NBER WORKING PAPER SERIES

LAGS, COSTS, AND SHOCKS:
AN EQUILIBRIUM MODEL OF THE OIL INDUSTRY

Gideon Bornstein
Per Krusell
Sergio Rebelo

Working Paper 23423
http://www.nber.org/papers/w23423

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2017, Revised April 2019

We thank Hilde Bjornland, Mario Crucini,Wei Cui, Jésus Fernández-Villaverde, Matteo Iacoviello, Ravi Jaganathan, Ryan Kellogg, Lutz Kilian, Markus Kirchner, Valerie Ramey, and Rob Vigfusson for their comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Lags, Costs, and Shocks: An Equilibrium Model of the Oil Industry
Gideon Bornstein, Per Krusell, and Sergio Rebelo
NBER Working Paper No. 23423
May 2017, Revised April 2019
JEL No. Q4,Q43

ABSTRACT

We use a new micro data set that covers all oil fields in the world to estimate a stochastic industry-equilibrium model of the oil industry with two alternative market structures. In the first, all firms are competitive. In the second, OPEC firms act as a cartel. 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.

Gideon Bornstein
2211 Campus Drive
Dept. of Economics, Northwestern University
3rd Floor
Evanston, IL 60208
United States
gideonbornstein2018@u.northwestern.edu

Per Krusell
Institute for International Economic Studies
Stockholm University
106 91 STOCKHOLM
SWEDEN
and NBER
per.krusell@iies.su.se

Sergio Rebelo
Northwestern University
Kellogg School of Management
Department of Finance
Leverone Hall
Evanston, IL 60208-2001
and CEPR
and also NBER
s-rebelo@northwestern.edu

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. Because 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 of these perils.

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, price, and investment outcomes as a function of these shocks.

In order to build our model on solid microeconomic foundations, we rely heavily on a new proprietary data set compiled by Rystad Energy that contains information on production, reserves, operational costs, and investment for all oil fields in the world. The data set includes information about roughly 14,000 oil fields operated by 3,200 companies. We use these data to guide the construction of our model in two ways. First, 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. Second, 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 estimate our structural model under two alternative market structures. In the first market structure, all firms are competitive. In the second market structure, OPEC firms act as a cartel and non-OPEC firms are a competitive fringe. We assume that the cartel has the ability to commit and we solve the model using the "timeless perspective" approach introduced by Woodford (1999) and Woodford (2011) in the context of monetary policy. The dynamics of the estimated competitive and cartel models are surprisingly similar. However, these models have different implications for the long-run level of oil prices. Using the estimates from our competitive model and allowing OPEC firms to act as cartel would raise the average long-run real oil price by 21 percent.

The two estimated versions of our model share three key 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: a lag between investment and production, convex costs of adjusting extraction rates, and decreasing returns to oil investment.

We use both versions of the model to measure the importance of demand and supply shocks in driving prices, production, and investment. Using our competitive model, 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. We obtain similar results for the cartel model

[1] While 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.

with one exception: the volatility of non-OPEC production is driven only by demand shocks.

One interesting property of our models is that they are consistent with the high correlation between real oil prices and real investment in the oil industry. In the literature on the cattle and hog cycles (e.g., Ezekiel (1938) and Nerlove (1958)) this positive correlation is often interpreted as resulting from backward-looking expectations. 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 advantage of the high oil prices. So, over time, the resulting supply expansion brings the oil price back to its steady state level.

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 both the competitive and cartel models 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.

Our work is closely related to research on the relation between fluctuations in oil prices and world real Gross Domestic Product. Examples include Backus and Crucini (2000), Leduc and Sill (2004), Blanchard and Gali (2007), Kilian (2009), 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), Anderson et al. (2017) and Bjørnland et al. (2017)).

The paper is organized as follows. We describe our model and the two market structures that we consider in Section 2. In Section 3, we present our parameter estimates obtained using both micro data and moments of key aggregate variables for the oil industry. We also discuss the properties of the estimated models. In Section 4, we estimate the properties of the fracking technology using micro data. We then study the implications of introducing competitive fracking

[2] Earlier 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)).

firms into both versions of our model. Section 5 concludes.

2 An industry-equilibrium model

In this section, we describe our equilibrium model of the oil industry. We consider two alternative market structures: (i) all firms are competitive; and (ii) non-OPEC firms are competitive and OPEC firms operate as a cartel with commitment.

Our goal is to design a parsimonious model consistent with the data that 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. To simplify the exposition, we summarize the demand for oil by consumers and firms around the world with a log-linear demand function.[3] In this simple demand specification, the short- and long-run price elasticities coincide. In subsection 3.2, we considered 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. Throughout the model description, we provide brief motivations behind our assumptions on technology. In Section 3, where we describe our rich data set, we discuss these assumptions more explicitly.

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

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

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

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

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 competitive non-OPEC firms. These firms maximize their value (V^N) which is given by

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

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

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 $\psi\left(\theta_t^N\right)^\eta K_t^N$ 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

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

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

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

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.[4]

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).

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

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

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, $\psi\eta\left(\theta_{t+1}^N\right)^{\eta-1}$, 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

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

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 additional extraction costs equal to $\psi\left(\theta_{t+1}^N\right)^{\eta}$. 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

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

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

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

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

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

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

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

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

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

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

Our motivation for assuming that only OPEC firms are subject to supply shocks is as follows. 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.[5]

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

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

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

2.1 Market structures

We consider two alternative market structures. In the first structure, both OPEC and non-OPEC firms are competitive. In the second structure, non-OPEC firms are competitive and OPEC firms operate as a cartel with commitment. We first discuss the problem of a competitive OPEC firm. Then, we consider the cartel case.

2.1.1 OPEC firms are competitive

The problem of the representative competitive 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).

[5]To investigate the potential importance of supply shocks to non-OPEC firms, we estimated a version of the competitive 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 times 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.

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

$$E_t\left(P_{t+1}e^{-u_{t+1}}\right) = \psi\eta\left(\theta_{t+1}^O\right)^{\eta-1} + E_t\left(\mu_{t+1}^O e^{-u_{t+1}}\right), \tag{11}$$

$$\mu_t^O = E_t\beta\left\{\left[P_{t+1}e^{-u_{t+1}}\theta_{t+1}^O - \psi\left(\theta_{t+1}^O\right)^\eta\right] + (1-e^{-u_{t+1}}\theta_{t+1}^O)\mu_{t+1}^O\right\}, \tag{12}$$

$$\nu_t^O = \lambda\mu_t^O + \beta(1-\lambda)E_t\nu_{t+1}^O, \tag{13}$$

$$1 = \alpha\left(I_t^O\right)^{\alpha-1}\left(L^O\right)^{1-\alpha}\nu_t^O. \tag{14}$$

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, $e^{-u_{t+1}}$.

Equilibrium and model solution. 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 (θ_t^O) and non-OPEC (θ_t^N),

$$P_t = p(d_t, u_t, \mathbf{K}_t^O, \mathbf{K}_t^N, \theta_t^O, \theta_t^N). \tag{15}$$

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 (15) 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

$$Q_t^N = \theta_t^N\mathbf{K}_t^N, \qquad Q_t^O = e^{-u_t}\theta_t^O\mathbf{K}_t^O.$$

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.

We solve the model using a second-order approximation around its non-stochastic steady state.

2.1.2 OPEC is a cartel

We now consider the case where OPEC operates as a cartel with commitment 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. The version of our model in which OPEC behaves as a cartel is more challenging to solve for two reasons. First, 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.

The problem of the OPEC cartel is to choose the stochastic sequences for its own quantities $\{I_t^O, \theta_{t+1}^O, K_{t+1}^O, X_{t+1}^O\}$, the quantities and Lagrange multipliers of the non-OPEC representative firm $\{I_t^N, \theta_{t+1}^N, K_{t+1}^N, X_{t+1}^N, \mu_t^N, \nu_t^N\}$, and the oil price, P_t, to maximize its value, (8). The constraints for this problem are the laws of motion for OPEC's exploration capital, (9), and reserves, (10), the demand equation

$$P_{t+1} = \exp(d_{t+1})(e^{-u_{t+1}}\theta_{t+1}^O K_{t+1}^O + \theta_{t+1}^N K_{t+1}^N)^{-1/\varepsilon}, \tag{16}$$

and the optimality conditions for non-OPEC firms. These optimality conditions, which together with equation (16) are the implementability conditions of the cartel problem, include the laws of motion for non-OPEC's exploration capital, (2), and reserves, (3), and the first-order conditions (4)-(7).

We set up the OPEC problem as a Lagrangian and use the following notation for the Lagrange multipliers. We denote by μ_t^O and ν_t^O the multipliers on constraints (9) and (10), respectively. We denote by φ_t^i, $i = 1, .., 7$ the multipliers on the implementability constraints (16), (2), (3), (4), (5), (6), (7), respectively.

Computing the first-order condition with respect to θ_{t+1}^O and dividing both sides by K_{t+1}^O we obtain:

$$E_t\left(P_{t+1}e^{-u_{t+1}}\right) = \psi\eta\left(\theta_{t+1}^O\right)^{\eta-1} + E_t\left(\mu_{t+1}^O e^{-u_{t+1}}\right) + \frac{1}{\varepsilon}E_t\left[\varphi_{t+1}^1 e^{-u_{t+1}}\exp(-\varepsilon d_{t+1})P_{t+1}^{1+\varepsilon}\right]. \tag{17}$$

Comparing with the analogous condition for the case in which OPEC firms are competitive, equation (11), we see that equation (17) has an additional term. This term represents the negative impact of an increase in the cartel's production on price.

The first-order condition with respect to K_{t+1}^O is:

$$\mu_t^O = E_t\beta\left\{\left[P_{t+1}e^{-u_{t+1}}\theta_{t+1}^O - \psi\left(\theta_{t+1}^O\right)^{\eta}\right] + (1 - e^{-u_{t+1}}\theta_{t+1}^O)\mu_{t+1}^O - \frac{\theta_{t+1}^O}{\varepsilon}E_t\left[\varphi_{t+1}^1\exp(-\varepsilon d_{t+1})P_{t+1}^{1+\varepsilon}\right]\right\}.$$

This equation also has an additional term relative to its competitive analogue, equation (12). An increase in OPEC reserves increases production, $e^{-u_{t+1}}\theta^O_{t+1}K^O_{t+1}$, for a given extraction rate, θ^O_{t+1}. The additional term represents the decline in price associated with this production increase.

The first-order conditions for X^O_{t+1} and I^O_t are the same as in the competitive case, equations (13) and (14), respectively.

The first-order condition with respect to P_{t+1} is

$$\varphi^1_{t+1} = e^{-u_{t+1}}\theta^O_{t+1}K^O_{t+1} - \varphi^4_{t+1} - \varphi^5_{t+1}\theta^N_{t+1}.$$

We can think of this equation as defining φ^1_{t+1}, the marginal value to the cartel of a higher oil price at $t+1$. The first component, $e^{-u_{t+1}}\theta^O_{t+1}K^O_{t+1}$, is the benefit for selling OPEC's production for a higher price. The second and third components, $-\varphi^4_{t+1}$ and $-\varphi^5_{t+1}\theta^N_{t+1}$, are the costs of tightening implementability constraints (4) and (5). This tightening occurs because a higher oil price incentivizes non-OPEC to increase extraction and oil reserves.

The first-order condition for K^N_{t+1} is

$$\varphi^2_t = \beta\frac{E_t\left[\varphi^1_{t+1}\exp(-\varepsilon d_{t+1})P^{1+\varepsilon}_{t+1}\right]}{\varepsilon}\theta^N_{t+1} + \beta\varphi^2_{t+1}(1-\theta^N_{t+1}).$$

The multiplier φ^2_t is the shadow cost of non-OPEC reserves for the OPEC cartel. This cost has two components. First, for a given value of the non-OPEC extraction rate, a rise in reserves increases non-OPEC production, lowering the price. Second, non-OPEC has more reserves at time $t+2$.

