

Lags, Costs, and Shocks: An Equilibrium Model of the Oil Industry*

Gideon Bornstein,[†] Per Krusell[‡] and Sergio Rebelo[§]

March 2018

Abstract

We use a new micro data set to estimate a stochastic industry-equilibrium model of the oil industry. This effort is a first step towards studying the importance of ongoing structural changes in the oil market in a general-equilibrium model of the world economy. We analyze the impact of the advent of fracking on the volatility of oil prices. Our model predicts a large decline in this volatility.

J.E.L. Classification: Q4

Keywords: oil, commodities, volatility.

*We thank Hilde Bjornland, Wei Cui, Matteo Iacoviello, Ravi Jaganathan, Ryan Kellogg, Lutz Kilian, Markus Kirchner, Valerie Ramey, and Rob Vigfusson for their comments.

[†]Northwestern University.

[‡]Institute for International Economic Studies, NBER, and CEPR.

[§]Northwestern University, NBER, and CEPR.

1 Introduction

How important is the oil market for the world economy? Although oil shocks are often viewed as responsible for the poor performance of many countries in the 1970s, these shocks have played a relatively minor role in leading macroeconomic models. Since oil represents a relatively small share of overall production costs, conventional models imply that oil shocks have a limited impact on aggregate output.

This conclusion has recently been challenged by Gabaix (2011), Acemoglu et al. (2012), and Baqaee and Farhi (2017). These authors argue that shocks to sectors with a small factor share that are highly complementary to other inputs can have a large impact on aggregate output. Baqaee and Farhi (2017) emphasize the perils of using linearization methods to analyze macroeconomic models with strong complementarities and use the impact of oil shocks in the 1970s as a leading example.

Motivated by this line of research, we revisit the workings of the oil market by proposing and estimating a stochastic industry-equilibrium model of the oil industry. This effort is important not only because oil shocks can be crucial determinants of macroeconomic outcomes, but also because there is ongoing structural change in this market that merits further study. While conventional oil production is characterized by long lags and various forms of adjustment costs, new forms of oil production, such as fracking, are much more nimble.

In this paper, we take a first step towards studying the importance of ongoing structural changes in the oil market in a general-equilibrium model of the world economy. To produce the first building block in this endeavor, we take as exogenous shocks to the demand for oil in the world economy as well as supply disruptions. We then derive quantity and price outcomes as a function of these shocks.

In order to build our model on solid microeconomic foundations, we rely heavily on a new micro data set that contains information on production, reserves, operational costs, and investment for all oil fields in the world. We use these data to guide the con-

struction of our model in three ways. First, we compile a set of facts about oil markets. Second, we produce micro estimates of two key model parameters: the average lag between investment and production and the elasticity of extraction costs with respect to production. Third, we use the generalized method of moments (GMM) to estimate the remaining model parameters, targeting a set of second moments for oil-related variables.

There is substantial heterogeneity across oil firms along various dimensions. We find two sources of heterogeneity that are particularly important. The first is the different behavior of firms that are part of the Organization of the Petroleum Exporting Countries (OPEC) and those that are not. The second is the difference between hydraulic fracturing (fracking) and conventional oil production. Our benchmark model features heterogeneity only in the OPEC/non-OPEC dimension. Our extended model includes firms that use conventional oil production methods as well as fracking.

We use our model to measure the importance of demand and supply shocks in driving prices, production, and investment. We find that supply and demand shocks contribute equally to the volatility of oil prices but that investment in the oil industry is driven mostly by demand shocks. The reason for this pattern is twofold. There is a long lag between investment and production and supply shocks are short-lived relative to demand shocks. So, investment responds much more to demand than to supply shocks. We also find that the volatility of OPEC production firms is driven primarily by supply shocks that disrupt the ability of these firms to extract oil. In contrast, both demand and supply shocks are important in explaining the volatility of non-OPEC production.

As discussed above, our data allows us to document two key differences between fracking and conventional oil production. First, it is less costly for fracking firms to adjust their level of production in the short run, so these firms are more responsive to changes in prices. Second, the lag between investment and production is much shorter in fracking operations than in conventional oil production.

We introduce fracking firms into our model to study their impact on the dynamics of the oil market. We find that their presence leads to a large decline in the volatility

of oil prices. The reason is simple: these firms are more nimble in adjusting production levels from existing fields and in starting production in new fields, so they can respond more quickly to price increases.


In Section 2 we review the following six facts about the oil market. First, the average rate of change of inflation-adjusted oil prices is not significantly different from zero. Second, oil prices are very volatile, much more volatile than returns to the stock market or exchange rates. Third, investment in the oil industry is very volatile, roughly seven times more volatile than U.S. aggregate investment. Fourth, investment in the oil industry shows a strong, positive correlation with the inflation-adjusted price of oil. Fifth, there is a low short-run elasticity of the supply of oil from individual oil fields with respect to price. Sixth, production by OPEC firms is more volatile and less correlated with oil prices than production by non-OPEC production firms.

We describe our model in Section 3. Section 4 is devoted to estimating model parameters using both micro data and moments of key aggregate variables for the oil industry. The estimated version of our model has four main features. First, demand is relatively inelastic. Second, supply is elastic in the long run because firms can invest in the discovery of new oil fields.¹ Third, supply is inelastic in the short run. This property results from several model features: there is a lag between investment and production, there are convex costs of adjusting extraction rates, and there are decreasing returns to oil investment. Fourth, there are shocks to demand, e.g., faster growth in China, and to the supply of OPEC firms, e.g., the Iran-Iraq war.

This paper is closely related to work on models of the response of oil prices and production to fluctuations in the world economy. Examples of this work, include Backus and Crucini (2000), Leduc and Sill (2004), Blanchard and Gali (2007), Kilian (2009),

¹While the amount of oil is ultimately finite, we can think about this investment process as including new ways of extracting oil as well as the development of oil substitutes, as in Adao et al. (2017). There has been a large expansion of oil reserves during our sample period. According to the U.S. Energy Information Administration, proved oil reserves measured in years of production have increased from roughly 30 years in 1980 to 52 years in 2015.

Figure 1: Real price of oil



Bodenstein et al. (2011), and Lippi and Nobili (2012).² Our paper is also related to a new, emerging literature that uses micro data to shed new light on key aspects of the oil industry (see, e.g., Kellogg (2014), Arezki et al. (2016), and Anderson et al. (2017)).

2 Key facts about the oil market

In this section, we review six facts about the oil market. Most of our analysis relies on proprietary data compiled by Rystad Energy that includes information on reserves, production, investment, and operational costs for the universe of oil fields. The data contains information about roughly 14,000 oil fields operated by 3,200 companies.

We focus our analysis on the period from 1970 to 2015 since until 1972, U.S. regulatory agencies sought to keep U.S. oil prices stable by setting production targets. (see Hamilton (1983) and Kilian (2014)).

²Earlier work on the impact of oil shocks on the economy generally treats oil prices as exogenous (see, e.g., Kim and Loungani (1992), Rotemberg and Woodford (1996), and Finn (2000)).

No significant time trend in real oil prices. In order to study long-term trends in real oil prices, we use a longer sample that covers the period from 1900 to 2015. Our data source for the period 1900-1947 is Harvey et al. (2010). After 1947, we use the price of West Texas Intermediate. We deflate the price of an oil barrel quoted in dollars by the U.S. consumer price index.

Figure 1 plots our measure of the real price of oil. The average annual growth rate in the real price of oil during this period is not statistically different from zero. For the period from 1900 to 2015 this average is 0.01 with a standard error of 0.21. For the period from 1970 to 2015 this average is 0.02 with a standard error of 0.27.

The property that average growth rates in real prices estimated over long time periods are close to zero is shared by many other commodities (Deaton and Laroque (1992), Harvey et al. (2010) and Chari and Christiano (2014)).

Oil prices are very volatile. Figure 1 shows that oil prices have been very volatile since the early 1970s. From 1970 to 2015, the volatility of oil prices is higher than that of returns to the stock market or exchange rates. The standard deviation of the annual percentage change in oil prices is 0.28 for nominal prices and 0.27 for real prices. In contrast, the standard deviation of nominal returns to the S&P 500 is 16 percent and the standard deviation of changes in exchange rates is roughly 10 percent.


These first two facts were aptly summarized by Deaton (1999) with the phrase “What commodity prices lack in trend, they make up for in variance.”

Investment in the oil industry is very volatile. The annual standard deviation of the growth rate of real world investment in the oil industry in the period 1970-2015 is 0.36. To put this number in perspective, this measure of volatility is 0.10 for U.S. manufacturing and 0.07 for U.S. aggregate investment.³

³The only major U.S. manufacturing sector with volatility of investment similar to the oil industry in the period 1970-2015 is Motor vehicle manufacturing, a sector that has struggled to compete with foreign manufacturers and had to be bailed out by the Federal government in 2009.

Figure 2 illustrates the high volatility of investment in the oil industry. This figure plots the linearly-detrended logarithm of two series: real world-wide investment in the oil industry and real aggregate investment in the U.S.

Figure 2: Detrended investment in the oil industry



Investment in the oil industry is positively correlated with oil prices. Figure 3 plots the logarithm of the inflation-adjusted price of oil and the logarithm of inflation-adjusted world investment in the oil industry. The correlation between the growth rate of the price of oil and the growth rate of investment is 0.51.


Table 1 reports the correlation between the growth rate of real investment and the growth rate of the real oil prices for each of the top twenty firms in the oil industry ranked according to their total oil production in 2015. We see that the correlation is high for every major oil firm with the exception of Iran’s NIOC, Total, and Eni, which show a correlation near zero.

Table 1: Investment and price correlation for top 20 firms

Firm	Headquarters	OPEC	$corr(\Delta p_t, \Delta i_t)$
Saudi Aramco	Saudi Arabia	✓	0.31
Rosneft	Russia	✗	0.34
PetroChina	China	✗	0.36
Kuwait Petroleum Corp (KPC)	Kuwait	✓	0.3
NIOC (Iran)	Iran	✓	0.06
Pemex	Mexico	✗	0.27
ExxonMobil	United States	✗	0.35
Lukoil	Russia	✗	0.41
Petrobras	Brazil	✗	0.3
PDVSA	Venezuela	✓	0.28
Abu Dhabi NOC	Abu Dhabi	✓	0.14
Chevron	United States	✗	0.43
Shell	Netherlands	✗	0.34
BP	United Kingdom	✗	0.35
Surgutneftegas	Russia	✗	0.26
South Oil Company (Iraq NOC)	Iraq	✓	0.19
Total	France	✗	0.04
CNOOC	China	✗	0.4
Statoil	Norway	✗	0.34
Eni	Italy	✗	0.04

Notes: x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$. P_t and I_t represent the real price of oil and the firm's real investment, respectively.

Figure 3: Oil prices and oil investment



The short-run elasticity of oil supply is very low. Oil producers can respond to an increase in the market price of oil in two ways. The first is to produce more oil from oil fields in operation by increasing the extraction rate. The second is to increase the number of oil fields in operation. We show that the short-run elasticity of the extraction rate (the ratio of production to reserves) with respect to an exogenous change in the price of oil is positive but small.⁴ We also show that the elasticity of response of the number of oil fields in operation to an exogenous change in the price of oil is not statistically different from zero.

Table 2 reports panel-data estimates of the elasticity of the extraction rate for a given oil field with respect to real oil prices.⁵ These estimates suggest that a rise in oil

⁴Anderson et al. (2017) estimate this elasticity to be close to zero. The difference between our results and theirs is likely to reflect differences in data frequency: their data is monthly while ours is annual.

⁵In our data, reserves are proven reserves, which measure the total amount of oil that can be produced from a given field. Reserves do not change in response to changes in oil prices, so there is no mechanical impact of oil prices on extraction rates.

prices leads to a very small increase in the supply of oil from a given oil field.⁶

Our estimates are obtained by running various versions of the following regression:

$$\ln \theta_{it} = \alpha_i + \beta \ln p_t + \gamma X_{it} + \varepsilon_{it},$$

where θ_{it} denotes the extraction rate of oil field i at time t , p_t is the real price of oil, and X_{it} represents other controls.⁷ These controls include a time trend, an oil-field fixed effect, and a fixed effect for year of operation to control for the life cycle of an oil field.⁸

Specification 1 in Table 2 is a simple OLS regression. The estimated slope coefficient that results from this regression can be biased if there is technical progress that lowers the cost of extraction, raising θ_{it} , increasing the supply of oil, and lowering p_t . To address this problem, we instrument the price of oil with our forecast of the cyclical component of world GDP. This forecast is correlated with aggregate demand and unaffected by aggregate cost shocks.

Our forecast of the cyclical component of world GDP is obtained as follows. We HP-filter the series for world real GDP and estimate an AR process for the resulting cyclical component. Choosing the number of lags using the Akaike information criterion resulted in an AR(2) process.

Specifications 2-5 use this instrument. Our benchmark specification is regression 2 which yields an estimate for β equal to 0.12.⁹ The following calculation is useful for evaluating the magnitude of this elasticity. The average extraction rate in our sample is 2.8 percent. A one standard deviation (27 percent) increase in the price of oil raises the extraction rate from 2.8 percent to 2.9 percent, resulting only in a 3.3 percent increase

⁶An oil field generally contains many oil rigs. Production increases can result from the intensive margin (higher production from existing oil rigs) or from the extensive margin (drilling new oil rigs). Anderson et al. (2017) use a sample of Texas oil rigs to show that the elasticity of the intensive margin is close to zero, so production increases result from the extensive margin.

