Supplementary Information to *The Influence of OPEC+ on Oil Prices: A Quantitative Assessment*

Dominic Quint and Fabrizio Venditti

1. DATA DESCRIPTION AND MANIPULATION

Our dataset is monthly and ranges from 1995m1-2020m1. Below we comment on the construction of the included variables.

Oil production is provided by the *International Energy Agency*. OPEC+ was founded in December 2016 by OPEC (at that time including 14 member states: Algeria, Angola, Ecuador, Gabon, Iran, Iraq, Kuwait, Libya, Nigeria, Qatar, Saudi Arabia, the United Arab Emirates and Venezuela) and 11 non-OPEC countries (Azerbaijan, Equatorial Guinea, Kazakhstan, Mexico, Oman, Russia, Bahrain, Brunei, Malaysia, Sudan and South Sudan). OPEC+ production is the sum of its members' oil production with the following adjustments:

- To obtain a consistent sample over time, we do not include all countries listed above as not all uninterruptedly participated in the cuts and some countries even left the alliance. We therefore exclude Iran, Libya, Nigeria and Venezuela as these countries did not participate in the agreed production cuts for an extended period of time.[1]
- We include Congo and keep Qatar and Ecuador in our sample. Congo joint OPEC in July 2018 and stated participating in the deal four month later. Qatar and Ecuador quit its OPEC membership in January 2019 and January 2020, respectively. We also include Equatorial Guinea in our sample which first participated in the deal as a non-OPEC member and continued as an OPEC member from June 2017.
- We only use crude oil supply in our data sample and do not include natural gas liquids (NGLs). Initially, the OPEC+ declaration of cooperation specified that production targets defined for OPEC countries only apply to crude oil production, whereas targets defined for non-OPEC members also include NGLs. However, the alliance later excluded NGLs also for non-OPEC members. For consistency, we thus only use crude oil data for the whole data sample.

To construct a counterfactual path of oil supply for the model without strategic interaction, described in Section 4.1 of the paper, we keep OPEC+ production constant at the reference level that the alliance used in 2016 to calculate the production cuts. When signing the deal in December 2016, the agreed production cuts were relative to the participating countries' production levels reached in October 2016, except for Angola, for which cuts were based on levels reached in September 2016. We finally construct a data series of aggregate production of for non-OPEC+ producers by subtracting aggregate OPEC+ production from global crude oil production.

Real economic activity are taken from Kilian (2009). As the data ranges only until June 2019, we extrapolate it for the remaining months using the growth rate of the Baltic Dry index. The index, which is available on a daily frequency, is a proxy for global dry bulk shipping stocks. The Baltic Dry index is downloaded as monthly averages from *Refinitiv Datastream*.

[1] Iran was exempted from the cuts in November 2018 following the U.S. sanctions imposed on the Iranian energy sector and Libya never participated in the cuts due to ongoing civil conflicts. Due to the dire economic situation in Nigeria and Venezuela, the former has started participating in the deal in November 2018 whereas the latter was exempted in November 2018 from the cuts.

Nominal oil prices are monthly averages of the Brent oil price taken from *Bloomberg*.
Inventory days-in-storage used in Section 5.2 of the paper is obtained by dividing total inventory levels available for OECD countries by the daily total oil consumption of these countries. The data is only available for total oil, including crude oil and NGLs and is provided by the *International Energy Agency*.

2. ALTERNATIVE COUNTERFACTUAL PATH FOR OPEC+ OIL PRODUCTION

For the model without strategic interaction described in Section 4.1 of the paper, we assume for the counterfactual path of oil production that OPEC+ would have kept production constant at the reference level it used for calculating the production cuts. Another plausible assumption would be that OPEC+ production would have continued to grow after December 2016 if the agreed production cuts had not taken place.

We derive such a plausible path by fitting an ARIMA model to the observed OPEC+ production and forecasting a hypothetical path after December 2016. Based on the AIC criteria, we find that an ARIMA(0,1,2) with drift is suited best to describe OPEC+ production between January 1995 and December 2016. Based on the estimated ARIMA model, Figure 1 (upper panel) plots the forecast for OPEC+ oil supply along with actual production and the baseline counterfactual path used before. Not surprisingly, the forecast differs greatly from the path based on the OPEC+ reference levels, especially at the end of the sample.

