Understanding the Estimation of Oil Demand and Oil Supply Elasticities

Lutz Kilian
Federal Reserve Bank of Dallas
CEPR

The identification problem in the global oil market

● What we care about in oil market VAR models is typically the magnitudes of the <u>one-month price elasticities</u> of oil supply and demand, which pin down the slopes of the short-run supply and demand curve.

● A structural model with a vertical short-run supply curve and a downward sloping short-run demand curve, for example, may explain the reduced-form VAR residuals as well as a structural model with an upward sloping supply curve and a downward sloping demand curve.

● Estimates of these elasticities based on extraneous data are crucial for the estimation of global oil market VAR models.

● The conventional approach of estimating global elasticities based on exogenous instruments is infeasible given the lack of suitable monthly instruments.

Microeconomic estimates of the oil supply elasticity

● Economic theory suggests that the short-run oil supply elasticity is zero if adjusting oil production is costly, as tends to be the case in practice (Anderson, Kellogg and Salant 2018).

● Newell and Prest (2019) use data from all major oil producing regions in the United States. Their preferred estimate of the one-quarter oil supply elasticity for <u>conventional crude</u> is 0.017 (with a standard error of 0.006).

● Even for <u>shale oil</u> their one-quarter supply elasticity is negligible, consistent with survey evidence that it typically takes at least four weeks to start production, if a producer wants to complete an existing shale well in response to higher oil prices (Golding 2019).

If we take these 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.

The Bjørnland et al. (2021) critique

● Bjørnland et al. define the oil supply elasticity as $\beta_1 + \beta_2$ in:

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

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. In contrast, Anderson, Kellogg and Salant (2018) using a similar model focus on β_1.

● Bjørnland et al. measure the response to the spot price, holding fixed the oil futures price. The more common thought experiment involves a persistent oil price increase that shifts up both the spot price and the oil futures price. This explains their higher estimates:

Elasticity definition:	Texas Conventional	Bakken Conventional	Bakken Shale
Bjørnland et al.	-	0.10	0.71
Anderson et al.	0.001	0.03	-0.12

Bounding the global oil supply elasticity

Kilian and Murphy (2012) proposed 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.

● This oil supply disruption 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.

<u>All else equal</u>, 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.

● Not all else was equal:

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.

The <u>simultaneous shift in Saudi Arabia's supply curve</u> 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 <u>upper bound</u> on the one-month price elasticity of oil supply rather than an estimate of this elasticity.

The Caldara et al. (2019) critique of this bound

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.

● Kilian and Murphy’s 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, 1990, threatening retribution if unspecified OPEC countries did not reduce their oil production.

If so, the decline in the UAE’s oil production in August 1990 has to be excluded when constructing the endogenous production response Δq, increasing the upper bound on the one-month oil supply elasticity from 0.026 to 0.045.

Why Caldara et al.'s view is not persuasive

1. 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 (e.g., Caldara et al. (2019), Appendix E, p. A25-A26).

2. At no point was there an immediate military threat to the UAE, which has no direct border with Iraq. Iraq lacked 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.

- ⇨ The UAE must be included in constructing the oil supply elasticity bound.

Panel IV estimates of the global oil supply elasticity

Caldara et al. (2019) focus on the response of oil production in a given country to <u>supply disruptions in other oil-producing countries</u>.

1. 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.

2. The fitted value from the first stage, $\widehat{\Delta p}_{i,t}$, is used in the second-stage IV regression to identify the one-month price elasticity of oil supply

$$\Delta q^s_{i,t} = \eta^s \widehat{\Delta p}_{i,t} + u^s_{i,t},$$

where $\Delta q^s_{i,t}$ denotes oil production growth in country i.

Panel IV estimates of the oil supply elasticity

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 η^s drops to 0.029 based on the narrow instrument and to 0.056 based on the broad instrument.

⇨ Even these estimates however, are problematic.

Problems with Caldara et al.'s IV estimator of η^s

1. Instrument relevance requires the F-statistic in the first-stage regression to exceed 10, when regressing the percent change in the real price of oil on an intercept and the instrument for 1985.1-2015.12.

⇨ Although the narrow instrument passes this test, the broad instrument (with or without the UAE) does not.

⇨ Excluding the August 1990 episode from the broad instrument, the F-statistic drops below 0.9, so only the 1990 episode matters.

<u>Remark</u>: Caldara et al. report larger F-statistics based on discarding all months when there is no oil supply shock. That estimator has much higher variance than the full sample estimator.

