

DISCUSSION PAPER SERIES

DP15244

**UNDERSTANDING THE ESTIMATION OF
OIL DEMAND AND OIL SUPPLY
ELASTICITIES**

Lutz Kilian

INTERNATIONAL MACROECONOMICS AND FINANCE



UNDERSTANDING THE ESTIMATION OF OIL DEMAND AND OIL SUPPLY ELASTICITIES

Lutz Kilian

Discussion Paper DP15244
Published 01 September 2020
Submitted 31 August 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- International Macroeconomics and Finance

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Lutz Kilian

UNDERSTANDING THE ESTIMATION OF OIL DEMAND AND OIL SUPPLY ELASTICITIES

Abstract

This paper examines the advantages and drawbacks of alternative methods of estimating oil supply and oil demand elasticities and of incorporating this information into structural VAR models. I not only summarize the state of the literature, but also draw attention to a number of econometric problems that have been overlooked in this literature. Once these problems are recognized, seemingly conflicting conclusions in the recent literature can be resolved. My analysis reaffirms the conclusion that the one-month oil supply elasticity is close to zero, which implies that oil demand shocks are the dominant driver of the real price of oil. The focus of this paper is not only on correcting some misunderstandings in the recent literature, but on the substantive and methodological insights generated by this exchange, which are of broader interest to applied researchers.

JEL Classification: Q43, Q41, C36, C52

Keywords: oil supply elasticity, Oil demand elasticity, IV estimation, structural VAR, Bayesian inference, oil price, gasoline price

Lutz Kilian - lkilian2019@gmail.com
Federal Reserve Bank of Dallas and CEPR

Acknowledgements

The views expressed in this paper are my own and should not be interpreted as reflecting the views of the Federal Reserve Bank of Dallas or any other member of the Federal Reserve System. I thank Ron Alquist, Ana Maria Herrera, Kunal Patel and Xiaoqing Zhou for helpful discussions. This paper builds on an earlier paper entitled "Facts and Fiction in Oil Market Modeling" by the same author. The discussion of oil market elasticities has been removed from that paper and in much expanded form is contained in the current paper, which replaces the earlier discussion.

Understanding the Estimation of Oil Demand and Oil Supply Elasticities

Lutz Kilian
Federal Reserve Bank of Dallas and CEPR

August 29, 2020

Abstract: This paper examines the advantages and drawbacks of alternative methods of estimating oil supply and oil demand elasticities and of incorporating this information into structural VAR models. I not only summarize the state of the literature, but also draw attention to a number of econometric problems that have been overlooked in this literature. Once these problems are recognized, seemingly conflicting conclusions in the recent literature can be resolved. My analysis reaffirms the conclusion that the one-month oil supply elasticity is close to zero, which implies that oil demand shocks are the dominant driver of the real price of oil. The focus of this paper is not only on correcting some misunderstandings in the recent literature, but on the substantive and methodological insights generated by this exchange, which are of broader interest to applied researchers.

JEL codes: Q43, Q41, C36, C52

Key words: Oil supply elasticity, oil demand elasticity, IV estimation, structural VAR, Bayesian inference, oil price, gasoline price.

Acknowledgments: The views expressed in this paper are my own and should not be interpreted as reflecting the views of the Federal Reserve Bank of Dallas or any other member of the Federal Reserve System. I thank Ron Alquist, Ana Maria Herrera, Kunal Patel and Xiaoqing Zhou for helpful discussions. This paper builds on an earlier paper entitled “Facts and Fiction in Oil Market Modeling” by the same author. The discussion of oil market elasticities has been removed from that paper and in much expanded form is contained in the current paper, which replaces the earlier discussion. The author has no conflict of interest to declare.

Correspondence to: Lutz Kilian, Federal Reserve Bank of Dallas, Research Department, 2200 N. Pearl St., Dallas, TX 75201, USA. E-mail: lkilian2019@gmail.com.

1. Introduction

Since Kilian and Murphy (2012), it has been known that the value of the impact price elasticity of oil supply in structural VAR models of the global oil market determines the ability of oil supply shocks to explain fluctuations in the real price of oil. In recent years there has been a heated debate about how to estimate this elasticity and the corresponding impact price elasticity of oil demand based on extraneous information. Some of this work has focused on estimating oil supply elasticities at the microeconomic level, while other studies have focused on estimating global or regional oil supply and oil demand elasticities (e.g., Newell and Prest 2019; Bjørnland, Nordvik and Rohrer 2019; Caldara, Cavallo and Iacoviello 2019).

At the same time, there is a closely related debate about how to utilize extraneous elasticity estimates in estimating structural VAR models of the global oil market. Approaches have ranged from imposing estimates of the price elasticity of oil supply directly in estimating the VAR model, as in Kilian (2009), to imposing bounds on the oil supply and/or oil demand elasticities, as in Kilian and Murphy (2012, 2014), or minimizing the distance between VAR elasticities and extraneous oil demand and oil supply elasticity estimates (Caldara et al. 2019), or explicitly specifying informative elasticity priors (Baumeister and Hamilton 2019). These competing approaches imply a wide range of oil demand and oil supply elasticity estimates and conflicting results about the relative importance of oil demand and oil supply shocks for the real price of oil.

It is important to note that what we care about in oil market VAR models is typically the magnitude of the one-month price elasticities of oil supply and demand. At longer horizons, the elasticities are left unconstrained. The purpose of this paper is to examine the advantages and drawbacks of alternative methods of estimating these one-month oil supply and oil demand

elasticities and of incorporating this information into structural VAR models. I not only summarize the state of the literature, but also draw attention to a number of econometric problems that appear to have been overlooked in this literature. I make the case that, once these problems are recognized, seemingly conflicting conclusions in the recent literature can be resolved. I also show how additional information can be used to assess the economic plausibility of elasticity estimates reported in the literature. My analysis reaffirms the conclusion that the one-month oil supply elasticity is low, which implies that oil demand shocks are the dominant driver of the real price of oil.

The remainder of this paper is organized as follows. Section 2 reviews the identification problem in the global market for crude oil and explains why extraneous elasticity estimates are crucial for estimating oil market models. Section 3 reviews the microeconomic evidence on the U.S. price elasticity of oil supply. While these estimates provide a useful benchmark, it is important to recognize that the supply elasticities of other oil producing countries such as Saudi Arabia, may be larger than the U.S. elasticity. Ultimately, we are interested in the global price elasticity of oil supply rather than in regional elasticities. Section 4 explains in detail the construction of the bound on the global oil supply elasticity originally proposed in Kilian and Murphy (2012) and addresses recent critiques of this approach. Section 5 examines the IV approach to estimating global oil supply and oil demand elasticities proposed by Caldara et al. (2019). Section 6 systematically evaluates the credibility of recent VAR estimates of oil demand and oil supply elasticities. Section 7 discusses subtle differences in the elasticity concept employed in recent VAR studies and their implications. The concluding remarks are in section 8.

