

The Role of Inventories and Speculative Trading in the Global Market for Crude Oil

Lutz Kilian

University of Michigan
CEPR

Daniel P. Murphy

University of Michigan

May 12, 2011

Abstract: We develop a structural model of the global market for crude oil that for the first time explicitly allows for shocks to the speculative demand for oil as well as shocks to flow demand and flow supply. The speculative component of the real price of oil is identified with the help of data on oil inventories. The model estimates rule out explanations of the 2003-08 oil price surge based on unexpectedly diminishing oil supplies and based on speculative trading. Instead, we find that this surge was caused by unexpected increases in world oil consumption driven by the global business cycle. There is evidence, however, that speculative demand shifts played an important role during earlier oil price shock episodes including 1979, 1986, and 1990. We also show that, even after accounting for the role of inventories in smoothing oil consumption, our estimate of the short-run price elasticity of oil demand is much higher than traditional estimates from dynamic models that do not account for price endogeneity. Our analysis implies that additional regulation of oil markets would not have prevented the 2003-08 oil price surge.

Key words: Oil market; speculation; fundamentals; peak oil; inventories; demand; supply; oil demand elasticity; gasoline demand elasticity; structural model; identification.

JEL: Q41, Q43, Q48, D84

Acknowledgements: We thank Christiane Baumeister, Domenico Giannone, Ana María Herrera, Ryan Kellogg, Chris Knittel, Harald Uhlig, Whitney Newey, the editor and three anonymous referees for helpful discussions.

Correspondence to: Lutz Kilian, Department of Economics, 611 Tappan Street, Ann Arbor, MI 48109-1220, USA. Email: lkilian@umich.edu.

1. Introduction

The dramatic price fluctuations in the global market for crude oil since 2003 have renewed interest in understanding the evolution of the real price of oil. There is much debate in policy circles about whether the surge in the real price of oil between 2003 and mid-2008 was caused by speculative trading. An alternative view is that unexpected reductions in oil supplies are to blame in the form of OPEC withholding oil supplies from the market or because global oil production has peaked, as predicted by the peak oil hypothesis. Yet another view is that this surge in the real price of oil was driven instead by unexpectedly strong economic growth in the global economy, in particular in emerging Asia.

The relative importance of these explanations is important to policymakers. To the extent that speculative trading is perceived to be the core of the problem, for example, there has been considerable political pressure to impose regulatory limits on trading in oil futures markets.¹ If dwindling global oil supplies are the problem, in contrast, there is little U.S. policymakers can do to avoid similar price surges but to promote energy conservation and the use of alternative sources of energy. Finally, if surges in the global business cycle are the chief cause of high oil prices, then efforts aimed at reviving the global economy after the financial crisis are likely to cause the real price of oil to recover as well, creating a policy dilemma.

This policy discussion has been accompanied by renewed debate among academic researchers about how much oil supply shocks matter for the real price of oil relative to speculative demand shocks and business-cycle driven demand shocks.² This debate has far-reaching implications for the specification of empirical models and for the design of theoretical models of the transmission of oil price shocks. Despite much progress in recent years, there is no consensus in the academic literature on how to model the global market for crude oil. One strand of the literature views oil as an asset, the price of which is determined by desired stocks. In this interpretation, shifts in the expectations of forward-looking traders are reflected in changes in the real price of oil and changes in oil inventories. The other strand of the literature views the price of oil as being determined by shocks to the flow supply of oil and flow demand for oil with little attention to the role of inventories. Much of the research on oil supply shocks is in that tradition, as are economic models linking the real price of oil to fluctuations in the global business cycle.

¹ For a discussion of the link between oil futures and oil spot markets see Alquist and Kilian (2010).

² See, e.g., Baumeister and Peersman (2010); Hamilton (2009a,b); Kilian (2008a, 2009a,b); Kilian and Murphy (2010).

Recently, there has been increasing recognition that both elements of price determination matter in modeling the market for oil (see, e.g., Dvir and Rogoff 2010; Frankel and Rose 2010; Hamilton 2009a,b; Kilian 2009a; Alquist and Kilian 2010). In sections 2 and 3, we propose a structural vector autoregressive (VAR) model of the global market for crude oil that explicitly nests these two explanations of the determination of the real price of oil. Drawing on insights from the economic theory for storable commodities, we design a set of identifying assumptions that allows us to estimate jointly the asset price component of the real price of oil and the components driven by flow demand and flow supply.

Constructing such an econometric model is nontrivial because the presence of an asset price component in the real price of oil considerably complicates the identification of the structural shocks. Traditional oil market VAR models implicitly equate market expectations with the econometric expectations formed on the basis of past data on oil production, global real activity, and the real price of oil. If traders respond to information about future demand and supply conditions not contained in the past data available to the econometrician, however, market expectations will differ from the expectations constructed by the econometrician, rendering these models invalid. We show that this problem can be overcome with the help of data on above-ground oil inventories. The intuition is that – unless the elasticity of oil demand is zero – any expectation of a shortfall of future oil supply relative to future oil demand not already captured by flow demand and flow supply shocks necessarily causes an increase in the demand for above-ground oil inventories and hence in the real price of oil. We refer to such a shock as a speculative demand shock in the spot market for crude oil. It is this type of shock that many pundits implicitly appeal to when attributing higher spot prices to rising oil futures prices. We show that speculative demand shocks can only be identified jointly with flow demand and flow supply shocks.

It may seem that speculative pressures on the spot price of oil could not be studied without including oil futures prices in the structural VAR model. This is not the case. The spot market for oil and the oil futures market are two distinct markets linked by an arbitrage condition which implies that changes in above-ground oil inventories are a sufficient statistic for the information in the oil futures spread (see the theoretical analysis in Alquist and Kilian 2010). This proposition can also be tested empirically, as discussed in Giannone and Reichlin (2006). Test results indicate that the information contained in oil futures prices is already contained in

our baseline model excluding oil futures prices.

Our analysis allows radically new insights into the genesis of historical oil price fluctuations, provides insights highly relevant for current policy debates, and overturns long-held views about the short-run price elasticity of oil demand. It is of particular relevance for recent policy discussions of the potential role of speculation in oil markets, as discussed in section 4. First, the model estimates rule out speculation as a cause of the surge in the real price of oil between 2003 and mid-2008. Given the absence of positive speculative demand shocks in the spot market, we may also infer from the arbitrage condition that speculation cannot explain the simultaneous surge in oil futures prices. Instead the model implies that both spot and futures prices during 2003-08 were driven by unexpected increases in world oil consumption. This finding implies that additional regulation of oil futures markets would not have stemmed the increase in the real price of oil. In fact, that conclusion would remain equally valid, if there were limits to arbitrage between spot and futures markets, as discussed in Singleton (2011).

Second, although speculative trading does not explain the recent surge in the real price of oil, we show that it played an important role in several earlier oil price shock episodes. For example, it was a central feature of the oil price shock of 1979, following the Iranian Revolution, consistent with the narrative evidence in Barsky and Kilian (2002), and it helps explain the sharp decline in the real price of oil in early 1986 after the collapse of OPEC. It also played a central role in 1990, following Iraq's invasion of Kuwait. Although neither negative flow supply shocks nor positive speculative demand shocks alone can explain the oil price spike and oil inventory behavior of 1990/91, their combined effects do.

Third, we document that unexpected fluctuations in global real activity explain virtually the entire surge in the real price of oil between 2003 and mid-2008, even acknowledging that negative oil supply shocks raised the real price of oil slightly. Business cycle factors were also responsible for the bulk of the 1979/80 oil price increase in conjunction with sharply rising speculative demand in the second half of 1979. In contrast, oil supply shocks played only a minor role in 1979. The continued rise in the real price of oil in 1980 reflected negative oil supply shocks (caused in part by the outbreak of the Iran-Iraq War) as much as continued (if slowing) global growth, amidst declining speculative demand. Finally, there is evidence that the recovery of the real price of oil starting in 1999, following an all-time low in post-war history, was aided by coordinated supply cuts. Although our analysis assigns more importance to oil

supply shocks than some previous studies, we conclude that, with the exception of 1990, the major oil price shocks were driven primarily by oil demand shocks.

Much of the *prima facie* case against an important role for speculative trading rests on the fact that there has been no noticeable increase in the rate of inventory accumulation in recent years. Recently, Hamilton (2009a) suggested that, as a matter of theory, speculative trading in oil futures markets may cause a surge in the real price of oil without any change in oil inventory holdings if the short-run price elasticity of demand for gasoline (and hence the short-run price elasticity of oil demand) is zero. Thus, it is essential that we pin down the value of these elasticities.

Hamilton observed that existing estimates of this elasticity in the literature are close to zero, lending some credence to this model of speculation. These estimates, however, are based on dynamic reduced-form regressions that ignore the endogeneity of the real price of oil. They have no structural interpretation and suffer from downward bias. In section 5, we address this limitation with the help of our structural VAR model. Not only do our response estimates show that speculative demand shocks are associated with systematic inventory building, but the econometric model allows us to construct a direct estimate of the elasticity parameter based on exogenous shifts of the oil supply curve along the oil demand curve. The model allows for arbitrarily low demand elasticities. Our median estimate of the short-run price elasticity of oil demand of -0.44 is seven times higher than standard estimates in the literature, but more similar in magnitude to recent estimates from alternative structural models.

This elasticity estimate of -0.44, however, like all existing estimates, is misleading because it ignores the role of oil inventories in smoothing oil consumption. Refiners hold crude oil inventories as insurance against unexpected disruptions of crude oil supplies or unexpected increases in the demand for refined products such as gasoline.³ The decline in crude oil inventories in response to a negative oil supply shock increases the magnitude of oil available to refiners and must be incorporated in constructing the elasticity of oil demand. Our structural model is designed to allow the estimation of this short-run oil demand elasticity in use. The median estimate is -0.26. Although much lower than the traditionally defined elasticity estimate from structural models, this estimate is still four times higher than the consensus view in the

³ For further discussion of this convenience yield provided by oil inventories see Alquist and Kilian (2010) and the references therein.

literature and much higher than zero.

To link our analysis to Hamilton's model, we provide a theoretical benchmark for thinking about the relationship between this short-run price elasticity of oil demand in use and the short-run price elasticity of gasoline demand. We show that under reasonable assumptions about the oil refining industry, the latter is about as high as the short-run price elasticity of oil demand in use, which we estimated to be a -0.26. That gasoline demand elasticity estimate is much larger than some recent estimates in the literature from reduced-form models, illustrating the importance of structural modeling in constructing estimates of the price elasticity of energy demand. Our gasoline demand elasticity estimate eliminates speculation as an explanation of the 2003-08 surge in the real price of oil. The concluding remarks are in section 6.