The first-order condition for X^N_{t+1} is

$$\varphi^3_t = \lambda\varphi^2_t + \beta(1-\lambda)\varphi^3_{t+1}.$$

The cost to OPEC of higher exploration capital for non-OPEC has two components. The first is that a fraction λ of the exploration capital turns into non-OPEC reserves. The second, is that a fraction $1-\lambda$ is available as exploration capital at time $t+2$.

The first-order condition for θ^N_{t+1} is

$$E_t\left[\varphi^1_{t+1}\exp(-\varepsilon d_{t+1})P^{1+\varepsilon}_{t+1}\right]\frac{K^N_{t+1}}{\varepsilon} = \varphi^2_{t+1}K^N_{t+1} + \varphi^4_{t+1}(\eta-1)\psi\eta\left(\theta^N_{t+1}\right)^{\eta-2}. \tag{18}$$

Increasing non-OPEC's extraction rate, θ^N_{t+1}, is costly as it lowers the price of oil, P_{t+1}. These costs are represented by the left-hand side of equation (18). There are also two benefits from increasing θ^N_{t+1}. The first is the decline in non-OPEC reserves at time $t+2$. The second is the relaxation of implementability constraint (4), associated with an increase in non-OPEC marginal extraction costs.

The first-order condition for μ_t^N is

$$-\varphi_t^5 = -\varphi_t^6\lambda - (1-\theta_t^N)\varphi_{t-1}^5 + \varphi_{t-1}^4. \tag{19}$$

The term $-\varphi_t^5$ is the shadow cost of tightening implementability constraint (5) by raising μ_t^N, the shadow value of reserves for non-OPEC. This rise increases the value of X_{t+1}^N to non OPEC, ν_t^N, tightening implementability constraint (6). An increase in μ_t^N also raises μ_{t-1}^N thus tightening implementability constraint (5) at time $t-1$. Finally, an increase in μ_t^N relaxes implementability constraint (4) at time $t-1$, because a higher marginal value of reserves at time t reduces the incentives of non-OPEC to extract oil at time $t-1$.

The first-order condition for ν_t^N is

$$-\varphi_t^6 = -\varphi_{t-1}^6(1-\lambda) - \varphi_t^7\alpha\left(I_t^N\right)^{\alpha-1}L_N^{1-\alpha}. \tag{20}$$

The term $-\varphi_t^6$ is the shadow cost of tightening implementability constraint (6) by raising ν_t^N, the shadow value of exploration capital for non-OPEC. An increase in ν_t^N also increases ν_{t-1}^N thus tightening implementability constraint (6) at time $t-1$. Finally, an increase in ν_t^N tightens implementability constraint (7) because it raises the marginal return on non-OPEC investment.

The first-order condition for I_t^N is

$$(1-\alpha)\,\varphi_t^7\nu_t^N = \varphi_t^3 I_t^N.$$

The multiplier φ_t^7 is the shadow value of relaxing implementability constraint (7). Relaxing this constraint allows OPEC to enforce lower investment of non-OPEC, which translates into a lower level of non-OPEC exploration capital.

Time inconsistency. Before describing the equilibrium, it is worth discussing the time inconsistent nature of the cartel's decision problem. The cartel has an incentive to promise that future equilibrium prices will be low to reduce the return to non-OPEC investment in exploration capital. However, once non-OPEC has low exploration capital, OPEC would like to renege on its promise and take actions that increase oil prices.

This time inconsistency is reflected in first-order conditions (19) and (20). When choosing the value of non-OPEC reserves and exploration capital to non-OPEC, μ_t^N and ν_t^N, OPEC must be consistent with the implementability constraints at time $t-1$. If OPEC could renege at time t on its commitments, its choices need not be consistent with the implementability constraints at time $t-1$, so the Lagrange multipliers φ_{t-1}^4, φ_{t-1}^5, and φ_{t-1}^6 in equations (19) and (20) would be zero.

Equilibrium and model solution. 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 (θ_t^O) and non-OPEC (θ_t^N):

$$P_t = p(d_t, u_t, \mathbf{K}_t^O, \mathbf{K}_t^N, \theta_t^O, \theta_t^N). \tag{21}$$

Non-OPEC firms maximize their value subject to the laws of motion for reserves and exploration capital. Each non-OPEC 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 non-OPEC firm.

OPEC firms maximize their value subject to the laws of motion for reserves and exploration capital and the implementability conditions discussed above. The aggregate state of this economy is given by $\{d_t, u_t, \mathbf{K}_t^O, \mathbf{K}_t^N, \mathbf{X}_t^O, \mathbf{X}_t^N, \theta_t^O, \theta_t^N, \varphi_{t-1}^4, \varphi_{t-1}^5, \varphi_{t-1}^6\}$. There are three new state variables, $\varphi_{t-1}^4, \varphi_{t-1}^5, \varphi_{t-1}^6$ relative to the model where OPEC firms are competitive.

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

$$Q_t^N = \theta_t^N \mathbf{K}_t^N, \qquad Q_t^O = e^{-u_t}\theta_t^O \mathbf{K}_t^O.$$

As in the competitive case, 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.

We solve the model using the "timeless perspective" approach introduced in Woodford (1999) and Woodford (2011). Under this approach, we solve for a non-stochastic steady state where the Lagrange multipliers $\varphi_{t-1}^4, \varphi_{t-1}^5, \varphi_{t-1}^6$ are constant, and linearize the model around this non-stochastic steady state in order to obtain a recursive, time-invariant system.

At time zero, the values of $\varphi_{-1}^4, \varphi_{-1}^5, \varphi_{-1}^6$ all equal to zero, so OPEC's choices at time zero do not have to be consistent with the implementability constraints at time -1. From time zero on, OPEC makes decisions under commitment. The timeless-perspective steady state is the rest point of this economy absent shocks.

2.2 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

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

The term $\beta E_t\left(\eta - 1\right)\psi\theta_{t+2}^{\eta}$ 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 θ_t and P_t are constant.[6] 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.[7]

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:

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

The long-run constancy of P_t is consistent with the fact that the average annual growth rate in

[6]When $\lambda = 0$ 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.

[7]One way to try to make the Hotelling model consistent with a constant real oil price is to assume that the marginal cost of extraction falls over time. However, this marginal cost has to eventually fall below zero for the price to remain constant.

the real price of oil is not statistically different from zero.[8] For the period 1900-2015, this average is 0.01 with a standard error of 0.21. For the period 1970-2015, 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)).

3 Estimation and quantitative analysis

In this section, we estimate the structural parameters for the two versions of our model using both micro and aggregate data. We then study the models' quantitative properties.

Our micro estimates are based on new proprietary data compiled by Rystad Energy that contains information on production, reserves, operational costs, and investment for all oil fields in the world. The data contains information on about roughly 14,000 oil fields operated by 3,200 companies.

We use these data to guide the construction of our model in two ways. First, we produce micro estimates of two key model parameters: the average lag between investment and production ($1/\lambda$) and the elasticity of extraction costs with respect to production (η). Second, we use the generalized method of moments (GMM) to estimate the remaining model parameters, targeting a set of second moments for oil-related variables.

Our analysis focuses 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), Kilian (2014), and Fernandez-Villaverde (2017)).

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.

[8] Between 1900 and 1947, our source for the price of oil is Harvey et al. (2010). After 1947, our measure of oil prices is the price per barrel of West Texas Intermediate. We deflate the price of oil using the U.S. consumer price index.

This estimate implies that the average lag between investment and production is 12 years. Figure 1 shows the empirical distribution of this lag together with the implied geometric distribution for our estimate of $\hat{\lambda}$.[9]

Figure 1: Empirical distribution of lags between investment and production

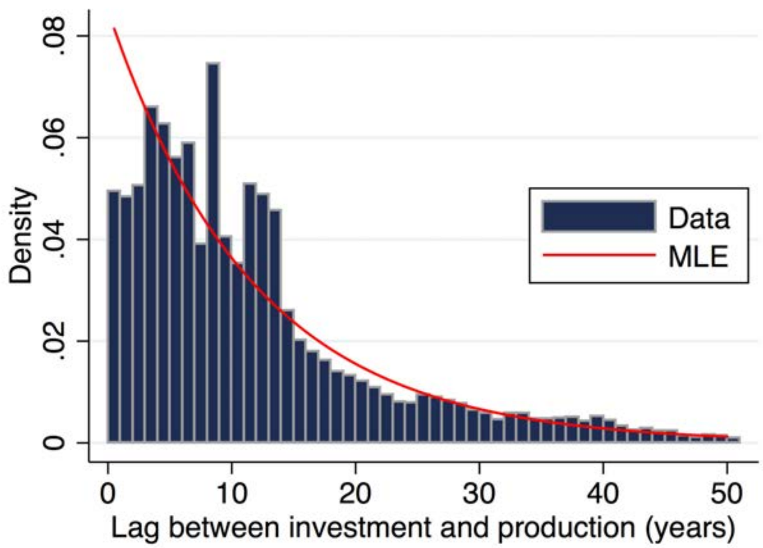

Notes: This figure presents the histogram of the lag between the first year of investment and first year of production across all oil fields. The MLE for λ, the Poisson arrival rate of production, is 0.085. The red line is the implied geometric distribution for the estimated λ.

We also use our micro data to estimate η, the parameter that controls the convexity of the extraction costs as displayed in (1). Our estimate is based on 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 our estimate of η. To address this problem, we instrument the extraction rate with the one-year-ahead forecast of detrended world real GDP. This forecast is correlated with aggregate demand and unaffected by field-specific cost shocks. Our forecast is computed by linearly detrending the time series for world real GDP and estimating an AR process for the

[9]Our 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.

Table 1: Extraction rate adjustment costs regression

Dep. variable: ln(prod. costs per barrel of oil reserves)

Variable	**(1)**	**(2)**	**(3)**
ln(extraction)	8.11***	7.69***	13.1
	(1.54)	(1.54)	(9.33)
Oil field FE	✓	✓	✓
Operation year FE	✓	✓	✓
Sample	All	Non-OPEC	OPEC
IV	✓	✓	✓
1^{st} stage F-stat	22	19	2
Clusters (oil fields)	11,527	9,969	1,558
Observations	174,339	146,879	27,460

Notes: This table presents the regression results for the adjustment-cost coefficient, η. Standard errors are clustered at the oil-field level. The instrument used is the one-year-ahead forecast of detrended world real GDP. *** - significant at the 1 percent level.

detrended data. Choosing the number of lags using the Akaike information criterion resulted in an AR(2) process.

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

Table 1 contains our slope estimates. Specifications 1 through 3 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. Perhaps as a result of this endogeneity problem, the point estimate is statistically insignificant. We use specification 2 as our benchmark and set η equal to the point estimate (7.7). 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 8).

Calibrated parameters. We choose the ratio L^O/L^N so that in the steady state the market share of OPEC production coincides with the average market share of OPEC in the data (45 percent). The total amount of land ($L^O + L^N$) is calibrated so that the steady-state extraction rate coincides with the average extraction rate in our data (2.8 percent). In the model where OPEC is competitive, the steady-state extraction rate is the same for OPEC and non-OPEC firms. In the version of the model where OPEC is a cartel, the steady-state extraction rate for OPEC firms is lower from that of non-OPEC firms. We calibrate α to match OPEC's average extraction rate (2.2 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.

GMM estimation In the model where OPEC firms are competitive, we estimate ϵ, α, the parameters of the AR(2) processes for demand and supply shocks, and the variance of the two shocks using GMM.[10] In the model where OPEC is a cartel, we estimate the same set of parameters with the exception of α, which is chosen to match OPEC's extraction rate.

The first column of Table 4 presents the moments targeted in our estimation. It is useful to highlight some salient facts about the oil market that are reflected in these moments. The first fact is 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. The high volatility of commodity prices in general was aptly summarized by Deaton (1999) with the phrase "What commodity prices lack in trend, they make up for in variance."

The second fact is that 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.[11]

The third fact is that investment in the oil industry is positively correlated with oil prices. The correlation between the growth rate of the price of oil and the growth rate of investment is 0.51. Table 2 reports the correlation between the growth rate of real investment and the growth rate of

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

[11]In the period 1970–2015, the only major U.S. manufacturing sector with investment volatility similar to that of the oil industry 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.