Despite the different hypothetical path in production, the implied path for the oil price is comparable to the benchmark results shown in Section 5.1 of the paper. Figure 1 (middle panel) compares the counterfactual path of the oil price based on the different assumption for production. The results are similar and are thus not too sensitive to the assumption we need to make regarding the counterfactual path of the oil supply. In addition, the sequence of supply shocks required to implement the counterfactual simulation based on the ARIMA forecast are also not unusually large by historical standards. This can be seen in Figure 1 (lower panel) plotting the simulated supply shocks and is confirmed by an F-test, which null hypothesis cannot be rejected with a p-value of 0.56.

3. SIMULATING OPEC+'S INVENTORY TARGET

Adding inventory days-in-storage to the model without strategic interaction described in Section 4.1 of the paper allows us to simulate how oil prices would have evolved if OPEC+ had focused on achieving its inventory target. Instead of assuming a counterfactual path for OPEC+ production, we construct a counterfactual path for our measure of inventories, back out the sequence of supply shocks that would have be needed for achieving the inventory target and derive the corresponding path of the oil prices–again holding constant the remaining structural shocks. As explained in the main text, we include days-in-storage as a more precise measure of market tightness.

With inventory levels and with it days-in-storage levels reaching historical highs in early 2017, we need to make an assumption on how quickly OPEC+ could have reached its target. We assume that production cuts would have been gradually, aiming at bringing days-in-storage within 12 months to the average level reached between 2010 and 2014. Figure 2 compares the actual path of days-in-storage with the assumed counterfactual path. After days-in-storage had actually fallen in 2017 and reached the desired level in early 2018, it started to deviate again from the target soon afterwards.

Structural identification is very similar to the model described in Section 4.1 of the paper, based on sign restrictions and bounds on the demand and supply elasticities. As before, we identify a supply shock and two demand shocks assuming that days-in-storage fall on impact following a negative supply shock and that they rise following a precautionary demand shock (Table 1). We leave the reaction of days-in-storage to a global demand shock unrestrained. The restrictions on the

Table 1: Sign Restrictions in Model with OECD Inventories

	oil supply	global oil demand	precautionary oil demand
global oil production	-	+	+
real economic activity	-	+	-
oil price	+	+	+
OECD inventories	-		+

Additional Restrictions
- Upper bound on elasticity of demand in production -0.1
- Upper bound on elasticity of supply 0.05

Table 2: Forecast Error Variance Decomposition

	Supply	Demand	Precautionary
Crude Production	0.57	0.16	0.27
Real activity	0.13	0.52	0.35
Real Oil Price	0.20	0.72	0.08

short-run price elasticity of oil supply and oil demand are similar to the previous model specifications.

4. MODEL SPECIFICATION, DIAGNOSTICS AND ADDITIONAL RESULTS

Both VAR models are estimated with four lags of the endogenous variables. The results of formal lag length selection criteria for the model without strategic interactions are shown in Figure 3. These results would call for a somewhat parsimonious model, favouring a specification of 2 lags. We prefer using more lags to ensure that any residual auto-correlation in the residuals is mopped up. As it is well known, overfitting a model only results in a loss of efficiency (i.e. wider standard errors) but ensures unbiased central estimates, whereas underfitting can severely bias the results (Kilian, 2001).

We apply some shrinkage on the model parameters through a standard Minnesota prior. Priors are implemented via dummy observations, see Del Negro and Schorfheide (2011, Part III, Section 2). Data are in levels, so the priors on the autoregressive coefficients are centered around 1 for first lag and 0 for the remaining lags. Let us collect the hyper-parameters in the vector $\lambda = [\lambda_1, \lambda_2, \lambda_3, \lambda_4, \lambda_5]$. The parameters λ_1 and λ_2 set the tightness of the prior on the coefficients for first and for the remaining lags, respectively. The prior for the covariance matrix is centered at a matrix that is diagonal with elements equal to the pre-sample variance of the data, with tightness λ_3. Finally, λ_4 and λ_5 regulate the tightness of the sum of coefficients prior and of the co-persistence priors. Values for these hyperparameters are set at conventional levels, $\lambda = [0.5, 1, 0.5, 0.5, 1]$. Importantly, a rejection sampler is employed, so that draws that imply explosive roots are discarded within the estimation algorithm. Table 2 shows the contribution of each shock to the forecast error variance of each variable at horizon 12 in the simpler model. In line with studies in the literature that have used this setup (Kilian and Murphy, 2014), we find that the contribution of supply shocks to oil price volatility stands at around 20 percent, whereas more than 70 percent of unpredictable changes in the price of oil come from demand shocks. Results for the 4 variables model, not shown for brevity, are very similar. Finally, IRFs for both models are reported in Figures 4 and 5.

