenous. The contemporaneous zero restrictions here have strong parallels to that approach. An identifying strategy based on exclusion restrictions assuming pre-determination of the oil price has been shown by Kilian and Vega (2011) to be sound.⁶

What needs to be made clear is that the choice of sign restrictions within this paper is not

⁶While Kilian and Vega only provide direct evidence of pre-determined oil prices in a monthly context, such restrictions have long been accepted to be tenable at quarterly and even annual frequencies as well. The reason this is not explicitly tested in their paper is due to the low power of the test beyond monthly frequency.

because the empirical analysis chooses to be “agnostic” about identification, as some authors profess to being when they argue for the use of sign restrictions instead of commonly used exclusion restrictions. The use of sign restrictions to identify the oil shock only stems from it being the only known means to disentangle oil shocks within the Kilian and Murphy (2013) framework.

The two residual shocks, labelled as other non-oil market shocks, are restricted to not have a contemporaneous impact on the oil price, inventories and production. The ideal scenario is to also restrict the residual shocks to have no impact on the aggregate demand index. This is certainly feasible if the VAR contains more than five variables, but it is easy to see this will violate the rank condition for identification in a five variable system. The imposition of a zero restriction on inventories, rather than the aggregate demand index, can be defended as follows. To the extent production and oil prices do not change contemporaneously with the residual shock, it is not sensible how inventory levels can change. The residual shocks, however, are outside the oil market and could be anything from monetary to fiscal shocks. It is not unreasonable for the demand for industrial commodities to change within the quarter. Combining using sign restrictions to identify oil shocks together with common oil price exclusion restrictions yields three advantages. First, identifying a full set of shocks uniquely identifies the three oil market shocks and distinguishes them from the other non-oil market shocks. Second, sign restrictions are considered only weak informational restrictions. As Paustian (2007) shows, many restrictions are needed for sign restrictions to be an informative exercise. The six zero restrictions aid in this respect. An alternative common approach is to allow the sign restrictions to hold across multiple periods. For example, the work of Peersman and Van Robays (2012) and Baumeister and Peersman (2013a) allow the sign restrictions to hold for four quarters. Perhaps the biggest controversy of using dynamic restrictions is the ruling out of possible general equilibrium effects after the initial period after the shock. If imposing zero restrictions tightens up the identification to the point where one learns sufficiently from the data without needing to impose dynamic restrictions to make the oil shock appear “reasonable”, this should be viewed favourably. Last, the approach can also be viewed as decomposing the underlying energy price shock of Clark and Terry (2010), who also utilise a TVP-VAR approach. Such an approach is consistent with their identification strategy to the extent that there is no contemporaneous feedback to the oil market from non-oil sources. If one can accept ordering oil shocks first in a structural VAR exercise, it is useful to think of this pre-determined oil shock as a combination of the three identified oil market shocks. By the same logic, if one rejects the exclusion restrictions here, then one should also similarly dismiss the validity of all VARs identified by ordering oil or energy prices first.

Table 1 summarises all the sign and zero restrictions for identification. All the oil shocks are normalised to increase the real oil price. The response of inflation is deliberately unrestricted, allowing the data and identification strategy to determine the inflationary impact of the oil market shocks.

2.4 Estimation

The TVP-VAR can be estimated using what are now standard Markov Chain Monte Carlo methods. Compared to these standard methods (i.e. Primiceri (2005)), distinctive features worth mentioning are the imposition of stable roots using the method of Koop and Potter (2011), and not using a training sample to calibrate the priors. 50,000 burn-in draws are taken. 1,000 draws are taken for inference by keeping one out of every 50 draws after the burn-in sample. The thinning is designed to reduce the autocorrelation within the Monte Carlo sample.⁷ At the same time, nonlinear impulse response functions, sometimes referred to within the class of Generalised impulse response functions, are used to account for the nonlinearity with the change in coefficients.

The details on the choice of priors, estimation strategy, implementation of the sign restrictions and computation of these nonlinear impulse response functions are relegated to the appendix.

3 Results and Discussion

The following subsection will first present the estimated impulse response functions. These will provide a base for a subsequent discussion on the role of oil shocks driving inflation during the 1970s and thereafter. As the core PCE index enters the estimation, any reference to inflation pertains to core PCE inflation.

3.1 Impulse Response Functions

Figures 2 and 3 present a time series of estimated impulse response functions taken at specific horizons. The impulse response functions will be from oil shocks scaled to raise the real oil price by 10% upon impact. This is a common normalisation for presenting responses to oil shocks with split samples or in TVP environments (e.g. Blanchard and Galí, 2009; Herrera and Pesavento, 2009; Baumeister and Peersman, 2013b). The bounds of the 16th and 84th percentile of the posterior distribution are presented to give a sense of the sampling uncertainty. The point estimate retains the orthogonal properties of the various identified shocks. They are the set of impulse response functions closest to the mean of the posterior distribution.⁸

⁷Convergence is monitored through visual inspection and the use of the Geweke (1992) Z statistic as a guide. The evidence by and large indicates convergence of the chain.

⁸The issues with the non-orthogonality of median impulse response functions are now well known (see Fry and Pagan, 2011). At this current point, there are no entirely satisfactory solutions to choosing a set of orthogonal decomposition in a TVP-VAR framework which utilises Bayesian methods. The compromise is to use Inoue and Kilian's (2013) suggestion of minimising the distance to the posterior mean, as unlike the median, the mean is a well defined statistical object. The minimisation approach is as per Fry and Pagan's (2011) proposal. To be clear, the posterior mean is taken to be from the *uncumulated and unnormalised* impulse response functions. These point estimates are then normalised to effect a 10% increase in the real oil price and then cumulated for the differenced variables.

Figure 2 displays responses for inflation and the cumulative impact on the price level on the left and right panel respectively. The choice of two quarters ahead for inflation is because it coincides with the estimated peak in the inflation response. Given that the impact of inflation generally dies out after four quarters, the four quarter ahead cumulative responses of the price level will give a sense how much overall inflation is induced by different oil shocks.

The point estimate generally indicates a dampening out of the inflation induced by flow supply and speculative demand shocks since the 1970s and early 1980s. Both the point estimate and the bounds of the credible sets indicate flow demand shocks are not inflationary. Regarding the precision of the estimates, the lower bound of the posterior credible set for the inflation response to both flow supply and speculative demand generally only clear zero during the 1979/80 oil shock and the subsequent inflation episode from 1979-1981.

Turning to the impact on the price level in the right panel, flow supply and speculative demand shocks in general induce an increase in the price level. Once again, there are no obvious patterns for flow demand. The lower bound of the credible set indicates that the impact of flow supply and speculative demand shocks on the price level often clear zero, which indicates a degree of inflation pass-through. The upper bound presents evidence consistent with the point estimates; inflation pass-through from flow supply and speculative demand shocks have dampened since the 1979/80 oil shock. Therefore, in general, the inflation and price level responses can be summarised as follows: flow demand shocks do not appear to induce any inflation pass-through. There is inflation pass-through from flow supply and speculative demand shocks, but this pass-through has dampened quite dramatically since the 1979/80 oil shock. The dampened pass-through coincides with the advent of the Great Moderation.

Figure 3 shows the four quarters ahead responses of oil production and inventories. The time horizon is chosen since the system generally settles at that horizon. Therefore, the final impact of the various oil shocks can be examined. All these responses are cumulated as they have been differenced prior to estimation. Given the units in estimating the model, oil production is in percentage change. Inventories are in level changes of hundreds of millions of barrels. All changes are relative to the time elapsed since the oil shock occur.

Once again, there is little systematic pattern in the responses of oil production and inventories from flow demand shocks. Turning to flow supply shocks, we can observe the quantity of oil produced four quarters after a flow supply shock reducing by a larger quantity before the mid-1980s compared to thereafter. This is clear from the point estimate, but also evident from the lower bound of the posterior credible set. The upper bound of the posterior credible set suggests, in a statistically significant sense, oil production falls after a flow supply shock before the mid-1980s, but not thereafter. Inventory accumulation also occurs after a flow supply shock. The increase in the quantity of inventories is much larger before the mid-1980s, whether one interprets the upper bound of the posterior credible set or the point estimate. The lower bound of the posterior credible

set also suggests inventory responses rose in a statistically significant sense before the mid-1980s, but not thereafter. It is important to point out that the raw inventory data is trending, implying the absolute level of inventory stocks was much smaller before the mid-1980s compared to thereafter. Cast in this light, the accumulation in the inventory stocks is even much larger in a relative sense. One can therefore deduce it is likely there is some structural change associated with how flow supply shocks interacted with inventory accumulation and oil production.

Turning to the responses of speculative demand, oil production is lower after four quarters. This is despite the sign restriction requiring the impact responses of oil production to be positive. Therefore, general equilibrium effects after initial impact of the oil shock must have reversed the directional response of oil production. This is a useful point to highlight because a common identification restriction requires the sign restrictions to hold for four quarters. Such dynamic identifying restrictions will *a priori* rule out possible general equilibrium impacts of the type reported here. If these restrictions are not valid, identification using such dynamic restrictions may be misleading.

In general, speculative demand shocks induce lower oil production levels and higher levels of inventory accumulation. While the point estimate only suggestively hint at this, the bounds of the credible sets present clearer evidence. In a statistically significant sense, oil production does fall four quarters after a speculative demand shock before the mid-1980s but not thereafter. The same point about statistical significance can be made about higher levels of inventory accumulation after four quarters before the mid-1980s but not thereafter. These changes in statistical significance before the mid-1980s and thereafter for both oil production and inventories respectively suggest a structural break around that time. In addition, once again recall inventories are in absolute levels, but trending over time from the raw data. Therefore inventory accumulation was much larger in a relative sense before the mid-1980s.