Table 2: 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: This table presents the correlation of investment changes and oil price changes for the top 20 oil producers, in descending order of production. 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.

the real oil prices for each of the top twenty firms in the oil industry ranked according to their total oil production in 2015. This table shows that, with a few exceptions, there is high correlation between oil prices and firm-level investment.

Finally, 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 be the result of supply shock to OPEC firms, such as the disruptions in oil markets associated with the Iranian revolution and the Iran-Iraq war.

Model estimates Table 3 reports our parameter estimates for both versions of the model. Our estimate for ϵ is 0.17 with a standard error of 0.02 for the competitive model, and 0.20 with a standard error of 0.05 for the cartel model.[12] These point estimates imply that demand is very inelastic. For $\epsilon = 0.17$, a 1 percent increase in production reduces the price by 5.9 percent.

To see why demand has to be inelastic to fit the data, it is useful to rewrite the demand function

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

Our models embed 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, d_t. The observed volatility and persistence of investment and production help determine the importance of the two channels that generate volatility in P_t and therefore identify ϵ. Volatile, persistent demand shocks are important for generating volatility and persistence in investment. A positive demand shock keeps the price of oil high for an extended period of time, creating incentives to increase investment.

Table 3 shows that the standard errors associated with our parameter estimates are generally small. The only parameter that is imprecisely estimated is α. In the competitive model this parameter has a point estimate of 0.44 with a standard error of 0.78. This imprecision results from the small impact of local changes in the value of α on the moments implied by the model. In the cartel case, α is calibrated to be 0.31 to match the average extraction rate of OPEC.

An important difference between the competitive and the cartel models is the estimated process for the supply shock. This shock is estimated to be persistent in the competitive model and to be i.i.d. in the cartel model.

Table 4 compares the estimated moments targeted by our GMM procedure with the population moments implied by the two versions of our model. We see that the fit of both models 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 both models and close to zero in the data.

As discussed in the introduction, both our models are consistent with the high correlation between the real price of oil and real investment.

Impulse response functions. Figure 2 depicts the impulse response function for a one standard deviation demand shock. The top and bottom row correspond to the competitive and cartel cases,

[12] These estimates are similar to the one obtained by Caldara et al. (2017).

Table 3: Estimated parameters

Parameter	**Competitive model**		**Cartel model**	
	Estimate	(s.e.)	Estimate	(s.e.)
ϵ	0.17	(0.02)	0.20	(0.05)
α	0.44	(0.78)	0.31	(-)
ρ_1^d	1.73	(0.09)	1.70	(0.06)
ρ_2^d	-0.75	(0.09)	-0.72	(0.06)
ρ_1^u	1.39	(0.16)	-0.06	(2.37)
ρ_2^u	-0.48	(0.13)	0.00	(2.14)
$var(e_t^d)$	0.014	(0.005)	0.023	(0.006)
$var(e_t^u)$	0.002	(0.001)	0.002	(0.001)

Notes: This table presents the benchmark GMM estimates of the structural parameters. Under the cartel specification, α is not estimated, but calibrated so that the model matches OPEC's average extraction rate.

respectively. The dynamics are similar in the two models. In both models the shock follows a hump-shaped pattern with a peak around year six. 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 production to take advantage of the extended period of high oil prices.

Figure 3 depicts the impulse response function for a one standard deviation supply shock. In the competitive model, the supply 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 production. But since the shock is short lived, the rise in investment is much smaller than the one 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 higher extraction costs than non-OPEC firms.

In the cartel model, the supply shock is i.i.d. Because firms cannot change their extraction rate within the period, the only impact on the quantity of oil produced is the direct, negative impact of the supply shock on OPEC production. This decline in production produces a temporary rise in the price of oil.

Table 4: Data and model moments

	Moment	Data	(s.e.)	Competitive model	Cartel model
(1)	std(Δp_t)	0.273	(0.028)	0.188	0.220
(2)	std(Δi_t^N)	0.192	(0.024)	0.218	0.212
(3)	std(Δi_t^O)	0.193	(0.027)	0.215	0.208
(4)	std(Δq_t^N)	0.022	(0.003)	0.026	0.024
(5)	std(Δq_t^O)	0.069	(0.011)	0.057	0.068
(6)	corr(Δp_t, Δi_t^N)	0.557	(0.147)	0.738	0.687
(7)	corr(Δp_t, Δi_t^O)	0.362	(0.109)	0.659	0.671
(8)	corr(Δp_t, Δq_t^N)	0.031	(0.069)	0.013	0.127
(9)	corr(Δp_t, Δq_t^O)	0.030	(0.122)	-0.585	-0.583
(10)	corr(Δi_t^N, Δi_t^O)	0.673	(0.096)	0.992	0.999
(11)	corr(Δi_t^N, Δq_t^N)	0.087	(0.094)	-0.041	-0.025
(12)	corr(Δi_t^N, Δq_t^O)	0.023	(0.112)	-0.150	-0.008
(13)	corr(Δi_t^O, Δq_t^N)	-0.034	(0.145)	-0.035	-0.027
(14)	corr(Δi_t^O, Δq_t^O)	-0.226	(0.153)	-0.044	0.013
(15)	corr(Δq_t^N, Δq_t^O)	-0.141	(0.125)	0.026	0.341
(16)	corr(Δp_t,Δp_{t-1})	-0.027	(0.088)	-0.043	-0.166
(17)	corr(Δi_t^N,Δi_{t-1}^N)	0.119	(0.135)	0.007	0.001
(18)	corr(Δi_t^O,Δi_{t-1}^O)	0.311	(0.096)	0.009	-0.001
(19)	corr(Δq_t^N,Δq_{t-1}^N)	0.643	(0.113)	0.291	0.416
(20)	corr(Δq_t^O,Δq_{t-1}^O)	0.213	(0.211)	0.314	-0.392
(21)	$\mathbb{E}\left[I_t/\left(P_t Q_t\right)\right]$	0.102	(0.016)	0.057	0.027

Notes: This table presents the targeted moments from the data and the model-implied moments under the benchmark specifications of the competitive and cartel models. Newey-West standard errors computed with 5-year lags in parenthesis. x_t represents the logarithm of X_t, Δx_t is equal to $x_t - x_{t-1}$.

Figure 2: Impulse response to a demand shock

Panel A - competitive model

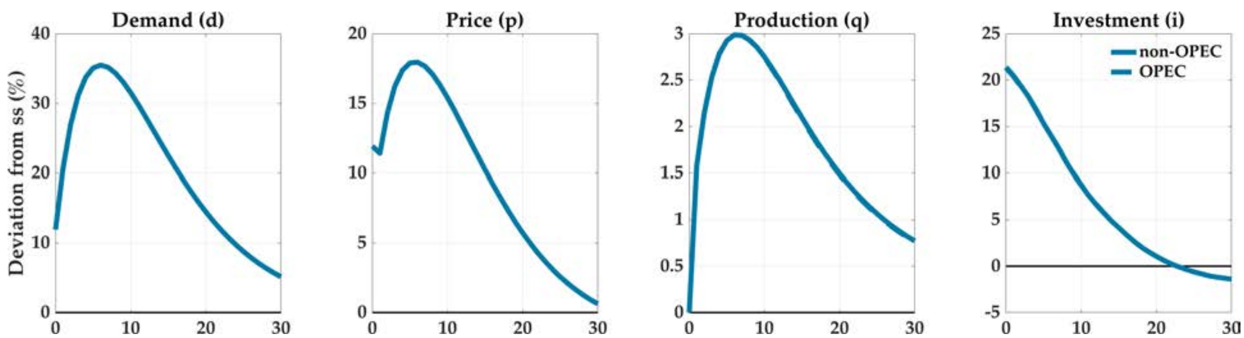

Panel B - cartel model

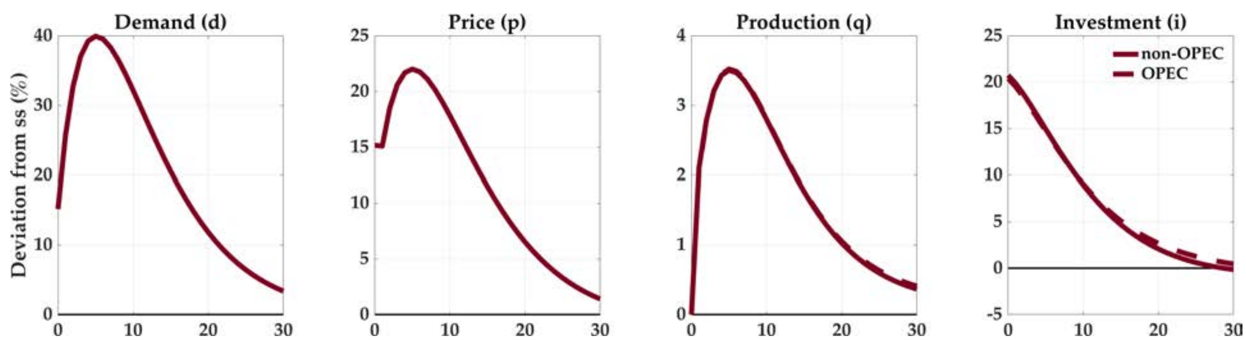

Notes: This figure presents the impulse response functions to a one standard deviation demand shock. The top panel corresponds to the benchmark specification of the competitive model. The bottom corresponds to the benchmark specification of the cartel model. In the two right columns, the solid line refers to non-OPEC quantities and the dashed line to OPEC quantities.

Figure 3: Impulse response to a supply shock

Panel A - competitive model

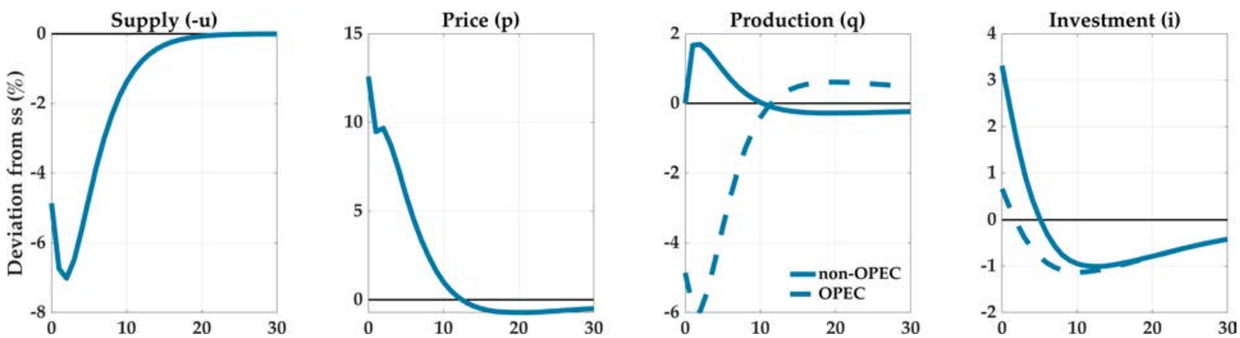

Panel B - cartel model

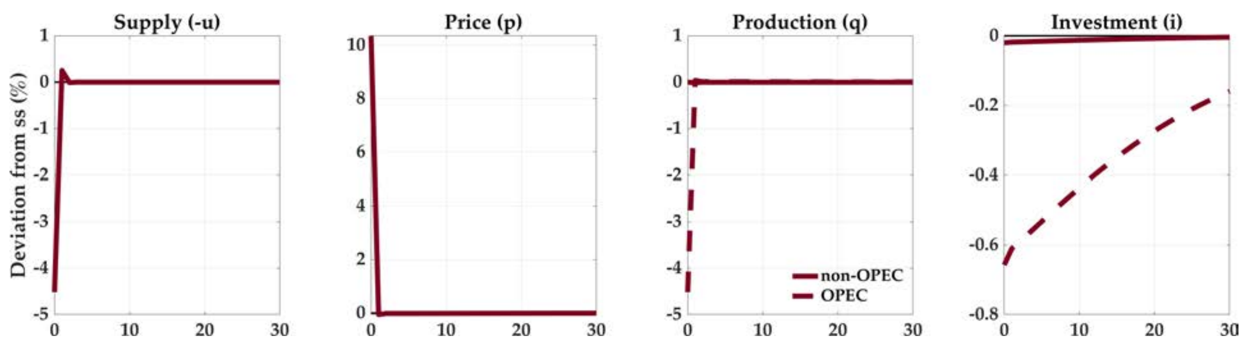

Notes: This figure presents the impulse response functions to a one standard deviation supply shock which reduces OPEC's oil production. The top panel corresponds to the benchmark specification of the competitive model. The bottom corresponds to the benchmark specification of the cartel model. In the two right columns, the solid line refers to non-OPEC quantities and the dashed line to OPEC quantities.