2. VAR Methodology

Our analysis is based on a four-variable dynamic simultaneous equation model in the form of a structural VAR. Let y_t be a vector of endogenous variables including the percent change in global crude oil production, a measure of global real activity expressed in percent deviations from trend, the real price of crude oil, and the change in oil inventories above the ground. All data are monthly and the sample period is 1973.2-2009.8. We remove seasonal variation by including seasonal dummies in the VAR model.

2.1. Data

Data on global crude oil production are available in the *Monthly Energy Review* of the Energy Information Administration (EIA). Our measure of fluctuations in global real activity is the dry cargo shipping rate index developed in Kilian (2009a). For more details on the rationale, construction and interpretation of this index the reader is referred to the related literature.⁴ While

⁴ The idea of using fluctuations in dry cargo freight rates as indicators of shifts in the global real activity dates back to Isserlis (1938) and Tinbergen (1959); also see Stopford (1997). One advantage of using this type of index is that it automatically accounts for any additional demand for industrial commodities generated by the depreciation of the U.S. dollar in recent years. The panel of monthly freight-rate data underlying the global real activity index was collected manually from *Drewry's Shipping Monthly* using various issues since 1970. The data set is restricted to dry cargo rates. The earliest raw data are indices of iron ore, coal and grain shipping rates compiled by *Drewry's*. The remaining series are differentiated by cargo, route and ship size and may include in addition shipping rates for oilseeds, fertilizer and scrap metal. In the 1980s, there are about 15 different rates for each month; by 2000 that number rises to about 25; more recently that number has dropped to about 15. The index was constructed by extracting the common component in the nominal spot rates. The resulting nominal index is expressed in dollars per metric ton and was deflated using the U.S. CPI and detrended to account for the secular decline in shipping rates. For this paper, this series has been extended based on the Baltic Exchange Dry Index, which is available from

there are other indices of global real activity available, none of these alternative proxies is as appropriate for our purpose of measuring shifts in the global demand for industrial commodities and none is available at monthly frequency for our sample period.

The real price of oil is defined as the U.S. refiners' acquisition cost for imported crude oil, as reported by the EIA, extrapolated from 1974.1 back to 1973.1 as in Barsky and Kilian (2002) and deflated by the U.S. consumer price index. We use the refiners' acquisition cost for imported crude oil because that price is likely to be a better proxy for the price of oil in global markets than the U.S. price of domestic crude oil which was regulated during the 1970s and early 1980s.

Given the lack of data on crude oil inventories for other countries, we follow Hamilton (2009a) in using the data for total U.S. crude oil inventories provided by the EIA. These data are scaled by the ratio of OECD petroleum stocks over U.S. petroleum stocks for each time period. That scale factor ranges from about 2.23 to 2.59 in our sample.⁵ We express the resulting proxy for global crude oil inventories in changes rather than percent changes. One reason is that the percent change in inventories does not appear to be covariance stationary, whereas the change in inventories does. The other reason is that the proper computation of the oil demand elasticity, as discussed below, requires an explicit expression for the change in global crude oil inventories in barrels. This computation is only possible if oil inventories are specified in changes rather than percent changes.⁶

2.2. A Model of the Global Market for Crude Oil

The reduced-form model allows for two years worth of lags. This approach is consistent with evidence in Hamilton and Herrera (2004) and Kilian (2009a) on the importance of allowing for long lags in the transmission of oil price shocks. The real activity index of Kilian (2009a) is

Bloomberg. The latter index, which is commonly discussed in the financial press, is essentially identical to the nominal index in Kilian (2009a), but only available since 1985.

⁵ Petroleum stocks as measured by the EIA include crude oil (including strategic reserves) as well as unfinished oils, natural gas plant liquids, and refined products. The EIA does not provide petroleum inventory data for non-OECD economies. We treat the OECD data as a proxy for global petroleum inventories. Consistent series for OECD petroleum stocks are not available prior to 1987.12. We therefore extrapolate the percent change in OECD inventories backwards at the rate of growth of U.S. petroleum inventories. For the period 1987.12-2009.8, the U.S. and OECD petroleum inventory growth rates are reasonably close with a correlation of about 80%.

⁶ Note that we focus on above-ground inventories in this paper. A different strand of the literature has considered oil below the ground as another form of oil inventories. We do not pursue this possibility, because oil below the ground is not fungible with oil above ground in the short run and because no reliable time series data exist on the quantity of oil below the ground. We do discuss, however, how speculation based on below-ground inventories would be recorded within our model framework and how it may be detected in section 4.3.

stationary by construction. We measure fluctuations in real activity in percent deviations from trend. Oil production is expressed in percent changes in the model, whereas oil inventories are expressed in differences. Preliminary tests provided no evidence of cointegration between oil production and oil inventories. Following Kilian (2009a), the real price of oil is expressed in log-levels.⁷ The corresponding structural model of the global oil market may be written as

$$B_0 y_t = \sum_{i=1}^{24} B_i y_{t-i} + \varepsilon_t, \quad (1)$$

where ε_t is a 4×1 vector of orthogonal structural innovations and $B_i, i = 0, \dots, 24$, denotes the coefficient matrices. The seasonal dummies have been suppressed for notational convenience. The vector ε_t consists of a shock to the flow of the production of crude oil (“flow supply shock”), a shock to the flow demand for crude oil and other industrial commodities (“flow demand shock”), a shock to the demand for above-ground oil inventories arising from forward-looking behavior (“speculative demand shock”), and a residual shock that captures all structural shocks not otherwise accounted for (such as weather shocks, unexpected non-speculative changes to strategic reserves, and shocks to oil companies’ preferences for inventory holdings or to inventory technology) and has no direct economic interpretation.

The flow supply shock corresponds to the classical notion of an oil supply shock, as discussed in the literature. It incorporates in particular supply disruptions associated with exogenous political events in oil producing countries as well as unexpected politically motivated supply decisions by OPEC members. The flow demand shock can be thought of as a demand shock reflecting the state of the global business cycle. The speculative demand shock is designed to capture innovations to the demand for crude oil that reflect revisions to expectations about future demand and future supply not captured by innovations to the current flow supply or flow demand. We specifically focus on speculation by oil traders in the spot market. An alternative view is that speculation may be conducted by oil producers that leave oil below the ground in anticipation of rising prices. The latter form of speculation would be associated with a negative

⁷ It is not clear a priori whether there is a unit root in the real price of oil. The level specification adopted in this paper has the advantage that the impulse response estimates are not only asymptotically valid under the maintained assumption of a stationary real price of oil, but robust to departures from that assumption, whereas incorrectly differencing the real price of oil would cause these estimates to be inconsistent. The potential cost of not imposing unit roots in estimation is a loss of asymptotic efficiency, which would be reflected in wider error bands. Since the impulse response estimates presented below are reasonably precisely estimated, this is not a concern in this study. It should be noted, however, that historical decompositions for the real price of oil rely on the assumption of covariance stationarity and would not be valid in the presence of unit roots.

flow supply shock in our framework rather than the building of above-ground inventories. Both forms of speculation are permitted in our model. We identify the structural shocks based on a combination of sign restrictions on the structural impulse response functions and other economically motivated restrictions.

Our model allows for alternative views of the determination of the real price of oil conditional on lagged data. One view is that the real price of crude oil is determined by the *current* flow supply of oil and the *current* flow demand for oil. The flow supply of crude oil is measured by the global production of crude oil. An unexpected disruption of that flow (embodied in a shift to the left of the contemporaneous oil supply curve along the oil demand curve) within the month will raise the real price of crude oil, and will lower global real activity, if it has any effect on real activity at all. The effect on oil inventories is ambiguous. On the one hand, a negative flow supply shock will cause oil inventories to be drawn down in an effort to smooth consumption. On the other hand, the same shock may raise demand for inventories to the extent that a negative flow supply shock triggers a predictable increase in the real price of oil. Which effect dominates is unclear a priori.

The flow demand for crude oil is driven by unexpected fluctuations in global real activity. These represent shifts in the demand for all industrial commodities including crude oil associated with the global business cycle.⁸ An unanticipated increase in global real activity (embodied in a shift to the right of the contemporaneous oil demand curve along the oil supply curve) within the month will raise the real price of oil, and will stimulate global oil production, if it has any effect on oil production at all. As in the case of a negative flow supply shock, the effect on inventories is ambiguous ex ante.

This standard view of the global crude oil market is incomplete, however. Given that crude oil is storable, it may also be viewed as an asset, the real price of which is determined by the demand for inventories. This means that we must allow the price of oil to jump in response to any news about *future* oil supply or *future* oil demand, including news that is not already embodied in flow supply and flow demand shocks. For example, upward revisions to expected future demand for crude oil (or downward revisions to expected future production of crude oil),

⁸ A well documented fact is that there is considerable comovement between the real price of crude oil and the real price of other industrial commodities during times of major fluctuations in global real activity (see, e.g., Kilian 2009b). The econometric model does not differentiate between shifts in the flow demand for oil and other industrial commodities arising from real economic growth and arising from changes in oil intensity for a given unit of real output.

all else equal, will increase the demand for crude oil inventories in the current period, resulting in an instantaneous shift of the contemporaneous demand curve for oil along the oil supply curve and an increase in the real price of oil. Such shifts could arise, for example, because of the anticipation of a political unrest in oil-producing countries in the Middle East, because of the anticipation of peak oil effects, because of the depletion of oil reserves, or because of the anticipation of faster growth in emerging Asia. Likewise, traders may anticipate a global recession in the wake of a financial crisis, may anticipate higher future oil production as new deep sea oil is discovered off the shores of Brazil, or may anticipate the resumption of oil production in Iraq, as the stability of that country improves.

Rather than being associated only with future oil supply conditions or only with future oil demand conditions, speculative demand shocks are associated with expected shortfalls of future oil supply relative to future oil demand. A positive speculative demand shock will shift the demand for above-ground oil inventories, causing the level of inventories and the real price of oil to increase on impact. The accumulation of inventories requires a reduction in oil consumption (reflected in lower global real activity) and/or an increase in oil production on impact.⁹ The feature that distinguishes flow demand shocks from speculative demand shocks is that positive flow demand shocks necessarily involve an increase in the demand for consumption in the current period, whereas speculative demand shocks do not.