⁷Our data includes all oil fields with a positive extraction rate for the period 1971-2015 excluding the last year of operation.

⁸See Arezki et al. (2016) and Anderson et al. (2017) for discussions of this life cycle.

⁹This estimate is similar to the one obtained by Caldara et al. (2017) by combining a narrative analysis of episodes of large drops in oil production with country-level instrumental variable regressions.

Table 2: Price elasticity of extraction rates

Variable	(1)	(2)	(3)	(4)	(5)
ln(price)	0.09*** (0.009)	0.12*** (0.012)	0.16*** (0.017)	0.17*** (0.023)	0.16*** (0.013)
ln(price) \times $\mathbb{1}_{\text{OPEC}}$			-0.2*** (0.069)	-0.19*** (0.069)	
ln(price) \times $\mathbb{1}_{\text{Big Firm}}$				-0.06 (0.054)	
ln(price) \times $\mathbb{1}_{ \Delta \ln(p) > 0.1}$					-0.17*** (0.009)
Oil field FE	✓	✓	✓	✓	✓
Operation year FE	✓	✓	✓	✓	✓
Year trend	✓	✓	✓	✓	✓
IV	✗	✓	✓	✓	✓
Clusters (oil fields)	12,187			11,479	
Observations	173,742			173,034	

Notes: oil fields with positive extraction in 1971-2015, excluding last year of operation. Standard errors in parenthesis. Instrument for price using the forecast of the cyclical component of world GDP. F-stat > 1000 in (2)-(5). (***) - significant at a 1% level.

in production.

Specification 3 includes the product of the logarithm of the price and an OPEC dummy. We see that OPEC firms respond less to changes in prices than non-OPEC firms. The point estimate for the response of OPEC to oil price shocks is negative (-0.04) and insignificant (P-value = 0.48).

Specification 4 includes also an interaction term that is the product of logarithm of price and a dummy for firm size. This dummy is equal to one for firms that produced more than 0.5 billion barrels of oil in 2015. Total production in 2015 was approximately 27.5 billion barrels of oil so that each large firm according to our definition has at least a 1.8 percent market share. The idea is to investigate whether large firms behave differently from small firms. We find no evidence of a firm-size effect in the response of extraction rates to changes in oil prices.

Specification 5 includes an interaction term that is the product of logarithm of price and a dummy for price changes that are larger than 10 percent in absolute value. The idea is to investigate whether firms react more to large oil price changes than to small price changes. We find that the coefficient on the interaction term is negative (-0.17) and statistically significant. This finding is consistent with the presence of convex adjustment costs in the extraction rate, so that the elasticity of response is higher for small price changes than for large price changes.

These results are robust when we extend the sample to start in 1900 (see Table 17 in the appendix).

Table 3 reports our time-series estimates of the elasticity of the number of oil fields in operation with respect to real oil prices. Specification 1 is a simple OLS regression where the dependent variable is the logarithm of oil fields in operation world wide and the independent variable is the logarithm of real oil prices. Specification 2 uses our forecast of the cyclical component of world GDP as an instrument for the logarithm of real oil prices. Specifications 3 and 4 report results for non-OPEC fields and OPEC fields, respectively. All four specifications yield elasticity estimates that are statistically

Table 3: Price elasticity of oil fields in operation

Variable	(1)	(2)	(3)	(4)
ln(price)	-0.05 (0.06)	-0.02 (0.08)	-0.01 (0.03)	0 (0.06)
Year trend	✓	✓	✓	✓
IV	✗	✓	✓	✓
Dep. variable	All fields	All fields	Non-OPEC fields	OPEC fields
Observations	45	45	45	45

Notes: Newey-West standard errors computed with 5-year lags in parenthesis.

insignificant. We also find an insignificant elasticity when we extend our sample to start in 1900 (see Table 16 in the appendix).

Taken together, these results suggest that the number of oil fields in operation does not respond in the short-run to changes in oil prices.

OPEC and non-OPEC firms behave differently. Table 4 shows that OPEC and non-OPEC firms differ in the volatility and persistence of production and investment, as well as in the correlation of these variables with real oil prices. The production of OPEC firms is more volatile and less persistent than that of non-OPEC firms. In addition, the correlation between investment and prices is higher for non-OPEC firms than for OPEC firms. These patterns are likely to result from supply shock to OPEC firms, such as the disruptions in oil markets associated with the Iranian revolution and the Iran-Iraq war.

Table 4: Correlation of production and investment for OPEC and non-OPEC

	Δp_t	Δi_t^n	Δi_t^o	Δq_t^n	Δq_t^o	st. dev.	auto-corr
Δp_t	1.00					0.27	-0.03
Δi_t^n	0.56	1.00				0.19	0.12 [†]
Δi_t^o	0.36	0.67	1.00			0.19	0.31
Δq_t^n	0.03 [†]	0.09 [†]	-0.03 [†]	1.00		0.02	0.64
Δq_t^o	0.03 [†]	0.02 [†]	-0.23 [†]	-0.14 [†]	1.00	0.07	0.21 [†]

Notes: † - not significant at a 5% level. See table 7 for standard errors. x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$. P_t , I_t^N , I_t^O , Q_t^N , and Q_t^O represent, respectively, the price of oil, investment by non-OPEC firms, investment by OPEC firms, quantity produced by non-OPEC firms, and quantity produced by OPEC firms. All variables defined in real terms.

3 An industry-equilibrium model

In this section we describe our equilibrium model of the oil industry. Our goal is to propose a parsimonious model consistent with the facts described in Section 2 which can ultimately serve as a building block in a model of the world economy.

To construct this building block, we take the world demand for oil as exogenous. We summarize the demand for oil by consumers and firms around the world with a log-linear demand function. In this simple demand specification, the short- and long-run price elasticities coincide. In subsection 4.2 we consider a generalization of this demand specification in which world demand can respond sluggishly to price changes, so that the short-run elasticity is lower than the long-run elasticity.¹⁰

In our benchmark specification, we assume that the world demand for oil is given by

$$P_t = \exp(d_t)Q_t^{-1/\varepsilon},$$

¹⁰To simplify, we also abstract from growth and consider a model where production and investment are constant in the non-stochastic steady state. We discuss in the Appendix a version of the model where production and investment grow at a constant rate in the non-stochastic steady state, while extraction rates and prices are constant.

where P_t is the real price of oil, Q_t is the quantity of oil consumed, and ε is the price elasticity of demand. The variable d_t is a stochastic demand shock that follows an AR(2) process in logarithms

$$\ln(d_t) = \rho_1^d \ln d_{t-1} + \rho_2^d \ln d_{t-2} + e_t^d.$$

We choose this AR(2) specification so that demand shocks can follow a hump-shaped pattern. This pattern allows an initial shock to contain news about a future rise in the demand for oil associated, for example, with faster growth in China.

Non-OPEC firms. There is a continuum of measure one of non-OPEC firms. These firms maximize their value (V^N) which is given by

$$V^N = E_0 \sum_{t=0}^{\infty} \beta^t [P_t \theta_t^N K_t^N - I_t^N - \psi (\theta_t^N)^\eta K_t^N]. \quad (1)$$

Here, I_t^N denotes investment, θ_t^N the extraction rate (the ratio of production to reserves), and K_t^N oil reserves.

The term $\psi (\theta_t^N)^\eta K_t^N$ represents the cost of extracting oil. We assume that this cost is linear in reserves so that aggregate production and aggregate extraction costs are invariant to the distribution of oil reserves across firms. This formulation allows us to use a representative firm to study production and investment decisions. We assume that $\eta > 1$, so that the costs of extraction are convex in the extraction rate. In addition, we assume that the time $t + 1$ extraction rate is chosen at time t .

We adopt a parsimonious way of modeling lags in investment by introducing *exploration capital*, which we denote by X_t . The timing of the realization of shocks and firm decisions is as follows. In the beginning of the period, the demand and supply shocks are realized, a fraction λ of the exploration capital materializes into new oil reserves, and production occurs according to the predetermined extraction rate. At the end of the period, the firm chooses its investment and its extraction rate for the next period.

The law of motion for exploration capital is as follows:

$$X_{t+1}^N = (1 - \lambda)X_t^N + (I_t^N)^\alpha (L^N)^{1-\alpha}. \quad (2)$$

Investment adds to the existing exploration capital, but only a fraction λ of the exploration capital materializes into oil reserves in every period. Investment requires land (L^N) and exhibits decreasing returns ($\alpha < 1$). Without this feature, investment would be extremely volatile, rising sharply when prices are high and falling deeply when prices are low.

One interpretation of equation (2) is as follows. Suppose each firm searches for oil on a continuum of oil fields containing X_t^N barrels of oil uniformly distributed across fields. The probability of finding oil is independent across oil fields and equal to λ . By the law of large numbers, each firm finds λX_t^N oil reserves at time t . We pursue this interpretation when we estimate λ using our micro data.

Oil reserves evolve as follows:

$$K_{t+1}^N = (1 - \theta_t^N)K_t^N + \lambda X_{t+1}^N. \quad (3)$$

Reserves fall with oil production ($\theta_t^N K_t^N$) and rise as exploration capital materializes into new reserves (λX_{t+1}^N).

The notion of exploration capital embodied in equations (2) and (3) is a tractable way of introducing time-to-build in investment that might be useful in other problems. This formulation allows us to introduce a lag between investment and production by adding only one state variable. The parameter λ allows us to smoothly vary the length of the lag.¹¹

The problem of the representative non-OPEC firm is to choose the stochastic sequences for I_t^N , θ_{t+1}^N , K_{t+1}^N , and X_{t+1}^N that maximize its value defined in equation (1), subject to constraints (2) and (3).

¹¹See Rouwenhorst (1991) for a discussion of the large state space and complex dynamics associated with time-to-build formulations.

The first-order conditions for θ_{t+1}^N is

$$E_t P_{t+1} = \psi \eta (\theta_{t+1}^N)^{\eta-1} + E_t \mu_{t+1}^N, \quad (4)$$

where μ_t^N is the Lagrange multiplier corresponding to equation (3). The extraction rate at time $t + 1$ is chosen at time t so as to equate the expected oil price to the sum of the marginal cost of extraction, $\psi \eta (\theta_{t+1}^N)^{\eta-1}$, and the expected value of a barrel of oil reserves at the end of time $t + 1$, $E_t \mu_{t+1}^N$.

The first-order condition for K_{t+1} is

$$\mu_t^N = E_t \beta \{ [P_{t+1} \theta_{t+1}^N - \psi (\theta_{t+1}^N)^\eta] + (1 - \theta_{t+1}^N) \mu_{t+1}^N \}. \quad (5)$$

For a given value of θ_{t+1}^N , each additional barrel of oil reserves results in additional revenue equal to $P_{t+1} \theta_{t+1}^N$ and increases extraction costs by $\psi (\theta_{t+1}^N)^\eta$. A fraction $1 - \theta_{t+1}^N$ of the barrel of reserves remains in the ground and has a value of $\beta E_t \mu_{t+1}^N$.

The first-order condition for X_{t+1} is

$$\nu_t^N = \lambda \mu_t^N + \beta (1 - \lambda) E_t \nu_{t+1}^N, \quad (6)$$

where ν_t^N is the Lagrange multiplier corresponding to equation (2). The value of increasing exploration capital by one unit, ν_t^N , has two components. A fraction λ materializes into oil reserves and has a value μ_t^N . A fraction $1 - \lambda$ remains as exploration capital and has an expected value $\beta E_t \nu_{t+1}^N$.

The first-order condition for I_t is

$$1 = \alpha (I_t^N)^{\alpha-1} (L^N)^{1-\alpha} \nu_t^N. \quad (7)$$

This condition equates the cost of investment (one unit of output) to the marginal product of investment in generating exploration capital, $\alpha (I_t^N)^{\alpha-1} (L^N)^{1-\alpha}$, evaluated at the value of exploration capital, ν_t^N .

OPEC firms. There is a continuum of measure one of OPEC firms. The problem for these firms is to maximize their value (V^O), which is given by

$$V^O = E_0 \sum_{t=0}^{\infty} \beta^t (P_t e^{-u_t} \theta_t^O K_t^O - I_t^O - \psi (\theta_t^O)^\eta K_t^O). \quad (8)$$

The key difference between OPEC and non-OPEC firms is that the former are subject to a supply shock, u_t . When this shock occurs, production falls for a given level of the extraction rate. We assume that supply shocks follow an AR(2) process:

$$\ln u_t = \rho_1^u \ln u_{t-1} + \rho_2^u \ln u_{t-2} + e_t^u,$$

and that innovations to demand (e_t^d) and supply (e_t^u) are uncorrelated.

The laws of motion for exploration capital and reserves are given by

$$X_{t+1}^O = (1 - \lambda)X_t^O + (I_t^O)^\alpha (L^O)^{1-\alpha}, \quad (9)$$

$$K_{t+1}^O = (1 - e^{-u_t} \theta_t^O)K_t^O + \lambda X_{t+1}^O. \quad (10)$$

The problem of the representative OPEC firm is to choose the stochastic sequences for I_t^O , θ_{t+1}^O , K_{t+1}^O , and X_{t+1}^O that maximize its value defined in equation (8), subject to constraints (9) and (10).