Table 5: Moments sensitivity to demand and supply shocks

Moment	**Data**	(s.e.)	**Competitive**			**Cartel**		
			Bench.	$std(d_t) = 0$	$std(u_t) = 0$	Bench.	$std(d_t) = 0$	$std(u_t) = 0$
$std(\Delta p_t)$	0.273	(0.028)	0.188	0.136	0.131	0.220	0.147	0.164
$std(\Delta i_t^n)$	0.192	(0.024)	0.218	0.036	0.214	0.212	0.000	0.212
$std(\Delta i_t^o)$	0.193	(0.027)	0.215	0.011	0.214	0.208	0.007	0.208
$std(\Delta q_t^n)$	0.022	(0.003)	0.026	0.018	0.018	0.024	0.000	0.024
$std(\Delta q_t^o)$	0.069	(0.011)	0.057	0.054	0.018	0.068	0.064	0.024
$corr(\Delta q_t^n, \Delta q_t^o)$	-0.141	(0.125)	0.026	-0.299	1.00	0.341	-0.551	1.000

Notes: This table presents key targeted data moments and the corresponding model-implied moments when we shut down the supply or demand shocks. Newey-West standard errors computed with 5-year lags are reported in parenthesis. x_t represents the logarithm of X_t, Δx_t is equal to $x_t - x_{t-1}$.

We can use our model to answer a classic question: what is the role of demand and supply shocks? Table 5 uses the two versions of our model to answer this question. Eliminating demand shocks lowers the volatility of prices but reduces the volatility of investment even more. Eliminating supply shocks reduces the volatility of prices, leads to a large decline in the volatility of OPEC production, and makes OPEC and non-OPEC production perfectly correlated.

Table 6 reports the variance decomposition of the key variables in our model to demand and supply shocks. We see that in both models demand and supply shocks contribute roughly equally to the variance of prices. This property results from demand and supply shocks having a similar short-run impact on the price of oil (see Figures 2 and 3) in both models. 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 6 also shows that in both models 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 2). In contrast, supply shocks have a much lower impact on investment because these shocks are less persistent than demand shocks in the competitive model and not persistent at all in the cartel model (see Figure 3).

The volatility of production by OPEC firms is dominated in both models by supply shocks. In the competitive model, demand and supply shocks contribute equally to the volatility of production by non-OPEC firms. In contrast, demand shocks account for all the volatility of non-OPEC production in the cartel model. This result reflects the i.i.d. nature of supply shocks in the cartel

Table 6: Variance decomposition

Moment	**Competitive**		**Cartel**	
	Demand	Supply	Demand	Supply
Δp_t	48.0%	52.0%	55.4%	44.6%
Δi_t^n	97.2%	2.8%	100%	0%
Δi_t^o	99.8%	0.2%	99.9%	0.1%
Δq_t^n	49.5%	50.5%	100%	0%
Δq_t^o	10.3%	89.7%	11.8%	88.2%

Notes: This table presents the variance decomposition of five key variables to demand and supply shocks in the two versions of the model. x_t represents the logarithm of X_t, Δx_t is equal to $x_t - x_{t-1}$.

model. Since extraction rates are chosen one period in advance, non-OPEC cannot respond to the supply shock.

Impact of market structure on steady-state oil price. While the dynamics of the competitive and cartel models are broadly similar, the market structure has a large impact on the steady state values of key variables. To study these effects, we take as given the structural parameters of the estimated competitive model and assume that OPEC acts as a cartel. We then solve for the steady state of this cartel model.

Table 7 compares the values of key variables in the competitive and cartel models. Because OPEC internalizes the effect of a production increase on the price, it cuts steady-state production relative to the competitive model by 29 percent. Non-OPEC firms increase their production by 18 percent but this increase is not large enough to compensate for the cutback in OPEC production. As a result, the price rises by 21 percent. OPEC's market share falls from 45 percent to 33 percent. In the competitive model the two groups of firms have the same extraction rate. In contrast, the extraction rate is lower in OPEC than in non-OPEC in the cartel model.

3.1 Other performance diagnoses

In this section, we evaluate whether our model is consistent with two sets of moments that were not targeted by our estimation. The first is our estimates of the short-run elasticity of supply obtained from micro data. The second is the volatility of one-year oil-price futures.

Table 7: Effect of market structure on the steady-state variables

Price change	21%
OPEC's production change	-29%
Non-OPEC's production change	18%
Competitive: OPEC's market share	45%
Cartel: OPEC's market share	33%
Competitive: extraction rate	2.79%
Cartel: OPEC's ext. rate	2.64%
Cartel: Non-OPEC's ext. rate	2.87%

Notes: This table compares the steady-state properties of a model calibrated with the benchmark parameters of the competitive model and solved under two alternative market structures: (i) all firms are competitive, and (ii) OPEC is a cartel and non-OPEC is a competitive fringe.

Estimating the short-run elasticity of oil supply. 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.[13] 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 statistically insignificant.

Table 8 reports panel-data estimates of the elasticity of the extraction rate for a given oil field with respect to real oil prices.[14] These estimates suggest that a rise in oil prices leads to only a slight increase in the supply of oil from a given oil field.[15]

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

$$\ln \theta_{it} = \alpha_i + \beta \ln p_t + \gamma X_{it} + \varepsilon_{it}, \tag{22}$$

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}

[13] Anderson 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.

[14] In 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.

[15] An 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.

Table 8: Price elasticity of extraction rates

Dep. variable: ln(extraction rate)

Variable	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**
ln(price)	0.09^{***}	0.18^{***}	0.20^{***}	0.22^{***}	0.14^{***}
	(0.009)	(0.030)	(0.036)	(0.042)	(0.029)
$\ln(\text{price}) \times \mathbb{1}_{\text{OPEC}}$			-0.14	-0.13	
			(0.087)	(0.088)	
$\ln(\text{price}) \times \mathbb{1}_{\text{Big Firm}}$				-0.04	
				(0.081)	
$\ln(\text{price}) \times \mathbb{1}_{\lvert\Delta \ln(p)\rvert>0.1}$					-0.013^{***}
					(0.002)
Oil field FE	✓	✓	✓	✓	✓
Operation year FE	✓	✓	✓	✓	✓
Year trend	✓	✓	✓	✓	✓
IV	✗	✓	✓	✓	✓
Clusters (oil fields)	12,187	11,479	11,479	11,479	11,479
Observations	173,742	173,034	173,034	173,034	173,034

Notes: This table presents estimates of the elasticity of extraction rates with respect to the price of oil. The data used includes all oil fields with positive extraction rates in 1971-2015, excluding the last year of operation. Standard errors, clustered at the oil-field level, are reported in parenthesis. The instrument for price is the one-year-ahead forecast of detrended world real GDP. $(***)$ - significant at a 1 percent level.

represents other controls.[16] 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.[17]

Specification 1 in Table 8 is a simple OLS regression. The estimated slope coefficient that results from this regression can be biased downwards 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 detrended world real GDP.

Specifications 2-5 use this instrument. Our benchmark specification is regression 2, which yields an estimate for β equal to 0.18.[18] The following calculation is useful for evaluating the mag-

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

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

[18]This 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 9: Price elasticity of extraction rates in both models

Dep. variable: ln(extraction rate)

Variable	**Data**	(s.e.)	**Competitive**	**Cartel**
ln(price)	0.18	(0.03)	0.16	0.16

Notes: This table presents estimates of the elasticity of extraction rates with respect to the price of oil using our data and simulated data from the two versions of our model. The instrument for price used in the data is the one-year-ahead forecast of detrended world real GDP. The instrument for price used in the model is the one-year-ahead forecast of demand, $\mathbb{E}_{t-1}[d_t]$.

nitude 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 5 percent production increase.

Specification 3 includes the product of the logarithm of the price and an OPEC dummy. We see that the response of OPEC and non-OPEC are not statistically different.

Specification 4 includes also an interaction term that is the product of the logarithm of the price and a dummy for firm size to investigate whether large and small firms behave differently. The 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. 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 the logarithm of the price and a dummy for price changes 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 (-0.01) 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 16 in the appendix).

We now discuss the implications of our two models for the analogue of the regressions reported in Table 8. We simulate data from both versions of our model and run regression 8 using the one-year-ahead forecast of the demand shock as an instrument for the price of oil. In both models, the elasticity of response of extraction rates to changes in oil prices is 0.16, which is within one standard error of the point estimate for the elasticity obtained using our benchmark specification.

Table 10: Price elasticity of oil fields in operation

Dep. variable: ln(number of operating oil fields)

Variable	**(1)**	**(2)**	**(3)**	**(4)**
ln(price)	−0.05	−0.19	−0.09	−0.10
	(0.06)	(0.13)	(0.08)	(0.07)
Year trend	✓	✓	✓	✓
IV	✗	✓	✓	✓
Dep. variable	All fields	All fields	Non-OPEC fields	OPEC fields
Observations	45	45	45	45

Notes: This table presents estimates of the elasticity of the number of oil fields in operation with respect to the price of oil. The dependent variable is the number of oil fields with positive extraction rates. Newey-West standard errors computed with 5-year lags are reported in parenthesis. The instrument for price is the one-year-ahead forecast of detrended world real GDP.

We now discuss the second channel through which production can increase, which is a rise in the number of oil fields in operation. Table 10 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 the number 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 insignificant. We also find an insignificant elasticity when we extend our sample to start in 1900 (see Table 15 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.

Oil price forecasts and oil futures. Both in the data and in our models 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 oil-price futures 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. [19] The volatility of one-year futures prices in the data is equal to 0.2, 26 percent lower than the volatility of spot oil prices in the data. In both version of our model, the volatility of one-year futures prices is equal to 0.16. In the competitive version of the model, the volatility of futures prices is 14 percent lower than the volatility of spot oil prices. In the cartel model, as in the data, futures volatility is 26 percent lower than spot prices volatility.

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

In this subsection, we present a variant of our models in which the short-run and long-run price elasticity of demand differ. This demand specification 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 substitution to occur.

Our dynamic demand specification takes the form:

$$P_t = \exp(d_t)\,(1-\phi)^{1/\xi}\,(Q_t - \phi Q_{t-1})^{-1/\xi} \;, \tag{23}$$

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 $\partial \ln Q_{ss}/\partial \ln P_{ss} = \xi$, where Q_{ss} and P_{ss} denote the steady-state quantity and price, respectively. 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 our models via GMM, including both ε and ϕ as parameters to be estimated. Table 11 reports our results. The estimated short-run demand elasticity, ε, is very similar to the

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

Table 11: Estimated parameters with dynamic demand elasticity

Parameter	**Competitive**		**Cartel**	
	Estimate	(s.e.)	Estimate	(s.e.)
ε	0.17	(0.03)	0.19	(0.01)
α	0.53	(0.72)	0.21	(-)
ρ_1^d	1.55	(0.14)	1.55	(0.003)
ρ_2^d	-0.60	(0.13)	-0.59	(0.004)
ρ_1^u	1.60	(0.17)	1.99	(0.008)
ρ_2^u	-0.64	(0.16)	-0.99	(0.01)
$var(e_t^d)$	0.018	(0.006)	0.037	(0.003)
$var(e_t^u)$	0.003	(0.001)	3×10^{-6}	(9×10^{-5})
ϕ	0.86	(0.15)	0.37	(0.01)

Notes: This table presents the GMM estimates of the structural parameters for the two versions of our model with dynamic demand. ϕ is the demand inertia parameter. Under the cartel specification, α is not estimated, but calibrated so that the steady state of the model matches the average extraction rate of OPEC.

estimated demand elasticity in our benchmark model. The estimated value of ϕ is 0.86 in the competitive model, indicating that the demand for oil features a high degree of inertia. The implied long-run demand elasticity, ξ, equals 1.21, approximately 7 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. Demand is less inertial in the cartel model, the estimated value of ϕ is 0.37.