2. The Identification Problem in Models of the Global Market for Crude Oil

As noted by Caldara et al. (2019), the same set of unexpected changes in oil production and in

the real price of oil can be explained by many different combinations of short-run oil supply and oil demand curves with different slopes. A model with a vertical short-run supply curve and a downward sloping short-run demand curve, for example, may explain the residuals of a reduced-form VAR for the global market as well as a model with an upward sloping supply curve and a downward sloping demand curve. These slopes are captured by the impact price elasticities of oil supply and oil demand. Being able to pin down at least one of these elasticities based on extraneous data helps estimate the other elasticity. Alternatively, one may be able to bring extraneous information on both of these impact elasticities to bear.

The conventional approach to estimating these elasticities is to rely on an exogenous instrument. For example, the identification of the impact price elasticity of oil demand (η^d) requires an instrument that shifts the global supply curve along the global demand curve. Such exogenous instruments have been developed by Kilian (2008a), for example.¹ However, evidence in Kilian (2008a,b) and Montiel Olea, Stock and Watson (2020) suggests that all existing oil supply shock instruments are weak in the econometric sense, which explains why this approach has been largely abandoned. Similarly, the identification of the impact price elasticity of oil supply (η^s) requires an instrument that shifts the global demand curve along the global supply curve. Since suitable demand instruments for the global oil market are hard to find, this approach has not received much attention.²

Given the difficulty of estimating global impact price elasticities directly, the literature has evolved in three directions, each of which is discussed below: (1) the estimation of the price

¹ An alternative ad hoc approach to this problem is discussed in Hamilton (2003).

² Examples of oil demand instruments are discussed in Kilian and Hicks (2013) and Känzig (2019). The former study focuses on exogenous changes in flow demand, whereas the latter focuses on exogenous changes in storage demand (see Kilian and Zhou 2020c). Neither of these instruments is available for extended time periods, however.

elasticity of oil supply using U.S. firm-level or well-level data; (2) bounds on the global price elasticity of oil supply constructed from natural experiments; and (3) the estimation of the global price elasticities of oil supply and oil demand using country-level panel data.

3. Microeconomic Estimates of the Oil Supply Elasticity

Microeconomic estimates of the oil supply elasticity are informative for the identification of oil market models because they constitute extraneous evidence. Whereas some earlier studies focused on U.S. oil producer data from selected regions such as Texas or North Dakota, Newell and Prest (2019) for the first time included data from all major oil producing regions in the United States, making it the most comprehensive study to date and a natural starting point.

The Evidence in Newell and Prest (2019)

Newell and Prests's preferred estimate of the one-quarter oil supply elasticity for conventional crude is 0.017 (with a standard error of 0.006) also provides an upper bound on the one-month elasticity. Their estimate is close to the benchmark provided by the theoretical analysis in Anderson, Kellogg and Salant (2018) who showed within an equilibrium model that the short-run oil supply elasticity is zero if adjusting oil production is costly, as tends to be the case in practice.

One might think that this conclusion would be changed when incorporating the one-month supply elasticity of shale oil producers. A common view is that the latter elasticity is at least as large as that for conventional crude oil. There are two points that must be kept in mind, however. One point is that even for shale oil the one-quarter supply elasticity is negligible. Newell and Prest (2019) report an estimate of -0.022 (with a standard error of 0.013) for U.S. shale oil producers. This estimate is close to zero and not statistically significant. This result does not contradict the widely held view that shale oil producers are nimbler in responding to market

conditions than conventional producers. It simply means that this response takes more than one month for shale oil producers and conventional producers alike. Even if a producer wants to complete an existing shale well in response to higher oil prices, it typically takes at least four weeks to start production.³ The view that the one-month shale oil supply elasticity is effectively zero is also consistent with survey evidence for oil producers in the 11th Federal Reserve District published by the Federal Reserve Bank of Dallas, as summarized in Golding (2019).⁴

The other point is that shale oil production did not exist for much of the estimation period considered in oil market studies. It only took off in late 2008 from a very low level and accounted only for a small share of world oil production even in the years after 2008, so we need to consider appropriately weighted averages in inferring the implied global oil supply elasticity. If we take the U.S. estimates as representative for oil producers in the world, given a share of 6% for shale oil production in global oil production in 2019, this implies a global one-month oil supply elasticity of under 0.016.⁵ We do not know the covariance of the elasticity estimates for conventional oil and for shale oil, but even after accounting for estimation uncertainty the upper bound of a confidence interval for this global oil supply elasticity estimate is unlikely to be much larger.

The Evidence in Bjørnland et al. (2019)

The evidence in Newell and Prest (2019) has been called into question by Bjørnland et al. (2019) and Bjørnland (2019), whose estimate of the one-month price elasticity of oil supply for

³ Note that what matters in measuring this elasticity is not how much time has passed before a new shale well is completed, but how much time it takes for a well to be completed in response to an unexpected change in the price of oil.

⁴ Golding (2019) reports that “the average horizontal well pad in the Permian Basin takes four to six months from the commencement of drilling to production coming online.” Moreover, even a “drilled-but-uncompleted well ... may take one to three months to go into production”. In other words, industry contacts are unanimous that the one-month price elasticity of shale oil supply is zero.

⁵ Specifically, $0.94 \times 0.017 + 0.06 \times 0 = 0.016$, where we treat the shale oil supply elasticity as zero rather than using the point estimate of -0.022.

conventional oil producers, obtained by regressing the change in conventional oil production in the Bakken on the change in the real price of oil and on the change in the oil futures spread, is 0.1 and statistically insignificant.⁶ Their corresponding baseline estimate of the one-month supply elasticity for shale oil is 0.7 (with additional estimates ranging from 0.3 to 0.9).

A closer look at their regressions, however, raises a number of econometric concerns.

The model is

$$\Delta q_{it} = \beta_1 \Delta p_t + \beta_2 \Delta(p_t - f_t^{(3)}) + \dots + e_{it}, \quad (1)$$

where q_{it} denotes the log of oil production, p_t is the log of the spot price of oil, and $f_t^{(3)}$ is the log of the 3-month oil futures price. Equation (1) may be equivalently rewritten as

$$\Delta q_{it} = \alpha_1 \Delta p_t + \alpha_2 \Delta f_t^{(3)} + \dots + e_{it}, \quad (2)$$

where $\alpha_1 = \beta_1 + \beta_2$. Bjørnland et al. suggest that $\hat{\alpha}_1$ represents the price elasticity of oil supply.