The first-order conditions for the problem of OPEC firms are as follows:

$$\begin{aligned} E_t (P_{t+1} e^{-u_{t+1}}) &= \psi \eta (\theta_{t+1}^O)^{\eta-1} + E_t (\mu_{t+1}^O e^{-u_{t+1}}), \\ \mu_t^O &= E_t \beta \{ [P_{t+1} e^{-u_{t+1}} \theta_{t+1}^O - \psi (\theta_{t+1}^O)^\eta] + (1 - e^{-u_{t+1}} \theta_{t+1}^O) \mu_{t+1}^O \}, \\ \nu_t^O &= \lambda \mu_t^O + \beta (1 - \lambda) E_t \nu_{t+1}^O, \\ 1 &= \alpha (I_t^O)^{\alpha-1} (L^O)^{1-\alpha} \nu_t^P. \end{aligned}$$

These first-order conditions are similar to those for non-OPEC firms. The key difference is that production of OPEC firms is scaled by the supply shock $e^{-u_{t+1}}$.

While there are supply shocks to non-OPEC producers (e.g., Canadian wildfires and Gulf of Mexico hurricanes), these shocks seem relative small compared to the

supply shocks to OPEC producers. This view is consistent with the higher volatility of production in OPEC relative to non-OPEC.

To investigate the potential importance of supply shocks to non-OPEC firms, we estimated a version of the model with uncorrelated supply shocks to both OPEC and non-OPEC firms. Despite having three more parameters, the statistical fit of this extended model is similar to that of the benchmark model. The variance of the supply shock innovation is roughly 500 lower for non-OPEC firms when compared to OPEC firms. This result suggests that it is empirically reasonable to abstract from non-OPEC supply shocks.

Equilibrium. In equilibrium, P_t is a function of demand and supply shocks, the aggregate level of reserves in OPEC (\mathbf{K}_t^O) and non-OPEC (\mathbf{K}_t^N), and the predetermined aggregate levels of extraction rates in OPEC ($\boldsymbol{\theta}_t^O$) and non-OPEC ($\boldsymbol{\theta}_t^N$):

$$P_t = p(d_t, u_t, \mathbf{K}_t^O, \mathbf{K}_t^N, \boldsymbol{\theta}_t^O, \boldsymbol{\theta}_t^N). \quad (11)$$

Firms maximize their value subject to the laws of motion for reserves and exploration capital. Each firm takes the law of motion for the aggregate levels of reserves, extraction rates, and exploration capital as given and so the price process is exogenous to each individual firm. These laws of motion for the aggregate variables depend on the six aggregate state variables included in equation (11) and the aggregate levels of exploration capital in OPEC (\mathbf{X}_t^O) and non-OPEC (\mathbf{X}_t^N).

The oil market clears, i.e., total oil production equals total oil demand:

$$Q_t^N + Q_t^O = Q_t,$$

where Q_t^N and Q_t^O are the aggregate quantities produced by non-OPEC and OPEC firms, respectively. These quantities are given by

$$Q_t^N = \boldsymbol{\theta}_t^N \mathbf{K}_t^N, \quad Q_t^O = e^{-u_t} \boldsymbol{\theta}_t^O \mathbf{K}_t^O.$$

There is a continuum of measure one of identical firms within the two groups, OPEC and non-OPEC. So, in equilibrium, the values of aggregate variables for each group coincide with the values of the corresponding variables for the representative firm in each group.

The Hotelling rule. The classic Hotelling (1931) rule emerges as a particular case of our model in which there are no OPEC firms, $\lambda = 0$, and $\eta = 1$. When $\lambda = 0$, investment does not result in more oil reserves, so oil is an exhaustible resource. Equation (6) implies that in this case the value of exploration capital is zero: $\nu_t^N = 0$. Combining equations (4) and (5) we obtain

$$E_t(P_{t+1} - \psi) = \beta E_t(P_{t+2} - \psi).$$

This equation is the Hotelling rule: the price of oil minus the marginal cost of production is expected to rise at the rate of interest in order to make oil producers indifferent between extracting oil at $t + 1$ and at $t + 2$.

For the general case where $\lambda \geq 0$ and $\eta \geq 1$, the marginal cost of production is $\eta\psi\theta_t^{\eta-1}$ and the difference between the price of oil and the marginal cost of production is given by

$$E_t(P_{t+1} - \eta\psi\theta_{t+1}^{\eta-1}) = \beta E_t(P_{t+2} - \eta\psi\theta_{t+2}^{\eta-1}) + \beta E_t(\eta - 1)\psi\theta_{t+2}^\eta.$$

The term $\beta E_t(\eta - 1)\psi\theta_{t+2}^\eta$ represents the marginal fall in production costs at time $t + 2$ from having an additional barrel of oil reserves. When $\eta = 1$, this term is zero and we recover the Hotelling rule.

When $\lambda = 0$ (no more oil can be found), there is no steady state in which P_t and θ_t are constant.¹² When the extraction rate is constant, production falls over time and, since demand is downward sloping, the price of oil rises over time. When the price is constant, production must also be constant and so the extraction rate must rise.

¹²When $\lambda = 0$ and $\eta > 1$ our economy resembles the model proposed by Anderson et al. (2017). Changes in the extraction rate in our model play a similar role to drilling new wells in their model.

In our model, $\lambda > 0$ and $\eta > 1$. Because it is feasible to find more oil, there is a steady state in which both P_t and θ_t are constant. Oil reserves are constant and so the quantity produced is also constant. In the steady state, the marginal decline in production costs from an additional barrel of oil is such that the difference between price and marginal cost remains constant:

$$\beta(\eta - 1)\psi\theta^\eta = (1 - \beta)(P - \eta\psi\theta^{\eta-1}).$$

4 Model solution and estimation

We solve the model using a second-order approximation around its non-stochastic steady state. In the non-stochastic steady state the ratio of production by OPEC and non-OPEC firms is equal to the ratio of the land available to each of the two groups:

$$\frac{L^O}{L^N} = \frac{\theta^O K^O}{\theta^N K^N}.$$

We calibrate some key parameters using our micro data set as well as information about the cost of capital in the oil industry. We estimate the remaining parameters using GMM.

We choose the ratio L^O/L^N so that in the steady state the ratio of OPEC to non-OPEC production is equal to the average of this ratio in our data (0.82). We calibrate the total amount of land ($L^O + L^N$) so that the steady-state extraction rate coincides with the average extraction rate in our data (2.8 percent).

The parameter ψ matters only for the level of oil prices, so we normalize it to one. We set $1/\beta - 1$, the real discount rate, to 8 percent which is the real cost of capital estimated by Damodaran (2017), which takes into account the systematic risk associated with investments in the oil industry.

4.1 Estimating λ and η with micro data


To estimate λ , we compute the lag between the first year of investment and first year of production (T_i) for every oil field in our data set. If the arrival of production occurs according to a Poisson process, the lag between investment and production follows a geometric distribution with mean λ . The maximum likelihood estimator for λ is:

$$\hat{\lambda} = \frac{N}{\sum_{i=1}^N T_i} = 0.085,$$

where N denotes the number of oil fields.

This estimate implies that the average lag between investment and production is 12 years. Figure 4 shows the empirical distribution of this lag together with the implied geometric distribution for our estimate of $\hat{\lambda}$.¹³

Figure 4: Empirical distribution of lags between investment and production



¹³Our estimate of the average production lag is higher than that reported in Arezki et al. (2016). This difference occurs because we estimate the lag between initial investment (which includes seismic analysis and drilling wells to discover and delineate oil fields) and production. Arezki et al. (2016) estimate the lag between oil discovery and production, which is shorter.

We also use our micro data to estimate η , the parameter that controls the convexity of the extraction costs. We run the following regression:

$$\ln \left[\frac{C(\theta_{it}, K_{it})}{K_{it}} \right] = \gamma_i + \eta \ln(\theta_{it}) + \varepsilon_{it},$$

where $C(\theta_{it}, K_{it})$ denotes extraction costs. The potential presence of cost shocks, either field specific or aggregate, creates an endogeneity problem. Suppose it becomes more costly to extract oil, so that firms reduce their extraction rates. This correlation between the cost and the rate of extraction biases downward in our estimate for η . To address this problem, we instrument the extraction rate with our forecast of the cyclical component of world GDP. This forecast is correlated with aggregate demand and unaffected by field-specific cost shocks.

Our data includes all oil fields with positive extraction rates between 1971 and 2015. We exclude the last year of oil field operation since the data for this year includes the costs of shutting down the field, which are not related to the rate of extraction.

Table 5 contains our slope estimates. Specifications 1 through 4 include fixed effects for oil field and operation year. Specification 1 includes all the oil fields in our sample. Specification 2 includes only non-OPEC firms. Specification 3 includes only OPEC firms. While our instrument is independent of oil-field-specific cost shocks, it may be correlated with aggregate supply shocks. The Iran-Iraq war, for example, may have caused a slowdown in world GDP at the same time as it disrupted the supply of oil in the warring countries. As a result, the point estimate is negative and statistically insignificant. We use specification 2 as our benchmark and set η equal to the point estimate (9.3). The fact that $\eta > 1$ is consistent with the fact that the elasticity of response of production to prices is higher for small price increases than for large price increases (see regression 3 in Table 2).

Table 5: Extraction rate adjustment costs regression

Variable	(1)	(2)	(3)
ln(extraction)	11.49*** (1.60)	9.30*** (1.14)	-36.65 (41.06)
Oil field FE	✓	✓	✓
Operation year FE	✓	✓	✓
Sample	All	Non-OPEC	OPEC
IV	✓	✓	✓
1 st stage F-stat	43	53	1
Clusters (oil fields)	11,527	9,969	1,558
Observations	174,339	146,879	27,460

GMM estimation. We estimate ϵ , α , the parameters of the AR(2) processes for demand and supply shocks, together with the variance of the two shocks using GMM.¹⁴ Table 6 reports our results. Our estimate for ϵ is 0.135 with a standard error of 0.015.¹⁵ This point estimate implies that demand is very inelastic. A 1 percent increase in production reduces the price by 7.4 percent.

To see why demand needs to be inelastic, it is useful to write the demand function:

$$\ln P_t = d_t - \frac{1}{\epsilon} \ln Q_t.$$

The model embeds two mechanisms that generate price volatility. The first mechanism is a low value of ϵ that makes low production volatility consistent with high price volatility. The second mechanism is volatile demand shocks. Since oil prices are highly persistent (changes in prices are roughly i.i.d.), demand shocks must also be highly persistent. Volatile, persistent demand shocks lead to high volatility in investment and production. The observed volatility of investment and production help determine the importance of the two channels and therefore identify ϵ .

Table 6 shows that the standard errors associated with our parameter estimates are

¹⁴Our weighting matrix is a diagonal matrix with diagonal elements equal to the inverse of the variance of the targeted moments.

¹⁵This estimate is similar to the one obtained by Caldara et al. (2017).

generally small. The only parameter that is imprecisely estimated is α . This imprecision results from the small impact of local changes in the value of α on the moments implied by the model. Our point estimate of α is close to zero. However, it is important that α be strictly positive since, when $\alpha = 0$, the model has the counterfactual implication that investment in the oil industry would be zero.

Table 6: Estimated parameters

Parameter	Estimate	(s.e.)
ϵ	0.135	(0.015)
α	3×10^{-6}	(0.978)
ρ_1^d	1.761	(0.075)
ρ_2^d	-0.775	(0.073)
ρ_1^u	1.428	(0.15)
ρ_2^u	-0.51	(0.11)
$var(e_t^d)$	0.02	(0.007)
$var(e_t^u)$	0.002	(0.001)

Table 7 compares the estimated moments targeted by our GMM procedure with the population moments implied by the model. We see that the fit of the model is good with most of the model population moments inside the 95 percent confidence interval of data moments. One exception is the correlation between changes in prices and change in quantities for OPEC firms which is negative in the model and close to zero in the data.