Table 12 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 competitive version of the model. Both the competitive and the cartel model come closer to reproducing the price volatility observed in the data. In addition, the competitive 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.

3.3 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

Table 12: Data and model moments with dynamic demand elasticity

	Moment	Data	(s.e.)	Competitive model		Cartel model	
				Benchm.	Dyn. demand	Benchm.	Dyn. demand
(1)	std(Δp_t)	0.273	(0.028)	0.188	0.219	0.220	0.233
(2)	std(Δi_t^N)	0.192	(0.024)	0.218	0.202	0.212	0.174
(3)	std(Δi_t^O)	0.193	(0.027)	0.215	0.213	0.208	0.217
(4)	std(Δq_t^N)	0.022	(0.003)	0.026	0.025	0.024	0.027
(5)	std(Δq_t^O)	0.069	(0.011)	0.057	0.070	0.068	0.061
(6)	corr(Δp_t, Δi_t^N)	0.557	(0.147)	0.738	0.644	0.687	0.820
(7)	corr(Δp_t, Δi_t^O)	0.362	(0.109)	0.659	0.357	0.671	0.515
(8)	corr(Δp_t, Δq_t^N)	0.031	(0.069)	0.013	0.024	0.127	0.232
(9)	corr(Δp_t, Δq_t^O)	0.030	(0.122)	-0.585	-0.274	-0.583	-0.366
(10)	corr(Δi_t^N, Δi_t^O)	0.673	(0.096)	0.992	0.917	0.999	0.641
(11)	corr(Δi_t^N, Δq_t^N)	0.087	(0.094)	-0.041	-0.076	-0.025	-0.036
(12)	corr(Δi_t^N, Δq_t^O)	0.023	(0.112)	-0.150	-0.053	-0.008	0.001
(13)	corr(Δi_t^O, Δq_t^N)	-0.034	(0.145)	-0.035	-0.074	-0.027	-0.062
(14)	corr(Δi_t^O, Δq_t^O)	-0.226	(0.153)	-0.044	0.232	0.013	-0.055
(15)	corr(Δq_t^N, Δq_t^O)	-0.141	(0.125)	0.026	0.054	0.341	0.375
(16)	corr(Δp_t,Δp_{t-1})	-0.027	(0.088)	-0.043	-0.194	-0.166	-0.291
(17)	corr(Δi_t^N,Δi_{t-1}^N)	0.119	(0.135)	0.007	-0.009	0.001	-0.008
(18)	corr(Δi_t^O,Δi_{t-1}^O)	0.311	(0.096)	0.009	-0.004	-0.001	0.014
(19)	corr(Δq_t^N,Δq_{t-1}^N)	0.643	(0.113)	0.291	0.544	0.416	0.638
(20)	corr(Δq_t^O,Δq_{t-1}^O)	0.213	(0.211)	0.314	0.622	-0.392	-0.198
(21)	$\mathbb{E}\left[I_t/\left(P_t Q_t\right)\right]$	0.102	(0.016)	0.057	0.069	0.027	0.019

Notes: This table presents the targeted moments from the data and the model-implied moments under both the benchmark and dynamic-demand specifications. Newey-West standard errors computed with 5-year lags are reported in parenthesis. x_t represents the logarithm of X_t, Δx_t is equal to $x_t - x_{t-1}$.

robustness results are included in the appendix.

In our benchmark results, we instrumented the price of oil with the forecast of detrended real 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 17 and 22 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 (0.18).

Tables 18 and 23 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 statistically insignificant.

Tables 19 and 24 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 7.7 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 20, 21, 25, and 26. 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, delivering lower price volatility than the benchmark versions. 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 in the data for these two countries. Table 27 shows the estimates of the elasticity of the extraction rate with respect to prices obtained using this restricted sample. This estimate (0.22) is similar to our benchmark estimate (0.18). Table 28 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 statistically insignificant. 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 30 and 31).

4 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.[20] 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 fracking operations are very flexible, comes from data compiled by Baker Hughes on the number of oil rigs in operation in the U.S. These data are depicted, together the nominal oil price in Figure 4. 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 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 5 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 13 reports our estimates of η for fields explored with fracking obtained using the forecast of real price of oil.[21] We see that for non-conventional oil fields our estimate of η is roughly 1.95.

[20] See 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.

[21] We use the real price of oil as an instrument because fracking firms can respond within the period to supply and demand shocks.

Figure 4: U.S. oil rigs in operation and the price of oil

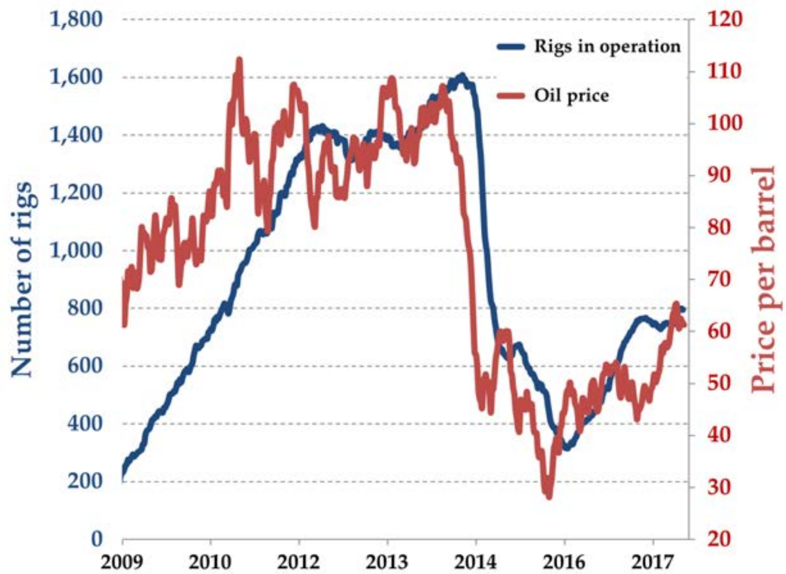

Notes: This figure presents the number of oil rigs in operation (blue line, left axis) and the nominal USD price of a barrel of oil (red line, right axis). Data source: Baker Hughes.

Figure 5: Empirical dist. of lags between investment to production - fracking fields

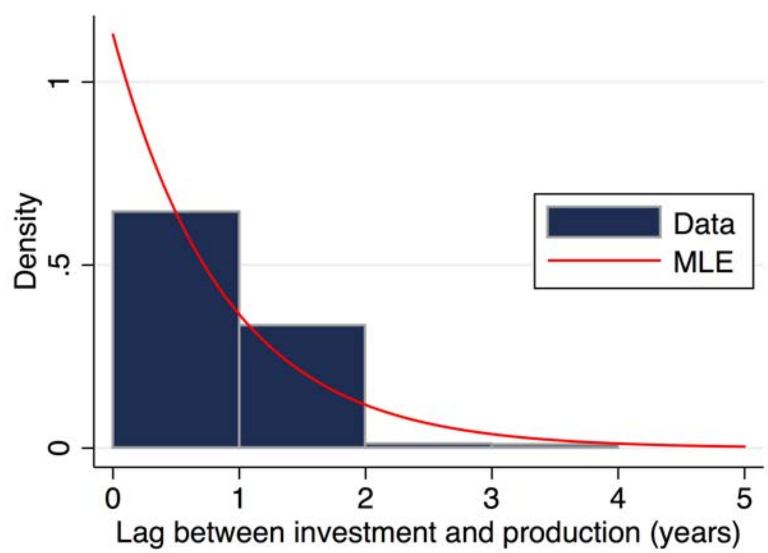

Notes: This figure presents the histogram of the lag between the first year of investment and first year of production across fracking oil fields. The MLE for λ, the Poisson arrival rate of production, is 1.13. The red line is the implied geometric distribution for the estimated λ.

Table 13: Extraction rate adjustment costs regression - fracking fields

Dep. variable: ln(prod. costs per barrel of oil reserves)

Variable	**(1)**	**(2)**
$\ln(\text{extraction})$	1.95^{***}	7.69^{***}
	(0.21)	(1.54)
Oil field FE	✓	✓
Year of operation FE	✓	✓
Sample	Fracking fields	All non-OPEC
IV	✓	✓
1^{st} stage F-stat	35	19
Clusters (oil fields)	952	9,969
Observations	4,940	146,879

Notes: This table presents estimates of the adjustment-cost coefficient, η. Specification (1) includes only fracking fields, and the instrument used is the real price of oil. Specification (2) includes all non-OPEC oil fields, and the instrument used is the one-year-ahead forecast of detrended world real GDP. Standard errors are clustered at the oil-field level. *** - significant at the 1 percent level.

In contrast, when we include all oil fields in our sample in the regression, we obtain an estimate of $\eta = 7.7$.

To study the impact of fracking, we include in both our models a third type of competitive firm that produces oil using fracking. These firms have extraction costs that are less convex ($\eta^F = 1.95 < 7.7$), 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 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):

$$\max_{\{I_t^F,\theta_{t+1}^F,K_{t+1}^F,X_{t+1}^F\}} V^F = E_0 \sum_{t=0}^{\infty} \beta^t \left[P_t \theta_t^F K_t^F - I_t^F - \psi^F \left(\theta_t^F\right)^{\eta^F} K_t^F \right], \tag{24}$$

subject to

$$X_{t+1}^F = (1-\lambda^F) X_t^F + \left(I_t^F\right)^{\alpha} \left(L^F\right)^{1-\alpha} \tag{25}$$

$$K_{t+1}^F = (1-\theta_t^F) K_t^F + \lambda^F X_{t+1}^F. \tag{26}$$

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:

$$Q_t = \theta_t^N \mathbf{K}_t^N + e^{-u_t}\theta_t^O \mathbf{K}_t^O + \theta_t^F \mathbf{K}_t^F,$$

where $\mathbf{K}_t^F$ and θ_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 the amount of land available for fracking so that in the steady state fracking represents 20 percent of global oil production. We calibrate ψ^F and ψ to be consistent with estimates of the difference between ψ^F and ψ obtained using our micro data given our point estimates of η^F and η.

Figures 6 and 7 show the impulse responses for demand and supply shocks, respectively. The first panel corresponds to the competitive model and the bottom panel to the cartel model. The panels for the price response also include the price response in our benchmark model without fracking firms. We see that in both the competitive and the cartel model, 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.

Table 14 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. In the competitive model, price volatility falls by 53 percent from 0.19 to 0.09. In the cartel model, price volatility falls by 50 percent from 0.22 to 0.11. There is a higher correlation between prices and quantities, and between investment and quantities in the models with fracking, reflecting the response of fracking firms to high-frequency movements in prices. Production is less volatile in the version of the models with fracking. This results reflects two opposing effects. The aggregate production response to demand shocks is higher than in the model without fracking, as fracking firms are more nimble. The aggregate production response to supply shocks is lower than in the model without fracking, because fracking firms can respond within the period to supply shocks, which smooths out the response of aggregate supply to these shocks.