We can thus summarise the evidence from the impulse response functions as follows. First, there has never been any systematic time-varying evidence from flow demand shocks. In particular, they do not appear to be a factor in driving core inflation. Second, the response of inflation to flow supply and speculative demand has dampened considerably since the 1970s, consistent with the view that oil shocks do not appear to be as inflationary since the Great Inflation. There is also similar evidence regarding structural breaks, looking at oil production and inventory responses associated with flow supply and speculative demand shocks. In particular, with these shocks, there appears a greater reduction in oil production and a larger degree of inventory accumulation before the mid-1980s compared with thereafter.

3.2 Were the 1970s Different?

The pass-through from oil shocks was large for both flow supply and speculative demand in the 1970s, but has been close to negligible in recent decades. Clark and Terry (2010) report evidence of

reduced energy price pass-through to core inflation. The results from the impulse response function suggest their result may be due to reduced pass-through of the flow supply and speculative demand shocks. Making the distinction between different oil shocks thus allows the presentation of evidence of possible sources of the reduction in the pass-through to core inflation.

Moving on, it is still clear flow supply shocks before the 1980s decreased oil production in the long-run, but this impact is more negligible thereafter. These results are similar to those obtained by Baumeister and Peersman's (2013a; 2013b) analysis of the oil demand curve despite differences in the identification strategy adopted within this paper. The oil demand curve has steepened over time as a smaller quantity adjustment is needed for a 10% increase in the real oil price arising from flow supply shocks. Their finding of a steepening of the oil supply curve is harder to reconcile within the construct of the empirical specification. We will revisit this point later. In Baumeister and Peersman (2013b), they argue smaller quantities is needed to clear the market due to a structural change in the oil market from long term contracting to a spot market based system, which coincide with the development of an oil futures market in the 1980s. This is a reasonable explanation for the smaller responses in oil production. While there is nothing within the analysis in this paper to contradict their argument, this leaves the door open for alternative interpretations of these patterns, especially in conjunction with inventory responses.

What is clear, however, is that there is evidence of structural change within the oil market and the oil-inflation relationship. This concludes the TVP-VAR is a suitable empirical model in uncovering this underlying relationship. The documented structural change within the analysis so far is entirely consistent with prior work. We now turn our attention to how structural changes within the oil market might coincide with a reduction in inflation pass-through.

3.2.1 Perception of Oil Shocks

The dampening of inflation responses to a generically identified oil shock is not a novel finding and is widely documented within the literature (e.g. Hooker, 2002; Blanchard and Galí, 2009; Bachmeier and Cha, 2011). Evidence pertaining to reduced inflation pass-through from flow supply and speculative demand shocks is noteworthy because such shocks were present in the 1979/80 oil shocks. To reiterate, the series of events in the 1979/80 oil shock include the Iranian revolution, Iran-Iraq War and Iranian hostage crisis, and these are obvious candidates for flow supply and speculative demand shocks.

Anecdotal evidence would suggest monetary policy had a role to play in the reduced pass-through of oil shocks since the 1970s. Without a monetary policy feedback rule estimated within the model, the analysis cannot rule in or rule out whether monetary policy is responsible for this change in pass-through. It is difficult to refute that an element of better monetary policy had a part to play in the dampening pass-through. This though, does not rule out the possibility that

alternative channels of the reduction in pass-through and the inclusion of inventories in a TVP-VAR model will help uncover further evidence. Such evidence is not intended to be competing, but rather complementary with existing explanation, like monetary policy.

What Do We Learn from Inventory Responses

As argued earlier in the paper, there are three main motives for holding inventories. These are to smooth consumption, to possess a convenience yield at hand to deal with uncertainty, and to arbitrage. While separately identifying these motives is not straightforward, smoothing in response to an increase in the real oil price can be inferred through a draw down in inventories. On the other hand, an increase in inventories can be interpreted as greater uncertainty and thus a corresponding increase in precautionary demand as per the mechanism in Alquist and Kilian's (2010) analysis. Alternatively, there could be an arbitrage motive. In which case, one can expect to see inventory movement reflecting an attempt to stock up cheaper oil today. We can thus indirectly infer expectations about the oil market in response to the different oil shocks by studying inventory responses from impulse response functions.

The responses of inventories to flow supply and speculative demand shocks in the 1970s appear to suggest that if inventories are used to smooth consumption, they certainly were not during the 1979/80 oil crisis. It is true one can make an argument there was no true flow supply shock during that period. This is because unlike the earlier OPEC episode in 1973/74, the Iranian revolution and the Iranian hostage crisis were accompanied by increases in oil production by Saudi Arabia in order to mitigate the production shortfalls caused by the supply disruptions triggered by those political events. Even if one ignores the possible links from flow supply disruption in the 1979/80 episode to fears of higher future oil prices or greater uncertainty, such market expectations are likely to still have been present as a response to speculative demand. In fact, the endogenous responses of inventories to speculative demand is similar to flow supply. Even if one could not separately identify one of the two, the presence of one of these shocks in 1979/80 is enough to trigger market expectations of higher future prices or more uncertainty. As already documented, conditional on a 10% increase in the real oil price, inventory stocking, in terms of the absolute number of barrels, was much larger for both flow supply and speculative demand shocks before the mid-1980s. Thereafter, there appears to have been a structural change in the oil market of the type which has earlier been documented by Baumeister and Peersman (2013b).

While the inventory behaviour to flow demand shocks does not systematically point to particular motives in their inventory management, they are sufficiently clear for flow supply and speculative demand. Before the 1980s, such shocks are associated with an expectation of higher future oil price or greater uncertainty given they dominate the motive to smooth oil consumption. Why then that such a break occur and how to explain the timing of it? There is perhaps a possibility where oil

shocks were new experiences in the 1970s, and thus such shocks back then were scarring to the extent they induced beliefs about greater future uncertainty or continuously higher prices. Another reason for the timing of the break may be due to the structure of the oil market. As Smith (2009) points out, if there is a rise in the oil price, refiners can either build up inventories or utilise the futures market in hedging their risk. Baumeister and Peersman's (2013b) analysis correctly point out the difference in the structure of the oil market in the 1970s is likely due to the lack of a futures market. The response of inventories strongly suggests the same mechanism is also likely at work here. Without the financial markets to hedge against price uncertainty, refiners have no choice but to increase their inventory holdings in response to large reductions in oil production. This is very clear from the evidence regarding flow supply and speculative demand shocks in the 1970s. With the development of the futures market in the 1980s, such behaviour is absent thereafter as risk may be better managed.

The analysis does not yield radically different insights from historical decompositions by Kilian and Murphy (2013). In particular, their findings of inventory build ups in the 1970s in response to flow supply and speculative demand are consistent with the findings here. Their model, however, is a constant parameter VAR which relies on historical decompositions to understand various oil shock episodes. Therefore, the outcome of the time series in particular historical episodes is due to the sequence of realised shocks. The TVP-VAR approach here allows the modelling of structural change, more in the tradition of Baumeister and Peersman (2013a; 2013b). The difference is while the analysis here is consistent with Kilian and Murphy's (2013) findings and interpretations, the approach here suggests these findings are more likely results of more deep seated structural change which a TVP-VAR is able to uncover.

The Role of Oil Producers

So far, the analysis only fleetingly mentions the role of producers. Consistency between different oil episodes has never been a feature of empirical analysis of OPEC (e.g., see Smith, 2005; Gately, 2007). This inconsistency suggests some form of time variation in their behaviour, which is clear from the oil production responses to the different oil shocks. Of course, world oil production does not merely consist of OPEC, but action relative to OPEC has often been important in understanding developments in the oil market.

Adelman (2002) argues one of OPEC's difficulties is price fixing. While they had been successful in fixing high prices, they have failed dismally with lowering prices because every member has an incentive to cheat on their production quotas. It is clear from the analysis that oil producers in the 1970s reduced their production by a greater magnitude in response to higher prices induced by flow supply and speculative demand. In other words, it appears they were speculating on higher prices. This is not entirely conclusive from the response of oil production from flow supply, simply because

the reduction may be due to sluggish adjustments to an agreed lower production level. However, it is conclusive for the responses to speculative demand. To the extent speculative demand implies an initial rightward shift of the oil demand curve, a reduction in production can only be achieved by either a leftward shift of the demand curve by a larger magnitude than the initial rightward shift, or a leftward shift of the supply curve. Given that inventories continuously rise in response to speculative demand, there is no reason to expect quantity to fall as a result of a reversal of the initial speculative demand shock. Therefore, it is quite likely oil producers were also speculating in response to speculative demand.

Why did this occur in the 1970s, but not thereafter? As Smith (2009) notes, a key development was OPEC's falling market share. Holding off production stimulates efforts to find new sources of oil and there is evidence OPEC's action in the 1970s reduced the world's dependence on them. It is quite possible OPEC's *modus operandi* had changed after the mid-1980s. Their behaviour likely changed after the mid-1980s because they could no longer afford to do as they did before.

3.2.2 Inflation-Pass Through

The uncovering of speculative motives on the demand side, and loose evidence on the supply side, suggests market participants' perception of the future oil price dynamics may have a role to play in the transmission of oil shocks. In particular, if oil prices are expected to be much higher in the future, inflation expectations will rise in response. This will eventually lead to higher realised inflation. The evidence uncovered within this paper suggests that the way oil price increases were viewed pre and post the 1980s may have had a way of feeding into realised inflation through expectations on future oil price dynamics. It is also likely oil shocks in the 1970s had a greater uncertainty effect compared to thereafter, and this stems from the market developments as described above. As far as uncertainty shocks go, this is a similar type of idea to that of Jo's (2013) analysis.