REFERENCES

Del Negro, M. and F. Schorfheide (2011). "Bayesian Macroeconometrics" in J. Geweke, G. Koop, and H. V. Dijk (editors), "The Oxford Handbook of Bayesian Econometrics," Oxford University Press, Oxford.

Kilian, L. (2001). "Impulse Response Analysis in Vector Autoregressions with Unknown Lag Order" *Journal of Forecasting* 20(3): 161–179.

Kilian, L. (2009). "Not All Oil Price Shocks are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market" *American Economic Review* 99(3): 1053–69 doi:https://doi.org/10.1257/aer.99.3.1053.

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(3): 454–478 doi:https://doi.org/10.1002/jae.2322.

Figure 1: Results Based on ARIMA Forecast

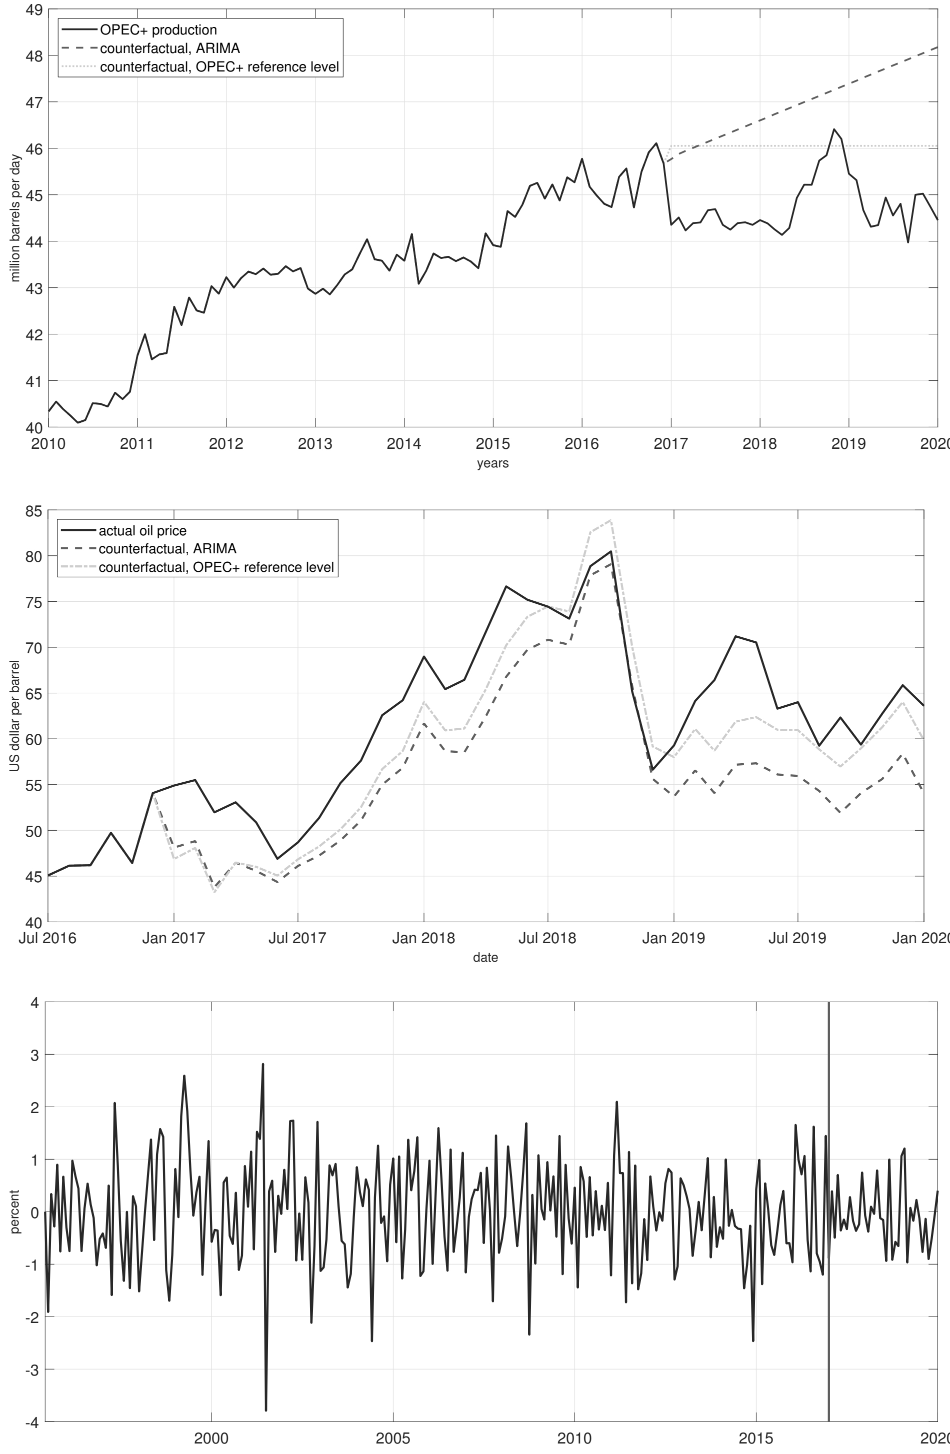