Figure 6: Impulse response to a demand shock - fracking

Panel A - competitive model

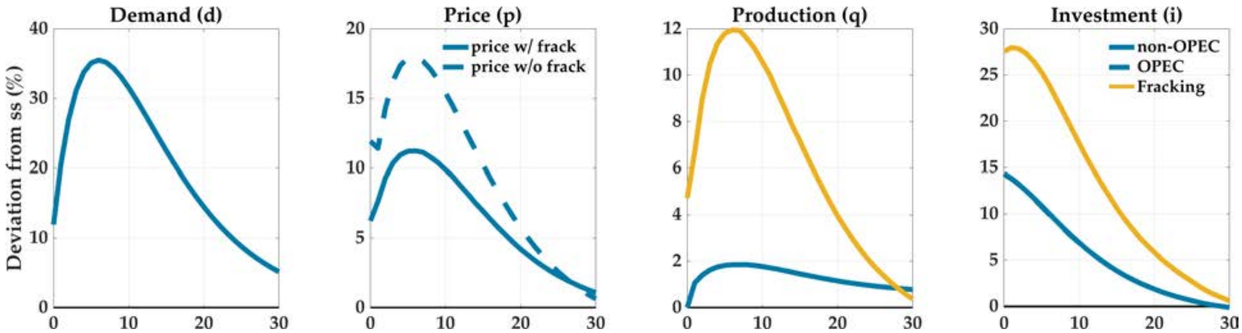

Panel B - cartel model

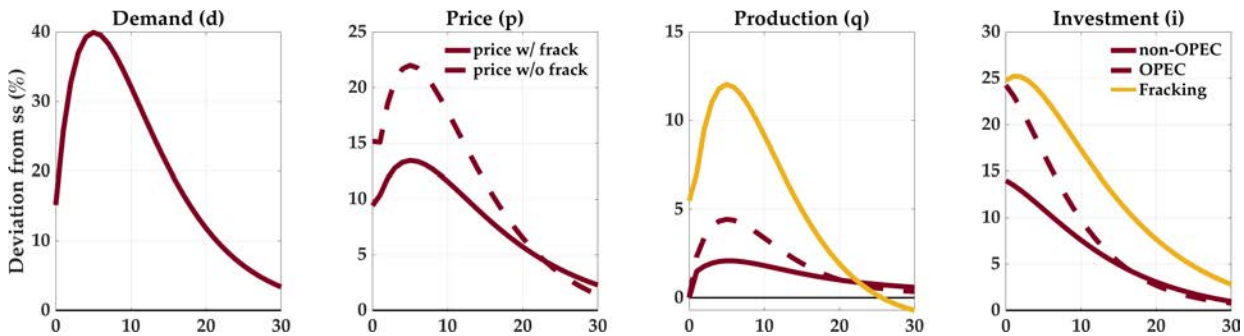

Notes: This figure presents the impulse response functions to a one standard deviation demand shock in models with competitive fracking firms. The top panel corresponds to the competitive model. The bottom panel corresponds to the cartel model. In the two right columns, the dashed line refers to OPEC quantities and the solid yellow line to fracking quantities. The dashed line on the second column refers to the price response under the benchmark specifications, which does not include fracking firms.

Table 14: Implication of fracking on key aggregate moments

	Moment	**Competitive**		**Cartel**	
		Benchm.	Fracking	Benchm.	Fracking
(1)	$std(\Delta p_t)$	0.19	0.09	0.22	0.11
(2)	$std(\Delta i_t)$	0.22	0.18	0.21	0.21
(3)	$std(\Delta q_t)$	0.03	0.02	0.04	0.03
(4)	$corr(\Delta p_t, \Delta i_t)$	0.70	0.79	0.69	0.85
(5)	$corr(\Delta p_t, \Delta q_t)$	-0.50	0.42	-0.44	0.34
(6)	$corr(\Delta q_t, \Delta i_t)$	-0.11	0.40	-0.02	0.41

Notes: This table presents second moments for our benchmark specifications and for versions of our models with fracking firms. x_t represents the logarithm of X_t, Δx_t is equal to $x_t - x_{t-1}$. I_t and Q_t represent aggregate investment and production at time t, respectively.

Figure 7: Impulse response to a supply shock - fracking

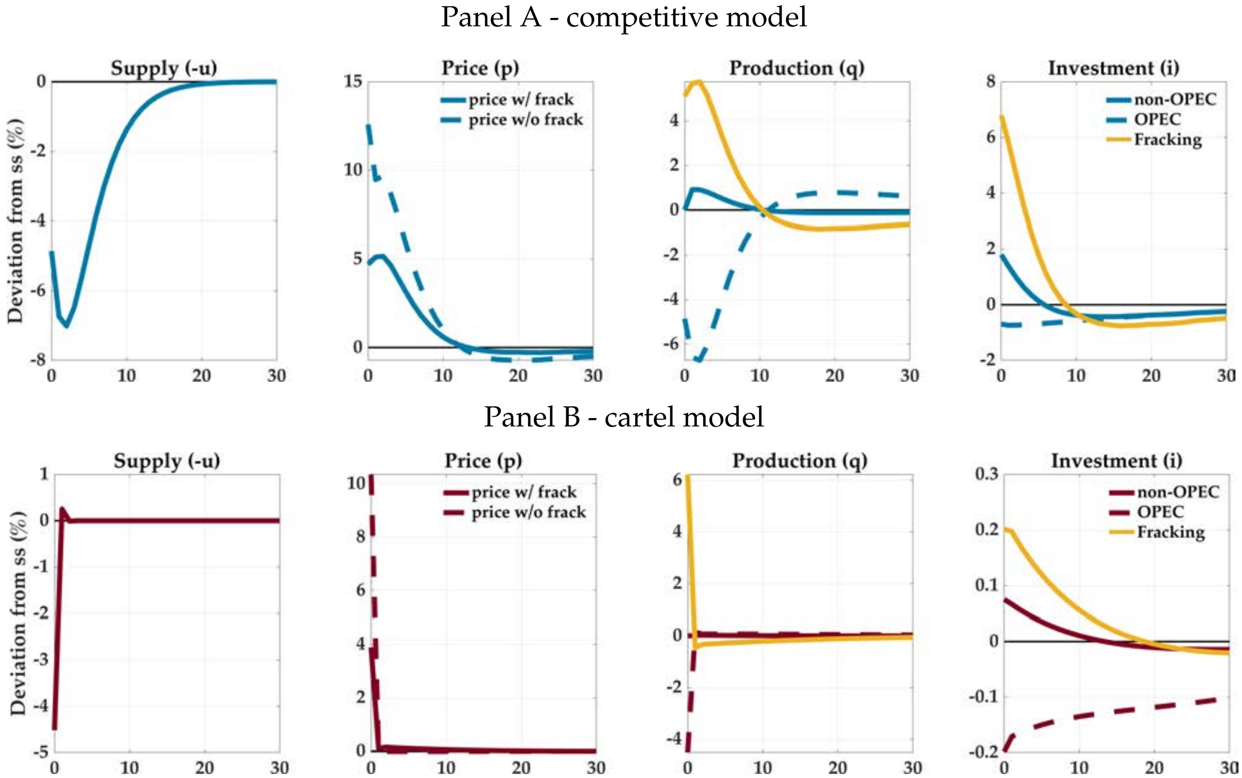

Notes: This figure presents the impulse response functions to a one standard deviation supply shock which reduces OPEC's oil production in models with competitive fracking firms. The top panel corresponds to the competitive model. The bottom panel corresponds to the cartel model. In the two right columns, the dashed line refers to OPEC quantities and the solid yellow line to fracking quantities. The dashed line on the second column refers to the price response under the benchmark specifications, which does not include fracking firms.

5 Conclusion

In this paper, we propose a parsimonious model consistent with micro data about the oil industry which can ultimately serve as a building block in a model of the world economy.

We leave three interesting projects for future research. The first is to develop a richer model of firm heterogeneity. In this paper we consider only two heterogeneity dimensions, OPEC versus non-OPEC, and conventional versus fracking producers. There are other heterogeneity dimensions which result in different choices of investment and extraction rates. The second 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.[22] The third is to combine our model of the oil market with a fully-fledged model of the 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., OZDAGLAR, A. and TAHBAZ-SALEHI, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, **80** 1977–2016.

ADAO, B., NARAJABAD, B. and TEMZELIDES, T. P. (2017). Renewable technology adoption and the macroeconomy.

ALQUIST, R., 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., KELLOGG, R. and SALANT, S. W. (2017). Hotelling under pressure. *Journal of Political Economy*.

AREZKI, R., 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.

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

BJØRNLAND, H. C., 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., ERCEG, C. J. and GUERRIERI, L. (2011). Oil shocks and external adjustment. *Journal of International Economics*, **83** 168–184.

CALDARA, D., 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.

FERNANDEZ-VILLAVERDE, J. (2017). Energy: The mover of output,in global economic history. Tech. rep., University of Pennsylvania.

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., 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., MADSEN, J. B. and WOHAR, M. E. (2010). The prebisch-singer hypothesis: four centuries of evidence. *The Review of Economics and Statistics*, **92** 367–377.

HASSLER, J., 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., 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., 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.

WOODFORD, M. (1999). Commentary: How should monetary policy be conducted in an era of price stability? *New challenges for monetary policy*, **277316**.

WOODFORD, M. (2011). *Interest and prices: Foundations of a theory of monetary policy*. princeton university press.

Appendix

A Robustness Tables

A.1 Sample period starting at 1900

Table 15: Price elasticity of extraction rates (1900–2015)

Dep. variable: ln(extraction rate)

Variable	(1)	(2)	(3)	(4)	(5)
$\ln(\text{price})$	0.06^{***}	0.09^{***}	0.12^{***}	0.14^{***}	0.16^{***}
	(0.006)	(0.008)	(0.009)	(0.01)	(0.011)
$\ln(\text{price}) \times \mathbb{1}_{\text{OPEC}}$			-0.21^{***}	-0.2^{***}	
			(0.022)	(0.023)	
$\ln(\text{price}) \times \mathbb{1}_{\text{Big Firm}}$				-0.06^{***}	
				(0.018)	
$\ln(\text{price}) \times \mathbb{1}_{\lvert\Delta \ln(p)\rvert>0.1}$					-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: This table presents estimates of the elasticity of extraction rates with respect to oil prices. The data includes oil fields with positive extraction rates in 1900-2015, excluding the last year of operation. Standard errors, clustered at the oil-field level, are reported in parenthesis. The instrument for price is the one-year-lagged price. $(***)$ - significant at a 1 precent level.

Table 16: Price elasticity of oil fields in operation (1900–2015)

Dep. variable: ln(number of operating oil fields)

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

Notes: This table presents estimates of the elasticity of the number oil fields in operation with respect to oil prices. The data used is for the period 1900-2015. The dependent variable is the number of oil fields with positive extraction rates. Newey-West standard errors computed with 5-year lags are reported in parenthesis. The instrument for the price is the one-year-lagged price.

A.2 Instrument with real copper prices

Table 17: Price elasticity of extraction rates

Dep. variable: ln(extraction rate)

Variable	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**
$\ln(\text{price})$	0.09^{***} (0.009)	0.26^{***} (0.018)	0.31^{***} (0.019)	0.34^{***} (0.021)	0.22^{***} (0.018)
$\ln(\text{price}) \times \mathbb{1}_{\text{OPEC}}$			-0.25^{***} (0.041)	-0.24^{***} (0.041)	
$\ln(\text{price}) \times \mathbb{1}_{\text{Big Firm}}$				-0.1^{***} (0.032)	
$\ln(\text{price}) \times \mathbb{1}_{\lvert\Delta \ln(p)\rvert>0.1}$					-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: This table presents estimates of the elasticity of the number oil fields in operation with respect to oil prices. The data used is for the period 1900-2015. The dependent variable is the number of oil fields with positive extraction rates. Newey-West standard errors computed with 5-year lags are reported in parenthesis. The instrument for the real oil price is the one-year-ahead forecast of the real price of copper. $(***)$ - significant at a one percent level.

Table 18: Price elasticity of oil fields in operation

Dep. variable: ln(number of operating oil fields)

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

Notes: This table presents estimates of the elasticity of the number oil fields in operation with respect to oil prices. The dependent variable is the number of oil fields with positive extraction rates. Newey-West standard errors computed with 5-year lags are reported in parenthesis. The instrument for the price is the one-year-ahead forecast of the real copper price.