Given that the oil futures spread is clearly stationary, it is unclear why the authors include the first difference of this spread in their regression.⁷ This specification choice has important consequences because it amounts to augmenting Δp_t in regression (2) by the regressor $\Delta f_t^{(3)}$. A high correlation between these regressors is problematic because it creates multi-collinearity and undermines the identification of α_1 . As Table 1 shows, for the estimation period of 1990.2-2017.6 used in this study., the correlation between Δp_t and $\Delta f_t^{(3)}$ is 98%. The high correlation

⁶ This estimate is higher than that obtained by Newell and Prest (2019), but also has a larger standard error, reflecting the smaller estimation sample and the existence of multicollinearity.

⁷ Bjørnland (2019) insists that the oil futures spread is integrated of order 1, making it necessary to difference the spread. This view is at odds with a large literature in empirical finance on modeling futures spreads. A plot of $f_t^{(3)} - p_t$ suggests that the three months oil futures spread is mean reverting. This visual impression is confirmed by a formal ADF test based on bootstrap critical values computed using the residual wild bootstrap to allow for conditional heteroskedasticity. Notwithstanding the low power of ADF tests in finite samples, the null of a unit root can be rejected at the 10% level.

between these regressors is the likely reason that including Δf_t^3 among the regressors changes the estimate of α_1 substantially compared to setting $\alpha_2 = 0$. Thus, the chief concern with regression specification (1) is not about using futures prices instead of spot prices. Table 1 makes it clear why it does not make much difference, in practice, whether we use Δp_t or $\Delta f_t^{(h)}$. The concern is that the elasticity should relate to the spot price or the futures price, but not to both prices.

Is there any justification for including the oil futures spread? One argument is that changes in the futures spread may affect oil production directly if producers respond to higher expected oil prices by storing oil below the ground rather than extracting it. For conventional oil, for technological reasons, this is not an option (see Newell and Prest 2019). For shale oil, producers could conceivably drill, but not complete a well in anticipation of rising prices. A simple testable implication of this argument is that the number of drilled, but not yet completed shale oil wells (known as DUCs) should increase, when the oil futures price rises above the spot price, making it more attractive to delay completions. However, the correlation between the growth in the number of DUCs in the Bakken, as reported by the EIA, and the 3-month WTI oil futures spread, defined as $f_t^{(3)} - p_t$, where lower case denotes logs, is only 0.03, suggesting that this effect is negligible. Using alternative data on the number of DUCs in the Bakken by Rystad Energy, the correlation in question is even of the wrong sign. As shown in Figure 1, during 2008.1-2017.6, the correlation is -0.15. As the estimation period is shortened, the correlation becomes more negative.

The fact that shale producers appear unable in the short run to adjust the number of DUCs in response to the futures spread is also consistent with the survey evidence in Golding (2019). Shale oil producers in this survey cite frictions in setting up well completions as the reasons for

not responding sooner to price signals. This evidence also suggests that the theoretical analysis in Anderson et al. (2008), which is based on adjustment costs to oil production and predicts a zero short-run oil supply elasticity, applies both to conventional oil and to shale oil.

4. Bounding the Global Oil Supply Elasticity

The fact that the U.S. one-month price elasticity of oil supply is near zero does not necessarily mean that the same is true of the corresponding global one-month price elasticity, which is of primary interest in oil market VAR models. Intuitively, it seems plausible that oil producing countries with spare capacity such as Saudi Arabia or Kuwait may have a higher oil supply elasticity than the United States.

Kilian and Murphy (2012) addressed this concern by constructing an upper bound on the global one-month price elasticity of oil supply based on the natural experiment of August 1990, when Iraq invaded Kuwait and oil production in these two countries ceased. Since there has been much confusion in the literature about the construction of this bound, it is useful to explain the rationale in more detail. The invasion of Kuwait was clearly an exogenous event with respect to the oil market and by all accounts unanticipated. The oil supply disruption of August 1990 boosted the demand for oil produced outside of Iraq and Kuwait. These countries' oil-demand curve was further shifted by a sharp rise in storage demand, reflecting expectations that Iraq would invade Saudi Arabia next. This reasoning may seem to suggest that the ratio of the percent change in oil production outside Iraq and Kuwait (Δq) to the percent change in the real price of oil (Δp) in August 1990 can be thought of as an estimate of the global one-month price elasticity of oil supply.

However, there is one complication. Saudi Arabia's supply also expanded in response to this geopolitical event, as part of Saudi Arabia's long-standing commitment to respond directly

to geopolitically driven oil supply disruptions in other OPEC member countries, possibly along with other OPEC oil producers.⁸ The simultaneous shift in Saudi Arabia's supply curve in August 1990 created an additional increase in Δq and a decline in Δp , causing the ratio $\Delta q / \Delta p$ to be larger than would have been the case in response to the demand shift only. Kilian and Murphy therefore interpreted the ratio $\Delta q / \Delta p = 0.026$ as an upper bound on the one-month price elasticity of oil supply rather than an estimate of this elasticity.

Not all OPEC oil producers outside of Iraq and Kuwait raised their production in August 1990. In the United Arab Emirates (UAE) notably, oil production fell. One view is that this production decline reflected increasing pressure from OPEC members for the UAE to adhere to its OPEC production quota and was unrelated to the invasion of Kuwait. An alternative view, suggested by Caldara et al. (2019), is that this decline was caused by a speech by Saddam Hussein on July 17, threatening retribution if unspecified OPEC countries did not reduce their oil production. The latter view implies that the decline in UAE's oil production in August 1990 was caused by an exogenous geopolitical event not unlike the invasion of Kuwait and hence has to be excluded when constructing the endogenous production response Δq , which would increase the upper bound on the one-month oil supply elasticity from 0.026 to 0.045.

This alternative bound, however, is not persuasive for two reasons. First, the UAE already agreed to lower its oil production at the OPEC meeting in Jeddah on July 11 several days before Saddam Hussein's speech, casting doubt on a causal link. Second, at no point was there an immediate military threat to the UAE, which has no direct border with Iraq. In fact, Iraq lacked

⁸ It is important to stress that this commitment does not mean that Saudi Arabia feels compelled to offset fluctuations in the demand of oil. For example, Saudi authorities made it clear in the 2000s that they would not respond to oil price increases driven by what they perceived to be shifts in speculative demand for oil, although they have always been willing to respond to exogenous oil supply disruptions driven by geopolitical events.

the ability to effectively project military force across the Persian Gulf to the UAE by air or sea, given the presence of U.S. and other opposing forces in the region. Thus, the UAE must be included in constructing the oil supply elasticity bound.