Note: The upper panel shows OPEC+ crude oil production together with the counterfactual path based on an ARIMA forecast and based on OPEC+ reference levels (data only plotted from 2010 onwards). The middle panel compares the counterfactual path of the oil price with the baseline result and the actual price of oil. The lower panel shows the sequence of supply shocks required to implement the counterfactual simulation. Until December 2016 (marked with the vertical line), the series corresponds to the estimated shocks.

Figure 2: OECD Inventory Days-In-Storage and Counterfactual Path

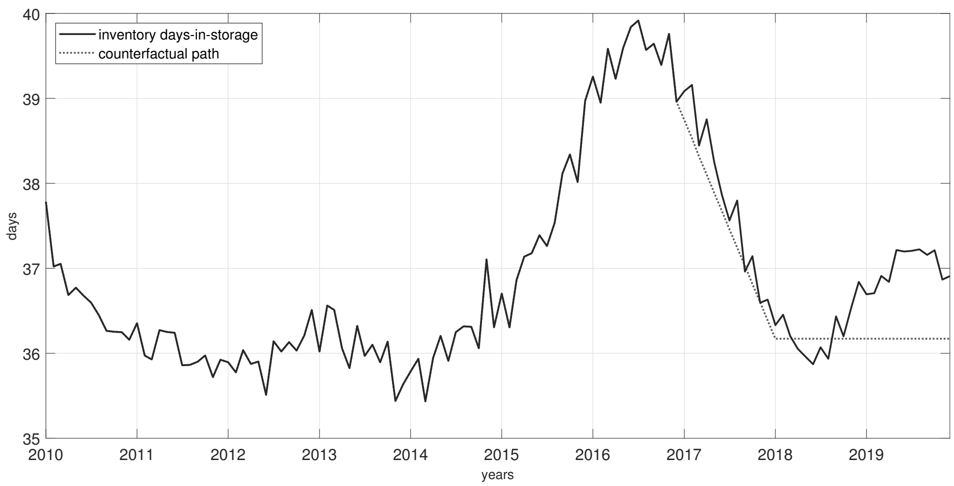

NOTE: The figure shows OECD inventory days-in-storage together with the counterfactual path based on the OPEC+ inventory target.