This of course does not rule out the role of better monetary policy. To the extent that better monetary policy anchors inflation expectations in response to oil shocks, poor monetary policy performance in the 1970s may still have played a role. This is possible given that oil shocks may be triggers in de-anchoring inflation expectations. If oil prices are expected to remain high, as is suggested by the speculative motives of market participants, then there is reason this can also trigger the wage-price spirals we observe in the 1970s. This type of wage-price spiral is an outcome of poor monetary policy.

3.2.3 Historical Decomposition

Figure 4 shows the historical decomposition of inflation to the three identified oil shocks.⁹ While the oil shocks play a larger role in the evolution of inflation pre-1980s, their impact are largely irrelevant post the 1980s. It is also clear a combination of flow supply and speculative demand raising oil prices in 1979/80, both had roles in raising inflation in the period 1980/81. The analysis also concurs with Baumeister and Peersman (2013a), who posit flow supply shocks had a minimal role in the 1979/80 inflation, as well as during 1990/91. Nonetheless, note the slight difference as Baumeister and Peersman (2013a) use the CPI, while the measure of inflation in this paper is core PCE inflation. This serves to assure the robustness of their findings are not sensitive to the inflation series chosen. It is also clear that flow supply and speculative demand had larger impacts pre-1980s, and their influence on the inflation series has moderated substantially since. Note this may be driven by inflation being more predictable (and thus shocks explaining very little) in an environment of low and stable inflation during the era of the Great Moderation. This is reflected by the relative magnitude of the reduced form residual from the inflation equation falling dramatically at the bottom panel of Figure 4.

3.3 A Caveat: Trend Inflation and Inflation Gaps

Due to the time-varying nature of parameters, including the constant, in TVP-VAR models, this implies the means of the time series are drifting over time. To see this, rewrite (1) in lag notation and take the expectation of the vector y_t

$$\begin{aligned}(I_5 - B_t(L))y_t &= c_t + \Gamma_t \epsilon_t \\ \mathbb{E}(y_t) &= (I_5 - B_t(L))^{-1}c_t\end{aligned}$$

It is clear the conditional mean is time-varying as well, implying trend inflation, taken as the mean inflation rate, is drifting as well. This is subtle, but reveals the impulse response functions pertain directly to an inflation gap rather than a level of inflation. This is an explicit point which Cogley et al.'s (2010) paper clarifies.

The subtlety is worth highlighting for a key reason. If trend inflation is constant, this subtlety within the model is essentially irrelevant as the impulse response functions will just model inflation moving around some constant steady state. However, there is sufficient evidence within the liter-

⁹The historical decomposition here are approximate. As the nonlinear impulse response functions are computationally very burdensome, it is impractical to constantly compute and integrate out impulse response functions of over extremely long horizons. Given the system settles within about eight quarters, computing using twenty quarters will not induce large truncation biases in the estimate. The historical decompositions presented are from the medians of all structural decompositions. Therefore, they are computed from orthogonal shocks and not mere decomposition of the median impulse responses which are not orthogonal.

ature that trend inflation is not likely to be constant (e.g. Stock and Watson, 2007; Cogley et al., 2010). In an empirical construct such as the TVP-VAR, shocks can only impact the cycle, or gap, and not the trend. Therefore, the finding where flow demand shocks are never inflationary may only pertain to it not affecting the inflation gap. They could well impact the trend inflation, but there is no mechanism to capture such features if this is indeed the case. This, of course, is not an issue just afflicting the analysis within this paper, as the same can be said about all existing TVP-VAR work.¹⁰ Even if the results in this paper are still informative, there is potential for future work to re-evaluate the conclusions in this paper. In particular, future advances in modelling identified shocks to trend variables should verify whether the results within the paper, as well as most of the TVP-VAR work from which this paper builds on and extends, are robust to modelling such features.

3.4 A Digression: Oil Supply and Demand Elasticities

It is worthwhile to revisit Baumeister and Peersman's (2013b) analysis given the possibility of computing time-varying demand and supply elasticities within the empirical framework. Figure 5 displays these estimated impact time-varying elasticities. Note the point estimate is once again chosen from the structural decomposition closest to the posterior mean as per much of the previous analysis.

It is worth first pointing out some differences in terms of comparing between these elasticities and the ones Baumeister and Peersman (2013b) estimate. First, given the inclusion of inventories, it is possible to compute both the oil demand elasticity in use and the oil demand elasticity in production like Kilian and Murphy (2013). Therefore, the conceptual difference may make comparison with Baumeister and Peersman (2013b) challenging given the underlying economic model differs. Second, the bound for the oil supply elasticity in the analysis is 0.3, which is half that of the bound used by Baumeister and Peersman (2013b). This admits the possibility of discarding structural decompositions which are valid in their analysis.

Nevertheless, there is some evidence that the oil demand elasticity in use has fallen over time. While this is not as clear from the point estimate, the bounds of the credible set would suggest so. The break around 1986 is consistent with that of Baumeister and Peersman's (2013b) analysis. This is less clear for the oil demand elasticity in production. Note the oil demand elasticity in use is more elastic than the oil demand elasticity in production, unlike the analysis by Kilian and Murphy (2013). The reason for this is the initial response of inventories within the analysis to a flow supply shock is to stock up, rather than draw down as in Kilian and Murphy's (2013) analysis. Therefore, the smoothing behaviour in Kilian and Murphy's (2013) analysis will by construction induce a

¹⁰Cogley et al. (2010) in fact mention this point in their paper, but acknowledge there are still no satisfactory approaches within the literature to circumvent this issue.

more inelastic response, given the greater unwillingness to change quantity consumed. Given the response in the results within this paper is to stock inventories with a flow supply shock, quantity consumed has to fall even further as some quantity of oil enters the inventory stock. Therefore, there is a greater response of oil in use, and a more elastic response.

The evidence for the oil supply elasticity, computed for both flow demand and speculative demand, do not appear to vary over time. This is even with a lower bound on the supply elasticity than the one Baumeister and Peersman (2013b) consider. Interestingly, if the model is estimated using the first difference, instead of the level, of the real oil price, the oil supply elasticity does steepen modestly. Differencing is the real oil price specification which Baumeister and Peersman (2013b) work with. Therefore, it appears using differences and levels, at least within the specification in this paper, matters with regards to whether oil supply elasticity has fallen over time.

While there are clearly elements within the analysis which are comparable to Baumeister and Peersman (2013b), such comparisons are only suggestive. The inclusion of inventories and inflation is meant to model the question of interest and is to some extent not entirely suitable for analysing the question which Baumeister and Peersman (2013b) pursue. In particular, there are differences in the model structure, the manner in which the priors are sourced and the identification scheme. Reconciling these differences is not within the scope of this paper and is best left for future research.

3.5 Sensitivity

The sensitivity of the results is tested on two dimensions. First, a suggestive bound of oil supply elasticity is reduced to 0.086. This is repeating the exercise by Kilian and Murphy (2012) in a quarterly context, and yields a bound about three and a half times larger than their bound of 0.025. The conclusions within the paper are not sensitive to this bound. This is not surprising as Figure 5 reveals that the large majority of model solutions would satisfy a bound as tight as this. Secondly, the real oil price enters the estimation as a first difference, instead of a level, like Baumeister and Peersman (2013b). This does not change the conclusions in the paper, except that the oil supply elasticity does fall. Therefore, the conclusions are fairly robust.¹¹

4 Conclusion

This paper studies inflation in a TVP-VAR framework with identified oil shocks utilising the Kilian and Murphy (2013) framework. The results indicate two possible explanations why the oil shocks of the 21st century have not been inflationary. First, the oil shocks of the 21st century are best described as being driven by the global business cycle, or flow demand shocks. The results suggest flow demand shocks have never been inflationary. Second, there has been a reduced pass-through

¹¹All these sensitivity results are available upon request.

from oil supply, or flow supply, and speculative oil demand shocks. These shocks certainly characterised the oil shocks during the 1970s Great Inflation and are different from the oil shocks of the recent decade. Even so, this is only a partial story. The analysis indicates a clear structural break with regard to flow supply and speculative demand shocks to inflation. Movements in oil production and inventories suggest oil market participants and oil producers were expecting higher oil prices to persist into the future, highlighting the importance of the role of modelling inventories. This essentially is not present post mid-1980s in response to flow supply and speculative demand shocks. With flow supply and speculative demand shocks no longer inducing expectations of continuously high oil prices, this may have contributed to a reduced pass-through into inflation.

This result is certainly positive. It suggests that even with another episode such as the 1970s oil shock occurring today, a similar outcome would probably not result due to structural change within the oil market and the oil-inflation relationship. Nevertheless, the reasons for this breakdown in the expectations channel of higher oil prices is not clear. Is this driven by a precautionary motive, where development of the future markets may have helped in hedging risk, as per the explanation suggested by Baumeister and Peersman (2013a)? Or has there been a structural shift where the diminishing market power of OPEC may have a role to play? Identifying this reason is an issue left for future research. What is clear however is the role of expectations in the determination of the oil price. It appears there are mechanisms where these expectation shifts in the oil market may spill over to macroeconomic outcomes like inflation. This hints that future work may have to account explicitly for the role of expectations in the oil market.