News about the level of future oil production and the level of future demand for crude oil are but one example of shocks to expectations in the global market for crude oil. An unexpected increase in the uncertainty about future oil supply shortfalls would have much the same effect. This point has been demonstrated formally in a general equilibrium model by Alquist and Kilian (2010). The main difference is that uncertainty shocks would not be associated with expected changes in future oil production or real activity. In this paper we refer to any oil demand shock that reflects shifts in expectations about future oil production or future real activity as a speculative oil demand shock. We do not take a stand on whether this economic activity is socially desirable; rather we are concerned with quantifying its empirical relevance.

⁹ Although oil producers could conceivably react to a speculative demand shock by lowering oil production in anticipation of predictable increases of the real price, there is no evidence that oil producers actually have responded systematically in this way. Instead, anecdotal evidence suggests that oil producers typically have increased their production levels following positive speculative demand shocks, consistent with the view that the expected impact response should be weakly positive. If there were on some occasion a reduction in oil production in response to a positive speculative demand shock, that cutback would be classified as a negative flow supply shock within our structural model.

The focus in this paper is on modeling the real price of oil in the spot market. We do not explicitly model the oil futures market. As shown in Alquist and Kilian (2010) and Hamilton (2009a), there is an arbitrage condition linking the oil futures market and the spot market for crude oil. To the extent that speculation drives up the price in the oil futures market, arbitrage will ensure that oil traders buy inventories in the spot market in response. Thus, we can focus on quantifying speculation in the spot market based on oil inventory data without loss of generality. The fact that the inclusion of oil inventory data makes oil futures prices redundant is particularly advantageous considering that oil futures markets were created only in the 1980s, and thus oil futures prices do not exist for a large part of our sample.

In fact, the analysis in Alquist and Kilian implies that data on oil futures prices are redundant in our structural VAR model given that we have already included changes in above-ground oil inventories. This implication of economic theory is testable. If there were additional information in the oil futures spread that is being omitted from the VAR model, then the structural shocks we recover would not be fundamental. Giannone and Reichlin (2006) proposed a statistical test of this type of proposition. They showed that nonfundamentalness of the structural shocks implies Granger causality from the oil futures spread to the remaining VAR model variables. We were unable to reject the null of no Granger causality at conventional significance levels for maturities of 1, 3, 6, 9, and 12 months, indicating that the VAR model is well specified.¹⁰ The next section discusses how we distinguish speculative demand shocks from flow demand and flow supply shocks in practice.

3. Identification

The structural VAR model is set-identified based on a combination of sign restrictions and bounds on the implied price elasticities of oil demand and oil supply.¹¹ Some of these restrictions are implied by the economic model discussed in section 2, while others can be motivated based on extraneous information. We impose four sets of identifying restrictions, each of which is discussed in turn.

¹⁰ Similar results also hold for ex-ante real interest rates.

¹¹ The use of sign restrictions in oil market VAR models was pioneered by Baumeister and Peersman (2010). For further discussion see Kilian and Murphy (2010). For a general exposition also see Uhlig (2005), Fry and Pagan (2011) and Inoue and Kilian (2011).

3.1. Impact Sign Restrictions

The sign restrictions on the impact responses of oil production, real activity, the real price of oil and oil inventories are summarized in Table 1. These restrictions directly follow from the model discussed in section 2. Implicitly, these restrictions also identify the fourth innovation. Given the difficulty of interpreting this residual economically, we do not report results for this fourth shock, but merely note that it is not an important determinant of the real price of oil.

Sign restrictions alone are typically too weak to be informative about the effects of oil demand and oil supply shocks. As demonstrated in Kilian and Murphy (2010) in the context of a simpler model, it is important to impose all credible identifying restrictions to allow us to narrow down the set of admissible structural models.¹² One such set of restrictions relates to bounds on impact price elasticities of oil demand and oil supply.

3.2. Bound on the Impact Price Elasticity of Oil Supply

The price elasticity of oil supply depends on the slope of the oil supply curve. A vertical short-run oil supply curve would imply a price elasticity of zero, for example. An estimate of the impact price elasticity of oil supply may be constructed from the dynamic simultaneous equation model (1) by evaluating the ratio of the impact responses of oil production and of the real price of oil to an unexpected increase in flow demand or in speculative demand. There is a consensus in the literature that this short-run price elasticity of oil supply is close to zero, if not effectively zero.¹³ This fact suggests the need for an upper bound on this elasticity in selecting the admissible models that allows for steep, but not quite vertical short-run oil supply curves (see Kilian and Murphy 2010). It is important to stress that this additional identifying restriction does not constrain the levels of the impact responses, but merely imposes a bound on their relative magnitude.

In our baseline model, we impose a fairly stringent bound of 0.025 on the impact price elasticity of oil supply. Because any such bound is suggestive only, we also experimented with higher bounds. It can be shown that doubling this bound, while increasing the number of

¹² For a similar point also see Canova and De Nicolo (2002), Uhlig (2005), and Canova and Paustian (2011).

¹³ For example, Hamilton (2009b, p. 25) observes that “in the absence of significant excess production capacity, the short-run price elasticity of oil supply is very low.” In practice, it often will take years for significant production increases. Kilian (2009a) makes the case that even in the presence of spare capacity, the response of oil supply within the month to price signals will be negligible because changing oil production is costly. Kellogg (2011) using monthly well-level oil production data from Texas finds essentially no response of oil production to either the spot price or the oil futures price.

admissible models, has little effect on the shape of the posterior distribution of the impulse responses. Even for a bound of 0.1 the 68% quantiles of the posterior distribution of the impulse responses remain qualitatively similar to the baseline model. Moreover, the estimates of the median price elasticity of oil demand reported in section 5 are remarkably robust to this change.

3.3. Bound on the Impact Price Elasticity of Oil Demand

A preliminary estimate of the impact price elasticity of oil demand may be constructed from the estimated model (1) by evaluating the ratio of the impact responses of oil production and of the real price of oil to an unexpected oil supply disruption. This *oil demand elasticity in production* corresponds to the standard definition of the oil demand elasticity used in the literature. It equates the production of oil with the consumption of oil. In the presence of changes in oil inventories that definition is inappropriate. The relevant quantity measure is instead the sum of the flow of oil production and the depletion of oil inventories triggered by an oil supply shock. To our knowledge, this distinction has not been discussed in the literature nor has there been any attempt in the literature to estimate this *oil demand elasticity in use*.

A natural additional identifying assumption is that the impact elasticity of oil demand in use, η_t^{Use} , must be weakly negative on average over the sample. We defer a formal definition of this elasticity to section 5.2.¹⁴ In addition to bounding the demand elasticity in use at zero from above, we also impose a lower bound. It is reasonable to presume that the impact price elasticity of oil demand is much lower than the corresponding long-run price elasticity of oil demand (see, e.g., Sweeney 1984). A benchmark for that long-run elasticity is provided by studies of nonparametric gasoline demand functions based on U.S. household survey data such as Hausman and Newey (1995) which have consistently produced long-run price elasticity estimates near -0.8. This estimate suggests a bound of $-0.8 \leq \eta_t^{Use} \leq 0$.¹⁵

3.4. Dynamic Sign Restrictions

Our final set of restrictions relates to the dynamic responses to an oil supply shock. This set of identifying restrictions is necessary to rule out structural models in which unanticipated oil supply disruptions cause a strong decline in the real price of oil below its starting level which is

¹⁴ Note that we do not need to restrict the oil demand elasticity in production. Our impact sign restrictions ensure that this elasticity is weakly negative on impact.

¹⁵ Schmalensee and Stoker (1999) question the reliability of the price data and the specification used in Hausman and Newey (1995), but Yatchew and No (2001) address those concerns using more detailed Canadian data and arrive at a gasoline demand elasticity estimate of -0.9, very close to Hausman and Newey's original estimate.

at odds with the standard view in the literature. Specifically, we impose the additional restriction that the response of the real price of oil to a negative oil supply shock must be positive for at least twelve months, starting in the impact period. Because the positive response of the real price of oil tends to be accompanied by a persistently negative response of oil production, once we impose this additional dynamic sign restriction, it furthermore must be the case that global real activity responds negatively to oil supply shocks. This is the only way for the oil market to experience higher prices and lower quantities in practice, because in the data the decline of inventories triggered by an oil supply disruption is much smaller than the shortfall of oil production. This implies a joint set of sign restrictions such that the responses of oil production and global real activity to an unanticipated oil supply disruption are negative for the first twelve months, while the response of the real price of oil is positive.

In contrast, we do not impose any dynamic sign restrictions on the responses to oil demand shocks. In this respect our approach differs from Baumeister and Peersman (2009), for example. In particular, we do not impose any dynamic sign restrictions on the responses of global real activity and oil production to speculative oil demand shocks. The reasoning is as follows. The speculative demand shock in our model is a composite of three distinct types of expectations shocks. For example, a pure uncertainty shock (such as a mean-preserving increase in the spread) that raises precautionary demand will be associated with an increase in the real price of oil and higher oil inventories without any change in expected oil production or expected real activity (see Alquist and Kilian 2010). As a result, one would expect no response in oil production or in real activity over time, except on to the extent that higher oil prices stimulate future oil production or reduce future real activity. Likewise, to the extent that a speculative shock involves upward revisions to expected future demand for crude oil, one would expect that shock to increase the real price of oil, oil inventories and future global real activity, but not to change oil production beyond the impact period, except to the extent that higher oil prices stimulate future global oil production. A classical example of such a situation would be speculation in anticipation of a booming world economy (or – with reverse signs – the anticipation of a major recession driven by a financial crisis).

In either of these first two cases, the response of oil production is expected to be weakly positive over time. On the other hand, a speculative shock involving downward revisions to expected future crude oil production would be expected to cause higher oil prices, higher

inventories, and lower future oil production, and possibly a decline in global real activity over time driven by higher oil prices. Thus, unlike in the earlier two cases, there is reason to expect the response of oil production to be negative over time. For example, it has been argued in the literature that the oil price spike of 1990, following the invasion of Kuwait by Saddam Hussein, reflected not only the oil supply shock associated with the cessation of oil production in Iraq and Kuwait, but also the specter of Iraq invading Saudi Arabia and future Saudi oil production ceasing (see Kilian 2008a). Likewise, Barsky and Kilian (2002) made the case that the increase in the price of oil following the Iranian Revolution of 1979 reflected in part the fear of a wider regional conflagration involving the U.S. and/or moderate Gulf oil producers. A third example would be speculation by traders driven by the anticipation of declining future oil supplies (as implied by the peak oil hypothesis). Such speculation, if it did occur, would have been associated with higher oil prices, higher inventories, and lower future oil production.¹⁶

This means that we cannot sign *ex ante* the response of oil production to a speculative demand shock beyond the impact period. Likewise, it is not clear whether the dynamic response of real activity should be positive, reflecting the realization of expectations of increased real activity, or negative, reflecting declines in real activity triggered by higher oil and other industrial commodity prices. Hence, we do not impose restrictions on the sign of the responses of real activity and oil production to speculative oil demand shocks beyond the impact period.