5. Macroeconomic Estimates of the Global Oil Demand and Oil Supply Elasticities

Motivated by the analysis in Kilian and Murphy (2012), Caldara et al. (2019) focus on the response of oil production in a given country to supply disruptions in other oil-producing countries, under the maintained assumption that all oil producers have the same supply elasticity. Their instrument for the real price of oil consists of a time series of oil supply disruptions in the United States, Mexico, Venezuela, Norway, Iran and various Arab oil producing countries that are classified as exogenous based on the authors' reading of the narrative evidence. In addition to an estimate of the global oil supply elasticity, the study also reports disaggregate oil supply elasticity estimates for Saudi Arabia, for OPEC excluding Saudi Arabia, and for non-OPEC countries.

It is useful to review this approach in more detail. Ignoring country-fixed effects, for expository purposes, the first-stage IV regression is

$$\Delta p_{i,t} = \gamma Z_{i,t} + \varepsilon_{i,t},$$

where $Z_{i,t}$ is the instrument for oil-producing country i constructed by interacting declines in oil production growth that take place in other oil-producing countries with a dummy indicating whether this decline is driven by exogenous events such as weather, strikes or wars. The fitted value from the first stage, $\widehat{\Delta p}_{i,t}$, is used in the second-stage IV regression

$$\Delta q_{i,t}^s = \eta^s \widehat{\Delta p}_{i,t} + u_{i,t}^s$$

to identify the one-month price elasticity of oil supply, where $\Delta q_{i,t}^s$ denotes oil production growth

in country i .

Caldara et al. report estimates for a narrow instrument including only oil supply disruptions of at least 2% of global oil production and a broad instrument including in addition a number of smaller oil supply disruptions. After excluding the August 1990 decline in UAE oil production from $Z_{i,t}$, their estimate of the one-month price elasticity of oil supply is 0.029 based on the narrow instrument and 0.056 based on the broad instrument. This compares to supply elasticity estimates of 0.054 and 0.081 for the narrow and the broad instrument, respectively, when the UAE is included in $Z_{i,t}$.

Even the estimates obtained after excluding the UAE from $Z_{i,t}$, however, are problematic. The first problem relates to the instrument relevance. Although the narrow instrument passes an F-test for weak instruments in the first stage, when regressing the percent change in the real price of oil on an intercept and the instrument for 1985.1-2015.12, the broad instrument (with or without the UAE) does not. In fact, after excluding the August 1990 episode from the broad instrument, the first-stage F-statistic drops below 0.9 (see Table 2).

The second problem relates to the exogeneity of the instrument. Since Saudi Arabia aims to directly offset geopolitical disruptions, both its oil supply curve and its oil demand curve shift in response to such an event, which violates the exclusion restriction required for IV estimation. The use of $Z_{i,t}$ is not only a problem when interpreting the Saudi elasticity estimate; the same concern also applies to other OPEC producers with spare capacity such as Kuwait and the UAE that have at times acted in line with Saudi Arabia in offsetting geopolitical oil supply disruptions. Not surprisingly, the elasticity of non-OPEC countries, which is not affected by this problem, is essentially zero (-0.004) with a standard error of 0.023.

A simple back-of-the-envelope calculation based on a specific episode of exogenous variation in the real price of oil helps illustrate this point. Between June 2014 and December 2014, the price of oil fell by 44%. There is a debate about the extent to which this price decline was caused by the unexpected rise of U.S. shale oil versus unexpected declines in global demand. Either way this unexpected decline was exogenous from Saudi Arabia's point of view. Given that Saudi Arabia did not respond to any exogenous geopolitical events during this half year, the Saudi production response can be used to identify the oil supply elasticity. Given the cumulative decline in Saudi oil production of 0.6%, the implied semi-annual Saudi oil supply elasticity is $-0.6/44 = 0.014$, which is much lower than the supply elasticity estimate of 0.212 for Saudi Arabia in Caldara et al (2019). The corresponding semi-annual oil supply elasticity estimate for OPEC is also zero for all practical purposes.

For the estimation of the corresponding one-month oil demand elasticity, Caldara et al. (2019) propose a similar IV approach. Consider an instrument $Z_{j,t}$ for the oil-consuming country j , consisting of exogenous foreign oil production disruptions. Events that are associated with an exogenous shift of both oil demand and oil supply in a given country such as Hurricanes along the U.S. Gulf coast that affect both refining and off-shore oil production are excluded. The first-stage regression

$$\Delta p_{j,t} = \gamma Z_{j,t} + \varepsilon_{j,t}$$

yields the fitted value $\widehat{\Delta p}_{j,t}$, which is used to estimate the price elasticity of oil demand from the second-stage IV regression

$$\Delta q_{j,t}^d = \eta^d \widehat{\Delta p}_{j,t} + \psi X_{j,t} + u_{j,t}^d,$$

where $\Delta q_{j,t}^d$ denotes petroleum consumption in country j and $X_{j,t}$ denotes controls for the state of the oil market and the global economy. It is assumed that all countries have the same demand

elasticity.

An additional problem with this elasticity estimator is that the second-stage regression for demand requires data for oil consumption. Such data do not exist. The “oil consumption” data reported by the International Energy Agency, for example, do not measure consumption of crude oil, but the consumption of refined products such as diesel, gasoline, jet fuel, bunker fuel and other residual products in refining. This means that the demand elasticity in the second-stage IV regression is not the own price elasticity of the demand for crude oil, but a cross-price elasticity of the demand for refined products.

How plausible are the IV elasticity estimates for the global oil market?

Table 3 summarizes the key oil supply and oil demand elasticity estimates generated by this IV approach. The authors’ preferred estimates based on the broad instrument are in the first row of Table 3. Since this instrument is weak, the most plausible estimates are those based on the narrow instrument excluding the UAE. The latter approach yields a one-month oil supply elasticity estimate of 0.03, close to the bound derived in Kilian and Murphy (2012).⁹

In assessing the oil demand elasticity estimates in Table 3, we can draw on additional evidence. There have been major methodological advances in estimating the one-month price elasticity of gasoline in recent years. State-of-the-art estimates based on data from the United States and Japan agree that this elasticity of demand is near -0.36 (e.g., Coglianesi, Davis, Kilian and Stock 2017; Levin, Lewis and Wolak 2017; Knittel and Tanaka 2019).¹⁰ Hamilton (2009) argues that the price elasticity of oil demand should be approximately half as large as the price

⁹ Including the UAE in the instrument inflates this estimate to near 0.05, which is close to the revised bound that can be derived under the same assumption.

¹⁰ Whereas the demand elasticity in Coglianesi et al. (2017) has a large standard error, the demand elasticity in the other two studies is precisely estimated. For example, Knittel and Tanaka (2019) report a standard error of only

elasticity of gasoline demand, given the cost share of oil in producing gasoline. If we take this argument at face value, the oil demand elasticities reported in Caldara et al. (2019) seem too low to be economically plausible.