APPENDIX

A Data

World crude oil production is measured in terms of the daily average number of barrels. The price of crude oil is measured using the U.S. crude oil imported acquisition cost by refiners (dollars per barrel). This is of indication for a world price for oil as it measures the cost of imported crude to the U.S. The U.S. Department of Energy supplies both series. The oil price is adjusted using the U.S. CPI for urban consumers to obtain the real oil price. Kilian's website provides the index for global economic activity. Kilian constructs the index using the cost of shipping freight. The idea is that freight rates provide sufficient proxy for global economic activity. In general, some of the main conceptual difficulty in constructing global economic activity indices are exchange rate and country weighting. Such issues are inherent in data like the OECD industrial production index, which, for example, does not sufficiently account for the rise of India and China. As freight rates are measured using a common currency, this eschews the issue of exchange rate weighting. In addition,

country weighting can also be avoided as freight rates reflect prices in a global market. The price is simply an indicator of demand pressures in the world, rendering the issue of country weighting moot. Kilian (2009) provides more details regarding the index. The index is monthly and thus averages are taken to convert to quarterly frequency.

Like Kilian and Murphy (2013), the OECD inventories stock acts as a proxy for world inventories stock and is constructed in the same way. This data is available from the EIA. The units in the paper are in hundreds of millions of barrels. Inventory is a stock variable. Differencing is thus taken of the stock at the end of the quarter. The OECD inventory data is incomplete before December 1987. This paper takes the same approach as Kilian and Murphy (2013) in reconstructing the data. As the ratio between OECD inventories and U.S. inventories is fairly stable, changes in U.S. inventories, where data is complete, is used to backcast the OECD inventories. The Personal Consumption Expenditure excluding food and energy is from the FRED databank.

B Constructing the Oil Demand Elasticity in Use

A key feature of the Kilian and Murphy (2013) model is that all oil produced is not necessarily consumed due to the presence of inventories. The relevant oil demand elasticity which is computed and bounded for the identification of flow supply shocks is the oil demand elasticity in use. This elasticity considers only crude oil consumption where derivations can be found in Kilian and Murphy (2013). Reproducing the key equation for computing oil demand elasticity in use from their paper using notation consistent within this paper

$$\eta_t^{use} = \frac{(P_{t-1} \times \tilde{\Gamma}_{11}) / 100 - \tilde{\Gamma}_{31}}{\tilde{\Gamma}_{21} / 100 - \overline{\Delta C}}$$

where η_t^{use} is oil demand elasticity in use. P_{t-1} is the quantity of oil production in period $t - 1$. $\tilde{\Gamma}_{11}$, $\tilde{\Gamma}_{31}$ and $\tilde{\Gamma}_{41}$ are the impact multipliers to a flow supply shock for the change in oil production (in percentage terms), the change real oil price (also in percentage terms) and the change in inventory stock (in level terms) respectively to an flow supply shock. $\overline{\Delta C}$ is the mean change in inventory stocks.

Note the elasticity is time-varying by construction even without using time-varying parameters as it is a function of oil production in the previous quarter. The key issue concerns the mean change in inventories. In a constant parameter VAR model setting in the original Kilian and Murphy (2013) paper, this is the unconditional mean of this variable. They thus take the mean change of inventories for the entire sample. However, things are not quite as straightforward in a TVP-VAR environment. In particular, it is possible for this mean to be time-varying as well. Plotting suggests that while volatility has dampened over time, the mean appears to be time-invariant. Fitting a

trend and/or quadratic trend term yields these terms as not statistically significant, which suggests the mean is not systematically moving in a certain direction. Of course, if there are drifts in the mean, this can only be estimated through a suitable state space specification. However, given there appear to be no systematic movements in the mean of the variable, taking the mean of the entire series will be a sufficiently close approximation of the conditional mean for each quarter. Therefore, the identification and construction of oil demand elasticity in use assumes the unconditional mean of the change in inventories as time-invariant.

C Estimation Details

C.1 Priors

The priors of R , W and blocks of S have an independent inverse-Wishart distribution. This places a normal distribution on the initial conditions of the time-varying coefficients B_0 , A_0 and $\log \sigma_0$.

Often, priors are constructed by designating a portion of the sample as a training sample. The training sample is then usually fitted using a constant parameter VAR. This is typically using least squares, though fitting Bayesian VARs with Minnesota priors are not uncommon. The parameters estimated from the training sample are then used to calibrate priors.

This typical choice of using a training sample is however not taken up here. There is difficulty getting reliable data before 1974. This will mean removing about ten years of data and no longer being able to study the interesting developments of the 1970s. Even though it is possible to reconstruct a training sample, as argued in the text, the regulatory structure means nominal oil price series (usually the West Texas Intermediary spot price which goes back as far as the 1940s) will resemble a step function. Deflating such a series through a price deflator results in most of the changes in the real oil price merely measuring inflation. Therefore, there are issues whether it is sensible fitting such a sample with a VAR to calibrate the priors.

In order to deal with this issue, the paper adopts a more Bayesian approach of trying to incorporate “reasonable” assumptions to discipline the priors. This takes the form of two steps. First, the idea is to source parameters as if there was a reasonable training sample which was fitted using least squares to calibrate the priors. This will mainly come from knowledge of what one can reasonably expect from oil market data. The next step is then to also incorporate features of the Minnesota priors. This is similar in spirit to Del Negro’s (2003) suggestion of fitting a Bayesian VAR with a Minnesota prior in order to calibrate priors for TVP-VARs. In doing so, the estimation procedure can utilise positive features of Del Negro’s suggestion. The key positive feature is that Minnesota priors have a reputation of forecasting well with small variances. Therefore, one can expect increased precision of the posterior and a greater practical ease of finding VARs with stable roots to feature in the estimation.

In the tradition of the Minnesota prior, the first lag of the real oil price and inflation in their respective equations is set to 0.8, reflecting the belief of persistence, but short of a unit root process. All other elements are set to zero. Denote the prior of these VAR coefficients as \hat{B} .

To elicit priors on the variance of the VAR equations, first notice the variance is merely a reflection of the spread of the unaccounted component of the respective VAR equations. Therefore, for the change in log oil production, the log real oil price, the change in inventories (in hundreds of millions of barrels), the global economic activity index and core PCE quarter on quarter inflation respectively, the priors on the standard deviation, denoted $\hat{\sigma} = [0.02, 0.2, 0.3, 5, 0.0025]$. For example, this reflects the belief that the standard deviation of an unaccounted for component of the percentage change in oil production is about 2%. These priors are reasonable, in so far as they are regarded as “typical” magnitudes.

The priors on the non zero elements of \hat{A} are constructed by simulating a data generating process implied by \hat{B} and $\hat{\sigma}$, but treating the residuals across equations as uncorrelated. Define the variance of the VAR coefficients, $V(\hat{B})$. In the tradition of the Minnesota prior, the variance of the VAR coefficients, $V(\hat{B})$ is diagonal. Given the large number of VAR coefficients, treating $V(\hat{B})$ as diagonal also makes it much easier to calculate inverses as the covariance matrix is better conditioned for computation. Let the subscript i and j represent the i^{th} equation in the VAR and the j^{th} variable in the VAR, then

$$V(\hat{B})_{i,j} = \begin{cases} \frac{0.2}{r^2}, & \text{for coefficients on their own lag, } r, r = 1, 2, 3, 4 \\ \frac{0.1}{r^2} \frac{\hat{\sigma}_i}{\hat{\sigma}_j}, & \text{for coefficients on lag, } r, i \neq j, r = 1, 2, 3, 4 \\ 100 \times \hat{\sigma}_i, & \text{for coefficients on the constants} \end{cases}$$

0.2, 0.1 and 100 represent the scaling factors which need to be chosen when specifying a Minnesota prior. The full priors are

$$\begin{aligned} B_0 &\sim N(\hat{B}, 4.V(\hat{B})) \\ A_0 &\sim N(\hat{A}, 4.V(\hat{A})) \\ \log \sigma_0 &\sim N(\log(\hat{\sigma}), I_N) \\ R &\sim IW(k_R.115.V(\hat{B}), 115) \\ W &\sim IW(k_W.6.I_N, 6) \\ S_i &\sim IW(k_S.(i+1).V(\hat{A}_i), i+1), \forall i \in \{1, 2, \dots, 5\} \end{aligned}$$

The scaling factors for the priors on the covariance matrices, k_R , k_S and k_W respectively, are similar to the ones Primiceri (2005) chooses. k_R , which governs the tightness of the prior on the covariance matrix of the time-varying VAR coefficients, is set to 0.0001. k_S and k_W are set to 0.01 and 0.0001 respectively and the degrees of freedom on the covariance matrix of the VAR coefficient is set to

115, modestly above the minimum degrees of freedom to ensure the prior is proper, 106. Such scaling factors are common choices within the TVP-VAR literature. The idea of using common scaling factors consistent with the TVP-VAR literature is in line with the idea of setting values which reflect a hypothetical and reasonable training sample estimated on least squares.