2. Since Saudi Arabia (along with selected other producers) aims to directly offset geopolitical disruptions, both its oil supply curve and its oil demand curve shift in response to such an event, which <u>violates the exclusion restriction required for IV estimation</u>.

Panel IV estimates of the global oil demand elasticity

● For the estimation of the corresponding one-month oil demand elasticity, Caldara et al. (2019) propose a similar IV approach, which is subject to the same problems as the oil supply elasticity estimator.

● An additional problem is the "oil consumption" data used by Caldara et al. refer to the consumption of refined products such as diesel, gasoline, jet fuel, bunker fuel.

This means that the demand elasticity in the second-stage IV regression is not the <u>own price elasticity</u> of the demand for crude oil, but a <u>cross-price elasticity</u> of the demand for refined products.

VAR estimates of global oil demand and supply elasticities

Four main approaches:

- Imposing estimates of the price elasticity of oil supply directly in estimating the VAR model (Kilian 2009)
- Imposing bounds on the oil supply and/or oil demand elasticities (Kilian and Murphy 2012, 2014).
- Minimizing the distance between VAR elasticities and extraneous oil demand and oil supply elasticity estimates (Caldara et al. 2019)
- Explicitly specifying informative elasticity priors (Baumeister and Hamilton 2019).

The last two approaches express the impact elasticities as a function of the parameters of B_0 in the structural VAR representation

$$B_0 y_t = B_1 y_{t-1} + \ldots + B_p y_{t-p} + w_t,$$

whereas the first two approaches express these elasticities as a function of the structural impact multiplier matrix, B_0^{-1}.

Alternative elasticity concepts

Baumeister and Hamilton (2019) 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, <u>holding constant all other variables in the model</u> such as global real economic activity and oil inventories.

In contrast, Kilian and Murphy (2014) define the 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, <u>allowing global real activity and oil inventories to respond contemporaneously</u> to the exogenous demand shift.

⇨ Clearly, these elasticity concepts in general are neither numerically nor conceptually equivalent.

Why both elasticity concepts are useful

● BH insist that their definition is the correct definition from a theoretical point of view, if we want to characterize the slope of demand and supply curves. There is no disagreement on this point.

● The fact is, however, that this is not the definition that has been used in the existing literature on estimating these elasticities from micro data or macro data. Kilian and Murphy (2012, 2014) did not invent a new elasticity concept, but used this term as it had been defined in the empirical micro literature for many years.

● The merits of this alternative definition depend on how we use the elasticity measure. For example, it makes a lot of sense to use this definition in restricting the magnitude of oil market VAR impulse responses because these responses do not hold constant the responses of other model variables.

● In contrast, Baumeister and Hamilton (2019) cannot appeal to these extraneous elasticity estimates (or elasticity bounds) that are inconsistent with their own elasticity definition. It is unclear how to design the elasticity priors within their framework or how to compare the Bayes estimate of the elasticity to extraneous estimates.

● Similarly, when Caldara et al. (2019) propose to minimize the distance between the impact elasticities in the structural VAR model and their extraneous IV elasticity estimates, there is a mismatch between the definition of their IV elasticity and their VAR elasticity.

● One could quibble that perhaps one should refer to this alternative definition by a name other than “elasticity”, but this question seems moot, given that the literature has chosen to call them elasticities. What matters is that we understand what concept we are using in a given application.

Pitfalls in implementing the BH approach

Example 1: Let $y_t = (\Delta q_t, a_t, p_t)'$, where Δq_t is the growth rate of global oil production, a_t is an appropriately chosen measure of global real economic activity, and p_t is the log real price of oil in global markets.

$$\Delta q_t = \alpha_{qa} a_t + \alpha_{qp} p_t + \ldots + w_{1t} \quad (1)$$

$$a_t = \alpha_{aq} \Delta q_t + \alpha_{ap} p_t + \ldots + w_{2t} \quad (2)$$

$$p_t = \alpha_{pq} \Delta q_t + \alpha_{pa} a_t + \ldots + w_{3t} \quad (3)$$

They interpret equation (3) as an inverted oil demand curve with α_{pq} denoting the reciprocal of the impact price elasticity of demand.

Implicit in this specification is the assumption that $Q_t = C_t \forall t$, where Q_t is global oil production and C_t is global oil consumption.

However, in reality, $Q_t = C_t + \Delta I_t$, where ΔI_t is the change in global oil inventories, so the model suffers from omitted variable bias.