6. Structural VAR Estimates of the Global Oil Supply and Oil Demand Elasticities

An alternative approach to estimating the impact price elasticities of oil supply and oil demand is by estimating structural VAR models. This requires restricting at least one of these elasticities. It is important to emphasize that pinning down one of these elasticities is not necessarily enough to infer the other, as illustrated by the example of Kilian (2009) who restricted the impact price elasticity of oil supply to zero in estimating a model of the global market for crude oil that includes global oil production growth, an index of cyclical variation in the global business cycle, and the real price of oil.¹¹ In that model, the impact price elasticity of oil demand cannot be defined explicitly. Attempts to compute this elasticity ignoring changes in inventories, as proposed by Baumeister and Hamilton (2019), will produce nonsensical estimates of the demand elasticity.

The reason is simple. In a global oil market model, the amount of oil produced in a given period may be consumed in a refinery or put into storage (see Kilian and Murphy 2014). Since oil is a storable commodity, defining the oil demand elasticity based on an accounting identity that equates oil production with oil consumption at each point in time is obviously incorrect. This problem arises even when using the alternative definition of the oil demand elasticity favored by Baumeister and Hamilton (2019), as discussed in section 7, because the structural model they use

0.03. An overview of this literature that explains why recent estimates of this elasticity differ dramatically from earlier IV estimates such as Hughes, Knittel and Sperling (2008) can be found in Kilian and Zhou (2020a).

¹¹ Kilian (2009) motivates this restriction by observing that OPEC historically has been slow to respond to cyclical demand fluctuations. Given that shifts in global oil demand are difficult to estimate reliably in real time and that changing oil production is costly, it makes sense for oil producers to delay production responses.

to interpret the Kilian (2009) estimates is misspecified due to the omission of oil inventories. Similar comments apply to Baumeister and Hamilton's reinterpretation of the sign-identified oil market models of Kilian and Murphy (2012) and Inoue and Kilian (2013) that bound the price elasticity of oil supply. It is easy to show that when augmenting this and similar structural models to include oil inventories, as in Kilian and Murphy (2014), the oil demand elasticity is much more reasonable than the estimates reported in Baumeister and Hamilton (2019).

Kilian and Murphy (2014)

The first VAR study to jointly estimate the price elasticities of oil supply and oil demand was Kilian and Murphy (2014). The approach taken by this paper was to bound the one-month price elasticity of oil supply by $0 < \eta^s < 0.026$. Kilian and Murphy also bounded the one-month price elasticity of oil demand to be lower than the long-run price elasticity of oil demand, which can be proxied for by microeconomic estimates of long-run price elasticity of gasoline demand in Hausman and Newey (1995) and Yatchew and No (2001) such that $-0.8 < \eta^d < 0$. These elasticity bounds in conjunction with sign and narrative inequality restrictions allow the estimation of the impact price elasticities of oil demand and oil supply using standard econometric methods for sign-identified models (see Antolin-Diaz and Rubio-Ramirez 2018).

Table 4 shows alternative elasticity estimates for this model based on different econometric methodologies and estimation periods. The correctly computed oil demand elasticity is inferred from the responses of the real price of oil and of oil consumption, defined as oil production adjusted for the change in oil inventories, to an oil supply shock. The incorrectly defined elasticity is based on equating oil consumption with oil production. As the first row in Table 4 shows, this model implies an estimate of the one-month oil supply elasticity of 0.01. More importantly, the correctly computed one-month oil demand elasticity is -0.26, whereas the

incorrectly computed demand elasticity is much higher. There are two key takeaways here. First, contrary to the claims in the recent literature, a low oil supply elasticity need not imply a high oil demand elasticity. Second, it is essential to define the price elasticity of oil demand correctly. The incorrect estimate is much larger than the corresponding extraneous estimates of the price elasticity of gasoline demand, whereas the correctly defined oil demand elasticity is much smaller and hence more consistent with the extraneous evidence.

The second row shows that this estimate is quite robust to changes in the estimation period and in the econometric methodology. It shows updated estimates of this model from Herrera and Rangaraju (2020), evaluated under Dirac delta loss (see Inoue and Kilian 2013, 2019). Using the same specification as Kilian and Murphy (2014) produces an estimate of 0.01 for the oil supply elasticity and a demand elasticity of -0.28 (correctly defined) and -0.51 (incorrectly defined), confirming the robustness of the original results.

The third row shows similarly updated estimates in Inoue and Kilian (2020) who recently re-estimated this model using a state-of-the-art Bayesian methodology that allows us to evaluate the posterior of the structural impulse response vector under absolute loss. They also relaxed the upper bound on the oil supply elasticity to 0.04 and replaced the indicator of global real economic activity, as proposed in Zhou (2020). The implied estimate of the price elasticity of oil supply is 0.01 and that for the (correctly defined) oil demand elasticity is -0.18, which is about half as large as extraneous estimates of the gasoline price elasticity. The estimate for the incorrectly defined elasticity of oil demand that equates oil production and oil consumption is -0.47. As Table 4 shows, none of these estimates are very far from the original estimates reported in Kilian and Murphy (2014).

Caldara, Cavallo and Iacoviello (2019)

An alternative approach proposed by Caldara et al. (2019) is to minimize the Euclidian distance between the VAR admissible impact elasticities and the elasticities estimated by the IV approach discussed in section 5. After expressing the demand elasticity as a function of the supply elasticity and the reduced-form error covariance matrix, Σ , this problem reduces to

$$\min_{\eta^s} \begin{bmatrix} \eta^s - \hat{\eta}^s \\ \eta^d(\eta^s, \Sigma) - \hat{\eta}^d \end{bmatrix} W^{-1} \begin{bmatrix} \eta^s - \hat{\eta}^s \\ \eta^d(\eta^s, \Sigma) - \hat{\eta}^d \end{bmatrix},$$

where $\hat{\eta}^s$ and $\hat{\eta}^d$ denote the IV elasticity estimates and W is a diagonal weighting matrix with the diagonal entries correspond to the variances of the IV elasticity estimates.

There are two caveats. First, all the concerns raised earlier about the IV estimator extend to the VAR elasticity estimator. Second, the validity of this estimator depends on the function $\eta^d(\eta^s, \Sigma)$ which reflects the VAR model structure. The structure of the oil market model used to derive $\eta^d(\eta^s, \Sigma)$ in Caldara et al. (2019) is inconsistent with the structure of the global oil market because their global market clearing condition equates oil production with oil consumption in every period, ignoring that oil is storable. The correct market clearing condition is that the quantity of oil produced in the world equals the quantity of oil consumed by refiners plus the accumulation of oil inventories. Since the structural VAR model does not account for the fact that countries can reduce oil stocks to deal with oil production shortfalls in addition to reducing oil consumption, there is an omitted variable problem (see Kilian and Murphy 2014).¹²

The estimates reported in Caldara et al. (2019) utilize the IV elasticity estimates obtained using the broad instrument that does not pass the weak instrument test, which suggests additional

¹² Caldara et al. (2019) acknowledge this problem, but insist that oil inventories are not important. As discussed in section 7, however, estimates of the supply and demand elasticities implied by structural models that account for oil inventories are inconsistent with the relationship $\eta^d(\eta^s, \Sigma)$ derived in Caldara et al.

caution in interpreting the estimates. As Table 5 shows, the implied VAR elasticity estimates are 0.1 for the oil supply elasticity and -0.14 for the oil demand elasticity.