One interesting property of the model is that it is consistent with the high correlation between prices and investment. In the literature on the cattle and hog cycles (e.g., Ezekiel (1938) and Nerlove (1958)) this positive correlation is often interpreted as suggesting that expectations have a backward-looking component. Investment rises when prices are high, sowing the seeds of a future fall in prices. In our model, the high correlation between the price of oil and investment results from the rational response of forward-looking firms. A positive demand shock raises the price of oil above its steady state level. As a result it is profitable to invest in oil to expand production and take

Table 7: Data and model moments

	Moment	Data	(s.e.)	Model
(1)	$\text{std}(\Delta p_t)$	0.273	(0.028)	0.227
(2)	$\text{std}(\Delta i_t^n)$	0.192	(0.024)	0.213
(3)	$\text{std}(\Delta i_t^o)$	0.193	(0.027)	0.211
(4)	$\text{std}(\Delta q_t^n)$	0.022	(0.003)	0.025
(5)	$\text{std}(\Delta q_t^o)$	0.069	(0.011)	0.054
(6)	$\text{corr}(\Delta p_t, \Delta i_t^n)$	0.557	(0.147)	0.709
(7)	$\text{corr}(\Delta p_t, \Delta i_t^o)$	0.362	(0.109)	0.657
(8)	$\text{corr}(\Delta p_t, \Delta q_t^n)$	0.031	(0.069)	0.049
(9)	$\text{corr}(\Delta p_t, \Delta q_t^o)$	0.030	(0.122)	-0.572
(10)	$\text{corr}(\Delta i_t^n, \Delta i_t^o)$	0.673	(0.096)	0.997
(11)	$\text{corr}(\Delta i_t^n, \Delta q_t^n)$	0.087	(0.094)	-0.001
(12)	$\text{corr}(\Delta i_t^n, \Delta q_t^o)$	0.023	(0.112)	-0.103
(13)	$\text{corr}(\Delta i_t^o, \Delta q_t^n)$	-0.034	(0.145)	0.003
(14)	$\text{corr}(\Delta i_t^o, \Delta q_t^o)$	-0.226	(0.153)	-0.036
(15)	$\text{corr}(\Delta q_t^n, \Delta q_t^o)$	-0.141	(0.125)	0.014
(16)	$\text{corr}(\Delta p_t, \Delta p_{t-1})$	-0.027	(0.088)	-0.018
(17)	$\text{corr}(\Delta i_t^n, \Delta i_{t-1}^n)$	0.119	(0.135)	0.008
(18)	$\text{corr}(\Delta i_t^o, \Delta i_{t-1}^o)$	0.311	(0.096)	0.009
(19)	$\text{corr}(\Delta q_t^n, \Delta q_{t-1}^n)$	0.643	(0.113)	0.334
(20)	$\text{corr}(\Delta q_t^o, \Delta q_{t-1}^o)$	0.213	(0.211)	0.326

Notes: x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$.

advantage of the high oil prices. This supply expansion brings the price back to its steady state level.

Oil price forecasts and oil futures. Both in the data and in the model annual changes in the price of oil are close to being i.i.d. So, in the short run, the stochastic process for oil prices is well approximated by a random walk.

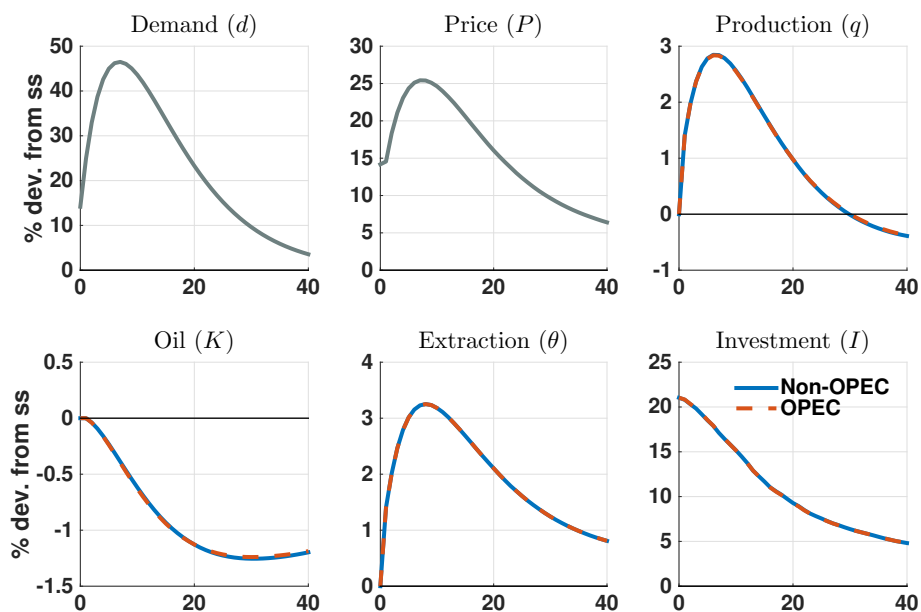
We compare the volatility of one-year futures prices in the data and in our model. Our data for oil futures covers the period from 1986 to 2015. We constructed annual real future oil prices by averaging all the future contracts with one-year maturity within year t and deflating the average by the time t consumer price index.¹⁶ We find that the volatility of one-year futures prices implied by our model coincides with the volatility of one-year futures in the data, 0.2.

Impulse response functions. Figure 5 depicts the impulse response function for a one standard deviation demand shock. The shock follows a hump-shaped pattern with a peak in year seven. On impact, firms cannot change their extraction rates so the price increases one to one with the demand shock. In the following periods, the price of oil increases but the magnitude of this increase is moderated by a rise in the extraction rate. Production rises and reserves are depleted. Since the shock is very persistent, investment rises to increase future reserves and production in order to take advantage of the extended period of high oil prices.

Figure 6 depicts the impulse response function for a one standard deviation supply shock. The shock follows a hump-shaped pattern with a peak in year 2. Compared to the demand shock, the supply shock is less persistent and smaller in magnitude. Extraction rates and production rise in non-OPEC firms and fall in OPEC firms. Non-OPEC firms increase their investment to boost reserves so they can raise production. But since the shock is short lived, the rise in investment is much smaller than the one

¹⁶See Alquist et al. (2013) for a discussion of the properties of oil price futures.

Figure 5: Impulse response to a demand shock

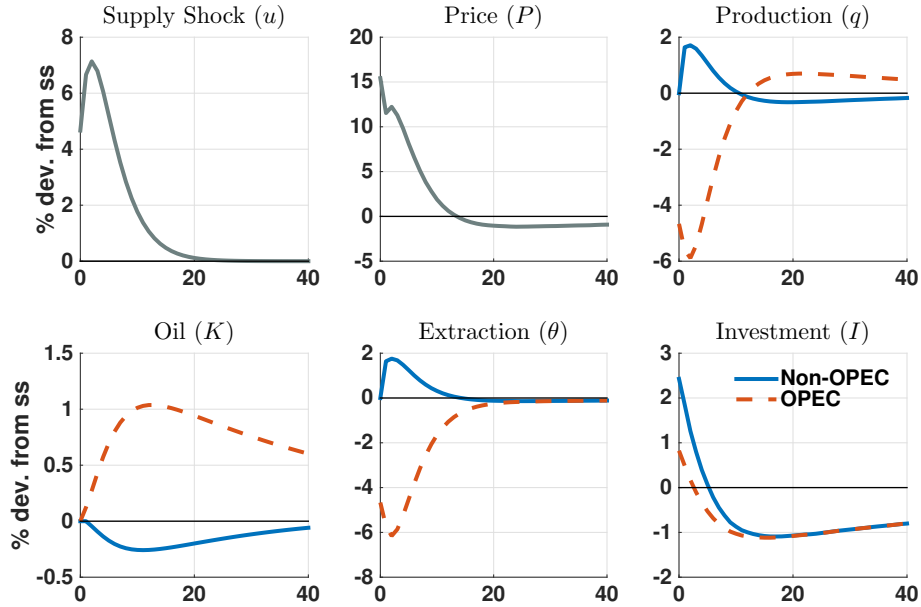


that occurs in response to a demand shock. OPEC firms also raise their investment but not as much as non-OPEC since in the short run OPEC firms have a lower extraction rate than non-OPEC firms.

We can use our model to answer a classic question: what is the role of demand and supply shocks? Table 8 illustrates the role played by demand and supply shocks in the performance of the model. Eliminating demand shocks leads to a small decline in the volatility of prices and to a large decline in the volatility of investment. Eliminating supply shocks leads to a small decline in the volatility of prices, to a large decline in the volatility of OPEC production, and implies a perfect correlation between OPEC and non-OPEC production levels.

Table 9 provides a decomposition of the annual standard deviation of the key variables in our model: price, production and investment. We see that demand and supply shocks contribute roughly equally to the volatility of prices. This result reflects the fact that demand and supply shocks tend to have a similar short-run impact on the price

Figure 6: Impulse response to a supply shock



of oil (see Figures 5 and 6). This result is consistent with the importance of macroeconomic performance in driving oil prices emphasized in Barsky and Kilian (2001) and Barsky and Kilian (2004).

Table 9 also shows that the volatility of investment is predominantly driven by demand shocks. These shocks are long-lived and so they elicit a large response of investment (see Figure 5). In contrast, supply shocks have a much lower impact on investment because these shocks are less persistent (see Figure 6).

Demand and supply shocks contribute equally to the volatility of production by non-OPEC firms. In contrast, the volatility of production by OPEC firms is dominated by supply shocks.

4.2 Dynamic demand

To simplify, we assume in our benchmark model that the demand for oil is static: the quantity demanded at time t depends only on the price at time t . As a result, the

Table 8: Moments sensitivity to demand and supply shocks

Moment	Data	(s.e.)	Bench.	No demand shocks	No supply shocks
$std(\Delta p_t)$	0.273	(0.028)	0.227	0.165	0.156
$std(\Delta i_t^n)$	0.192	(0.024)	0.213	0.028	0.211
$std(\Delta i_t^o)$	0.193	(0.027)	0.211	0.011	0.211
$std(\Delta q_t^n)$	0.022	(0.003)	0.025	0.018	0.017
$std(\Delta q_t^o)$	0.069	(0.011)	0.054	0.051	0.017
$corr(\Delta q_t^n, \Delta q_t^o)$	-0.141	(0.125)	0.01	-0.303	1.00

Notes: x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$.

Table 9: Decomposition of annual standard deviation

Moment	Shocks	
	Demand	Supply
$std(\Delta p_t)$	47.2%	52.8%
$std(\Delta i_t^n)$	98.3%	1.7%
$std(\Delta i_t^o)$	99.7%	0.3%
$std(\Delta q_t^n)$	48.3%	51.7%
$std(\Delta q_t^o)$	10.1%	89.9%

Notes: x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$.

short-run and the long-run elasticities of demand coincide.

In this subsection, we present a variant of the model in which the short-run and long-run price elasticity of demand are different. This model is consistent with the notion that when oil prices are high, households and firms substitute towards other forms of energy, but it might take time for this to occur.

Our dynamic demand specification takes the form:

$$P_t = \exp(d_t) (1 - \phi)^{1/\xi} (Q_t - \phi Q_{t-1})^{-1/\xi} , \quad (12)$$

where ϕ is the inertia parameter.

This demand specification is similar to the one derived in the literature on deep habits (Ravn et al. (2006) and Binsbergen (2016)). Differences between short- and long-run elasticities of demand also emerge in models with endogenous technology choice, as in Leon-Ledesma and Satchi (2016).

All other model equations remain unchanged. The multiplicative term $(1 - \phi)^{1/\xi}$ ensures that the non-stochastic steady state of the model does not depend on ϕ . The short-run demand elasticity is given by

$$\frac{\partial \ln Q_t}{\partial \ln P_t} = \frac{Q_t - \phi Q_{t-1}}{Q_t} \xi .$$

The long-run demand elasticity is equal to ξ , as $\frac{\partial \ln Q_{ss}}{\partial \ln P_{ss}} = \xi$. So, the short-run demand elasticity is lower than the long-run demand elasticity. The local short-run demand elasticity around the non-stochastic steady state equals $(1 - \phi)\xi$. We denote this short-run demand elasticity by $\varepsilon \equiv (1 - \phi)\xi$.

We re-estimate the model via GMM, including both ε and ϕ as parameters to be estimated. Table 10 reports our results. The estimated short-run demand elasticity, ε , is very similar to the estimated demand elasticity in our benchmark model. The estimated value of ϕ is 0.836, indicating that the demand for oil features a high degree of inertia. The implied long-run demand elasticity, ξ , equals 0.746, approximately 6 times higher than the short-run demand elasticity. This finding suggests that the parameter ε estimated in our benchmark model is the short-run demand elasticity.

Table 10: Estimated parameters with dynamic demand elasticity

Parameter	Estimate	(s.e.)
ε	0.122	(0.020)
α	3×10^{-8}	(0.522)
ρ_1^d	1.693	(0.132)
ρ_2^d	-0.716	(0.115)
ρ_1^u	1.696	(0.139)
ρ_2^u	-0.717	(0.136)
$var(e_t^d)$	0.017	(0.008)
$var(e_t^u)$	0.002	(0.001)
ϕ	0.836	(0.105)

Table 11 presents the empirical and model-implied moments. Allowing the model the flexibility of having a different short-run and long-run demand elasticities substantially improves the fit of the model. The model comes very close to reproducing the price volatility observed in the data. In addition, the model does a better job than the benchmark at fitting the volatility of oil produced by OPEC firms and the correlation of prices with investment of OPEC and non-OPEC firms.