C.2 Simulation of the Posterior Distribution

A Metropolis within Gibbs Sampler is used to simulate the posterior distribution. The outline of the steps are

1. Draw from $p(B_t^{(i)}|Y, A_t^{(i-1)}, \Sigma_t^{(i-1)}, V^{(i-1)})$
2. Draw from $p(A_t^{(i)}|Y, B_t^{(i)}, \Sigma_t^{(i-1)}, V^{(i-1)})$
3. Draw from $p(\Sigma_t^{(i)}|Y, B_t^{(i)}, A_t^{(i)}, V^{(i-1)})$
4. Draw from $p(V^i|Y, B_t^{(i)}, A_t^{(i)}, \Sigma_t^{(i)})$
5. Repeat steps 1 to 4

where the superscript denotes the i^{th} draw of the simulation. Steps 1 to 4 consist of time-varying parameters and are estimated through the standard Kalman filtering and smoothing techniques for time-varying parameter models as per Carter and Kohn (1994), except for imposing a restriction requiring the VAR roots to lie within the unit circle in step 1. The posterior simulator achieves this by utilising an algorithm developed by Koop and Potter (2011).

The index for the order of the simulation draw will now be dropped and replaced with a superscript T to denote the parameters are time-varying.

Drawing elements of B_t . The density of the time-varying parameters can be factored

$$p(B^T|Y^T, A^T, \Sigma^T, V) = \prod_{t=1}^T p(B_t|B_{t-1}, Y^T, A^T, \Sigma^T, V) \mathbb{1}(B_t \in \Phi)$$

where $\mathbb{1}(\cdot)$ is an indicator function and Φ denotes the entire set of VAR coefficients which satisfy the stability condition by having roots lying within the unit circle. Without imposing stable roots, the time-varying VAR coefficients can normally be drawn using a Kalman filtering and smoothing technique like Carter and Kohn (1994) where

$$B_t|(B_{t+1}, Y^T, A^T, \Sigma^T, V) \sim N(B_{t|t+1}, P_{t|t+1})$$

$P_{t|t+1}$ is the precision matrix of the Kalman filter. Filtering forward of the Kalman filter and smoothing backwards allows computation of $B_{t|t+1}$ and $P_{t|t+1}$.

As Koop and Potter (2011) demonstrate, an accept/reject algorithm like Cogley and Sargent (2005) is not exactly correct due to ignoring the integrating constant. To account for this, the usual step of taking a draw by using the Kalman filter is done first. This will be a proposal draw for the metropolis step. If any of the time-varying VAR coefficients imply an unstable system, the algorithm immediately rejects this draw and stays with the previous draw. The algorithm is allowed to switch between the multi and single move algorithm in Koop and Potter (2011). This is particularly helpful if the multi move algorithm is rejecting a long sequence of draws which contain explosive roots. Otherwise, if the condition of stable VAR roots is achieved, an acceptance probability as per equations (9) and (10) or (16) in Koop and Potter (2011) is computed, depending on whether the algorithm is a single or multi move one.

Drawing elements of A_t . Conditional on Y^T , B^T and Σ^T . Rewriting $A_t \hat{Y}_t \equiv A_t(Y_t - X_t' B_t) = \Sigma_t \epsilon_t = u_t$ where

$$\begin{aligned}\hat{Y}_{1,t} &= u_{1,t} \\ \hat{Y}_{2,t} &= -\alpha_{21,t} \hat{Y}_{1,t} + u_{2,t} \\ \hat{Y}_{3,t} &= -\alpha_{31,t} \hat{Y}_{1,t} - \alpha_{32,t} \hat{Y}_{2,t} + u_{3,t} \\ \hat{Y}_{4,t} &= -\alpha_{41,t} \hat{Y}_{1,t} - \alpha_{42,t} \hat{Y}_{2,t} - \alpha_{43,t} \hat{Y}_{3,t} + u_{4,t} \\ \hat{Y}_{5,t} &= -\alpha_{51,t} \hat{Y}_{1,t} - \alpha_{52,t} \hat{Y}_{2,t} - \alpha_{53,t} \hat{Y}_{3,t} - \alpha_{54,t} \hat{Y}_{4,t} + u_{5,t}\end{aligned}$$

and $\hat{Y}_t \equiv [\hat{Y}_{1,t} \hat{Y}_{2,t} \hat{Y}_{3,t} \hat{Y}_{4,t} \hat{Y}_{5,t}]$. The model has now been transformed into a series of time-varying regression models.

The α 's can now be drawn using the Kalman filter algorithm similar to the one described in drawing the time-varying VAR coefficients one equation at a time.

Drawing elements of Σ_t . Consider $A_t(y_t - X_t' B_t) = y^* = \Sigma_t \epsilon_t$. Taking the square and log of every element leads to

$$y_t^{**} = 2h_{i,t} + e_{i,t}$$

where $y_t^{**} = \log(y^* + c)^2$ where c is a small constant set to 0.0001 to prevent taking the log of zero, $e_{i,t} = \log(\epsilon_t^2)$ and $h_{i,t} = \log \sigma_{i,t}$. A mixture of normals as described by Kim et al. (1998) is used to approximate the transformed innovations, which are distributed $\log \chi^2(1)$. First, the expression is demeaned by utilising the relationship $\mathbb{E}(e_{i,t}) = -1.2704$. To define a mixture of seven normal probabilities with component probabilities q_j mean $m_j - 1.2704$ and variances ν_j^2 , $j \in (1, 2, \dots, 7)$, let $s^T = [s_1 \dots s_7]'$, a matrix of indicator variables selecting the mixtures of normal approximation at every point of time. We can sample the h 's as per the Kalman filter described previously and

$$Pr(s_{i,t} = j | y_t^{**}, h_{i,t}) \propto q_j f_N(y_t^{**} | 2h_{i,t} + m_j - 1.2704, \nu_j^2), j \in (1 \dots 7), i \in (1 \dots n)$$

Kim et al. (1998) compute the weights of the seven component mixtures. These same weights are also used in the estimation.

Drawing elements of V . R , W and the blocks of S have an inverse Wishert distribution. They can be drawn as the innovations are observable.

D Implementing Sign Restrictions With Nonlinear Impulse Response Functions

To implement sign restrictions, recall the impact matrix Γ_t in (1). The objective is to construct Γ_t where $\Gamma_t\Gamma_t' = \Omega_t$ and Ω_t is the reduced form covariance matrix from (2). Suppose there exists a matrix Q which is orthonormal, satisfying $QQ' = I$. Then from (2), it must be true that $\Omega_t = A_t^{-1}\Sigma_tQQ'\Sigma_t'A_t^{-1'} = \Gamma_t\Gamma_t'$, the last equality holding as long as $\Gamma_t = A_t^{-1}\Sigma_tQ$. Therefore identification is through finding a (a number of) Q matrix which rotates a base set of orthogonalised shocks, in this case $A_t^{-1}\Sigma_t$, in order to satisfy a set of pre-specified conditions. In order to impose the zero restrictions, Q matrices have to be constructed for two columns to have the relevant zero restrictions, so that a sign restriction algorithm will merely search over if the other three columns contain the identified oil shocks. To do this, an algorithm which Binning (2013) proposes, which is a variant of the ones in Rubio-Ramírez et al. (2010), is used to construct these orthonormal rotation matrices. Briefly the orthonormal matrix in the QR decomposition of a 5×5 matrix populated randomly by $N(0,1)$ elements is first used to randomise a base set of shocks formed by a Cholesky decomposition. Following which, the block of zero restrictions are imposed by using a sequence of QR decompositions to construct the two columns of zero restrictions as per the steps proposed by Rubio-Ramírez et al. (2010) before constructing the rest of the matrix used to create a candidate set of shocks.

The impulse response function computed is similar in spirit to the generalised impulses of Koop et al. (1996) and Kilian and Vigfusson (2011). Let k denote the desired horizon of the impulse response function to be computed at time t for variable y , then

$$IRF_{y,t+k} = \mathbb{E}(y_{t+k}|e, \theta^t) - \mathbb{E}(y_{t+k}|\theta^t)$$

where e is a one time shock to be perturbed at time t and θ^t is the information set available at time t . To construct the nonlinear impulse response functions, first sample a draw of the posterior which gives a draw of B_t , α_t and Σ_t at time t . Generate Γ by rotating a lower triangular decomposition of the covariance matrix by generating a Q matrix constructed using the algorithm Binning (2013) suggests and which is described above. If this rotation of the shocks satisfies the sign and elasticity restrictions, repeat the following steps 100 times. Otherwise, discard and take another rotation of

the shocks.

- Stochastically simulate a random path of length k (the desired horizon for the impulse response function) for B_t , α_t and Σ using equations (4) to (6). This will require knowledge of the covariance matrices R , S and W to generate sequences of length k for v , ζ and η .
- Take N random draws from a $N(0, 1)$. Call this $\epsilon_t \equiv (\epsilon_1, \epsilon_2 \dots \epsilon_N)$. Create the reduced form residual u_t through $u_t = \Gamma_t \epsilon_t$. This will be for a baseline path. Without any loss of generality, suppose we desire to have a structural shock form the first variable, generate an alternate set of reduced form residual where 1 is added to the desired shock. Call this $\tilde{\epsilon}_t \equiv (\epsilon_1 + 1, \epsilon_2 \dots \epsilon_N)$. Similarly the reduced form residual for the alternate path perturbed by the shock can be computed using $\tilde{u}_t = \Gamma_t \tilde{\epsilon}_t$.
- Generate the two paths from the two reduced form series in the previous step. Along the path of length k , simulate the same reduced form shocks hitting both paths from $t + 1$ onwards and using the stochastic path of time-varying coefficients generated. This allows the system to be hit by different shocks along the time path. Both series thus only differ with respect to the shock in period t .
- Find the difference between both series. This represents one possible path.

Taking an average across the 100 simulations allows for Monte Carlo integration over all possible paths. This represents one set of impulse response functions for constructing the posterior density.