Baumeister and Hamilton (2019)

An alternative approach in the literature has been to rely on Bayesian methods that allow us to specify explicit priors on the one-month price elasticities of oil demand and oil supply (see Baumeister and Hamilton 2019).¹³ Baumeister and Hamilton argue that the magnitude of the oil supply elasticity estimates in the literature and the uncertainty surrounding these estimates justifies their choice of a diffuse oil supply elasticity prior which allows for oil supply elasticity values in the range $[0, \infty]$. Specifically, they appeal to Bjørnland et al.'s (2019) regional shale oil supply elasticity estimates of 0.3-0.9. They also appeal to the global supply elasticity estimate in Caldara et al. (2019) based on the broad instrument. The drawbacks of these estimates have been discussed in detail in earlier sections.

Baumeister and Hamilton make the case that we also need to account for uncertainty about extraneous elasticity estimates. This argument is sensible, but does not justify an unbounded prior distribution for η^s . For example, if we add two standard errors to Caldara et al.'s (2019) estimate of the supply elasticity bound based on the narrow instrument including the UAE, we obtain an upper bound of about 0.09, which is far from ∞ and well below the posterior supply elasticity estimate of 0.15 reported by Baumeister and Hamilton (see Table 5). Herrera and Rangaraju (2020) show that bounding the oil supply elasticity by 0.1 reduces the posterior median estimate of the supply elasticity in Baumeister and Hamilton's model from 0.15 to 0.08

¹³ As shown in Inoue and Kilian (2020), this approach implies highly unrealistic priors on the impulse response functions. For a more detailed review of the econometric approach of Baumeister and Hamilton (2019) and its drawbacks see Kilian and Zhou (2020c) and Kilian (2020). For a discussion of the related debate about how to measure oil inventories and global real activity the reader is referred to Kilian and Zhou (2018, 2019), Funashima (2020) and Kilian (2019, 2020).

and raises the oil demand elasticity estimate to -0.53. This bound is still conservative in that the UAE has not been excluded from the instrument.

If, instead, we start with the conventional view that the one-month oil supply elasticity is positive, but close to zero, a much tighter bound is obtained. For example, the posterior estimate of 0.15 in Table 5 is 22 standard errors above the point estimate for the United States in Newell and Prest (2019). As observed by Zhou (2020), even adding four standard errors to the Newell and Prest (2019) estimate results in a supply elasticity bound of only 0.04, which would reduce Baumeister and Hamilton's supply elasticity estimate from 0.15 to 0.03 (see Herrera and Rangaraju 2020).¹⁴

In short, there is no justification for the supply elasticity priors underlying the analysis in Baumeister and Hamilton (2019) and their posterior supply elasticity estimate is highly sensitive to imposing tighter bounds based on extraneous information. Likewise, the corresponding posterior elasticity estimates reported in Table 5 are difficult to reconcile with extraneous gasoline demand elasticity estimates. For example, the estimate of -0.35 for the oil demand elasticity is at odds with Hamilton's (2009) rule of thumb that the oil demand elasticity should be half as large as extraneous gasoline demand elasticity estimates, regardless of what extraneous estimate one appeals to. Section 7 discusses in more detail why coming up with elasticity priors that are informed by extraneous elasticity estimates is particularly difficult in Baumeister and Hamilton's framework.

7. Alternative Definitions of the Global Price Elasticities

¹⁴ In related work, Bornstein et al. (2019) estimate the one-year global price elasticity of oil supply. Since the one-month elasticity cannot be larger than the one-year elasticity, their estimate of 0.18 is fully consistent with a one-month elasticity near zero, but casts doubt on the view that the one-month global oil supply elasticity could be as large as 0.15.

An interesting point that calls for discussion is the difference in the definition of the elasticity concept used in Baumeister and Hamilton (2019) and related studies, on the one hand, and that used by studies building on Kilian and Murphy (2014), on the other. My discussion of this point focuses on the one-month price elasticity of oil supply for expository purposes.

Baumeister and Hamilton define the oil supply elasticity as the impact response of oil production to an increase in the real price of oil triggered by an exogenous demand shift, holding constant not only the remaining structural shocks, but also all other variables in the model such as global real economic activity and oil inventories. In contrast, Kilian and Murphy (2014) define the one-month price elasticity of oil supply as the ratio of the impact response of oil production to the impact response in the real price of oil triggered by an exogenous demand shift, with all other structural shocks set to zero. This allows global real activity and oil inventories to respond contemporaneously to the exogenous demand shift. The latter changes in turn may affect the quantity produced. Clearly, these elasticity concepts in general are neither numerically nor conceptually equivalent.¹⁵

The key difference is that Baumeister and Hamilton's elasticity definition is a theoretical construct that one is not likely to observe in reality since both global real activity and oil inventories will in general move on impact in response to a demand shock. Most extraneous elasticity estimates reported in the literature are not constructed to be consistent with this elasticity definition. For example, the elasticity bound derived by Kilian and Murphy (2012) or the IV supply elasticity estimator in Caldara et al. (2019) do not hold constant the remaining

¹⁵ A practical difference is that Baumeister and Hamilton's approach ensures by construction a unique estimate of the oil supply elasticity, whereas Kilian and Murphy's approach produces two estimates of the oil supply elasticity that need not be identical, one in response to the flow demand shock and one in response to the storage demand shock. Given that these estimates in practice tend to differ only by one second decimal point, however, there is little loss in generality in reporting an average estimate as in Table 4.

model variables. The same is true for the microeconomic estimates of the oil supply elasticity in Newell and Prest (2019). Thus, it makes sense to choose an elasticity definition that corresponds to empirical elasticity estimates in the literature, which is what Kilian and Murphy (2014), along with many other researchers, have done. This approach ensures the consistency of the internal and the extraneous elasticity estimate. In contrast, Baumeister and Hamilton cannot appeal to these extraneous elasticity estimates (or elasticity bounds) to motivate their prior specification because these estimates are inconsistent with their own elasticity definition. It is unclear how to work around this problem when designing the elasticity priors within their framework.