3.5. Implementation of the Identification Procedure

Given the set of identifying restrictions and consistent estimates of the reduced-form VAR model, the construction of the set of admissible structural models follows the standard approach in the literature on VAR models identified based on sign restrictions (see, e.g., Canova and De Nicolò 2002; Uhlig 2005). Consider the reduced-form VAR model $A(L)y_t = e_t$, where y_t is the N -dimensional vector of variables, $A(L)$ is a finite-order autoregressive lag polynomial, and e_t is the vector of white noise reduced-form innovations with variance-covariance matrix Σ_e . Let ε_t denote the corresponding structural VAR model innovations. The construction of structural

¹⁶ In contrast, speculation by oil producers, who choose to leave oil below the ground in anticipation of rising prices, as discussed in Hamilton (2009b), in our modeling framework would be classified as a shock to the flow supply of crude oil.

impulse response functions requires an estimate of the $N \times N$ matrix \tilde{B} in $e_t = \tilde{B}\varepsilon_t$.¹⁷ Let $\Sigma_e = P\Lambda P'$ and $B = P\Lambda^{0.5}$ such that B satisfies $\Sigma_e = BB'$. Then $\tilde{B} = BD$ also satisfies $\tilde{B}\tilde{B}' = \Sigma_e$ for any orthogonal $N \times N$ matrix D . One can examine a wide range of possibilities for \tilde{B} by repeatedly drawing at random from the set \mathbf{D} of orthogonal matrices D . Following Rubio-Ramirez, Waggoner and Zha (2010) we construct the set $\tilde{\mathbf{B}}$ of admissible models by drawing from the set \mathbf{D} and discarding candidate solutions for \tilde{B} that do not satisfy a set of a priori restrictions on the implied impulse response functions.

The procedure consists of the following steps:

- 1) Draw an $N \times N$ matrix K of $NID(0,1)$ random variables. Derive the QR decomposition of K such that $K = Q \cdot R$ and $QQ' = I_N$.
- 2) Let $D = Q'$. Compute impulse responses using the orthogonalization $\tilde{B} = BD$. If all implied impulse response functions satisfy the identifying restrictions, retain D . Otherwise discard D .
- 3) Repeat the first two steps a large number of times, recording each D that satisfies the restrictions and record the corresponding impulse response functions.

The resulting set $\tilde{\mathbf{B}}$ comprises the set of admissible structural VAR models.

4. Estimation Results

The identifying restrictions described in section 3 do not yield point-identified structural impulse responses, but a range of models consistent with the identifying assumptions. We generated 5 million rotations based on the least-squares reduced-form VAR estimate. 14 candidate models satisfied all identifying restrictions. The results are robust to changes in the random seed. For expository purposes, in the analysis below, we focus on the structural model that yields an impact price elasticity of oil demand in use (defined in detail in section 5.2) closest to the posterior median of this elasticity. We also conducted the same analysis with every one of the other 13 admissible structural models. Our main results are robust to the choice of admissible model. The only difference is that some admissible models assign even more explanatory power to flow demand shocks than the benchmark model at the expense of speculative demand shocks.

¹⁷ For a review of the construction of these structural impulse responses the reader is referred to Fry and Pagan (2011), for example.

4.1. Responses to Oil Supply and Oil Demand Shocks

Figure 1 plots the responses of each variable in this benchmark model to the three oil supply and oil demand shocks along with the corresponding pointwise 68% error bands.¹⁸ All shocks have been normalized such that they imply an increase in the real price of oil. In particular, the flow supply shock refers to an unanticipated oil supply disruption. Figure 1 illustrates that the role of storage differs depending on the nature of the shock. A flow supply disruption causes inventories to be drawn down in an effort to smooth production of refined products. A positive flow demand shock causes almost no change in oil inventories on impact, followed by a temporary drawdown of oil inventories. After one year, oil inventories reach a level in excess of their starting level. A positive speculative demand shock causes a persistent increase in oil inventories.

A negative flow supply shock is associated with a reduction in global real activity and a persistent drop in oil production, but much of the initial drop is reversed within the first half year. The real price of oil rises only temporarily. It peaks after three months. After one year, the real price of oil falls below its starting value, as global real activity drops further. A positive shock to the flow demand for crude oil, in contrast, is associated with a persistent increase in global real activity. It causes a persistent hump-shaped increase in the real price of oil with a peak after one year. Oil production also rises somewhat, but only temporarily. Finally, a positive speculative demand shock is associated with an immediate jump in the real price of oil. The real price response overshoots, before declining gradually. The effects on global real activity and global oil production are largely negative, but small. These estimates imply a larger role for flow supply shocks than the structural VAR model in Kilian (2009a), for example, illustrating the importance of explicitly modeling speculative demand shocks and oil inventories.

4.2. What Drives Fluctuations in Oil Inventories and in the Real Price of Oil?

It can be shown that in the short run, 29% of the variation in crude oil inventories is driven by speculative demand shocks, followed by oil supply shocks with 26%. Flow demand shocks have a negligible impact with 2%. At long horizons, in contrast, the explanatory power of speculative demand shocks declines to 27% and that of flow supply shocks to 24%, while the explanatory power of flow demand shocks increases to 15%. This evidence suggests that, on average,

¹⁸ As is standard in the literature, inference about the structural impulse responses is based on Bayesian methods. Following Uhlig (2005), we specify a diffuse Gaussian-inverse Wishart prior distribution for the reduced-form parameters and a Haar distribution for the rotation matrix, which ensures that all admissible models are equally likely a priori.

fluctuations in oil inventories mainly reflect speculative trading as well as production smoothing by refiners in response to oil supply shocks. This contrasts with a much larger role of flow demand shocks in explaining the variability of the real price of oil. For example, in the long run, 87% of the variation in the real price of oil can be attributed to flow demand shocks, compared with 9% due to speculative demand shocks and 3% due to flow supply shocks.

Impulse responses and forecast error variance decompositions are useful in studying average behavior. To understand the historical evolution of the real price of oil, especially following major exogenous events in oil markets, it is more useful to compute the cumulative effect of each shock on the real price of oil and on the change in oil inventories. Figure 2 allows us to answer not only the question of how important the speculative component in the real price of oil was between 2003 and mid-2008, but we can assess the quantitative importance of speculation at each point in time since the late 1970s.¹⁹

4.3. Did Speculators Cause the Oil Price Shock of 2003-2008?

A common view in the literature is that speculators caused part or all of the run-up in the real price of oil between 2003 and mid-2008. Especially the sharp increase in the real price of oil in 2007/08 has been attributed to speculation. The standard interpretation is (a) that there was an influx of “hedge funds” (defined loosely as investors not customarily in the oil business) into the oil futures market, (b) that this influx drove up oil futures prices, and (c) that the increase in oil futures prices somehow spilled over into the spot market for crude oil, causing the spot price to increase as well.

Regarding (a) there is evidence of increased participation of hedge funds in the oil futures market, but it is easy to overlook that oil companies could speculate in the oil futures market just as easily as anybody else, so the distinction between hedge funds and other market participants in the oil futures market is not economically meaningful. Regarding (b) there is no independent evidence supporting this claim. Recent academic studies such as Singleton (2011) have documented a predictive correlation, but stopped short of establishing causality. Let us nevertheless suppose that, for the sake of argument, hedge funds indeed drove up the oil futures price. Then the key question becomes *how* this increase in the price of oil futures contracts spilled over to the spot price of oil, as asserted in (c).

¹⁹ We do not include the contribution of the residual inventory shock because that shock makes no large systematic contribution to the evolution of the real price of oil.

Alquist and Kilian (2010) showed that there are two ways for market participants to “speculate” in general. One involves buying an oil futures contract; the other involves buying oil inventories and holding them with the intent of selling them at a higher price in the future. Economically, these strategies are equivalent. Hence, there is an arbitrage condition that ensures that – if there is speculation in the oil futures market – the speculative demand for inventories in the spot market must also rise, all else equal causing an increase in the spot price and in oil inventories. This same point has been reiterated more recently by Hamilton (2009a). Unless the price elasticity of oil demand is zero – a possibility that we discuss below – a necessary implication of speculation driving up oil futures prices is that above-ground oil inventories must increase.

Notwithstanding the popular perception that speculative demand helped cause the run-up in the real price of oil after 2003 and in particular in 2007/08 – either in anticipation of stronger economic growth or in anticipation of declining oil production as predicted by the peak oil hypothesis – there is no indication in Figure 2 that this oil price surge had much to do with speculative demand shocks. In light of the arbitrage condition, the absence of speculation in the spot market also implies the absence of speculation in oil futures markets and refutes the conventional wisdom in the press about the causes of this oil price surge.²⁰ Instead, the model supports the substantive conclusions in Kilian (2009a) that the surge in the real price of oil between 2003 and mid-2008 was mainly caused by shifts in the flow demand for crude oil driven by the global business cycle. This finding is important because it tells us that further regulation of these markets would have done nothing to stem the increase in the real price of oil.²¹

An alternative view of speculation is that OPEC in anticipation of even higher oil prices held back its production after 2001, using oil below ground effectively as inventories (see, e.g.,

²⁰ If this arbitrage were impeded or broke down completely, our analysis of the spot market would remain valid, but the oil futures price would become disconnected from the spot price. This is a possibility recently discussed in Singleton (2011). In the limiting case of no arbitrage we would be unable to infer from our model whether there is speculation in oil futures markets, although we could still infer whether there is speculation taking place in the spot market. Clearly, a situation in which arbitrage breaks down, is not consistent with the scenario envisioned by the pundits who attribute rising spot prices to speculation by hedge funds, because in that case speculation-driven increases in the oil futures price could not possibly be transmitted to the spot price of oil.

²¹ One may object that the structural VAR model implies repeated positive flow demand shocks during 2003-08 which may seem unlikely a priori, but this feature is consistent with an in-depth analysis of the record of professional real GDP forecasts in Kilian and Hicks (2010). Even professional forecasters persistently underestimated global growth during 2003-08, especially growth in emerging Asia.