The above steps are repeated until each of the quarters in the sample have 2,000 sets of impulse response functions which satisfy the sign and elasticity restrictions. The zero restrictions are already satisfied by the manner in which the orthonormal rotation matrices are constructed. These 2,000 impulse response functions are taken together to construct the posterior density of the impulse response functions.

References

- Adelman MA. 2002. World oil production & prices 1947-2000. *The Quarterly Review of Economics and Finance* **42**: 169–191, doi: 10.1016/S1062-9769(02)00129-1.
- Alquist R, Kilian L. 2010. What do we learn from the price of crude oil futures? *Journal of Applied Econometrics* **25**: 539–573, doi: 10.1002/jae.1159.
- Alquist R, Kilian L, Vigfusson R. 2011. Forecasting the price of oil. In Graham E, Timmermann A (eds.) *Handbook of Economic Forecasting*, volume 2.
- Bachmeier LJ, Cha I. 2011. Why don't oil shocks cause inflation? *Journal of Money, Credit, and Banking* **43**: 1165–1183, doi: 10.1111/j.1538-4616.2011.00421.x.

- Barsky RB, Kilian L. 2002. Do we really know that oil caused the great stagflation? A monetary alternative. In Bernanke BS, Rogoff K (eds.) *NBER Macroeconomics Annual 2001*, volume 16. NBER, 137–183.
- Baumeister C, Peersman G. 2013a. The role of time-varying price elasticities in accounting for volatility changes in the crude oil market. *Journal of Applied Econometrics* **forthcoming**, doi: 10.1002/jae.2283.
- Baumeister C, Peersman G. 2013b. Time-varying effects of oil supply shocks on the U.S. economy. *American Economic Journal: Macroeconomics* **forthcoming**.
- Binning A. 2013. Underidentified SVAR models: A framework for combining short and long-run restrictions with sign-restrictions. *Norges Bank Working Paper 2013/14* .
- Blanchard O, Galí J. 2009. The macroeconomic effects of oil price shocks: Why are the 2000s so different from the 1970s? In Galí J, Gertler M (eds.) *International Dimensions of Monetary Policy*. NBER, 373–421.
- Blinder AS, Rudd JB. 2012. The supply-shock explanation of the great stagflation revisited. In Bordo MD, Orphanides A (eds.) *The Great Inflation: The Rebirth of Modern Central Banking*. NBER, 119–180.
- Carter C, Kohn R. 1994. On Gibbs sampling for state space models. *Biometrika* **81**: 541–553, doi: 10.1093/biomet/81.3.541.
- Chen SS. 2009. Revisiting the inflationary effects of oil prices. *The Energy Journal* **30**: 141–154, doi: 10.5547/ISSN0195-6574-EJ-Vol30-No4-5.
- Clark TE, Terry SJ. 2010. Time variation in the inflation passthrough of energy prices. *Journal of Money, Credit and Banking* **42**: 1419–1433, doi: 10.1111/j.1538-4616.2010.00347.x.
- Cogley T, Primiceri GE, Sargent TJ. 2010. Inflation-gap persistence in the U.S. *American Economic Journal: Macroeconomics* **2**: 43–69, doi: 10.1257/mac.2.1.43.
- Cogley T, Sargent TJ. 2001. Evolving post-World War II U.S. inflation dynamics. In Bernanke BS, Rogoff K (eds.) *NBER Macroeconomics Annual 2001*, volume 16. NBER, 331–388.
- Cogley T, Sargent TJ. 2005. Drift and volatilities: Monetary policies and outcomes in the post WWII U.S. *Review of Economic Dynamics* **8**: 262–302, doi: 10.1016/j.red.2004.10.009.
- Del Negro M. 2003. Discussion of Cogley and Sargent’s “Drifts and volatilities: Monetary policy and outcomes in the Post WWII U.S.”. *Federal Reserve Bank of Atlanta Working Paper 2003-26* .
- Fry R, Pagan A. 2011. Sign restrictions in structural vector autoregressions: A critical review. *Journal of Economic Literature* **49**: 938–960, doi: 10.1257/jel.49.4.938.
- Gately D. 2007. What oil export levels should we expect from OPEC. *The Energy Journal* **28**: 151–173, doi: 10.5547/ISSN0195-6574-EJ-Vol28-No2-7.
- Geweke J. 1992. Evaluating the accuracy and sampling-based approaches to the calculation of posterior moments. In Bernardo JM, Berger JO, Dawid P, Smith AF (eds.) *Bayesian Statistics*. Oxford University Press, 169–193.
- Hamilton JD. 1983. Oil and the macroeconomy since World War II. *Journal of Political Economy* **91**: 228–248, doi: 10.1086/261140.
- Hamilton JD. 2009. Causes and consequences of the oil shock of 2007-08. *Brookings Papers on Economic Activity* **40**: 215–261.
- Hamilton JD, Herrera AM. 2004. Oil shocks and aggregate macroeconomic behavior: The role of monetary policy: Comment. *Journal of Money, Credit and Banking* **36**: 265–286, doi: 10.1353/mcb.2004.0012.
- Hausman JA, Newey WK. 1995. Nonparametric estimation of exact consumers surplus and deadweight loss. *Econometrica* **63**: 1445–1476, doi: 10.2307/2171777.

- Herrera AM, Pesavento E. 2009. Oil price shocks, systematic monetary policy, and the great moderation. *Macroeconomic Dynamics* **13**: 107–137, doi: 10.1017/S1365100508070454.
- Hooker MA. 2002. Are oil shocks inflationary? Asymmetric and nonlinear specifications versus changes in regime. *Journal of Money, Credit, and Banking* **34**: 540–561, doi: 10.1016/S0304-3932(96)01281-0.
- Inoue A, Kilian L. 2013. Inference on impulse response functions in structural VAR models. *Journal of Econometrics* **forthcoming**, doi: 10.1016/j.jeconom.2013.02.009.
- Jo S. 2013. The effects of oil price uncertainty on the macroeconomy. *Journal of Money, Credit and Banking* **forthcoming**.
- Kilian L. 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review* **99**: 1053–1069, doi: 10.1257/aer.99.3.1053.
- Kilian L. 2010. Explaining fluctuations in gasoline prices: A joint model of the global crude oil market and the U.S. retail gasoline market. *The Energy Journal* **31**: 87–112, doi: 10.5547/ISSN0195-6574-EJ-Vol31-No2-4.
- Kilian L, Lee TK. 2013. Quantifying the speculative component in the real price of oil: The role of global oil inventories. *Journal of International Money and Finance* **forthcoming**.
- Kilian L, Murphy DP. 2012. Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market VAR models. *Journal of European Economic Association* **10**: 1166–1188, doi: 10.1111/j.1542-4774.2012.01080.x.
- Kilian L, Murphy DP. 2013. The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics* **forthcoming**, doi: 10.1002/jae.2322.
- Kilian L, Park C. 2009. The impact of oil price shocks on the U.S. stock market. *International Economic Review* **50**: 1267–1287, doi: 10.1111/j.1468-2354.2009.00568.x.
- 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**: 660–671, doi: 10.1162/REST_a.00086.
- Kilian L, Vigfusson RJ. 2011. Are the responses of the U.S. economy asymmetric in energy price increases and decreases? *Quantitative Economics* **2**: 419–453, doi: 10.3982/QE99.
- Kim S, Shephard N, Chib S. 1998. Stochastic volatility: Likelihood inference and comparison with ARCH models. *Review of Economic Studies* **3**: 361–93, doi: 10.1111/1467-937X.00050.
- Koop G, Pesaran MH, Potter SM. 1996. Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics* **74**: 119–147, doi: 10.1016/0304-4076(95)01753-4.
- Koop G, Potter S. 2011. Time varying VARs with inequality restrictions. *Journal of Economic Dynamics and Control* **35**: 1126–1138, doi: 10.1016/j.jedc.2011.02.001.
- Paustian M. 2007. Assessing sign restrictions. *The B.E. Journal of Macroeconomics* **7**, doi: 10.2202/1935-1690.1543.
- Peersman G, Van Robays I. 2012. Cross-country differences in the effects of oil shocks. *Energy Economics* **34**: 1532–1547, doi: 10.1016/j.eneco.2011.11.010.
- Primiceri G. 2005. Time varying structural vector autoregressive and monetary policy. *Review of Economic Studies* **72**: 821–852, doi: 10.1111/j.1467-937X.2005.00353.x.
- Rubio-Ramírez JF, Waggoner DF, Zha T. 2010. Structural vector autoregressions: Theory of identification and algorithms for inference. *Review of Economic Studies* **77**: 665–696, doi: 10.1111/j.1467-937X.2009.00578.x.
- Smith JL. 2005. Inscrutable OPEC? Behavioral tests of the cartel hypothesis. *The Energy Journal* **26**: 51–82, doi: 10.5547/ISSN0195-6574-EJ-Vol26-No1-3.

Smith JL. 2009. World oil: Market or mayhem? *Journal of Economic Perspectives* **23**: 145–164, doi: 10.1257/jep.23.3.145.

Stock JH, Watson MW. 2007. Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking* **39**: 3–33, doi: 10.1111/j.1538-4616.2007.00090.x.

Table 1: Table of Identifying Restrictions

	Flow Shock	Supply Shock	Flow Demand Shock	Speculative Oil Demand Shock	Other Shocks	
Oil Production	—		+	+	0	0
Real Oil Price	+		+	+	0	0
Inventories	x		x	+	0	0
Global Economic Activity	—		+	—	x	x
Inflation	x		x	x	x	x

Sign indicates restriction on the directional response to the respective shocks.

0 indicates no response to respective shocks.

An x indicates no restrictions imposed.