Table 19: Extraction rate adjustment costs regression

Dep. variable: ln(prod. costs per barrel of oil reserves)

Variable	**(1)**	**(2)**	**(3)**
$\ln(\text{extraction})$	4.54^{***}	4.05^{***}	15.19^{**}
	(0.22)	(0.18)	(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

Notes: This table presents estimates of the elasticity of extraction rates with respect to oil prices, η. Standard errors are clustered at the oil-field level. The instrument used is the one-year-ahead forecast of the real price of copper. *** - significant at the 1 percent level.

Table 20: Estimated parameters

Parameter	**Competitive**		**Cartel**	
	Estimate	(s.e.)	Estimate	(s.e.)
ϵ	0.19	(0.04)	0.20	(0.06)
α	0.42	(0.85)	0.34	(-)
ρ_1^d	1.86	(0.07)	1.79	(0.07)
ρ_2^d	-0.86	(0.07)	-0.80	(0.08)
ρ_1^u	1.56	(0.14)	-0.25	(0.79)
ρ_2^u	-0.61	(0.13)	-0.02	(2.30)
$var(e_t^d)$	0.004	(0.002)	0.011	(0.004)
$var(e_t^u)$	0.001	(0.000...)	0.002	(0.001)

Notes: This table presents the GMM estimates of the structural parameters when η is obtained using the real price of copper as an instrument. Under the cartel specification, α is not estimated, but calibrated so that the steady-state extraction rate matches the average extraction rate for OPEC.

Table 21: Data and model moments

	Moment	Data	(s.e.)	Competitive model		Cartel model	
				Benchm.	Copper IV	Benchm.	Copper IV
(1)	std(Δp_t)	0.273	(0.028)	0.188	0.112	0.220	0.175
(2)	std(Δi_t^N)	0.192	(0.024)	0.218	0.220	0.212	0.215
(3)	std(Δi_t^O)	0.193	(0.027)	0.215	0.215	0.208	0.209
(4)	std(Δq_t^N)	0.022	(0.003)	0.026	0.027	0.024	0.026
(5)	std(Δq_t^O)	0.069	(0.011)	0.057	0.044	0.068	0.066
(6)	corr(Δp_t, Δi_t^N)	0.557	(0.147)	0.738	0.723	0.687	0.602
(7)	corr(Δp_t, Δi_t^O)	0.362	(0.109)	0.659	0.589	0.671	0.594
(8)	corr(Δp_t, Δq_t^N)	0.031	(0.069)	0.013	-0.008	0.127	0.079
(9)	corr(Δp_t, Δq_t^O)	0.030	(0.122)	-0.585	-0.491	-0.583	-0.680
(10)	corr(Δi_t^N, Δi_t^O)	0.673	(0.096)	0.992	0.979	0.999	1.000
(11)	corr(Δi_t^N, Δq_t^N)	0.087	(0.094)	-0.041	-0.004	-0.025	0.005
(12)	corr(Δi_t^N, Δq_t^O)	0.023	(0.112)	-0.150	-0.151	-0.008	0.003
(13)	corr(Δi_t^O, Δq_t^N)	-0.034	(0.145)	-0.035	0.002	-0.027	0.003
(14)	corr(Δi_t^O, Δq_t^O)	-0.226	(0.153)	-0.044	0.004	0.013	0.012
(15)	corr(Δq_t^N, Δq_t^O)	-0.141	(0.125)	0.026	-0.075	0.341	0.333
(16)	corr(Δp_t,Δp_{t-1})	-0.027	(0.088)	-0.043	-0.122	-0.166	-0.333
(17)	corr(Δi_t^N,Δi_{t-1}^N)	0.119	(0.135)	0.007	0.028	0.001	0.018
(18)	corr(Δi_t^O,Δi_{t-1}^O)	0.311	(0.096)	0.009	0.031	-0.001	0.017
(19)	corr(Δq_t^N,Δq_{t-1}^N)	0.643	(0.113)	0.291	0.434	0.416	0.533
(20)	corr(Δq_t^O,Δq_{t-1}^O)	0.213	(0.211)	0.314	0.523	-0.392	-0.384
(21)	$\mathbb{E}\left[I_t/\left(P_tQ_t\right)\right]$	0.102	(0.016)	0.057	0.049	0.027	0.029

Notes: This table presents key targeted moments from the data and the model-implied moments under the benchmark specification and when η is estimated using the real price of copper instrument. Newey-West standard errors computed with 5-year lags are reported in parenthesis. x_t represents the logarithm of X_t, Δx_t is equal to $x_t - x_{t-1}$.

A.3 Instrument with real metal prices

Table 22: Price elasticity of extraction rates

Dep. variable: ln(extraction rate)

Variable	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**
$\ln(\text{price})$	0.09^{***}	0.20^{***}	0.24^{***}	0.26^{***}	0.22^{***}
	(0.009)	(0.014)	(0.015)	(0.016)	(0.015)
$\ln(\text{price}) \times \mathbb{1}_{\text{OPEC}}$			-0.23^{***}	-0.22^{***}	
			(0.035)	(0.035)	
$\ln(\text{price}) \times \mathbb{1}_{\text{Big Firm}}$				-0.07^{**}	
				(0.027)	
$\ln(\text{price}) \times \mathbb{1}_{\lvert\Delta \ln(p)\rvert>0.1}$					0.04^{***}
					(0.005)
Oil field FE	✓	✓	✓	✓	✓
Operation year FE	✓	✓	✓	✓	✓
Year trend	✓	✓	✓	✓	✓
IV	✗	✓	✓	✓	✓
Clusters (oil fields)	12,187	11,479	11,479	11,479	11,479
Observations	173,742	173,034	173,034	173,034	173,034

Notes: This table presents estimates of the elasticity of extraction rates with respect to oil prices. The data includes oil fields with positive extraction rates in 1971-2015, excluding the last year of operation. Standard errors, clustered at the oil-field level, are reported in parenthesis. The instrument for the real oil price is the one-year-ahead forecast of the real price of metal. $(***)$ [$**$] - significant at a 1 percent level [5 percent].

Table 23: Price elasticity of oil fields in operation

Dep. variable: ln(number of operating oil fields)

Variable	**(1)**	**(2)**	**(3)**	**(4)**
$\ln(\text{price})$	-0.05	-0.23^{***}	-0.08^{**}	-0.15^{***}
	(0.06)	(0.07)	(0.03)	(0.04)
Year trend	✓	✓	✓	✓
IV	✗	✓	✓	✓
Dep. variable	All fields	All fields	Non-OPEC fields	OPEC fields
Observations	45	45	45	45

Notes: This table presents estimates of the elasticity of the number of oil fields in operation with respect to oil prices. The dependent variable is the number of oil fields with positive extraction rates. Newey-West standard errors computed with 5-year lags are reported in parenthesis. The instrument for the price is the one-year-ahead forecast of real price of metals.

Table 24: Extraction rate adjustment costs regression

Dep. variable: ln(prod. costs per barrel of oil reserves)

Variable	**(1)**	**(2)**	**(3)**
$\ln(\text{extraction})$	4.97^{***}	4.35^{***}	33.0
	(0.26)	(0.21)	(37.6)
Oil field FE	✓	✓	✓
Operation year FE	✓	✓	✓
Sample	All	Non-OPEC	OPEC
IV	✓	✓	✓
1^{st} stage F-stat	243	275	0.72
Clusters (oil fields)	11,527	9,969	1,558
Observations	174,339	146,879	27,460

Notes: This table presents estimates the adjustment-cost coefficient, η. Standard errors are clustered at the oil-field level. The instrument used is the one-year-ahead forecast of the real price of metals. *** - significant at the 1 percent level.

Table 25: Estimated parameters

Parameter	Competitive		Cartel	
	Estimate	(s.e.)	Estimate	(s.e.)
ϵ	0.19	(0.04)	0.21	(0.07)
α	0.42	(0.84)	0.33	(-)
ρ_1^d	1.84	(0.08)	1.78	(0.06)
ρ_2^d	-0.85	(0.08)	-0.79	(0.06)
ρ_1^u	1.54	(0.14)	-0.16	(0.75)
ρ_2^u	-0.60	(0.13)	0.00	(2.08)
$var(e_t^d)$	0.005	(0.003)	0.012	(0.004)
$var(e_t^u)$	0.001	(0.000...)	0.002	(0.001)

Notes: This table presents the GMM estimates of the structural parameters when η is estimated using the real price of metals as an instrument. Under the cartel specification, α is not estimated, but calibrated so that the steady state of the moel matches the average extraction rate of OPEC.

Table 26: Data and model moments

	Moment	**Data**	(s.e.)	**Competitive model**		**Cartel model**	
				Benchm.	Metal IV	Benchm.	Metal IV
(1)	std(Δp_t)	0.273	(0.028)	0.188	0.120	0.220	0.175
(2)	std(Δi_t^N)	0.192	(0.024)	0.218	0.220	0.212	0.215
(3)	std(Δi_t^O)	0.193	(0.027)	0.215	0.215	0.208	0.209
(4)	std(Δq_t^N)	0.022	(0.003)	0.026	0.027	0.024	0.026
(5)	std(Δq_t^O)	0.069	(0.011)	0.057	0.046	0.068	0.067
(6)	corr(Δp_t, Δi_t^N)	0.557	(0.147)	0.738	0.726	0.687	0.617
(7)	corr(Δp_t, Δi_t^O)	0.362	(0.109)	0.659	0.600	0.671	0.608
(8)	corr(Δp_t, Δq_t^N)	0.031	(0.069)	0.013	-0.008	0.127	0.089
(9)	corr(Δp_t, Δq_t^O)	0.030	(0.122)	-0.585	-0.506	-0.583	-0.660
(10)	corr(Δi_t^N, Δi_t^O)	0.673	(0.096)	0.992	0.981	0.999	1.000
(11)	corr(Δi_t^N, Δq_t^N)	0.087	(0.094)	-0.041	-0.009	-0.025	0.002
(12)	corr(Δi_t^N, Δq_t^O)	0.023	(0.112)	-0.150	-0.153	-0.008	0.002
(13)	corr(Δi_t^O, Δq_t^N)	-0.034	(0.145)	-0.035	-0.003	-0.027	-0.000
(14)	corr(Δi_t^O, Δq_t^O)	-0.226	(0.153)	-0.044	-0.003	0.013	0.011
(15)	corr(Δq_t^N, Δq_t^O)	-0.141	(0.125)	0.026	-0.061	0.341	0.350
(16)	corr(Δp_t,Δp_{t-1})	-0.027	(0.088)	-0.043	-0.113	-0.166	-0.297
(17)	corr(Δi_t^N,Δi_{t-1}^N)	0.119	(0.135)	0.007	0.026	0.001	0.017
(18)	corr(Δi_t^O,Δi_{t-1}^O)	0.311	(0.096)	0.009	0.028	-0.001	0.015
(19)	corr(Δq_t^N,Δq_{t-1}^N)	0.643	(0.113)	0.291	0.414	0.416	0.528
(20)	corr(Δq_t^O,Δq_{t-1}^O)	0.213	(0.211)	0.314	0.495	-0.392	-0.374
(21)	$\mathbb{E}\left[I_t/\left(P_t Q_t\right)\right]$	0.102	(0.016)	0.057	0.050	0.027	0.028

Notes: This table presents key targeted moments from the data and the model-implied moments under the benchmark specification and when η is estimated using the real price of metals as an instrument. Newey-West standard errors computed with 5-year lags are reported in parenthesis. x_t represents the logarithm of X_t, Δx_t is equal to $x_t - x_{t-1}$.