Table 11: Data and model moments with dynamic demand elasticity

	Moment	Data	(s.e.)	Benchm. model	Dynamic demand elasticity
(1)	$\text{std}(\Delta p_t)$	0.273	(0.028)	0.227	0.260
(2)	$\text{std}(\Delta i_t^n)$	0.192	(0.024)	0.213	0.196
(3)	$\text{std}(\Delta i_t^o)$	0.193	(0.027)	0.211	0.198
(4)	$\text{std}(\Delta q_t^n)$	0.022	(0.003)	0.025	0.023
(5)	$\text{std}(\Delta q_t^o)$	0.069	(0.011)	0.054	0.070
(6)	$\text{corr}(\Delta p_t, \Delta i_t^n)$	0.557	(0.147)	0.709	0.585
(7)	$\text{corr}(\Delta p_t, \Delta i_t^o)$	0.362	(0.109)	0.657	0.353
(8)	$\text{corr}(\Delta p_t, \Delta q_t^n)$	0.031	(0.069)	0.049	0.108
(9)	$\text{corr}(\Delta p_t, \Delta q_t^o)$	0.030	(0.122)	-0.572	-0.276
(10)	$\text{corr}(\Delta i_t^n, \Delta i_t^o)$	0.673	(0.096)	0.997	0.945
(11)	$\text{corr}(\Delta i_t^n, \Delta q_t^n)$	0.087	(0.094)	-0.001	-0.008
(12)	$\text{corr}(\Delta i_t^n, \Delta q_t^o)$	0.023	(0.112)	-0.103	-0.074
(13)	$\text{corr}(\Delta i_t^o, \Delta q_t^n)$	-0.034	(0.145)	0.003	-0.011
(14)	$\text{corr}(\Delta i_t^o, \Delta q_t^o)$	-0.226	(0.153)	-0.036	0.153
(15)	$\text{corr}(\Delta q_t^n, \Delta q_t^o)$	-0.141	(0.125)	0.014	-0.032
(16)	$\text{corr}(\Delta p_t, \Delta p_{t-1})$	-0.027	(0.088)	-0.018	-0.152
(17)	$\text{corr}(\Delta i_t^n, \Delta i_{t-1}^n)$	0.119	(0.135)	0.008	0.011
(18)	$\text{corr}(\Delta i_t^o, \Delta i_{t-1}^o)$	0.311	(0.096)	0.009	0.014
(19)	$\text{corr}(\Delta q_t^n, \Delta q_{t-1}^n)$	0.643	(0.113)	0.334	0.652
(20)	$\text{corr}(\Delta q_t^o, \Delta q_{t-1}^o)$	0.213	(0.211)	0.326	0.691

Notes: x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$.

4.3 Modeling OPEC as a cartel

Our model assumes that oil firms are competitive. It would be interesting to study the properties of a version of our model in which OPEC firms behave as a cartel and non-OPEC firms are a competitive fringe. Stiglitz (1976) and Hassler et al. (2010) solve for an equilibrium in which the oil market is controlled by a monopolist that faces a demand with constant elasticity. A version of our model in which OPEC behaves as a cartel is significantly more challenging to solve for three reasons. First, there is a sizable competitive fringe, so the cartel faces a residual demand that is endogenous and

does not have constant elasticity. Second, in our model the extraction decision has a dynamic element because the marginal cost function of oil at time t is a function of all past investment decisions. Third, and most importantly, a version of our model in which OPEC is a cartel involves a complex dynamic game. In particular, the cartel can influence its future competition by manipulating today's oil prices in order to influence the investment decisions of its competitors. It is challenging to solve the resulting game even if we assume perfect commitment on the part of the cartel. We leave this task for future research.

Here, we take two modest steps towards investigating the potential impact of non-competitive behavior by OPEC on our results. First, we solved a version of the model in which OPEC firms behave as a cartel for one period and behave competitively thereafter. Non-OPEC act as a competitive fringe in all periods. We find that the response of this economy to demand and supply shocks is similar quantitatively and qualitatively to the impulse response functions we discuss in the previous section.

Second, we re-estimate the model allowing demand and supply shocks to be correlated. One interpretation of this correlation is that it might be induced by non-competitive behavior on the part of OPEC firms; for example, OPEC firms can be thought to change their supply so as to smooth oil prices. We present the results in two tables included in the Appendix. Table 14 reports our parameter estimates and Table 15 the implications of this version of the model for the moments targeted in the estimation. The estimated correlation between innovations to demand and supply (ε^d and ε^u) is -0.64 . This negative correlation is consistent with the smoothing hypothesis: OPEC firms cut production when demand is low. The fit of this version of the model is very similar to the fit of the version with independent shocks.

4.4 Robustness

In this section we discuss the results of two robustness exercises. First we consider two alternative instruments for the price of oil. Second, we exclude from the sample two

countries for which the data might have larger measurement error (Saudi Arabia and Venezuela). The tables with our robustness results are included in the appendix.

In our benchmark results, we instrumented the price of oil with the forecast of the cyclical component of world GDP. Here we consider two alternative instruments for oil prices: copper prices, as in Newell et al. (2016), and the IMF's metals price index. We deflate both indexes by the U.S. consumer price index. Tables 18 and 23 show our estimates of the elasticity of the extraction rate with respect to prices. These estimates are still quite low (0.26 and 0.20 instrumenting with real copper and metals prices, respectively) but are higher than our benchmark estimate (0.12).

Tables 19 and 24 report the elasticity of the extensive margin (number of oil fields in operation) with respect to the real price of oil. As in our benchmark results we find that this elasticity is not significantly different from zero.

Tables 20 and 25 contain our estimates of η obtained using real copper prices and real metals prices as instruments. Both instruments yield a lower value of η . Our estimates of η for non-OPEC are 9.3 in the benchmark case, 4.1 with real copper prices as an instrument, and 4.4 with real metals prices. We re-estimated the model with these values of η and report the results in Tables 21, 22, 26, and 27. The fit of the two alternative models with $\eta = 4.1$ and $\eta = 4.4$ is only slightly worse than the fit of the benchmark model. The parameter estimates are similar across the three models. The main difference is that demand shocks are more persistent in the models with $\eta = 4.1$ or $\eta = 4.4$ than in the benchmark model.

Next we redo our analysis excluding Saudi Arabia and Venezuela from the sample. Our motivation is the possibility of larger measurement error for these two countries. Table 28 shows the estimates of the elasticity of the extraction rate with respect to prices obtained using this restricted sample. This estimate (0.15) is similar to our benchmark estimate (0.12). Table 29 reports the elasticity of the extensive margin (number of oil fields in operation) with respect to the real price of oil. As in our benchmark results we find that this elasticity is not significantly different from zero. We re-estimated

our structural model excluding Saudi Arabia and Venezuela from the countries used to compute the data moments. The point estimates of the data moments are very similar to those of the full sample and, as a result, the estimated parameters and model fit are quite similar to those obtained in the benchmark specification (see Tables 31 and 32).

5 The impact of fracking


The advent of fracking is transforming the oil industry, making the U.S. once again one of the world's top oil producers.¹⁷ We study the quantitative impact of fracking using an extended version of our model that incorporates fracking firms. Since our dataset includes the universe of oil fields in operation, it allows us to compare the properties of conventional oil fields with the properties of oil fields explored using fracking. We find that fracking operations differ greatly from conventional oil production in terms of production flexibility and lags between investment and production. Our model implies that an expansion in the share of fracking production in total oil production will result in a sizable decline in the volatility of oil prices.

There are two important differences between fracking and conventional forms of oil production. First, the lag between investment and production is much shorter for fracking operations. Second, it is much less costly to adjust the extraction rate in fracking operations than in conventional oil operations.

Our results are consistent with the findings of Bjørnland et al. (2017). These authors use monthly data for North Dakota to show that oil production from shale wells is much more flexible than conventional oil production. Additional evidence consistent with the notion that that fracking operations are very flexible, comes from data compiled by Baker Hughes on the number of oil rigs in operation in the U.S. Between January 2009 and September 2014, oil prices rose from 42 to 93 dollars per barrel. During this period, the number of oil rigs in operation increased from 345 to 1,600. Most of the new rigs are

¹⁷See Kilian (2016), Gilje et al. (2016) and Melek, Plante and Yucel (2017) for a discussion of the impact of fracking on oil and gasoline markets.

Figure 7: Empirical dist. of lags between investment to production - fracking fields



likely to have been used in fracking operations. Between September 2014 and February 2016, oil prices plummeted from 93 to 30 dollars per barrel. During this period, the number of oil rigs in operation fell from 1,600 to 400.

Figure 7 shows the distribution of the lag between investment and the first year of production in non-conventional oil fields. Our maximum likelihood estimator of λ is 1.13, so the average lag between investment and production is about one year. Recall that this lag is 12 years for conventional production.

Table 11 reports our estimates of η for fields explored with fracking obtained using the cyclical component of world real GDP as an instrument. We see that for non-conventional oil fields our estimate of η is roughly 1.7. In contrast, when we include all oil fields in our sample in the regression, we obtain an estimate of $\eta = 9.3$.

To study the impact of fracking, we include in our model a third type of firm that produces oil using fracking. These firms have extraction costs that are less convex ($\eta^F = 1.7 < 9.3$), no lag between investment and production ($\lambda^F = 1$), and no lag in the adjustment of the extraction rate. We also assume that fracking firms are not

Table 12: Extraction rate adjustment costs regression - fracking fields

Variable	(1)	(2)
ln(extraction)	1.68*** (0.34)	9.30*** (1.14)
Oil field FE	✓	✓
Year of operation FE	✓	✓
Sample	Fracking fields	All non-OPEC
IV	✓	✓
1 st stage F-stat	9.1	53
Clusters (oil fields)	952	9,969
Observations	4,940	146,879

subject to supply shocks.

There is a continuum of measure one of fracking firms. The problem of the representative firm is to maximize its value (V^F):

$$\max_{\{I_t^F, \theta_{t+1}^F, K_{t+1}^F, X_{t+1}^F\}} V^F = E_0 \sum_{t=0}^{\infty} \beta^t \left[P_t \theta_t^F K_t^F - I_t^F - \psi^F (\theta_t^F)^{\eta^F} K_t^F \right], \quad (13)$$

subject to

$$X_{t+1}^F = (1 - \lambda^F) X_t^F + (I_t^F)^\alpha (L^F)^{1-\alpha} \quad (14)$$

$$K_{t+1}^F = (1 - \theta_t^F) K_t^F + \lambda^F X_{t+1}^F. \quad (15)$$

Here, I_t^F denotes investment, θ_t^F the extraction rate, X_t^F exploration capital, K_t^F oil reserves, and L^F the land available to the fracking firm.

We assume that the convexity of extraction costs is lower for fracking firms: $\eta^F < \eta$. This assumption does not imply that the average and marginal cost of extraction for fracking firms are different than those of the traditional firms, as this comparison also depends on ψ^F and on the equilibrium levels of reserves and extraction rates. The optimality conditions for fracking firms are identical to those for non-OPEC firms.

Total production is now given by:

$$\boldsymbol{\theta}_t^N \mathbf{K}_t^N + e^{-u_t} \boldsymbol{\theta}_t^O \mathbf{K}_t^O + \boldsymbol{\theta}_t^F \mathbf{K}_t^F = q_t,$$

where \mathbf{K}_t^F and $\boldsymbol{\theta}_t^F$ denote the aggregate reserves and aggregate extraction rate of fracking firms, respectively.

These new forms of oil production are only feasible in some parts of the globe. Rystad estimates that they will represent 20 percent of oil production by 2050. We calibrate ψ^F and the amount of land available for fracking to satisfy two properties. First, in the steady state fracking represents 20 percent of total oil production. Second, the steady-state extraction rate for fracking coincides with the one in our data (1.2 percent).

In equilibrium, the marginal return to investing in fracking is equal to that of investing in conventional oil production. Our calibration implies that the steady state marginal cost of extracting a barrel of oil is 19 percent higher for fracking firms than for conventional production. This higher production cost is compensated by the shorter lag between investment and production associated with fracking.

Figures 8 and 9 show the impulse response for a demand and supply shock, respectively. The panels for the price response also include the price response in our benchmark model without fracking firms. We see that fracking firms respond much more to the shock, in terms of production, extraction and investment, than non-fracking firms. The result is a much lower increase in the price of oil.

Comparing the response of conventional firms in a world with (Figures 8 and 9) and without fracking (Figures 5 and 6), we see that conventional firms respond much less to shocks in a world with fracking firms. Since fracking firms are more nimble, they respond rapidly to shocks, moderating the equilibrium price increase and making it less profitable for conventional firms to change their investment and extraction rates.