The identifying restrictions to hold only upon impact.

Figure 1: Real Oil Price Inflation and Core PCE Inflation

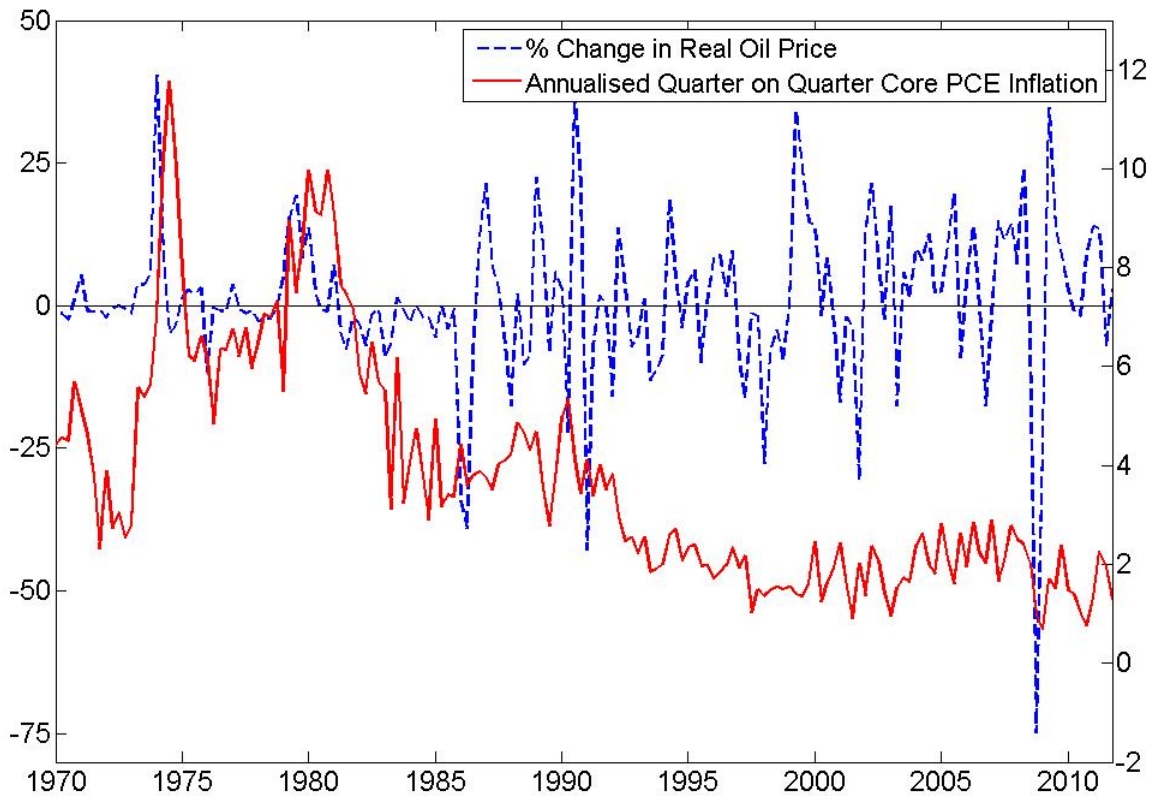
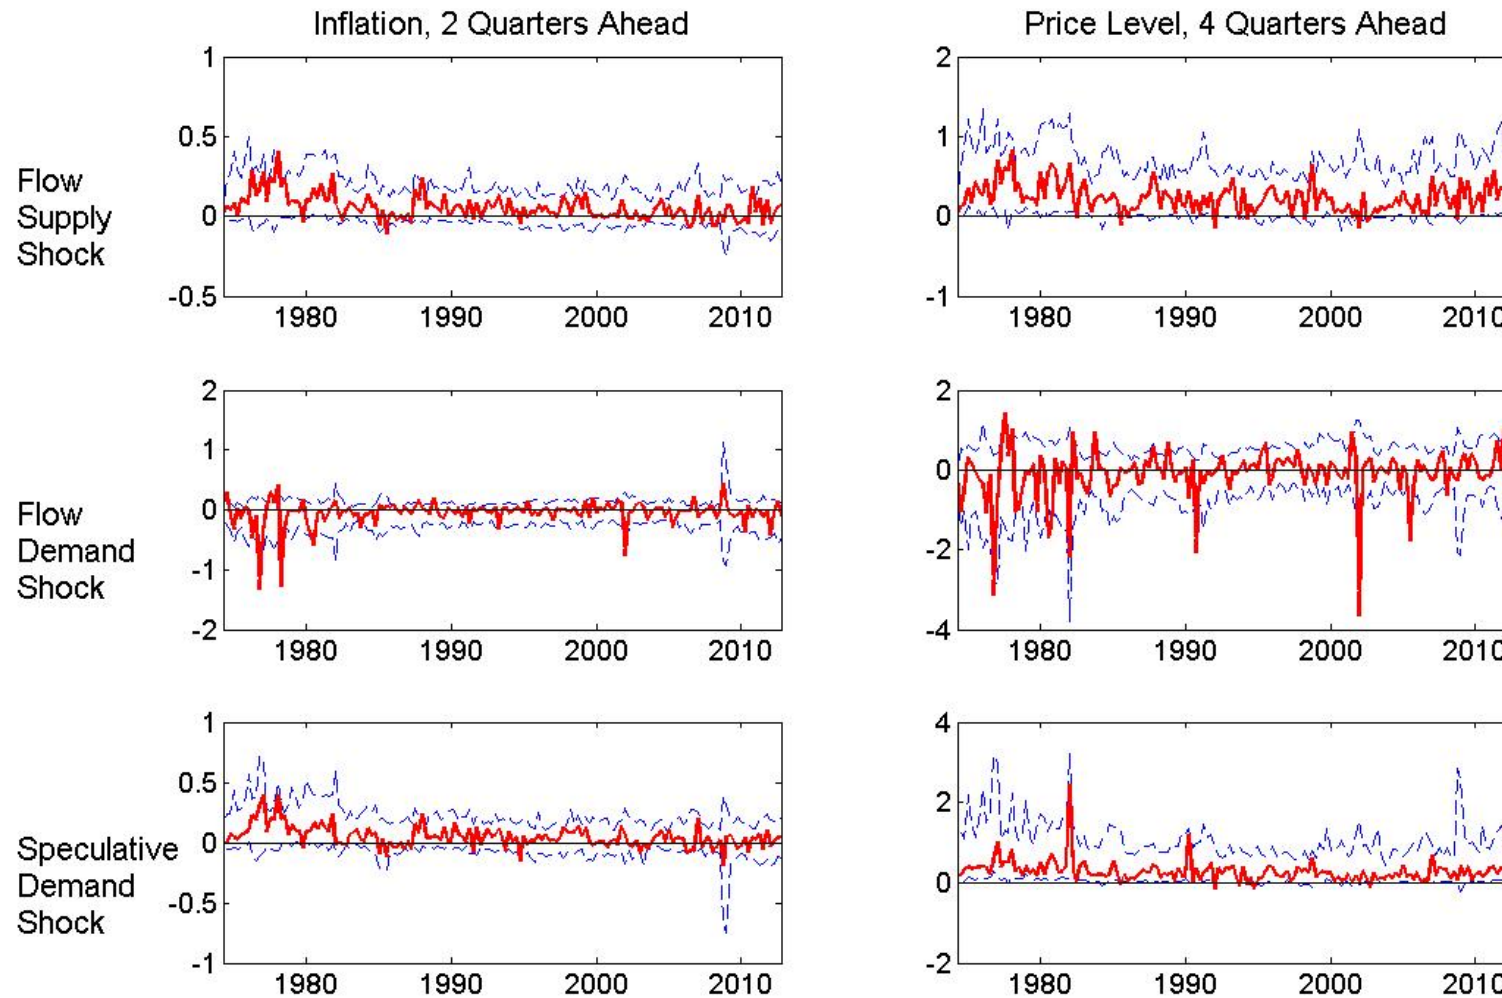