A.4 Excluding Saudi Arabia and Venezuela from the sample

Table 27: Price elasticity of extraction rates

Dep. variable: ln(extraction rate)

Variable	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**
ln(price)	0.11^{***}	0.22^{***}	0.19^{***}	0.16^{***}	0.19^{***}
	(0.009)	(0.032)	(0.036)	(0.42)	(0.03)
$\ln(\text{price}) \times \mathbb{1}_{\text{OPEC}}$			0.16^{*}	0.16^{*}	
			(0.10)	(0.10)	
$\ln(\text{price}) \times \mathbb{1}_{\text{Big Firm}}$				0.11	
				(0.09)	
$\ln(\text{price}) \times \mathbb{1}_{\lvert\Delta \ln(p)\rvert>0.1}$					-0.01^{***}
					(0.002)
Oil field FE	✓	✓	✓	✓	✓
Operation year FE	✓	✓	✓	✓	✓
Year trend	✓	✓	✓	✓	✓
IV	✗	✓	✓	✓	✓
Clusters (oil fields)	11,893	11,195	11,195	11,195	11,195
Observations	167,848	167,150	167,150	167,150	167,150

Notes: This table presents estimates of the elasticity of extraction rates with respect to oil prices. The data includes oil fields with positive extraction rates in 1971-2015, excluding the last year of operation and oil fields in Saudi Arabia and Venezuela. Standard errors, clustered at the oil-field level, are reported in parenthesis. The instrument for the price is the one-year-ahead forecast of detrended world real GDP. $(***)$ $[*]$ - significant at a 1 percent $[10percent]$ level.

Table 28: Price elasticity of oil fields in operation

Dep. variable: ln(number of operating oil fields)

Variable	**(1)**	**(2)**	**(3)**	**(4)**
ln(price)	−0.03 (0.08)	−0.22 (0.18)	−0.09 (0.08)	−0.13 (0.11)
Year trend	✓	✓	✓	✓
IV	✗	✓	✓	✓
Dep. variable	All fields	All fields	Non-OPEC fields	OPEC fields
Observations	45	45	45	45

Notes: This table presents estimates for the elasticity of the number of oil fields in operation when we exclude oil fields in Saudi Arabia and Venezuela. The dependent variable is the number of oil fields with positive extraction rates. Newey-West standard errors computed with 5-year lags are reported in parenthesis. The instrument for the price is the one-year-ahead forecast of detrended world real GDP.

Table 29: Extraction rate adjustment costs regression

Dep. variable: ln(prod. costs per barrel of oil reserves)

Variable	**(1)**	**(2)**	**(3)**
ln(extraction)	6.88^{***} (1.09)	7.69^{***} (1.54)	4.21^{***} (0.83)
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

Notes: This table presents estimates of the adjustment-cost coefficient, η, when we exclude oil fields in Saudi Arabia and Venezuela. Standard errors are clustered at the oil-field level. The instrument used is the one-year-ahead forecast of detrended world real GDP. *** - significant at the 1% level.

Table 30: Estimated parameters

Parameter	Competitive		Cartel	
	Estimate	(s.e.)	Estimate	(s.e.)
ϵ	0.16	(0.03)	0.21	(0.08)
α	0.43	(0.77)	0.29	(-)
ρ_1^d	1.74	(0.09)	1.70	(0.06)
ρ_2^d	-0.76	(0.09)	-0.72	(0.06)
ρ_1^u	1.50	(0.10)	-0.10	(2.48)
ρ_2^u	-0.56	(0.10)	0.00	(1.84)
$var(e_t^d)$	0.016	(0.006)	0.023	(0.006)
$var(e_t^u)$	0.002	(0.001)	0.002	(0.002)

Notes: This table presents the GMM estimates of the structural parameters when the targeted moments and η are computed excluding observations on Saudi Arabia and Venezuela. Under the cartel specification, α is not estimated, but calibrated so that steady state extraction rate matches OPEC's average extraction rate.

Table 31: Data and model moments

	Moment	Data	(s.e.)	Competitive model	Cartel model
(1)	std(Δp_t)	0.273	(0.028)	0.188	0.218
(2)	std(Δi_t^N)	0.192	(0.024)	0.226	0.217
(3)	std(Δi_t^O)	0.208	(0.031)	0.220	0.212
(4)	std(Δq_t^N)	0.022	(0.003)	0.026	0.024
(5)	std(Δq_t^O)	0.076	(0.016)	0.052	0.073
(6)	corr(Δp_t, Δi_t^N)	0.557	(0.147)	0.800	0.690
(7)	corr(Δp_t, Δi_t^O)	0.360	(0.120)	0.708	0.672
(8)	corr(Δp_t, Δq_t^N)	0.031	(0.069)	0.030	0.158
(9)	corr(Δp_t, Δq_t^O)	-0.025	(0.104)	-0.540	-0.569
(10)	corr(Δi_t^N, Δi_t^O)	0.626	(0.108)	0.986	0.999
(11)	corr(Δi_t^N, Δq_t^N)	0.087	(0.094)	-0.041	-0.024
(12)	corr(Δi_t^N, Δq_t^O)	-0.065	(0.112)	-0.202	-0.007
(13)	corr(Δi_t^O, Δq_t^N)	-0.067	(0.138)	-0.036	-0.026
(14)	corr(Δi_t^O, Δq_t^O)	-0.357	(0.246)	-0.068	0.015
(15)	corr(Δq_t^N, Δq_t^O)	-0.126	(0.089)	0.012	0.328
(16)	corr(Δp_t,Δp_{t-1})	-0.027	(0.088)	-0.037	-0.133
(17)	corr(Δi_t^N,Δi_{t-1}^N)	0.119	(0.135)	0.008	0.001
(18)	corr(Δi_t^O,Δi_{t-1}^O)	0.252	(0.110)	0.010	-0.002
(19)	corr(Δq_t^N,Δq_{t-1}^N)	0.643	(0.113)	0.339	0.420
(20)	corr(Δq_t^O,Δq_{t-1}^O)	0.349	(0.154)	0.395	-0.401
(21)	$\mathbb{E}\left[I_t/\left(P_t Q_t\right)\right]$	0.102	(0.016)	0.056	0.026

Notes: This table presents key targeted moments from the data and the model-implied moments when observations from Saudi Arabia and Venezuela are excluded. Newey-West standard errors computed with 5-year lags are reported in parenthesis. x_t represents the logarithm of X_t, Δx_t is equal to $x_t - x_{t-1}$.

B Balanced growth path

In the non-stochastic steady state of our models, oil production is constant. In this section, we extend the competitive version of our 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 growth in oil demand.

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

$$X_{t+1}^i = (1-\lambda)X_t^i + A_t \left(I_t^i\right)^{\alpha} \left(L^i\right)^{1-\alpha} , \qquad i \in \{O, N\} . \tag{27}$$

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

$$P_t = \exp\left(d_t^{AR2}\right) \exp\left(d_t^g\right) Q_t^{-\frac{1}{\epsilon}} , \tag{28}$$

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 the existence of 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

$$V\left(K, X, \theta, A, u; \Omega\right) = \max_{\{K', X', \theta', I\}} P\theta K e^{-u} - I - \psi \left(\frac{\theta}{\theta_{ss}}\right)^{\eta} K + \beta \mathbb{E} V\left(K', X', \theta', A', u'; \Omega'\right) , \tag{29}$$

$$\text{s.t.} \quad K' = (1-\theta e^{-u})K + \lambda X' ,$$

$$X' = (1-\lambda)X + A I^{\alpha} L^{1-\alpha} ,$$

where $u = 0$ for non-OPEC firms.[23] 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

$$\tilde{V}\left(\tilde{K}, \tilde{X}, \theta, u; \tilde{\Omega}\right) = \max_{\{\tilde{K}', \tilde{X}', \theta', \tilde{I}\}} P\theta\tilde{K}e^{-u} - \tilde{I} - \psi\left(\frac{\theta}{\theta_{ss}}\right)^{\eta}\tilde{K} + e^{g}\beta\mathbb{E}\tilde{V}\left(\tilde{K}', \tilde{X}', \theta', u'; \tilde{\Omega}'\right) ,$$

$$\text{s.t.} \quad \tilde{K}' = e^{-g}(1 - e^{-u}\theta)\tilde{K} + \lambda\tilde{X}' ,$$

$$\tilde{X}' = e^{-g}(1-\lambda)\tilde{X} + e^{-g}\tilde{I}^{\alpha}L^{1-\alpha} ,$$

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.

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

$$[\tilde{K}'] \qquad \mu_1 = \beta e^{g}\mathbb{E}\left[P'\theta' e^{-u'} + e^{-g}(1-\theta' e^{-u'})\mu_1' - \psi\left(\frac{\theta'}{\theta_{ss}}\right)^{\eta}\right] , \tag{30}$$

$$[\tilde{X}'] \qquad \mu_2 = \lambda\mu_1 + \beta(1-\lambda)\mathbb{E}_t\mu_2' , \tag{31}$$

$$[\tilde{I}] \qquad 1 = e^{-g}\alpha\left(L/\tilde{I}\right)^{1-\alpha}\mu_2 , \tag{32}$$

$$[\theta'] \qquad \eta\psi\left(\theta'\right)^{\eta-1}(\theta_{ss})^{-\eta}\tilde{K}' = \mathbb{E}\left[(P'e^{-u'} - e^{-g}e^{-u'}\mu_1')\right]\tilde{K}' , \tag{33}$$

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 FOCs, we have

$$\mu_1 = \beta e^{g}\frac{\theta P - \psi}{1 - \beta(1-\theta)} . \tag{34}$$

[23] For 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 second FOC can be written as

$$\mu_2 = \frac{\lambda}{1-\beta(1-\lambda)}\mu_1 \; . \tag{35}$$

Substituting into the third FOC, we have

$$\left(\frac{\tilde{I}}{L}\right)^{1-\alpha} = e^{-g}\alpha\frac{\lambda}{1-\beta(1-\lambda)}\mu_1 \; . \tag{36}$$

The last FOC simplifies to

$$\eta\psi\left(\theta\right)^{-1} = P - e^{-g}\mu_1 \; . \tag{37}$$

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

$$P = e^{d_t^g}\left(\theta^N K^N + \theta^O K^O\right)^{-\frac{1}{\epsilon}} \; . \tag{38}$$

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

$$P = \left(\theta^N \tilde{K}^N + \theta^O \tilde{K}^O\right)^{-\frac{1}{\epsilon}} \; . \tag{39}$$

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

$$\tilde{X} = \frac{1}{e^g-(1-\lambda)}\left(\tilde{I}\right)^{\alpha}(L)^{1-\alpha} \; , \tag{40}$$

$$\tilde{K} = \frac{e^g}{e^g-(1-\theta)}\frac{\lambda}{e^g-(1-\lambda)}\left(\tilde{I}\right)^{\alpha}(L)^{1-\alpha} \; . \tag{41}$$

Combining equations (34) and (37) we get

$$P - \eta\psi\left(\theta\right)^{-1} = \beta\frac{\theta P-\psi}{1-\beta(1-\theta)} \; . \tag{42}$$

Rearranging we get

$$(1-\beta)P = \left[\left(\frac{1-\beta}{\theta}+\beta\right)\eta-\beta\right]\psi \; , \tag{43}$$

or simply

$$P = \eta\psi\left(\theta\right)^{-1} + \frac{\beta}{1-\beta}(\eta-1)\psi \; . \tag{44}$$

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 (34) and (36), we can pin down the investment to land ratios for OPEC and non-OPEC firms.

$$\left(\frac{\tilde{I}}{L}\right)^{1-\alpha} = \frac{\alpha\lambda}{1-\beta(1-\lambda)}\frac{\beta}{1-\beta}(\eta-1)\psi \; . \tag{45}$$

Substituting into the demand function we get

$$\left[\eta\psi\left(\theta\right)^{-1}+\frac{\beta}{1-\beta}(\eta-1)\psi\right]^{-\epsilon}=\frac{e^{g}\theta}{e^{g}-(1-\theta)}\frac{\lambda}{e^{g}-(1-\lambda)}\left[\frac{\alpha\lambda}{1-\beta(1-\lambda)}\frac{\beta}{1-\beta}(\eta-1)\psi\right]^{\frac{\alpha}{1-\alpha}}\left(L^{N}+L^{O}\right)\ . \tag{46}$$

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 and the levels of reserves, exploration capital, and production, as well as investment and production costs grow at a constant rate, $g = \frac{1}{1-\alpha}g_A$.