Table 13 compares the implications of versions of the model with and without fracking for some key moments. We find that the main impact of fracking is to reduce the volatility of oil prices by 65 percent from 0.23 to 0.08. Investment and production

Figure 8: Impulse response to a demand shock - fracking

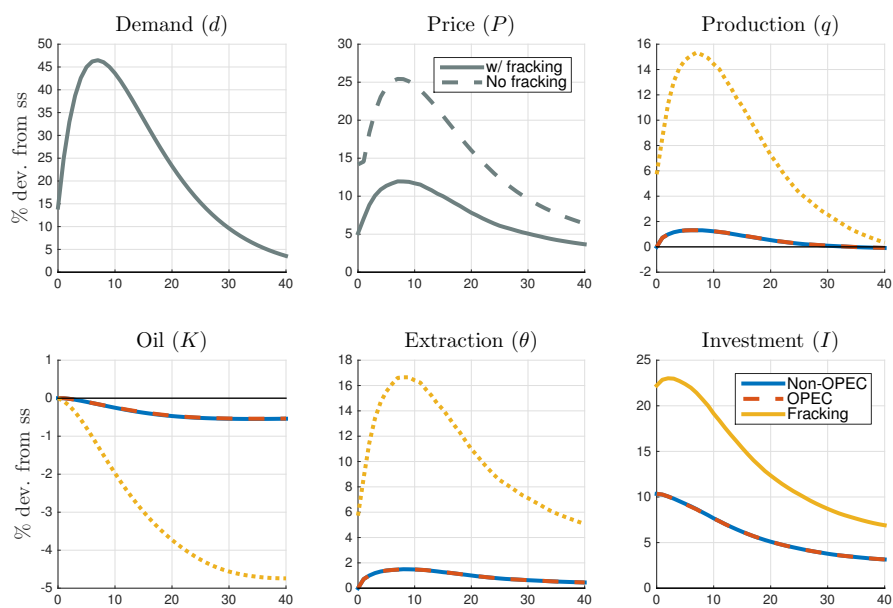


Figure 9: Impulse response to a supply shock - fracking

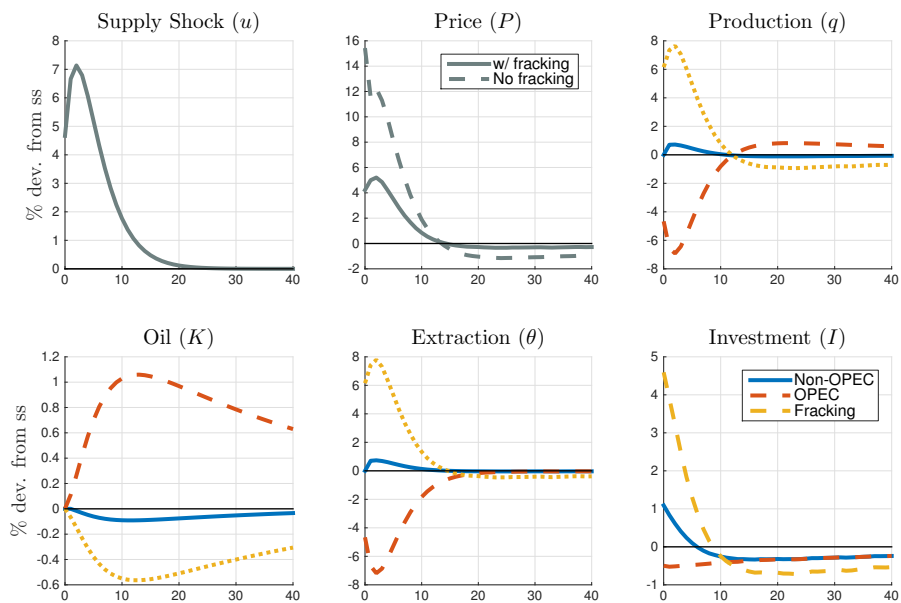


Table 13: Implication of fracking on key aggregate moments

	Moment	Benchm. model	Model w/ fracking
(1)	$std(\Delta p_t)$	0.23	0.08
(2)	$std(\Delta i_t)$	0.21	0.12
(3)	$std(\Delta q_t)$	0.03	0.02
(4)	$corr(\Delta p_t, \Delta i_t)$	0.69	0.7
(5)	$corr(\Delta p_t, \Delta q_t)$	-0.47	0.6
(6)	$corr(\Delta q_t, \Delta i_t)$	-0.06	0.53

Notes: x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$.

are also less volatile in the version of the model with fracking. In addition, there is also a higher correlation between prices and quantities, and between investment and quantities, reflecting the response of fracking firms to high-frequency movements in prices.

6 Conclusion

Our paper reviews some key facts about the oil market and proposes a simple industry equilibrium model that is consistent with these facts.

We leave four interesting projects for future research. The first is to develop a richer model of firm heterogeneity. In our model there are only two types of firms, OPEC and non-OPEC. In practice, firms in both groups differ in their attributes which results in different choices of investment and extraction rates. The second is to study a version of our model in which OPEC firms act as a cartel while non-OPEC firms are a competitive fringe. The third is to introduce the possibility of above-ground inventories that can be used by commodity speculators to respond to high-frequency changes in oil prices.¹⁸ The fourth is to combine our model of the oil market with a fully-fledged model of the

¹⁸See Deaton and Laroque (1992), Deaton and Laroque (1996), Kilian and Murphy (2014), and Olovsson (2016) for a discussion.

world economy. Such a combined model would allow us to study the effect of energy-saving technical change and evaluate the impact of solar, wind, and other alternative energy sources on the dynamics of the oil markets and the world economy.

References

- ACEMOGLU, D., CARVALHO, V. M. and OZDAGLAR, A. and TAHBAZ-SALEHI, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, **80** 1977–2016.
- ADAO, B. and NARAJABAD, B. and TEMZELIDES, T. P. (2017). Renewable technology adoption and the macroeconomy.
- ALQUIST, R. and KILIAN, L. and VIGFUSSON, R. J. (2013). Forecasting the price of oil. In *Handbook of economic forecasting*, vol. 2. Elsevier, 427–507.
- ANDERSON, S. T. and KELLOGG, R. and SALANT, S. W. (2017). Hotelling under pressure. *Journal of Political Economy*.
- AREZKI, R. and RAMEY, V. A. and SHENG, L. (2016). News shocks in open economies: Evidence from giant oil discoveries. *The Quarterly Journal of Economics*.
- BACKUS, D. K. and CRUCINI, M. J. (2000). Oil prices and the terms of trade. *Journal of international Economics*, **50** 185–213.
- BAQAEE, D. R. and FARHI, E. (2017). The macroeconomic impact of microeconomic shocks: Beyond hulten’s theorem. Tech. rep., National Bureau of Economic Research.
- BARSKY, R. B. and KILIAN, L. (2001). Do we really know that oil caused the great stagflation? a monetary alternative. *NBER Macroeconomics annual*, **16** 137–183.
- BARSKY, R. B. and KILIAN, L. (2004). Oil and the macroeconomy since the 1970s. *The Journal of Economic Perspectives*, **18** 115–134.

- BINSBERGEN, J. H. (2016). Good-specific habit formation and the cross-section of expected returns. *The Journal of Finance*, **71** 1699–1732.
- BJØRNLAND, H. C. and NORDVIK, F. M. and ROHRER, M. (2017). Supply flexibility in the shale patch: Evidence from north dakota. *manuscript, Norges Bank*.
- BLANCHARD, O. J. and GALI, J. (2007). The macroeconomic effects of oil shocks: Why are the 2000s so different from the 1970s? Tech. rep., National Bureau of Economic Research.
- BODENSTEIN, M. and ERCEG, C. J. and GUERRIERI, L. (2011). Oil shocks and external adjustment. *Journal of International Economics*, **83** 168–184.
- CALDARA, D. and CAVALLO, M. and IACOVIELLO, M. M. (2017). Oil price elasticities and oil price fluctuations.
- CHARI, V. and CHRISTIANO, L. (2014). The optimal extraction of exhaustible resources. Tech. rep., Federal Reserve Bank of Minneapolis.
- DAMODARAN, A. (2017). Cost of equity and capital. Tech. rep., New York University.
- DEATON, A. (1999). Commodity prices and growth in africa. *The Journal of Economic Perspectives*, **13** 23–40.
- DEATON, A. and LAROQUE, G. (1992). On the behaviour of commodity prices. *The Review of Economic Studies*, **59** 1–23.
- DEATON, A. and LAROQUE, G. (1996). Competitive storage and commodity price dynamics. *Journal of Political Economy*, **104** 896–923.
- EZEKIEL, M. (1938). The cobweb theorem. *The Quarterly Journal of Economics*, **52** 255–280.

- FINN, M. G. (2000). Perfect competition and the effects of energy price increases on economic activity. *Journal of Money, Credit and banking* 400–416.
- GABAIX, X. (2011). The granular origins of aggregate fluctuations. *Econometrica*, **79** 733–772.
- GILJE, E. and READY, R. and ROUSSANOV, N. (2016). Fracking, drilling, and asset pricing: Estimating the economic benefits of the shale revolution. Tech. rep., National Bureau of Economic Research.
- HAMILTON, J. D. (1983). Oil and the macroeconomy since world war ii. *Journal of Political Economy*, **91** 228–248.
- HARVEY, D. I., KELLARD, N. M. and MADSEN, J. B. and WOHR, M. E. (2010). The prebisch-singer hypothesis: four centuries of evidence. *The Review of Economics and Statistics*, **92** 367–377.
- HASSLER, J. and KRUSELL, P. and OLOVSSON, C. (2010). Oil monopoly and the climate. *The American Economic Review*, **100** 460–464.
- HOTELLING, H. (1931). The economics of exhaustible resources. *Journal of Political Economy*, **39** 137–175.
- KELLOGG, R. (2014). The effect of uncertainty on investment: evidence from texas oil drilling. *The American Economic Review*, **104** 1698–1734.
- KILIAN, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *The American Economic Review*, **99** 1053–1069.
- KILIAN, L. (2014). Oil price shocks: causes and consequences. *Annu. Rev. Resour. Econ.*, **6** 133–154.

- KILIAN, L. (2016). The impact of the shale oil revolution on us oil and gasoline prices. *Review of Environmental Economics and Policy*, **10** 185–205.
- KILIAN, L. and MURPHY, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, **29** 454–478.
- KIM, I.-M. and LOUNGANI, P. (1992). The role of energy in real business cycle models. *Journal of Monetary Economics*, **29** 173–189.
- LEDUC, S. and SILL, K. (2004). A quantitative analysis of oil-price shocks, systematic monetary policy, and economic downturns. *Journal of Monetary Economics*, **51** 781–808.
- LEON-LEDESMA, M. A. and SATCHI, M. (2016). Appropriate technology and balanced growth. *The Review of Economic Studies*.
- LIPPI, F. and NOBILI, A. (2012). Oil and the macroeconomy: a quantitative structural analysis. *Journal of the European Economic Association*, **10** 1059–1083.
- NERLOVE, M. (1958). Adaptive expectations and cobweb phenomena. *The Quarterly Journal of Economics*, **72** 227–240.
- NEWELL, R. G. and PREST, B. C. and VISSING, A. (2016). Trophy hunting vs. manufacturing energy: The price-responsiveness of shale gas. Tech. rep., National Bureau of Economic Research.
- OLOVSSON, C. (2016). Oil prices in a real-business-cycle model with precautionary demand for oil.
- RAVN, M. and SCHMITT-GROHÉ, S. and URIBE, M. (2006). Deep habits. *The Review of Economic Studies*, **73** 195–218.

- ROTEMBERG, J. J. and WOODFORD, M. (1996). Imperfect competition and the effects of energy price increases on economic activity. *Journal of Money, Credit and Banking*, **28** 549–577.
- ROUWENHORST, K. G. (1991). Time to build and aggregate fluctuations: A reconsideration. *Journal of Monetary Economics*, **27** 241–254.
- STIGLITZ, J. E. (1976). Monopoly and the rate of extraction of exhaustible resources. *The American Economic Review*, **66** 655–661.

Appendix

A Additional tables

Table 14: Estimated parameters when shocks are correlated

Parameter	Benchm. model	(s.e.)	Model w/ corr. shocks	(s.e.)
ϵ	0.135	(0.015)	0.084	(0.013)
α	3×10^{-6}	(0.978)	1×10^{-7}	(0.674)
ρ_1^d	1.761	(0.075)	1.759	(0.089)
ρ_2^d	-0.775	(0.073)	-0.773	(0.09)
ρ_1^u	1.428	(0.15)	1.719	(0.072)
ρ_2^u	-0.51	(0.11)	-0.733	(0.063)
$var(e_t^d)$	0.02	(0.007)	0.05	(0.018)
$var(e_t^u)$	0.002	(0.001)	0.002	(0.001)
$corr(e_t^d, e_t^u)$	-	-	-0.639	(0.105)

Table 15: Data and model moments when shocks are correlated

	Moment	Data	(s.e.)	Benchm. model	Model w/ corr. shocks
(1)	$\text{std}(\Delta p_t)$	0.273	(0.028)	0.227	0.232
(2)	$\text{std}(\Delta i_t^n)$	0.192	(0.24)	0.213	0.205
(3)	$\text{std}(\Delta i_t^o)$	0.193	(0.27)	0.211	0.216
(4)	$\text{std}(\Delta q_t^n)$	0.022	(0.003)	0.025	0.022
(5)	$\text{std}(\Delta q_t^o)$	0.069	(0.011)	0.054	0.072
(6)	$\text{corr}(\Delta p_t, \Delta i_t^n)$	0.557	(0.147)	0.709	0.749
(7)	$\text{corr}(\Delta p_t, \Delta i_t^o)$	0.362	(0.109)	0.657	0.53
(8)	$\text{corr}(\Delta p_t, \Delta q_t^n)$	0.031	(0.069)	0.049	0.003
(9)	$\text{corr}(\Delta p_t, \Delta q_t^o)$	0.030	(0.122)	-0.572	-0.418
(10)	$\text{corr}(\Delta i_t^n, \Delta i_t^o)$	0.673	(0.096)	0.997	0.946
(11)	$\text{corr}(\Delta i_t^n, \Delta q_t^n)$	0.087	(0.094)	-0.001	0.008
(12)	$\text{corr}(\Delta i_t^n, \Delta q_t^o)$	0.023	(0.112)	-0.103	0.027
(13)	$\text{corr}(\Delta i_t^o, \Delta q_t^n)$	-0.034	(0.145)	0.003	0
(14)	$\text{corr}(\Delta i_t^o, \Delta q_t^o)$	-0.226	(0.153)	-0.036	0.251
(15)	$\text{corr}(\Delta q_t^n, \Delta q_t^o)$	-0.141	(0.125)	0.014	-0.221
(16)	$\text{corr}(\Delta p_t, \Delta p_{t-1})$	-0.027	(0.088)	-0.018	-0.125
(17)	$\text{corr}(\Delta i_t^n, \Delta i_{t-1}^n)$	0.119	(0.135)	0.008	0.025
(18)	$\text{corr}(\Delta i_t^o, \Delta i_{t-1}^o)$	0.311	(0.096)	0.009	0.022
(19)	$\text{corr}(\Delta q_t^n, \Delta q_{t-1}^n)$	0.643	(0.113)	0.334	0.46
(20)	$\text{corr}(\Delta q_t^o, \Delta q_{t-1}^o)$	0.213	(0.211)	0.326	0.671

Notes: x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$.