Hamilton 2009a, p. 239).²² One way of testing this hypothesis is through the lens of our structural model. In the model, OPEC holding back oil production in anticipation of rising oil prices would be classified as a negative flow supply shock. Figure 2 provides no indication that negative flow supply shocks played an important role between 2003 and mid-2008. Nor does the evidence provide any support for the notion that negative oil supply shocks associated with the peak oil hypothesis explain the surge in the real price of oil after 2003. Any unanticipated exogenous decline in oil supplies associated with that hypothesis likewise would be captured by an unexpected flow supply reduction, yet there is no indication of such shocks persistently driving up the real price of oil toward the end of the sample. What evidence there is of a small supply-side driven increase in the real price of oil is dwarfed by the price increases associated with flow demand shocks.

Moreover, the sharp V-shaped dip in the real price of oil in late 2008 is unambiguously driven by a similar dip in the global real activity measure associated with the global financial crisis. A similar, if less pronounced, dip followed the Asian crisis in 1997. Whereas the recovery from the all-time low in the real price of oil in 1999-2000 resulted from a combination of coordinated OPEC oil supply cuts, a gradual increase in global real activity (often associated with the U.S. productivity boom) and increased speculative demand in anticipation of increased future real activity and/or further oil supply reductions, the resurgence of the real price of oil starting in early 2009 reflected primarily a recovery of global real activity (see Figure 2).

4.4. The Inventory Puzzle of 1990

Although speculative motives played no important role after 2003, there are other oil price shock episodes when they did, suggesting that our model has the ability to detect speculative demand shocks when they exist. One particularly interesting example is the oil price shock associated with the Persian Gulf War of 1990/91. In related work, Kilian (2009a) presented evidence based on a model without oil inventories that the 1990 oil price increase was driven mainly by a shift in speculative demand (reflecting the uncertainty about future oil supplies from neighboring Saudi Arabia) rather than the physical reduction in oil supplies associated with the war. As noted by Hamilton (2009a), this result is puzzling upon reflection because oil inventories moved little and,

²² We do not consider oil below the ground as part of oil inventories in this paper. Unlike above-ground oil inventories that can be drawn down at short notice, oil below the ground is inaccessible in the short run and not available for consumption smoothing. Thus, it must be differentiated from inventories in the usual sense.

if anything, slightly declined following the invasion of Kuwait. This observation prompted Hamilton to reject the hypothesis that shifts in speculative demand were behind the sharp increase in the real price of oil in mid-1990 and its fall after late 1990. Given the consensus that flow demand did not move sharply in mid-1990, Hamilton suggested that perhaps this price increase must be attributed to flow supply shocks after all. The inventory data, however, seem just as inconsistent with this alternative hypothesis. Inventories declined in August of 1990, but only by one third of a standard deviation of the change in inventories. Given one of the largest unexpected oil supply disruptions in history, one would have expected a much larger decline in oil inventories given the impulse response estimates in Figure 1. In light of this evidence, neither the supply shock explanation nor the speculative demand shock explanation by itself seems compelling.

Our econometric model resolves this inventory puzzle. The explanation is that the invasion of Kuwait in August of 1990 represented two shocks that occurred simultaneously. On the one hand it involved an unexpected flow supply disruption and on the other an unexpected increase in speculative demand. Whereas the flow supply shock caused a decline in oil inventories, increased speculative demand in August caused an increase in oil inventories, with the net effect being a slight decline in oil inventories. At the same time, the observed increase in the real price of oil was caused by both shocks working in the same direction.

The historical decomposition in Figure 3 contrasts the price and inventory movements caused by each shock during 1990/91. It shows that about one third of the price increase in August of 1990 was caused by speculative demand shocks and two thirds by flow supply shocks. This result is in sharp contrast to the estimates in Kilian (2009a) who found no evidence of oil supply shocks contributing to this increase, illustrating again that the inclusion of inventories in the structural model matters.

Figure 3 also highlights that the decline in the real price of oil, starting in November of 1990 when the threat of Saudi oil fields being captured by Iraq had been removed by the presence of U.S. troops, was almost entirely caused by a reduction in speculative demand rather than increased oil supplies. The latter observation is consistent with evidence in Kilian (2008a) that it is difficult to reconcile the sharp decline in the real price of oil starting in late 1990 with data on oil production. The evidence of a sharp decline in speculative demand in late 1990 raises the obvious question of when and why speculative demand had surged in the first place. The

bottom panel of Figure 2 reveals there was in fact a substantial increase in speculative demand already in the months leading up to the invasion. This result is consistent with a sharp increase in oil inventories in the months leading up to the invasion. One interpretation is that the invasion was anticipated by informed oil traders or, more likely, that traders responded to evidence of increased political tension in the Middle East.²³ The reason that this increase in the speculative component of the real price of oil went unnoticed by the general public was a simultaneous substantial increase in oil production with offsetting effects on the real price of oil in early 1990, as shown in the top panel of Figure 2. In fact, that expansion of oil production served as the motivation for Iraq's increasing hostility to neighboring countries such as Kuwait which it accused of undermining the price of oil, making it more difficult for Iraq to service its foreign debt. Taken in conjunction our evidence implies a much larger role for speculative demand in 1990/91 than the data for the month of August alone would suggest.

4.5. What Caused the 1979 and 1980 Oil Price Shocks?

Speculative demand played an even more important role in 1979. The traditional interpretation is that this oil price increase was driven by flow supply disruptions associated with the Iranian Revolution of late 1978 and early 1979. Much of the observed increase in the real price of oil, however, only occurred after Iranian oil production had resumed later in 1979. Barsky and Kilian (2002) therefore attribute the price increase starting in May of 1979 and extending into 1980 in part to increased flow demand for oil and in part to a substantial increase in speculative demand for oil, consistent with anecdotal evidence from oil industry sources and with the perception of a noticeable increase in the risk of an oil supply disruption in the Persian Gulf in 1979.²⁴

This hypothesis is testable in our model. Figure 2 shows that not only was there a dramatic and persistent increase in the real price of oil driven by positive flow demand shocks in 1979 and 1980 (not unlike the persistent price increase after 2003), but that increase was reinforced after May of 1979 by a sharp increase in speculative demand, exactly as conjectured in Barsky and Kilian (2002). Whereas flow demand pressures on the real price of oil gradually

²³ In this regard, Gause (2002) notes a shift in Iraqi foreign policy toward a more aggressive stance in early 1990.

²⁴ For example, Terzian (1985, p. 260) writes that in 1979 "spot deals became more and more infrequent. The independent refineries, with no access to direct supply from producers, began to look desperately for oil on the so-called 'free market'. But from the beginning of November, most of the big oil companies invoked *force majeure* and reduced their oil deliveries to third parties by 10% to 30%, when they did not cut them off altogether. Everybody was anxious to hang on to as much of their own oil as possible, until the situation had become clearer. The shortage was purely psychological, or 'precautionary' as one dealer put it."

receded starting in 1981, speculative demand pressures on average remained relatively high until the collapse of OPEC in late 1985. In contrast, there is little evidence of flow supply shocks being responsible for the oil price surge of 1979, consistent with the fact that overall global oil production increased in 1979, reflecting additional oil production outside of Iran. Only in late 1980 and early 1981 is there a moderate spike in the real price of oil driven by flow supply shocks, in part associated with the outbreak of the Iran-Iraq War (see Figure 2).

It is useful to explore the price and inventory dynamics in 1979 in more detail. The historical decompositions in Figure 4 show that indeed flow supply shocks caused a temporary drop in oil inventories in December of 1978 and January of 1979, but for the next half year positive flow supply shocks increased oil inventories. This result is also consistent with the fact that global oil production starting in April of 1979 exceeded its level prior to the Iranian Revolution. At the same time, after March of 1979, repeated speculative demand shocks caused a persistent accumulation of inventories, while driving up the real price of oil. The inventory accumulation continues into 1980. There is no indication that flow supply shocks played an important role in the oil price surge of 1979.

It was not until September of 1980 when the Iran-Iraq War broke out that the oil market experienced another major disruption of flow supply. This event is once again associated with declining oil inventories initially and subsequently rising inventories driven by unexpected flow supply increases, reflecting in part the growing importance of new non-OPEC oil producers. As Figure 5 shows, the increase in the real price of oil in response to this flow supply shock is larger than the price response to the 1979 shock. There is also evidence of a slow resurgence of speculative demand following the outbreak of the war, reflected in rising inventories and a higher real price of oil.

4.6. The Collapse of OPEC in 1986

In late 1985, Saudi Arabia decided that it would no longer attempt to prop up the price of oil by reducing its oil production, creating a major positive shock to the flow supply of oil. The same event also markedly changed market perceptions about OPEC's market power. Figure 6 shows that, as expected, the positive flow supply shock in early 1986 drove down the real price of oil, while oil inventories rose in response. Simultaneously, a drop in speculative demand reinforced the decline in the real price of oil, while lowering inventory holdings. This pattern is similar to the pattern in Figure 3, except in reverse. Although OPEC attempted to reunite and control

production in 1987, amidst increased speculation, as shown in Figure 2, these attempts proved unsuccessful in the long run.

4.7. The Venezuelan Crisis and Iraq War of 2002/03

Figure 7 focuses on the flow supply shock of 2002/2003 when within months first Venezuelan oil production slowed considerably at the end of 2002 and then Iraqi oil production ceased altogether in early 2003. The combined cutback in oil production was of similar magnitude to the oil supply shocks of the 1970s (see Kilian 2008a). Figure 7 shows that this event reflected a combination of negative flow supply shocks and positive speculative demand shocks.

The Venezuelan oil supply crisis of late 2002 was associated with declining oil inventories, consistent with an unexpected oil supply disruption, but it also was associated with an increase in speculative demand in anticipation of the Iraq War that dampened the decline in inventories, while reinforcing the increase in the real price of oil. The military conflict in Iraq lasted from late March 2003 until the end of April 2003. Despite the additional loss of Iraqi output in early 2003, global oil production unexpectedly increased. The production shortfalls in Iraq and Venezuela were more than offset at the global level by increased oil production elsewhere. These positive flow supply shocks lowered the real price of oil starting in early 2003 and resulted in positive inventory accumulation. At the same time, as early as March 2003, lower speculative demand caused the real price of oil to drop and oil inventories to fall. Again the effect of the two shocks on inventories was offsetting, whereas the effect on the price worked in the same direction. This last example again underscores that geopolitical events in the Middle East matter not merely because of the disruptions of the flow supply of oil they may create, but also because of their effect on speculative demand.