The distinction between these elasticity concepts also matters for understanding the relationship between oil demand and oil supply elasticities. For example, if we adopt the elasticity concept of Baumeister and Hamilton (2019), the relationship between the oil supply and the oil demand elasticity may be derived analytically from the underlying structural oil market model. In general, small positive oil supply elasticities are associated with large negative oil demand elasticities. This relationship forces the value of the oil supply elasticity to be quite high, if we want to avoid generating large negative values for the oil demand elasticity.¹⁶ In contrast, given the elasticity concept in Kilian and Murphy (2014), oil supply elasticities close to zero (say, 0.01) may coexist with quite small oil demand elasticities (say, -0.18), after taking account of the response of oil inventories (see Table 4). The latter values can also be matched against extraneous elasticity estimates in the literature, whereas extraneous elasticity estimates that satisfy Baumeister and Hamilton's elasticity definition are harder to come by, making it

¹⁶ The precise tradeoff is model-specific. For example, according to the theoretical relationship between η^s and η^d derived in Caldara et al. (2019, Figure 2), which underlies their VAR minimum distance elasticity estimate, an oil supply elasticity of 0.15 should be associated with an oil demand elasticity of about -0.1. In contrast, Baumeister and Hamilton's estimate of the oil demand elasticity in Table 5, which is based on a model that includes oil inventories, with -0.35 is much higher than -0.1, illustrating that the relationship derived in Caldara et al. (2019) does not hold in general.

difficult to evaluate their elasticity estimates.¹⁷

In the end what matters for choosing between these elasticity definitions is what question we are interested in. For example, if the question is how much oil consumption responds to an exogenous supply disruption raising the real price of oil, the approach of Kilian and Murphy (2014) provides the correct answer. If we are interested in how oil consumption responds to the same shock, holding constant global real activity, Baumeister and Hamilton's definition is appropriate (with the important caveat that applications of their definition to oil market models excluding oil inventories are invalid).¹⁸

Either approach, in principle, may be used to constrain the VAR impulse responses to ensure the economic plausibility of the VAR impulse response estimates. What matters is that the elasticity concept used in estimating the structural model matches that in constructing the extraneous estimates used to constrain the model estimates. At this point, the existence of extraneous elasticity estimates and bounds that are estimated without holding constant the other oil market model variables makes it easier to apply and evaluate the Kilian and Murphy (2014) approach.

8. Conclusion

The value of the price elasticities of oil supply and oil demand in structural VAR models of the global market for crude oil largely determine the relative importance of oil supply and oil

¹⁷ In this context, it should be also noted that Caldara et al.'s claim that the relative contribution of oil demand and oil supply shocks to the variation the real price of oil sharply varies, as the oil supply elasticity is increased modestly, is at odds with the findings in a range of other studies including Kilian and Murphy (2012) based on their elasticity definition.

¹⁸ Baumeister and Hamilton maintain that this is not a problem since inventories in their view are quantitatively unimportant. This view is at odds with the evidence presented in Kilian and Murphy (2014) and several related studies including Kilian and Lee (2014), Kilian and Zhou (2020b), and Cross, Nguyen and Tran (2019). Indeed, even the evidence in Baumeister and Hamilton (2019) suggests that accounting for oil inventories has a substantial effect on the oil demand elasticity, controlling for the definition of the oil demand elasticity.

demand shocks. In this paper I examined alternative econometric approaches to estimating one-month price elasticities of oil supply and oil demand. The discussion focused on microeconomic estimators of the U.S. price elasticity of oil supply, the construction of bounds on the global price elasticity of oil supply, and IV estimators and VAR estimators of the global oil supply and oil demand elasticities.

I showed that some of these methodologies suffer from drawbacks that call into question the estimates they generate. I also illustrated how extraneous information may be used to judge the economic plausibility of oil supply and oil demand elasticity estimates. My analysis suggests that recent findings of rather large one-month oil supply elasticities are misleading, which implies that oil demand shocks are the dominant driver of the real price of oil and that the recessionary effect of oil supply shocks is more modest than suggested by some recent VAR studies. My analysis also raises questions about many estimates of the one-month price elasticity of oil demand reported in the literature. Some of these estimates are shown to be implausibly low in absolute terms, while others are implausibly high. Finally, I explained the rationale for the use of alternative elasticity definitions in the literature and discussed the trade-off between these definitions.

The insights of this paper are by no means restricted to the oil market. Indeed, the global oil market model in many ways resembles a traditional textbook model of demand and supply. Very similar econometric problems arise in many applications in empirical macroeconomics and related fields. For example, elasticity estimates and elasticity bounds may be used in constructing fiscal multipliers. They also naturally arise in labor economics and may be used in studying labor market responses to demand and supply shocks. Another example are macroeconomic studies of the effects of demand and supply shocks at the aggregate level.

References

- Anderson, S.T., Kellogg, R., and S.W. Salant (2018), “Hotelling Under Pressure,” *Journal of Political Economy*, 126, 984-1026.
- Antolin-Diaz, J., and J.F. Rubio-Ramirez (2018), “Narrative Sign Restrictions for SVARs,” *American Economic Review*, 108, 2802-2839.
- Baumeister, C., and J.D. Hamilton (2019), “Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Oil Demand Shocks,” *American Economic Review*, 109, 1873-1910.
- Bjørnland, H.C., Nordvik, F.M., and M. Rohrer (2019), “Supply Flexibility in the Shale Patch: Evidence from North Dakota,” manuscript, Norwegian Business School.
- Bjørnland, H.C. (2019), “Supply Flexibility in the Shale Patch: Facts, no Fiction,” manuscript, BI Norwegian Business School.
- Bornstein, G., Krusell, P., and S. Rebelo (2019), “Lags, Costs, and Shocks: An Equilibrium Model of the Oil Industry,” manuscript, Northwestern University.
- Caldara, D., Cavallo, M, and M. Iacoviello (2019), “Oil Price Elasticities and Oil Price Fluctuations,” *Journal of Monetary Economics*, 103, 1-20.
- Coglianese, J., Davis, L.W, Kilian, L., and J.H. Stock (2017), “Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand,” *Journal of Applied Econometrics*, 32, 1-15.
- Cross, J.L., Nguyen, B.H., and T.D. Tran (2019), “The Role of Precautionary and Speculative Demand in the Global Market for Crude Oil,” manuscript, Norwegian School of Business.
- Funashima, Y. (2020), “Global Economic Activity Indexes Revisited,” forthcoming: *Economics Letters*.
- Golding, G. (2019), “Don’t Expect U.S. Shale Producers to Respond Quickly to Geopolitical Supply Disruption,” <https://www.dallasfed.org/research/economics/2019/1003>.
- Hamilton, J.D. (2003), “What Is an Oil Shock?,” *Journal of Econometrics*, 113, 363-398.
- Hamilton, J.D. (2009). “Causes and Consequences of the Oil Shock of 2007-08,” *Brookings Papers on Economic Activity*, 40, 215-283.