B Robustness Tables

B.1 Sample period starting at 1900

Table 16: Price elasticity of extraction rates (1900–2015)

Variable	(1)	(2)	(3)	(4)	(5)
$\ln(\text{price})$	0.06*** (0.006)	0.09*** (0.008)	0.12*** (0.009)	0.14*** (0.01)	0.16*** (0.011)
$\ln(\text{price}) \times \mathbb{1}_{\text{OPEC}}$			-0.21*** (0.022)	-0.2*** (0.023)	
$\ln(\text{price}) \times \mathbb{1}_{\text{Big Firm}}$				-0.06*** (0.018)	
$\ln(\text{price}) \times \mathbb{1}_{ \Delta \ln(p) > 0.1}$					-0.15*** (0.01)
Oil field FE	✓	✓	✓	✓	✓
Operation year FE	✓	✓	✓	✓	✓
Year trend	✓	✓	✓	✓	✓
IV	✗	✓	✓	✓	✓
Clusters (oil fields)	13,811			13,372	
Observations	351,442			351,003	

Notes: oil fields with positive extraction in 1901-2015, excluding last year of operation. Standard errors in parenthesis. Instrument for price using one-year lagged price (F-stat > 1000 in (2)-(5)). (***) - significant at a 1% level.

Table 17: Price elasticity of oil fields in operation (1900–2015)

Variable	(1)	(2)	(3)	(4)
ln(price)	−0.4 (0.27)	−0.51 (0.32)	−0.18 (0.13)	−0.32 (0.19)
Year trend	✓	✓	✓	✓
IV	✗	✓	✓	✓
Dep. variable	All fields	All fields	Non-OPEC fields	OPEC fields
Observations	114	114	114	114

Notes: Newey-West standard errors computed with 5-year lags in parenthesis.

B.2 Instrument with real copper prices

Table 18: Price elasticity of extraction rates

Variable	(1)	(2)	(3)	(4)	(5)
ln(price)	0.09*** (0.009)	0.26*** (0.018)	0.31*** (0.019)	0.34*** (0.021)	0.22*** (0.018)
ln(price) \times $\mathbb{1}_{\text{OPEC}}$			-0.25*** (0.041)	-0.24*** (0.041)	
ln(price) \times $\mathbb{1}_{\text{Big Firm}}$				-0.1*** (0.032)	
ln(price) \times $\mathbb{1}_{ \Delta \ln(p) > 0.1}$					-0.03*** (0.005)
Oil field FE	✓	✓	✓	✓	✓
Operation year FE	✓	✓	✓	✓	✓
Year trend	✓	✓	✓	✓	✓
IV	✗	✓	✓	✓	✓
Clusters (oil fields)	12,187			11,479	
Observations	173,742			173,034	

Notes: oil fields with positive extraction in 1971-2015, excluding last year of operation. Standard errors in parenthesis. Instrument for price using real copper price prediction (F-stat > 1000 in (2)-(5)). (***) - significant at a 1% level.

Table 19: Price elasticity of oil fields in operation

Variable	(1)	(2)	(3)	(4)
ln(price)	-0.05 (0.06)	-0.44 (0.29)	-0.13 (0.10)	-0.31 (0.19)
Year trend	✓	✓	✓	✓
IV	✗	✓	✓	✓
Dep. variable	All fields	All fields	Non-OPEC fields	OPEC fields
Observations	45	45	45	45

Notes: Newey-West standard errors computed with 5-year lags in parenthesis.

Table 20: Extraction rate adjustment costs regression

Variable	(1)	(2)	(3)
ln(extraction)	4.54*** (0.22)	4.05*** (0.18)	15.19** (7.72)
Oil field FE	✓	✓	✓
Operation year FE	✓	✓	✓
Sample	All	Non-OPEC	OPEC
IV	✓	✓	✓
1 st stage F-stat	276	198	3.4
Clusters (oil fields)	11,527	9,969	1,558
Observations	174,339	146,879	27,460

Table 21: Estimated parameters

Parameter	Estimate	(s.e.)
ϵ	0.077	(0.014)
α	0.062	(0.079)
ρ_1^d	1.991	(0.003)
ρ_2^d	-0.991	(0.003)
ρ_1^u	1.729	(0.077)
ρ_2^u	-0.748	(0.072)
$var(e_t^d)$	0.001	(0.000...)
$var(e_t^u)$	0.001	(0.000...)

Table 22: Data and model moments

	Moment	Data	(s.e.)	Benchm. model	Model w/ copper IV
(1)	$std(\Delta p_t)$	0.273	(0.028)	0.227	0.195
(2)	$std(\Delta i_t^n)$	0.192	(0.24)	0.213	0.228
(3)	$std(\Delta i_t^o)$	0.193	(0.27)	0.211	0.217
(4)	$std(\Delta q_t^n)$	0.022	(0.003)	0.025	0.025
(5)	$std(\Delta q_t^o)$	0.069	(0.011)	0.054	0.034
(6)	$corr(\Delta p_t, \Delta i_t^n)$	0.557	(0.147)	0.709	0.621
(7)	$corr(\Delta p_t, \Delta i_t^o)$	0.362	(0.109)	0.657	0.484
(8)	$corr(\Delta p_t, \Delta q_t^n)$	0.031	(0.069)	0.049	-0.051
(9)	$corr(\Delta p_t, \Delta q_t^o)$	0.030	(0.122)	-0.572	-0.357
(10)	$corr(\Delta i_t^n, \Delta i_t^o)$	0.673	(0.096)	0.997	0.977
(11)	$corr(\Delta i_t^n, \Delta q_t^n)$	0.087	(0.094)	-0.001	0.228
(12)	$corr(\Delta i_t^n, \Delta q_t^o)$	0.023	(0.112)	-0.103	-0.066
(13)	$corr(\Delta i_t^o, \Delta q_t^n)$	-0.034	(0.145)	0.003	0.238
(14)	$corr(\Delta i_t^o, \Delta q_t^o)$	-0.226	(0.153)	-0.036	0.095
(15)	$corr(\Delta q_t^n, \Delta q_t^o)$	-0.141	(0.125)	0.014	-0.198
(16)	$corr(\Delta p_t, \Delta p_{t-1})$	-0.027	(0.088)	-0.018	-0.108
(17)	$corr(\Delta i_t^n, \Delta i_{t-1}^n)$	0.119	(0.135)	0.008	0.295
(18)	$corr(\Delta i_t^o, \Delta i_{t-1}^o)$	0.311	(0.096)	0.009	0.324
(19)	$corr(\Delta q_t^n, \Delta q_{t-1}^n)$	0.643	(0.113)	0.334	0.591
(20)	$corr(\Delta q_t^o, \Delta q_{t-1}^o)$	0.213	(0.211)	0.326	0.609

Notes: x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$.

B.3 Instrument with real metal prices

Table 23: Price elasticity of extraction rates

Variable	(1)	(2)	(3)	(4)	(5)
ln(price)	0.09*** (0.009)	0.20*** (0.016)	0.24*** (0.018)	0.26*** (0.02)	0.07*** (0.014)
ln(price) \times $\mathbb{1}_{\text{OPEC}}$			-0.22*** (0.04)	-0.21*** (0.041)	
ln(price) \times $\mathbb{1}_{\text{Big Firm}}$				-0.06* (0.032)	
ln(price) \times $\mathbb{1}_{ \Delta \ln(p) > 0.1}$					-0.12*** (0.006)
Oil field FE	✓	✓	✓	✓	✓
Operation year FE	✓	✓	✓	✓	✓
Year trend	✓	✓	✓	✓	✓
IV	✗	✓	✓	✓	✓
Clusters (oil fields)	12,187			11,479	
Observations	173,742			173,034	

Notes: oil fields with positive extraction in 1971-2015, excluding last year of operation. Standard errors in parenthesis. Instrument for price using real metal price index prediction (F-stat > 1000 in (2)-(5)). *** [*] - significant at a 1% [10%] level.

Table 24: Price elasticity of oil fields in operation

Variable	(1)	(2)	(3)	(4)
ln(price)	-0.05 (0.06)	-0.08 (0.07)	-0.02 (0.02)	-0.06 (0.06)
Year trend	✓	✓	✓	✓
IV	✗	✓	✓	✓
Dep. variable	All fields	All fields	Non-OPEC fields	OPEC fields
Observations	45	45	45	45

Notes: Newey-West standard errors computed with 5-year lags in parenthesis.

Table 25: Extraction rate adjustment costs regression

Variable	(1)	(2)	(3)
ln(extraction)	4.72*** (0.22)	4.26*** (0.18)	11.73** (4.15)
Oil field FE	✓	✓	✓
Operation year FE	✓	✓	✓
Sample	All	Non-OPEC	OPEC
IV	✓	✓	✓
1 st stage F-stat	271	286	3.5
Clusters (oil fields)	11,527	9,969	1,558
Observations	174,339	146,879	27,460

Table 26: Estimated parameters

Parameter	Estimate	(s.e.)
ϵ	0.077	(0.013)
α	0.067	(0.087)
ρ_1^d	1.991	(0.003)
ρ_2^d	-0.991	(0.003)
ρ_1^u	1.728	(0.076)
ρ_2^u	-0.746	(0.071)
$var(e_t^d)$	0.001	(0.000...)
$var(e_t^u)$	0.001	(0.000...)

Table 27: Data and model moments

	Moment	Data	(s.e.)	Benchm. model	Model w/ metal IV
(1)	std(Δp_t)	0.273	(0.028)	0.227	0.199
(2)	std(Δi_t^n)	0.192	(0.24)	0.213	0.229
(3)	std(Δi_t^o)	0.193	(0.27)	0.211	0.217
(4)	std(Δq_t^n)	0.022	(0.003)	0.025	0.025
(5)	std(Δq_t^o)	0.069	(0.011)	0.054	0.035
(6)	corr($\Delta p_t, \Delta i_t^n$)	0.557	(0.147)	0.709	0.621
(7)	corr($\Delta p_t, \Delta i_t^o$)	0.362	(0.109)	0.657	0.484
(8)	corr($\Delta p_t, \Delta q_t^n$)	0.031	(0.069)	0.049	-0.041
(9)	corr($\Delta p_t, \Delta q_t^o$)	0.030	(0.122)	-0.572	-0.372
(10)	corr($\Delta i_t^n, \Delta i_t^o$)	0.673	(0.096)	0.997	0.977
(11)	corr($\Delta i_t^n, \Delta q_t^n$)	0.087	(0.094)	-0.001	0.244
(12)	corr($\Delta i_t^n, \Delta q_t^o$)	0.023	(0.112)	-0.103	-0.076
(13)	corr($\Delta i_t^o, \Delta q_t^n$)	-0.034	(0.145)	0.003	0.238
(14)	corr($\Delta i_t^o, \Delta q_t^o$)	-0.226	(0.153)	-0.036	0.087
(15)	corr($\Delta q_t^n, \Delta q_t^o$)	-0.141	(0.125)	0.014	-0.202
(16)	corr($\Delta p_t, \Delta p_{t-1}$)	-0.027	(0.088)	-0.018	-0.111
(17)	corr($\Delta i_t^n, \Delta i_{t-1}^n$)	0.119	(0.135)	0.008	0.293
(18)	corr($\Delta i_t^o, \Delta i_{t-1}^o$)	0.311	(0.096)	0.009	0.323
(19)	corr($\Delta q_t^n, \Delta q_{t-1}^n$)	0.643	(0.113)	0.334	0.587
(20)	corr($\Delta q_t^o, \Delta q_{t-1}^o$)	0.213	(0.211)	0.326	0.599

Notes: x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$.

B.4 Excluding Saudi Arabia and Venezuela from the sample

Table 28: Price elasticity of extraction rates

Variable	(1)	(2)	(3)	(4)	(5)
ln(price)	0.11*** (0.009)	0.15*** (0.012)	0.17*** (0.015)	0.18*** (0.02)	0.18*** (0.013)
ln(price) \times $\mathbb{1}_{\text{OPEC}}$			-0.17*** (0.07)	-0.17*** (0.041)	
ln(price) \times $\mathbb{1}_{\text{Big Firm}}$				-0.05 (0.05)	
ln(price) \times $\mathbb{1}_{ \Delta \ln(p) > 0.1}$					-0.18*** (0.001)
Oil field FE	✓	✓	✓	✓	✓
Operation year FE	✓	✓	✓	✓	✓
Year trend	✓	✓	✓	✓	✓
IV	✗	✓	✓	✓	✓
Clusters (oil fields)	11,893			11,195	
Observations	167,848			167,150	

Notes: oil fields with positive extraction in 1971-2015, excluding last year of operation. Standard errors in parenthesis. Instrument for price using real metal price index prediction (F-stat > 1000 in (2)-(5)). *** [*] - significant at a 1% [10%] level.