5. Can Speculation Occur Without a Change in Oil Inventories?

Our structural model of the oil market not only sheds light on the historical evolution of prices and quantities, but it may be used to obtain direct estimates of the short-run price elasticities of oil supply and oil demand. The elasticity of oil demand in particular plays a central role in assessing the empirical content of a recently proposed alternative model of speculation in oil markets. Specifically, Hamilton (2009a) shows that speculation in oil futures markets could result in a surge in the real price of oil without any additional oil inventory accumulation, provided the short-run price elasticity of gasoline demand is zero. As shown below, that

elasticity is closely related to the short-run price elasticity of oil demand, estimates of which can be obtained from our structural model.

5.1. The Short-Run Oil Demand Elasticity in Production

While there is little doubt that the price elasticity of oil supply is near zero in the short run, the literature does not offer much direct evidence on the short-run price elasticity of oil demand.²⁵ The identification of this demand elasticity requires an exogenous shift of the contemporaneous oil supply curve along the contemporaneous oil demand curve within the context of a structural model. In contrast, much of the existing literature on estimating oil demand elasticities from dynamic models has been based on models that do not distinguish between oil demand and oil supply shocks (see, e.g., Dahl 1993; Cooper 2003).²⁶

In addition, existing estimates of the oil demand elasticity in the literature have equated the percent change in quantity with the percent change in the production of crude oil. In this paper, we refer to the resulting elasticity measure as an oil demand elasticity in production, denoted by $\eta^{O, Production}$. In model (1), this elasticity can be estimated as the ratio of the impact response of oil production to an oil supply shock relative to the impact response of the real price of oil. Much of the recent literature builds on the consensus that the short-run price elasticity of oil demand in production is very low. For example, Hamilton (2009b) concludes based on a review of other studies that this elasticity is -0.06. For example, Dahl (1993) and Cooper (2003) report estimates between -0.05 and -0.07. Our posterior median estimate of the oil demand elasticity estimate, as shown in the first column of the upper panel of Table 2, is -0.44. This estimate is seven times higher than typical conjectures in the recent literature. It is also much higher than conventional regression estimates of this elasticity.

One reason for this difference is that standard econometric estimates of the crude oil demand elasticity fail to account for the endogeneity of the price of crude oil. Standard concerns about price endogeneity with respect to quantity suggest that these elasticity estimates are biased toward zero. For example, if we employed the conventional log-level specification used in some the earlier literature on our data the implied elasticity estimate would be only -0.02. The use of a fully structural econometric model allows us to overcome this bias problem. In fact, a number of

²⁵ Our benchmark estimate of the impact price elasticity of oil supply is only about 0.02, consistent with the conventional view that the short-run oil supply curve is nearly vertical.

²⁶ A recent exception is Baumeister and Peersman (2010) who proposed an alternative structural VAR model of the oil market and used this model to estimate the oil demand elasticity in production.

recent studies relying on alternative structural models have obtained similarly large oil demand elasticity estimates ranging from -0.35 to -0.41 (see, e.g., Serletis et al. 2010; Baumeister and Peersman 2010; Bodenstein and Guerrieri 2011). Moreover, accounting for estimation uncertainty does not overturn our result. The first column in the upper panel of Table 2 shows the posterior median of the oil demand elasticity in production with the 68% posterior error band. The model assigns substantial probability mass to values between -0.80 and -0.23 and very little probability mass to values of 0.06 or lower.

5.2. The Short-Run Oil Demand Elasticity in Use

The initial estimate of -0.44 is based on the change in oil production associated with a flow supply shock. This is the conventional definition in empirical work. It is not the definition of the price elasticity of oil demand that matters for policy questions, however. The latter elasticity is based on the change in the use of oil, defined as the sum of the change in oil production and of the depletion of oil inventories, which more accurately captures the behavior of oil consumers. To our knowledge, this *oil demand elasticity in use* has not been discussed in the literature nor has there been any attempt to estimate it. Doing so requires a structural model that explicitly includes oil inventories. In this section we show how the oil demand elasticity in use can be approximated with the help of our structural dynamic model of the oil market. By construction, allowing for inventory responses will tend to lower the price elasticity of oil demand. An important question is how much of a difference the inclusion of oil inventories makes for the elasticity estimate.

The amount of oil used in period t , denoted by U_t , equals the quantity of oil produced in that period (Q_t) minus the oil that is added to the stock of inventories (ΔS_t):

$$U_t = Q_t - \Delta S_t.$$

The change in oil used over time therefore equals the change in oil produced minus the change in the addition to inventory stocks: $\Delta U_t = \Delta Q_t - \Delta^2 S_t$. The price elasticity of oil demand in use is defined as:

$$\eta_t^{Use} \equiv \frac{\% \Delta U_t}{\% \Delta P_t} = \frac{\frac{\Delta Q_t - \Delta^2 S_t}{Q_{t-1} - \Delta S_{t-1}}}{\% \Delta P_t},$$

where Δ represents changes and $\% \Delta$ indicates percent changes in response to an oil supply shock in period t , and P_t denotes the real price of oil. Denote by \tilde{B}_{11} the impact response of the percent change in oil production to an oil supply shock, where \tilde{B}_{ij} refers to the ij th element of \tilde{B} . Then the implied change in oil production is $\Delta Q_t = Q_{t-1} \times (1 + \tilde{B}_{11} / 100) - Q_{t-1} = Q_{t-1} \times \tilde{B}_{11} / 100$. Moreover, $\Delta^2 S_t = \Delta S_t - \Delta S_{t-1} = \overline{\Delta S} + \tilde{B}_{41} - \overline{\Delta S} = \tilde{B}_{41}$, where the change in oil inventories in response to the oil supply shock equals the impact response \tilde{B}_{41} and, prior to the shock, the change in crude oil inventories is equal to its mean $\overline{\Delta S}$, which is observable. Finally, the impact percent change in the real price of oil in response to an oil supply shock is \tilde{B}_{31} . Hence, the demand elasticity in use can be expressed equivalently as

$$\eta_t^{Use} = \frac{(Q_{t-1} \times \tilde{B}_{11} / 100) - \tilde{B}_{41}}{\frac{Q_{t-1} - \overline{\Delta S}}{\tilde{B}_{31} / 100}}.$$

Note that by construction η_t^{Use} depends on Q_{t-1} and hence will be time-varying even though the oil demand elasticity in production is not. We therefore report the average oil demand elasticity in use over the sample period throughout this paper, denoted by $\eta^{O,Use}$.

The second column in the upper panel of Table 2 shows that, as expected, the demand elasticity in use is much lower than the elasticity in production. The median estimate is only -0.26 compared with an estimate of -0.44 for the demand elasticity in oil production. The 68% posterior error bands range from -0.54 to -0.09. This finding is important in that it suggests that even the inclusion of inventories does not overturn our findings that the short-run price elasticity of oil demand is much higher than previously thought. The lower panels of Table 2 show that these estimates are quite robust to relaxing the upper bound on the impact price elasticity of oil supply. Relaxing this bound to 0.050 or even to 0.100 raises the median oil demand elasticity in use slightly without affecting the substance of the conclusions.

5.3. Bounding the Short-Run Gasoline Demand Elasticity

The magnitude of the short-run price elasticity of oil demand has important implications for theoretical models of speculative demand. All else equal, standard models of speculation imply that oil inventories must increase to enable the price of oil to increase. Recently, it has been

suggested that under certain conditions speculation in oil futures markets may drive up the real price of oil without any change in oil inventories (see Hamilton 2009a). Specifically, this may occur if refiners are able to pass on fully to gasoline consumers an exogenous increase in the real price of oil driven by speculation. This requires the demand for gasoline to be completely price-inelastic. Whether the alternative model of speculation in Hamilton (2009a) could explain our data depends on the magnitude of the short-run price elasticity of gasoline demand.

That magnitude is directly related to the impact price elasticity of oil demand in use. Rather than extend the econometric model to permit a joint analysis of the crude oil and gasoline markets, which does not seem feasible given the limited degrees of freedom, in this section we derive an explicit relationship between consumers' demand for gasoline and refiners' demand for crude oil in a model in which refiners are allowed to, but not required to have market power in the gasoline market. Refiners are treated as price-takers in the crude oil market. Our analysis is strictly short-term, as is appropriate in constructing impact price elasticities. In the interest of tractability, we abstract from the fact that gasoline is only one of several refined products jointly produced from crude oil.

We postulate that gasoline is produced according to a Leontief production function over capital, labor, and oil, $G = \min(K, L, \alpha O)$. If capital is fixed in the short run and refiners' labor input can be varied on the intensive margin, which seems plausible in practice, refiners produce gasoline in fixed proportion to the quantity of oil consumed, $G = \alpha O$, and pay a marginal cost equal to the price of oil, P_o , plus the marginal cost of labor, $MC = P_o + c$. P_G denotes the price of gasoline.

Consumers demand $G(P_G) = XP_G^{-\sigma} / P^{1-\sigma}$, where X is the expenditure on gasoline, P is the consumer price index, and the price elasticity of demand for gasoline, η^G , equals $-\sigma$. The inverse demand function is $P_G(G) = \omega G^{-1/\sigma}$ where $\omega \equiv X / P^{1-\sigma}$. In the Cournot-Nash equilibrium, each of J identical refinery firms will choose its own quantity of gasoline output, g_j , $j = 1, \dots, J$, given the outputs of other firms, to maximize profits $\pi_j = P_G(G)g_j - (c + P_o)g_j$ with respect to g_j , where $G = \sum_j g_j$. This yields the first-order conditions

$$\omega G^{-1/\sigma} - \omega g_j G^{(-\sigma-1)/\sigma} / \sigma - (c + P_o) = 0 \quad j = 1, \dots, J.$$

Summing over j and solving for the market price and gasoline production yields

$$P_G = \frac{J\sigma(P_o + c)}{J\sigma - 1} \quad G = \left(\frac{\omega(J\sigma - 1)}{J\sigma(P_o + c)} \right)^\sigma.$$

Given $G = \alpha O$, we obtain

$$\alpha O = \left(\frac{\omega(J\sigma - 1)}{J\sigma(P_o + c)} \right)^\sigma.$$

Log-linearization yields

$$\eta^{O,Use} \approx \frac{P_o}{P_o + c} \eta^G,$$

where $\eta^{O,Use}$ denotes the price elasticity of demand for crude oil in use.