Example 2: BH's preferred model is:

$$\Delta q_t = \alpha_{qp} p_t + \ldots + w_{1t} \tag{7}$$

$$a_t = \alpha_{ap} p_t + \ldots + w_{2t} \tag{8}$$

$$\Delta q_t - \Delta i_t = \beta_{ca} a_t + \beta_{cp} p_t + \ldots + w_{3t} \tag{9}$$

$$\Delta i_t = \alpha_{iq} \Delta q_t + \alpha_{ia} a_t + \alpha_{ip} p_t + \ldots + w_{4t}, \tag{10}$$

were $\Delta i_t \equiv 100 \Delta I_t / Q_{t-1}$ with ΔI_t denoting the change in global oil inventories and Q_t denoting global oil production and Δq_t is the growth rate of global oil production.

BH interpret equation (9) as an oil demand curve with β_{cp} representing the impact price elasticity of oil demand.

Key assumption: $\Delta q_t - \Delta i_t$ in equation (9) is the monthly growth rate of global oil consumption.

Problem: $\Delta q_t - \Delta i_t = 100(C_t - Q_{t-1}) / Q_{t-1} \neq 100 \Delta c_t$.

Measurement Error in Baumeister and Hamilton's Oil Consumption Growth Measure

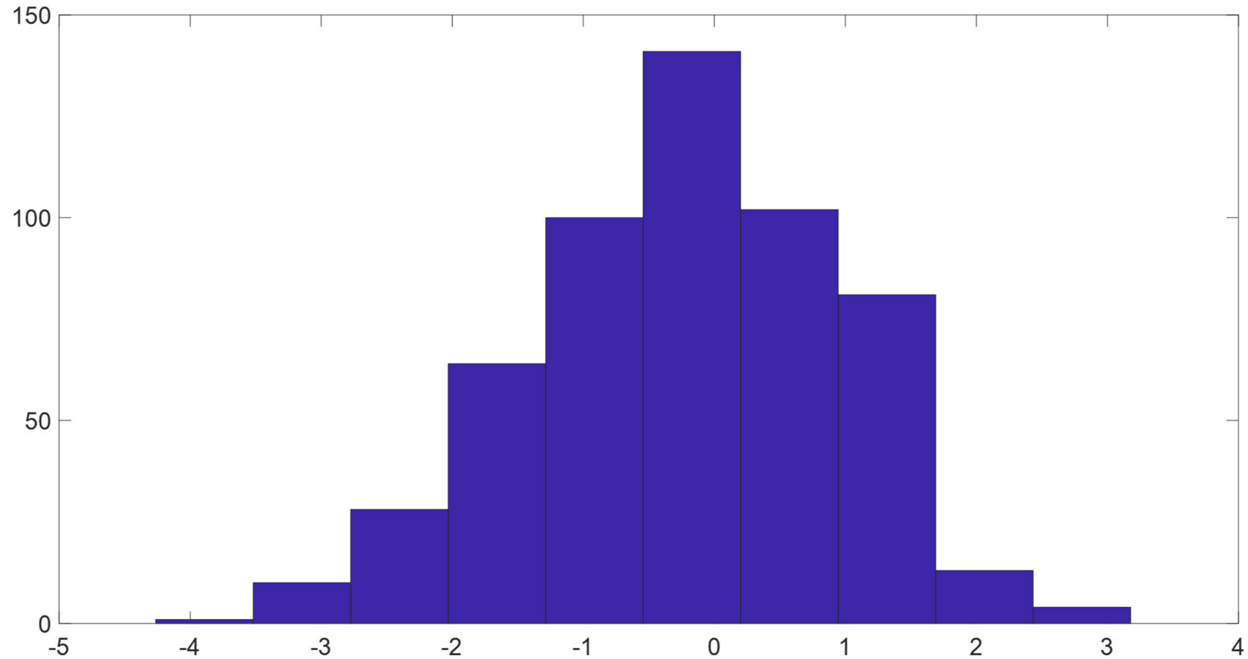

NOTES: Monthly growth rates, not annualized.

⇨ This undermines the credibility of the implied oil demand elasticity estimate in Baumeister and Hamilton (2019).

Conclusions

- Recent estimates of one-month oil demand elasticities close to zero and one-month oil supply elasticities of between 0.1 and 0.9 are misleading.

- These estimates were derived under questionable econometric and economic assumptions.

- My analysis reaffirms the conclusion that the one-month oil supply elasticity is low (and much lower than the corresponding oil demand elasticity), which implies that oil demand shocks are the dominant driver of the real price of oil.