Table 29: Price elasticity of oil fields in operation

Variable	(1)	(2)	(3)	(4)
ln(price)	-0.03 (0.08)	0.01 (0.1)	-0.01 (0.03)	0.02 (0.09)
Year trend	✓	✓	✓	✓
IV	✗	✓	✓	✓
Dep. variable	All fields	All fields	Non-OPEC fields	OPEC fields
Observations	45	45	45	45

Notes: Newey-West standard errors computed with 5-year lags in parenthesis.

Table 30: Extraction rate adjustment costs regression

Variable	(1)	(2)	(3)
ln(extraction)	9.97*** (1.20)	9.3*** (0.18)	19.77 (12.03)
Oil field FE	✓	✓	✓
Operation year FE	✓	✓	✓
Sample	All	Non-OPEC	OPEC
IV	✓	✓	✓
1 st stage F-stat	56	53	2.5
Clusters (oil fields)	11,243	9,969	1,274
Observations	168,388	146,879	21,509

Table 31: Estimated parameters

Parameter	Estimate	(s.e.)
ϵ	0.122	(0.017)
α	7×10^{-7}	(0.930)
ρ_1^d	1.765	(0.077)
ρ_2^d	-0.779	(0.075)
ρ_1^u	1.529	(0.094)
ρ_2^u	-0.585	(0.092)
$var(e_t^d)$	0.002	(0.001)
$var(e_t^u)$	0.022	(0.008)

Table 32: Data and model moments

Moment	Benchmark			No Saudi Arabia and Venezuela		
	Data	(s.e.)	Model	Data	(s.e.)	Model
(1) $std(\Delta p_t)$	0.273	(0.028)	0.227	0.273	(0.028)	0.226
(2) $std(\Delta i_t^n)$	0.192	(0.024)	0.213	0.192	(0.024)	0.221
(3) $std(\Delta i_t^o)$	0.193	(0.027)	0.211	0.208	(0.031)	0.218
(4) $std(\Delta q_t^n)$	0.022	(0.003)	0.025	0.273	(0.028)	0.024
(5) $std(\Delta q_t^o)$	0.069	(0.011)	0.054	0.076	(0.016)	0.049
(6) $corr(\Delta p_t, \Delta i_t^n)$	0.557	(0.147)	0.709	0.557	(0.147)	0.767
(7) $corr(\Delta p_t, \Delta i_t^o)$	0.362	(0.109)	0.657	0.36	(0.12)	0.708
(8) $corr(\Delta p_t, \Delta q_t^n)$	0.031	(0.069)	0.049	0.031	(0.069)	0.061
(9) $corr(\Delta p_t, \Delta q_t^o)$	0.030	(0.122)	-0.572	-0.025	(0.104)	-0.53
(10) $corr(\Delta i_t^n, \Delta i_t^o)$	0.673	(0.096)	0.997	0.626	(0.108)	0.995
(11) $corr(\Delta i_t^n, \Delta q_t^n)$	0.087	(0.094)	-0.001	0.087	(0.094)	0.001
(12) $corr(\Delta i_t^n, \Delta q_t^o)$	0.023	(0.112)	-0.103	-0.065	(0.112)	-0.145
(13) $corr(\Delta i_t^o, \Delta q_t^n)$	-0.034	(0.145)	0.003	-0.067	(0.138)	0.004
(14) $corr(\Delta i_t^o, \Delta q_t^o)$	-0.226	(0.153)	-0.036	-0.357	(0.246)	-0.062
(15) $corr(\Delta q_t^n, \Delta q_t^o)$	-0.141	(0.125)	0.014	-0.126	(0.089)	0.002
(16) $corr(\Delta p_t, \Delta p_{t-1})$	-0.027	(0.088)	-0.018	-0.027	(0.088)	-0.019
(17) $corr(\Delta i_t^n, \Delta i_{t-1}^n)$	0.119	(0.135)	0.008	0.119	(0.135)	0.009
(18) $corr(\Delta i_t^o, \Delta i_{t-1}^o)$	0.311	(0.096)	0.009	0.252	(0.11)	0.01
(19) $corr(\Delta q_t^n, \Delta q_{t-1}^n)$	0.643	(0.113)	0.334	0.643	(0.113)	0.381
(20) $corr(\Delta q_t^o, \Delta q_{t-1}^o)$	0.213	(0.211)	0.326	0.349	(0.154)	0.401

Notes: x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$.

C A balanced growth path

In the non-stochastic steady state of our benchmark model, oil production is constant. In this section, we extend the benchmark model so that it features a balanced growth path. Along the non-stochastic balanced growth path, the levels of production, reserves, and exploration capital, as well as investment and production costs, grow at a constant rate, while prices and extraction rates are constant. The model features investment-specific technological progress as well as a growing demand for oil.

The difference between the balanced growth model and our benchmark model narrows down to two equations, the law of motion for exploration capital and the demand function. The law of motion for exploration capital is given by

$$X_{t+1}^i = (1 - \lambda)X_t^i + A_t (I_t^i)^\alpha (L^i)^{1-\alpha} , \quad i \in \{O, N\} . \quad (16)$$

In our benchmark model $A_t = 1$ for all t . Here, instead, we assume that A_t grows at a constant rate g_A so that

$$\ln A_{t+1} = \ln A_t + g_A ,$$

for all t . The demand function is given by

$$P_t = \exp(d_t^{AR2}) \exp(d_t^g) Q_t^{-\frac{1}{\epsilon}} , \quad (17)$$

where d_t^{AR2} is a stochastic demand component that follows an AR(2) process, which corresponds to d_t in our benchmark model. The second demand component, d_t^g , grows at a constant rate g_d so that

$$d_{t+1}^g - d_t^g = g_d ,$$

for all t . To ensure a balanced growth path we assume that $g_d = \frac{1}{\epsilon(1-\alpha)}g_A$. As we show below, this assumption implies that prices and extraction rates are constant along the non-stochastic balanced growth path.

The Bellman equation of each firm is

$$\begin{aligned}
V(K, X, \theta, A, u; \Omega) &= \max_{\{K', X', \theta', I\}} P\theta K e^{-u} - I - \psi \left(\frac{\theta}{\theta_{ss}} \right)^\eta K + \beta EV(K', X', \theta', A', u'; \Omega') , \\
\text{s.t.} \quad K' &= (1 - \theta e^{-u})K + \lambda X' , \\
X' &= (1 - \lambda)X + AI\alpha L^{1-\alpha} ,
\end{aligned} \tag{18}$$

where $u = 0$ for non-OPEC firms.¹⁹ The variable Ω denotes the aggregate state of the economy, which consists of the aggregate level of reserves and exploration capital of non-OPEC and OPEC firms, the current extraction rates chosen by OPEC and non-OPEC firms, the state of the stochastic supply and demand shocks, and the level of investment technology A and demand d^g . To make the value function stationary over time we need to normalize it.

Let the scaled variable $\tilde{Z} = e^{gt}Z$, where $g = \frac{1}{1-\alpha}g_A$, and let $\tilde{V}(\cdot) = e^{gt}V(\cdot)$. The scaled Bellman equation is given by

$$\begin{aligned}
\tilde{V}(\tilde{K}, \tilde{X}, \theta, u; \tilde{\Omega}) &= \max_{\{\tilde{K}', \tilde{X}', \theta', \tilde{I}\}} P\theta\tilde{K}e^{-u} - \tilde{I} - \psi \left(\frac{\theta}{\theta_{ss}} \right)^\eta \tilde{K} + e^g\beta E\tilde{V}(\tilde{K}', \tilde{X}', \theta', u'; \tilde{\Omega}') , \\
\text{s.t.} \quad \tilde{K}' &= e^{-g}(1 - e^{-u}\theta)\tilde{K} + \lambda\tilde{X}' , \\
\tilde{X}' &= e^{-g}(1 - \lambda)\tilde{X} + e^{-g}\tilde{I}\alpha L^{1-\alpha} ,
\end{aligned}$$

where $\tilde{\Omega}$ is the aggregate state of the economy, which contains the scaled aggregate levels of OPEC and non-OPEC reserves and exploration capital, the extraction rates of OPEC and non-OPEC, and information about the stochastic processes of the AR(2) demand and supply shocks. As we confirm below, the scaled value function above and the scaled variables are constant in the non-stochastic steady state.

¹⁹For ease of notation we have omitted the superscripts for OPEC and non-OPEC firms, and use primed variables to denote a variable in the following period.

The first-order conditions (FOCs) for the firms' problem are given by:

$$[\tilde{K}'] \quad \mu_1 = \beta e^g E \left[P' \theta' e^{-u'} + e^{-g} (1 - \theta' e^{-u'}) \mu_1' - \psi \left(\frac{\theta'}{\theta_{ss}} \right)^\eta \right], \quad (19)$$

$$[\tilde{X}'] \quad \mu_2 = \lambda \mu_1 + \beta (1 - \lambda) E_t \mu_2', \quad (20)$$

$$[\tilde{I}] \quad 1 = e^{-g} \alpha \left(L / \tilde{I} \right)^{1-\alpha} \mu_2, \quad (21)$$

$$[\theta'] \quad \eta \psi (\theta')^{\eta-1} (\theta_{ss})^{-\eta} \tilde{K}' = E \left[(P' e^{-u'} - e^{-g} e^{-u'} \mu_1') \right] \tilde{K}', \quad (22)$$

These are the first-order conditions for both OPEC and non-OPEC firms, where $u = u' = 0$ for non-OPEC firms.

Non-stochastic steady state. Consider the non-stochastic steady state. Rearranging the first FOC we have

$$\mu_1 = \beta e^g \frac{\theta P - \psi}{1 - \beta(1 - \theta)}. \quad (23)$$

The second FOC can be written as

$$\mu_2 = \frac{\lambda}{1 - \beta(1 - \lambda)} \mu_1. \quad (24)$$

Substituting into the third FOC, we have

$$\left(\frac{\tilde{I}}{L} \right)^{1-\alpha} = e^{-g} \alpha \frac{\lambda}{1 - \beta(1 - \lambda)} \mu_1. \quad (25)$$

The last FOC simplifies to

$$\eta \psi (\theta)^{-1} = P - e^{-g} \mu_1. \quad (26)$$

The aggregate demand equation in the non-stochastic steady state is

$$P = e^{d_t^g} (\theta^N K^N + \theta^O K^O)^{-\frac{1}{\epsilon}}. \quad (27)$$

Since $g_d = \frac{1}{\epsilon} g$, we can simplify the aggregate demand equation to

$$P = \left(\theta^N \tilde{K}^N + \theta^O \tilde{K}^O \right)^{-\frac{1}{\epsilon}}. \quad (28)$$

Finally, the laws of motion for productive and exploration capital imply

$$\tilde{X} = \frac{1}{e^g - (1 - \lambda)} \left(\tilde{I} \right)^\alpha (L)^{1-\alpha} , \quad (29)$$

$$\tilde{K} = \frac{e^g}{e^g - (1 - \theta)} \frac{\lambda}{e^g - (1 - \lambda)} \left(\tilde{I} \right)^\alpha (L)^{1-\alpha} . \quad (30)$$

Combining equations (23) and (26) we get

$$P - \eta\psi(\theta)^{-1} = \beta \frac{\theta P - \psi}{1 - \beta(1 - \theta)} . \quad (31)$$

Rearranging we get

$$(1 - \beta)P = \left[\left(\frac{1 - \beta}{\theta} + \beta \right) \eta - \beta \right] \psi , \quad (32)$$

or simply

$$P = \eta\psi(\theta)^{-1} + \frac{\beta}{1 - \beta}(\eta - 1)\psi . \quad (33)$$

Given P , there is a single value of θ that solves the equation above. This property implies that in the non-stochastic steady state, the extraction rate is the same ($\theta^N = \theta^O$) for OPEC and non-OPEC firms. Using the equation above together with equations (23) and (25), we can pin down the investment to land ratios for OPEC and non-OPEC firms.

$$\left(\frac{\tilde{I}}{L} \right)^{1-\alpha} = \frac{\alpha\lambda}{1 - \beta(1 - \lambda)} \frac{\beta}{1 - \beta}(\eta - 1)\psi . \quad (34)$$

Substituting into the demand function we get

$$\left[\eta\psi(\theta)^{-1} + \frac{\beta}{1 - \beta}(\eta - 1)\psi \right]^{-\epsilon} = \frac{e^g\theta}{e^g - (1 - \theta)} \frac{\lambda}{e^g - (1 - \lambda)} \left[\frac{\alpha\lambda}{1 - \beta(1 - \lambda)} \frac{\beta}{1 - \beta}(\eta - 1)\psi \right]^{\frac{\alpha}{1 - \alpha}} (L^N + L^O) . \quad (35)$$

This equation implicitly pins down the equilibrium rate of extraction along the non-stochastic balanced growth path. Using the equilibrium extraction rate, we can find all other endogenous variables using the equations above. Thus, there exists a balanced growth path equilibrium along which oil prices and extraction rates are constant over time while the levels of reserves, exploration capital, and production, as well as investment and production costs are growing at a constant rate $g = \frac{1}{1 - \alpha}g_A$.