The marginal cost estimates in Considine (1997) suggest that $c \approx 0$, which implies $\eta^{O,Use} \approx \eta^G$. Hence, we conclude based on Table 2 that the posterior median estimate of the short-run price elasticity of gasoline demand is -0.26. There is little probability mass on elasticity values closer to zero. With 84% probability the elasticity exceeds -0.09 in magnitude, allowing us to discount the possibility discussed in Hamilton (2009a) that speculation in oil futures markets may drive up the spot price without any change in oil inventories and to rule out the hypothesis that the 2003-08 oil price surge was driven by speculation.²⁷

6. Conclusion

Standard structural VAR models of the market for crude oil implicitly equate oil production with oil consumption and ignore the role of oil inventories. Traditionally these models have focused on shocks to the contemporaneous flow supply of oil and the contemporaneous flow demand for oil. In this paper we augmented the structural model to include shocks to inventory demand

²⁷ It is worth noting that our estimate of the price elasticity of gasoline demand is much larger than some estimates in the literature. For example, Hughes, Knittel and Sperling (2008) in an influential recent study report a baseline elasticity estimate of only -0.04 for the United States based on data for 2001-06. Their estimate, however, like earlier studies, fails to account for price endogeneity and cannot be interpreted as an elasticity in the textbook sense. Moreover, it is sensitive to the sample period. The average of the least-squares estimates in Hughes et al. for 1975-80 and for 2001-06 is -0.19, similar to estimates in Dahl and Sterner (1991), which is much closer to our estimate. Burger and Kaffine (2009) report estimates as high as -0.29. On the other hand, our estimate is smaller than the instrumental variable regression estimate of the gasoline tax elasticity of gasoline demand of -0.47 reported in Davis and Kilian (2011) with a standard error of 0.23.

reflecting shifts in expectations about future oil supply and future oil demand. Such speculative demand shocks must be represented as shifts of the contemporaneous oil demand curve rather than the contemporaneous oil supply curve, even if the shift in expectations is about a cut in future oil supplies rather than an increase in future oil demand. The reason is that traders in anticipation of the expected oil shortage will buy and store crude oil now with the expectation of selling later at a profit. We proposed a dynamic simultaneous equation model including oil inventories that allows the simultaneous identification of all three types of shocks.

The inclusion of oil inventories matters. The structural model proposed in this paper implies a larger role for flow supply shocks in explaining fluctuations in the real price of oil than previous estimates. The added explanatory power of oil supply shocks in explaining fluctuations in the real price of oil, especially in 1990, comes at the expense of the explanatory power of speculative demand shocks. We showed that the largest and most persistent fluctuations in the real price of oil since the 1970s have been driven primarily by business cycle fluctuations affecting the demand for crude oil. Of particular interest in this paper has been the increase in the real price of oil from 2003 until mid-2008. We were able to provide direct evidence against the popular view that this increase was driven by speculation among oil traders. This is true even for the 2007/08 period. Shifts in speculative demand played a more important role during several earlier oil price shock episodes, however, notably in 1979, 1986 and in 1999/2000. We showed that, without accounting for shifts in the speculative demand for oil, it is not possible to understand the evolution of the real price of oil during these episodes.

Our analysis also suggests that there is no evidence that peak oil or that deliberate production cutbacks by oil producers had much bearing on the recent oil price surge. Rather our results support recent findings in the literature that the sustained run-up in the real price of oil between 2003 and mid-2008 was caused primarily by shifts in the global flow demand for oil. This implies that the real price of oil is expected to rise, as the global economy recovers from the financial crisis, creating a policy dilemma, unless energy consumption can be reduced or new energy sources can be found. Indeed, the recovery of the real price of oil since early 2009 has been primarily driven by increased flow demand. By contrast, additional regulation of oil traders is not likely to prevent the price of oil from rising again in the future nor can increased domestic oil production in the U.S. be expected to have much of an effect on the real price of oil, given the small magnitude of the production increments involved on a global

scale (also see Baumeister and Kilian 2011).

Hamilton (2009a,b) recently has cast doubt on explanations of major oil price increases based on shifts in speculative demand during previous oil price shock episodes. He observed in particular that following the outbreak of the Persian Gulf War in August 1990, oil inventories did not increase as one would have expected in response to a positive speculative demand shock. At the same time, the absence of a sharp decline in oil inventories in August of 1990 is inconsistent with the view that the price increase reflected a negative oil supply shock. We demonstrated that this inventory puzzle can be resolved with the help of a structural oil market model. Our analysis showed that the price and inventory data can be explained only based on a combination of these two shocks. Because the implied inventory responses are of opposite sign, the net effect in inventories is close to zero, where the sharp price increase reflects the fact that the implied price responses are of the same sign. Similar relationships were shown to hold during other key historical episodes. These examples illustrate that it is essential to rely on structural models rather than reduced form evidence in interpreting the price and quantity data.

The use of a structural regression model also is important for the construction of the short-run price elasticity of oil demand. For example, Hamilton (2009a,b) suggests that 1978-81 is one episode where one might clearly and without a regression model attribute cumulative changes in the price of oil to exogenous oil supply shifts only, allowing one to construct a demand elasticity estimate from the ratio of cumulative changes in quantities and prices for that period. The structural model we have analyzed suggests otherwise. We showed that oil demand shocks were the main cause of the observed oil price increase in 1978-81. Oil supply shocks played a small role only, violating the premise of Hamilton's calculations.

We observed that traditional estimates of the short-run price elasticity of oil demand are not credible. One problem is that conventional estimates of this elasticity from dynamic reduced form regressions, as in Dahl (1993) and Cooper (2003), have ignored the endogeneity of the real price of oil, causing the elasticity estimate to be downward biased. Moreover, all existing estimates, including the structural estimates recently provided by Baumeister and Peersman (2009), have ignored the role of inventories in smoothing oil consumption in response to oil supply shocks. We provided a model that allows the estimation of both the traditional oil demand elasticity in production and of the more relevant oil demand elasticity in use which incorporates changes in oil inventories. Our short-run elasticity estimates are substantially higher than

standard estimates cited in the literature, and rule out recently proposed models of speculative trading based on a zero short-run price elasticities of oil demand and of gasoline demand.

Our structural analysis hinged on the use of a proxy for global crude oil inventories. This allowed us to model speculative demand shocks directly rather than absorbing them into a residual shock as in Kilian (2009). Although our inventory data are only an approximation, they are likely to be informative. For example, much has been made of media reports that some speculators in 2007/08 have used oil tankers on the high seas to store oil. Such storage is not covered by national statistics. This is less of a concern than it may seem in that one would expect speculation, if widespread, to result in a systematic increase of all forms of oil inventories. Moreover, the extent to which tankers have been used for storage appears insignificant and is in any case limited to the very end of our sample. Likewise, there has been concern about the expansion of strategic reserves in non-OECD countries such as China. Non-OECD strategic reserves are not included in our inventory data set. However, the construction of the expanded Chinese oil storage facilities was only completed in early 2009, so this fact cannot help explain the surge in the real price of oil from 2003 until mid-2008.

References

- Alquist, R., and L. Kilian (2010), "What Do We Learn from the Price of Crude Oil Futures?" *Journal of Applied Econometrics*, 25, 539-573.
- Barsky, R.B., and L. Kilian (2002), "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," in: *NBER Macroeconomics Annual 2001*, B.S. Bernanke and K. Rogoff (eds.), MIT Press: Cambridge, MA, 137-183.
- Baumeister, C., and L. Kilian (2011), "Real-Time Forecasts of the Real Price of Oil," mimeo, University of Michigan.
- Baumeister, C., and G. Peersman (2010), "Sources of the Volatility Puzzle in the Crude Oil Market," mimeo, Ghent University.
- Bodenstein, M., and L. Guerrieri (2011), "On the Sources of Oil Price Fluctuations," mimeo, Federal Reserve Board.
- Burger, N.E., and D.T. Kaffine (2009), "Gas Prices, Traffic, and Freeway Speeds in Los Angeles," *Review of Economics and Statistics*, 91, 652-657.
- Canova, F. and D. De Nicrolo (2002), "Monetary Disturbances Matter for Business Cycle Fluctuations in the G-7," *Journal of Monetary Economics*, 49, 1131-1159.
- Canova, F., and M. Paustian (2011), "Measurement with Some Theory: Using Sign Restrictions to Evaluate Business Cycle Models," mimeo, Pompeu Fabra.
- Considine, T.J. (1997), "Inventories under Joint Production: An Empirical Analysis of Petroleum Refining," *Review of Economics and Statistics*, 79, 493-502.
- Cooper, J.C.B. (2003), "Price Elasticity of Demand for Crude Oil: Estimates for 23 Countries," *OPEC Review*, 27, 1-8.
- Dahl, C.A. (1993), "A Survey of Oil Demand Elasticities for Developing Countries," *OPEC Review*, 17, 399-419.
- Dahl, C.A., and T. Sterner (1991), "Analyzing Gasoline Demand Elasticities: A Survey," *Energy Economics*, 13, 203-210.
- Davis, L.W., and L. Kilian (2011), "Estimating the Effect of a Gasoline Tax on Carbon Emissions," forthcoming: *Journal of Applied Econometrics*.
- Dvir, E., and K. Rogoff (2010), "Three Epochs of Oil," mimeo, Harvard University.
- Edelstein, P., and L. Kilian (2009), "How Sensitive Are Consumer Expenditures to Retail Energy Prices?" *Journal of Monetary Economics*, 56, 766-779.

- Frankel, J., and A. Rose (2010), “Determinants of Agricultural and Mineral Commodity Prices,” in: Fry, R., Jones, C., and C. Kent (eds.), *Inflation in an Era of Relative Price Shocks*, Reserve Bank of Australia, Sydney, 9-51.
- Fry, R. and A.R. Pagan (2011), “Sign Restrictions in Structural Vector Autoregression: A Critical Review,” forthcoming: *Journal of Economic Literature*.
- Gause, G.F. (2002), “Iraq’s Decision to Go to War, 1980 and 1990,” *Middle East Journal*, 56, 47-71.
- Giannone, D., and L. Reichlin (2006), “Does Information Help Recover Structural Shocks from Past Observations,” *Journal of the European Economic Association*, 4, 455-465.
- Hamilton, J.D. (2009a), “Causes and Consequences of the Oil Shock of 2007-08,” *Brookings Papers on Economic Activity*, 1 (Spring), 215-261.
- Hamilton, J.D. (2009b), “Understanding Crude Oil Prices”, *Energy Journal*, 30, 179-206.
- Hamilton, J.D., and A.M. Herrera (2004), “Oil Shocks and Aggregate Economic Behavior: The Role of Monetary Policy,” *Journal of Money, Credit and Banking*, 36, 265-286.
- Hausman, J.A., and W.K. Newey (1995), “Nonparametric Estimation of Exact Consumers’ Surplus and Deadweight Loss,” *Econometrica*, 63, 1445-1476.
- Hughes, J.E., Knittel, C.R., and D. Sperling (2008), “Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand,” *Energy Journal*, 29, 113-134.
- Inoue, A., and L. Kilian (2011), “Inference on Impulse Response Functions in Structural VAR Models,” mimeo, University of Michigan.
- Isserlis, L. (1938), “Tramp Shipping Cargoes and Freights,” *Journal of the Royal Statistical Society*, 101(1), 53-134.
- Kellogg, R. (2011), “Short-Run Responses of Oil Production to Spot and Future Oil Prices,” mimeo, University of Michigan.
- Kilian, L. (2008a), “Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?” *Review of Economics and Statistics*, 90, 216-240.
- Kilian, L. (2008b), “The Economic Effects of Energy Price Shocks,” *Journal of Economic Literature*, 46, 871-909.
- Kilian, L. (2009a), “Not all Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market,” *American Economic Review*, 99, 1053-1069.