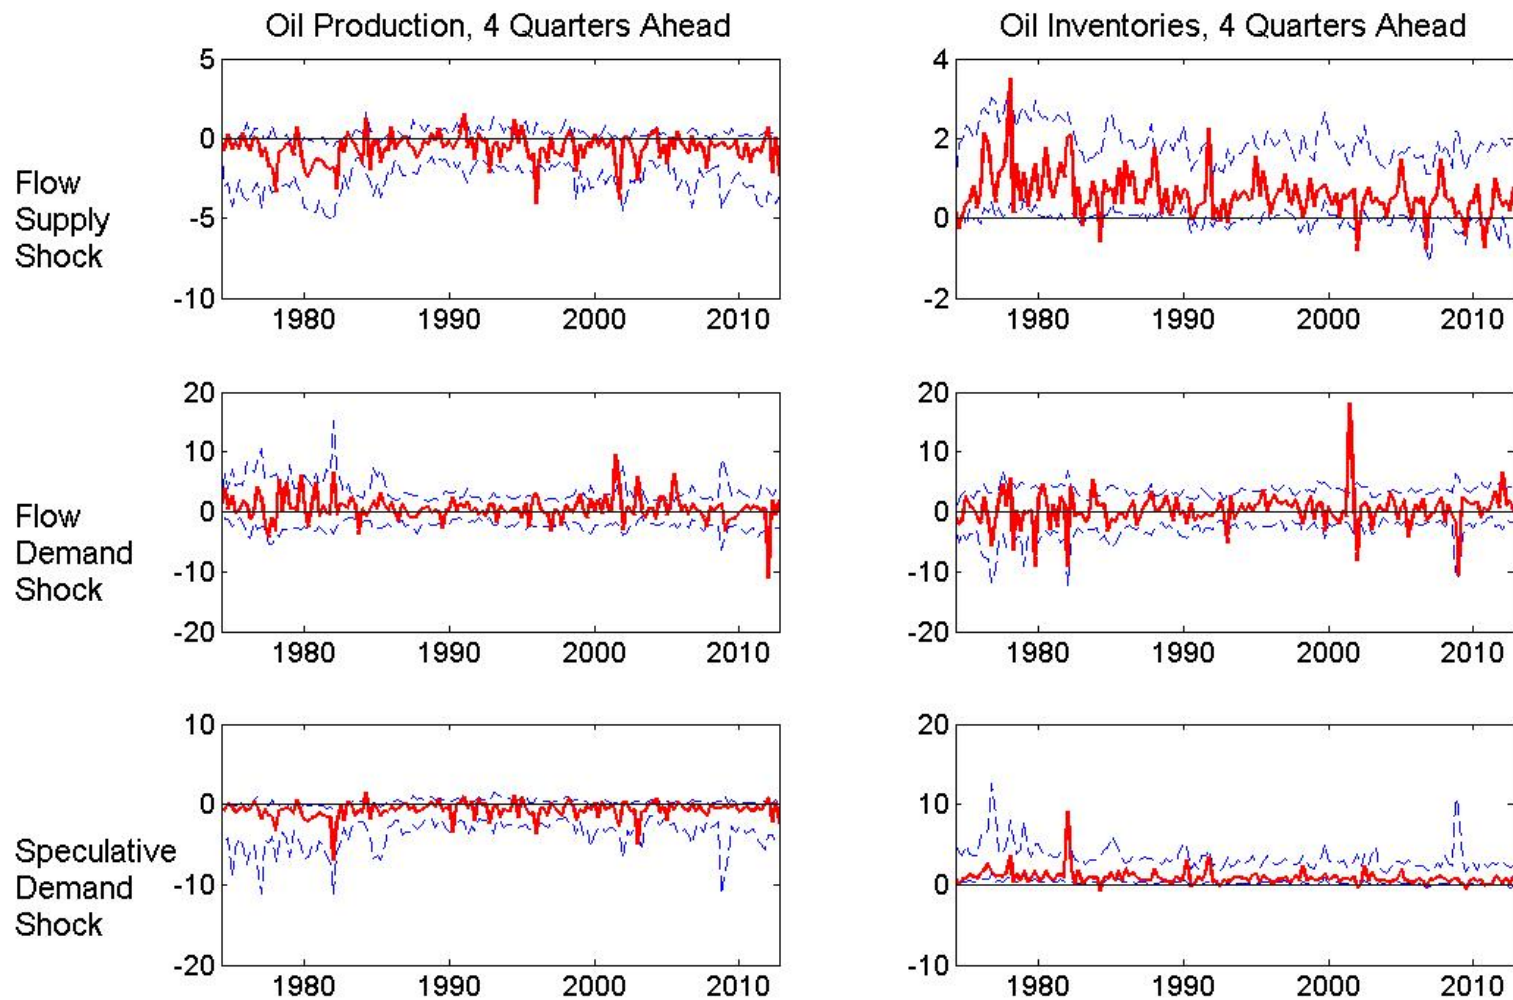
Figure 2: Impulse Response Function to Respective Oil Shocks
 Oil Shocks Normalised to Raise Real Oil Price By 10% Upon Impact



33

*16th and 84th Percentile.
 Point Estimate from Orthogonal Decomposition Closest to the Posterior Mean.
 Inflation is non-annualised quarter on quarter inflation.
 Price Level in Cumulative Percentage Change.*

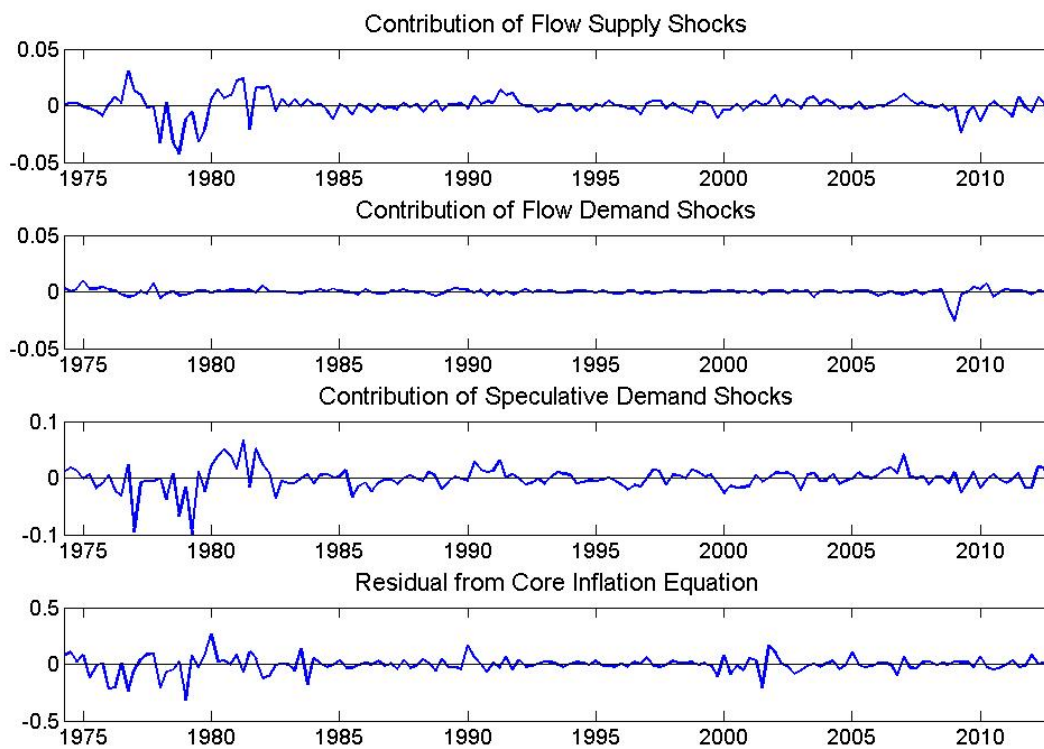
Figure 3: Impulse Response Function to Respective Oil Shocks
 Oil Shocks Normalised to Raise Real Oil Price By 10% Upon Impact



34

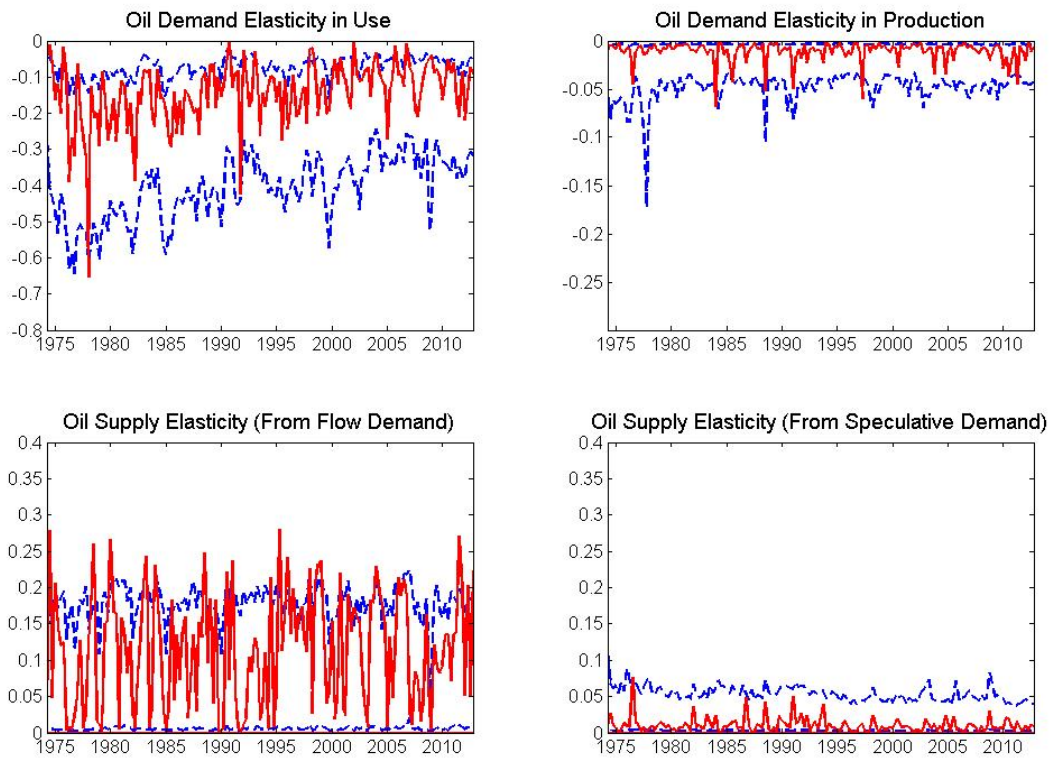
*16th and 84th Percentile.
 Point Estimate from Orthogonal Decomposition Closest to the Posterior Mean.
 Oil Production in Cumulative Percentage Change.
 Inventories in hundreds of millions barrels.*

Figure 4: Historical Decomposition of Inflation



Historical decomposition done for all admissible structural decompositions with the medians of the distribution being reported.

Figure 5: Elasticities



16th and 84th Percentile.

Point Estimate from Orthogonal Decomposition Closest to the Posterior Mean.