Figure 3: Lag Length Selection Criteria

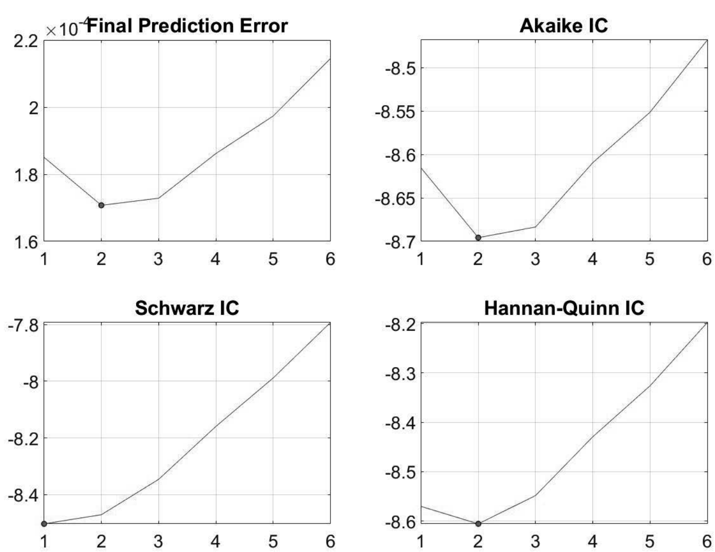

Figure 4: Impulse Responses Model Without Strategic Interaction

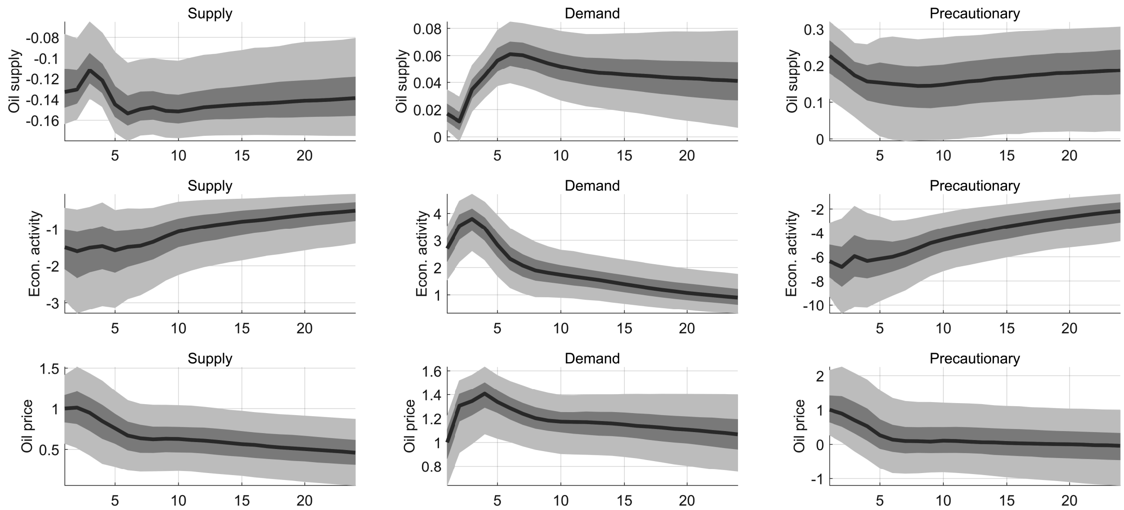

Note: Impulse response functions are standardized so as to lead to a 1 percent increase in the price of oil with variables in logs (oil supply and oil price) being expressed as percentage. The shaded areas are 65 and 84 percent confidence bands.

Figure 5: Impulse Responses Model With Strategic Interaction

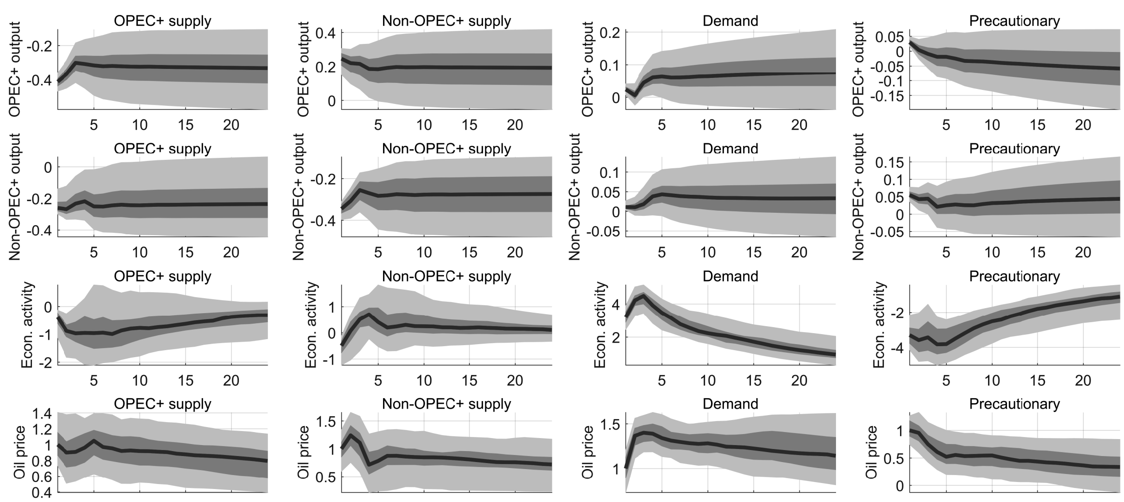

Note: Impulse response functions are standardized so as to lead to a 1 percent increase in the price of oil with variables in logs (OPEC+ and Non-OPEC+ output and oil price) being expressed as percentage. The shaded areas are 65 and 84 percent confidence bands.