- Hausman, J.A., and W.K. Newey (1995), “Nonparametric Estimation of Exact Consumers’ Surplus and Deadweight Loss,” *Econometrica*, 63, 1445-1476.
- Herrera, A.M., and S.K. Rangaraju (2020), “The Effect of Oil Supply Shocks on U.S. Economic Activity: What Have We Learned?” *Journal of Applied Econometrics*, 35, 141-159.
- 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 (2013), “Inference on Impulse Response Functions in Structural VAR Models,” *Journal of Econometrics*, 177, 1-13.
- Inoue, A., and L. Kilian (2019), “Corrigendum to ‘Inference on Impulse Response Functions in Structural VAR Models’ [J. Econometrics 177 (2013), 1-13],” *Journal of Econometrics*, 209, 139-143.
- Inoue, A., and L. Kilian (2020), “The Role of the Prior in Estimating VAR Models with Sign Restrictions,” manuscript, Federal Reserve Bank of Dallas.
- Känzig, D. (2019), “The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements,” manuscript, London Business School.
- 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(2), 216-240.
- Kilian, L. (2008b), “The Economic Effects of Energy Price Shocks,” *Journal of Economic Literature*, 46, 871-909.
- Kilian, L. (2009), “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market”, *American Economic Review*, 99, 1053-1069.
- Kilian, L. (2019), “Measuring Global Real Economic Activity: Do Recent Critiques Hold Up to Scrutiny?,” *Economics Letters*, 178, 106-110.
- Kilian, L. (2020), “Facts and Fiction in Oil Market Modeling,” manuscript, Federal Reserve Bank of Dallas.
- Kilian, L., and B. Hicks (2013), “Did Unexpectedly Strong Economic Growth Cause the Oil Price Shock of 2003-2008?,” *Journal of Forecasting*, 32, 385-394.

- Kilian, L., and T.K. Lee (2014), “Quantifying the Speculative Component in the Real Price of Oil: The Role of Global Oil Inventories,” *Journal of International Money and Finance*, 42, 71-87.
- Kilian, L., and D.P. Murphy (2012), “Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models,” *Journal of the European Economic Association*, 10, 1166-1188.
- Kilian, L., and D.P. Murphy (2014), “The Role of Inventories and Speculative Trading in the Global Market for Crude Oil,” *Journal of Applied Econometrics*, 29, 454-478.
- Kilian, L., and X. Zhou (2018), “Modeling Fluctuations in the Global Demand for Commodities,” *Journal of International Money and Finance*, 88, 54-78.
- Kilian, L., and X. Zhou (2019), “Oil Supply Shock Redux?” manuscript, Federal Reserve Bank of Dallas.
- Kilian, L., and X. Zhou (2020a), “Gasoline Demand More Responsive to Price Changes than Economists Once Thought,” <https://www.dallasfed.org/research/economics/2020/0616>.
- Kilian, L., and X. Zhou (2020b), “Does Drawing Down the U.S. Strategic Petroleum Reserve Help Stabilize Oil Prices?,” forthcoming: *Journal of Applied Econometrics*.
- Kilian, L., and X. Zhou (2020c), “The Econometrics of Oil Market VAR Models,” manuscript, Federal Reserve Bank of Dallas.
- Knittel, C.R., and S. Tanaka (2019), “Driving Behavior and the Price of Gasoline: Evidence from Fueling Micro Data,” CEEPR Working Paper 2019-019.
- Newell, R., and B. Prest (2019), “The Unconventional Oil Supply Boom: Aggregate Price Response from Microdata,” *Energy Journal*, 40, 1-30.
- Montiel Olea, J.L., Stock, J.H., and M.W. Watson (2020), “Inference in SVARs Identified with an External Instrument,” forthcoming: *Journal of Econometrics*.
- Levin, L., Lewis, M.S., and F.A. Wolak (2017), “High Frequency Evidence on the Demand for Gasoline,” *American Economic Journal: Economic Policy*, 9, 314-347.
- Yatchew, A., and J.A. No (2001), “Household Gasoline Demand in Canada,” *Econometrica*, 69, 1697-1709.

Zhou, X. (2020), "Refining the Workhorse Oil Market Model," *Journal of Applied Econometrics*, 35, 130-10.

Table 1: Correlations between Δp_t and $\Delta f_t^{(h)}$, 1990.2-2017.6

h	$\text{Corr}(\Delta p_t, \Delta f_t^{(h)})$
1	0.999
2	0.989
3	0.977
4	0.967

NOTES: All data are from the U.S. Energy Information Administration and transformed to logs. The estimation period matches that in Bjørnland et al. (2019).

Table 2: Weak instrument diagnostics for Caldara et al. (2019), 1985.1-2015.12

		First-stage F-statistic
Narrow instrument	Original	13.93
	UAE excluded	13.08
Broad instrument	Original	8.83
	UAE excluded	7.87
	August 1990 shock excluded	0.87

NOTES: Based on regressions of the percent change in the WTI spot oil price on an intercept and each of the instruments provided in Caldara et al. (2019, Table 2). Given one endogenous variable, a standard rule of thumb is that the weak instrument null can be rejected if the first-stage F-statistic exceeds 10.

Table 3: IV Global Elasticity Estimates

	Instrument	η^s	η^d (incorrectly defined)
Caldara et al. 2019	Broad ^{a,c}	0.08	-0.08
Caldara et al. 2019	Narrow ^a	0.05	-0.03
Caldara et al. 2019	Narrow ^b	0.03	Not reported

^a Including UAE in instrument

^b Excluding UAE from instrument

^c Based on instrument that does not pass weak IV test


Table 4: VAR Global Elasticity Estimates based on Kilian-Murphy (2014) Framework

	η^s	η^d (correctly defined)	η^d (incorrectly defined)
Kilian and Murphy 2014	0.01	-0.26	-0.44
Herrera and Rangaraju 2020	0.01	-0.28	-0.51
Inoue and Kilian 2020	0.01	-0.18	-0.47

Table 5: VAR Global Elasticity Estimates Based on Other Frameworks

	η^s	η^d
Caldara et al. 2019 ^a	0.10	-0.14
Baumeister and Hamilton 2019	0.15	-0.35

^a Using IV elasticity estimates based on broad instrument that fails weak IV test

Figure 1: Growth in the Number of Drilled, but Uncompleted Shale Oil wells (DUCs) in the Bakken and the WTI Oil Futures Spread, 2008.1-2017.6

NOTES: The price data were obtained from the U.S. Energy Information Administration. The Bakken DUC counts are courtesy of Rystad Energy. The estimation period ends in June 2017, as in Bjørnland et al. (2019) and starts in January 2008, before the start of the U.S. shale oil boom.