- Kilian, L. (2009b), "Comment on 'Causes and Consequences of the Oil Shock of 2007-08' by James D. Hamilton," *Brookings Papers on Economic Activity*, 1 (Spring), 267-278.
- Kilian, L., and B. Hicks (2010), "Did Unexpectedly Strong Economic Growth Cause the Oil Price Shock of 2003-2008?" mimeo, University of Michigan.
- Kilian, L., and D. Murphy (2010), "Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models," forthcoming: *Journal of the European Economic Association*.
- Rubio-Ramirez, J., Waggoner, D., and T. Zha. (2010) "Markov-Switching Structural Vector Autoregressions: Theory and Application", *Review of Economic Studies*, 77, 665-696.
- Schmalensee, R., and T. Stoker (1999), "Household Gasoline Demand in the United States," *Econometrica*, 67, 645-662.
- Serletis, A., G. Timilsina, and O. Vasetsky (2010), "International Evidence on Sectoral Interfuel Substitution," *Energy Journal*, 31, 1-29.
- Sweeney, J.L. (1984), "The Response of Energy Demand to Higher Prices: What Have We Learned?" *American Economic Review*, 74, 31-37.
- Terzian, P. 1985. *OPEC. The Inside Story*. Zed Books: London.
- Singleton, K.J. (2011), "Investor Flows and the 2008 Boom/Bust Cycle in Oil Prices," mimeo, Stanford University.
- Stopford, M. (1997), *Maritime Economics*. 2nd ed., London, U.K.: Routledge.
- Tinbergen, J. (1959). "Tonnage and Freight" in: *Jan Tinbergen Selected Papers*, Amsterdam: North Holland, 93-111.
- Uhlig, H. (2005) "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure," *Journal of Monetary Economics*, 52, 381-419.
- Yatchew, A., and J.A. No (2001), "Household Gasoline Demand in Canada," *Econometrica*, 69, 1697-1709.

Table 1: Sign Restrictions on Impact Responses in VAR Model

| | Flow supply shock | Flow demand shock | Speculative demand shock |
|-------------------|-------------------|-------------------|--------------------------|
| Oil production | - | + | + |
| Real activity | - | + | - |
| Real price of oil | + | + | + |
| Inventories | | | + |


Note: All shocks have been normalized to imply an increase in the real price of oil. Missing entries mean that no sign restriction is imposed.

Table 2: Posterior Distribution of the Short-Run Price Elasticity of Demand for Crude Oil

| | | $\eta^{O, Production}$ | $\eta^{O, Use}$ |
|------------------------------|-----------------------------|------------------------|-----------------|
| $\eta_t^{Supply} \leq 0.025$ | 16 th Percentile | -0.80 | -0.54 |
| | 50 th Percentile | -0.44 | -0.26 |
| | 84 th Percentile | -0.23 | -0.09 |
| $\eta_t^{Supply} \leq 0.050$ | 16 th Percentile | -0.80 | -0.57 |
| | 50 th Percentile | -0.45 | -0.27 |
| | 84 th Percentile | -0.29 | -0.09 |
| $\eta_t^{Supply} \leq 0.100$ | 16 th Percentile | -0.76 | -0.61 |
| | 50 th Percentile | -0.47 | -0.30 |
| | 84 th Percentile | -0.24 | -0.10 |


Note: Based on 112 draws from the reduced form posterior with 5 million rotations each. $\eta^{O, Production}$ refers to the impact price elasticity of oil demand in production and $\eta^{O, Use}$ to the impact price elasticity of oil demand in use. The latter definition accounts for the role of inventories in smoothing oil consumption. η^{Supply} refers to the impact price elasticity of oil supply.

Figure 1: Structural Impulse Responses: 1973.2-2009.8




Note: Results for the admissible structural model with an impact price elasticity of oil demand in use closest to the posterior median of that elasticity. Dashed lines indicate pointwise 16% and 84% posterior quantiles based on 112 draws from the reduced form posterior with 5 million rotations each. Oil production refers to the cumulative percent change in oil production and inventories to cumulative changes in inventories.

Figure 2: Historical Decompositions for 1978.6-2009.8




Note: Based on benchmark estimate as in Figure 1. The vertical bars indicate major exogenous events in oil markets, notably the outbreak of the Iranian Revolution in 1978.9 and of the Iran-Iraq War in 1980.9, the collapse of OPEC in 1985.12, the outbreak of the Persian Gulf War in 1990.8, the Asian Financial Crisis of 1997.7, and the Venezuelan crisis in 2002.11, which was followed by the Iraq War in early 2003. In constructing the historical decomposition we discard the first five years of data in an effort to remove the transition dynamics.

Figure 3: Historical Decompositions for the Persian Gulf War Episode of 1990/91




Note: Based on benchmark estimate of structural model (1) on data for 1973.2-2009.8

Figure 4: Historical Decompositions for the Iranian Revolution of 1978/79




Note: See Figure 3.

Figure 5: Historical Decompositions for the Outbreak of the Iran-Iraq War in 1980



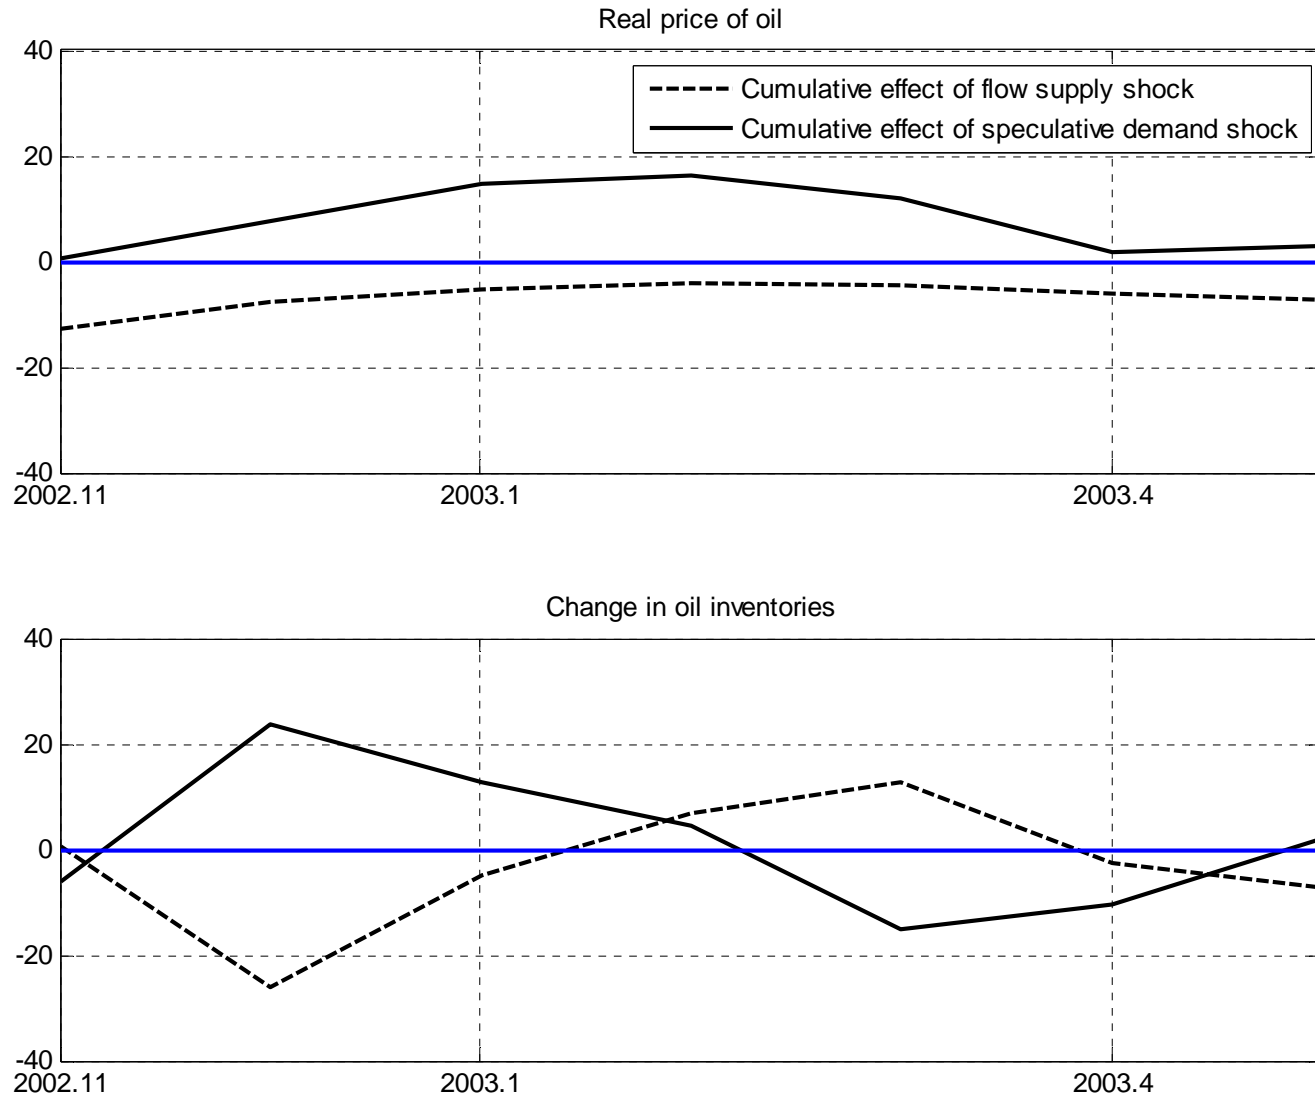
Note: See Figure 3.

Figure 6: Historical Decompositions for the Collapse of OPEC in 1986



Note: See Figure 3.

Figure 7: Historical Decompositions for Venezuelan Crisis and Iraq War in 2002/03



Note: See Figure 3.