

NBER WORKING PAPER SERIES

A WORLD EQUILIBRIUM MODEL OF THE OIL MARKET

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 January 2021

An earlier version of this paper was circulated under the title "Lags, Costs, and Shocks: An Equilibrium Model of the Oil Industry." We thank Hilde Bjornland, Craig Burnside, Mario Crucini, Wei Cui, Jesús Fernández-Villaverde, Matteo Iacoviello, Ravi Jaganathan, Ryan Kellogg, Lutz Kilian, Markus Kirchner, Laura Murphy, 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.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Gideon Bornstein, Per Krusell, and Sergio Rebelo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A World Equilibrium Model of the Oil Market
Gideon Bornstein, Per Krusell, and Sergio Rebelo
NBER Working Paper No. 23423
May 2017, Revised January 2021
JEL No. Q4,Q43

ABSTRACT

We use new, comprehensive micro data on oil fields to build and estimate a structural model of the oil industry embedded in a general equilibrium model of the world economy. In the model, firms that belong to OPEC act as a cartel. The remaining firms are a competitive fringe. We use the model to study the macroeconomic impact of the advent of fracking. Fracking weakens the OPEC cartel, leading to a large long-run decline in oil prices. Fracking also reduces the volatility of oil prices in the long run because fracking firms can respond more quickly to changes in oil demand.

Gideon Bornstein
20 Washington Rd
Economics Department
JRR Building
Princeton, NJ 08540
United States
bornstein@princeton.edu

Per Krusell
Institute for International Economic Studies
Stockholm University
106 91 STOCKHOLM
SWEDEN
and CEPR
and also 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

In this paper, we build and estimate a structural model of the oil industry embedded in a general equilibrium model of the world economy. Our modeling of production and investment in the oil industry relies 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. The data include information about roughly 14,000 oil fields operated by 3,200 companies across 109 countries.

These data guide the construction of our model in two ways. First, we use it to obtain 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 employ the data to obtain a set of second moments for oil-related aggregate variables. We use the generalized method of moments (GMM) to estimate the remaining model parameters, targeting these moments.

We use our model to study the macroeconomic consequences of the ongoing large structural changes in the oil industry associated with the advent of hydraulic fracturing (fracking). This production technique involves pumping a mixture of water, sand, and chemicals at high pressure into shale rock formations to open up small fissures that release oil and gas. Combined with the ability to drill horizontally through shale layers over long distances, fracking has transformed the U.S. from a top oil importer to a top oil exporter. The expansion of fracking continues, not just in the U.S. but also in countries such as Argentina, China, Mexico, and Russia.

We document two important differences between fracking and conventional oil production. First, adjusting production in the short run is cheaper for fracking firms, so these firms are more responsive to changes in oil prices. The elasticity of operational costs with respect to the extraction rate is about five times lower in fracking than in conventional oil production. Second, the average lag between investment and production is much shorter in fracking (one year) than in conventional oil production (12 years).

Our model features two types of firms: those that belong to the Organization of the Petroleum Exporting Countries (OPEC) and those that do not. OPEC firms act as a cartel. Non-OPEC firms are a competitive fringe.

This structural model implies that the expansion of fracking results in large long-run declines in the level and volatility of oil prices. The main reason for the decline in the level of oil prices is that fracking weakens OPEC's market power. Even though OPEC manages to retain its market share, it does so by expanding production, contributing to the fall in oil prices. The main reason

for the decline in the volatility of oil prices is that fracking firms can respond more quickly to changes in oil demand.

Our model's implication that OPEC's market share remains relatively stable after the advent of fracking is consistent with the data. Most of the rise in the market share of fracking firms has been compensated by a fall in the market share of non-OPEC conventional oil producers.

Our estimated model has 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 three model features: a lag between investment and production, convex costs of adjusting extraction rates, and decreasing returns to oil investment.

One interesting property of our model is that it accounts for 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. Under this interpretation, when prices rise, firms expect prices to remain high, so they increase their investments. This rise in investment expands supply and causes prices to fall, so the expectation of high prices is irrational. In our model, the high correlation between the price of oil and investment follows naturally from the rational response of forward-looking firms. A persistent, 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.

Our work relates to four strands of research. The first strand studies the relation between oil price fluctuations and business cycles. Examples include [Backus and Crucini \(2000\)](#), [Leduc and Sill \(2004\)](#), [Blanchard and Gali \(2007\)](#), [Kilian \(2009\)](#), [Bodenstein, Erceg, and Guerrieri \(2011\)](#), [Lippi and Nobili \(2012\)](#), and [Arezki, Ramey, and Sheng \(2016\)](#).² Relative to this literature, our main contribution is to construct a model of the world economy in which OPEC acts as a cartel with the oil production technology parameters estimated using micro data.

¹While the amount of oil is ultimately finite, we can think about this investment process as including new ways of extracting oil as well as the development of oil substitutes, as in [Adao, Narajabad, and Temzelides \(2017\)](#) and [Stuermer and Schwerhoff \(2020\)](#). A large expansion of oil reserves took place 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.

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

The second strand is a new, emerging literature that uses micro data to shed new light on key aspects of the oil industry. Examples include Kellogg (2014), Anderson, Kellogg, and Salant (2017), Bjornland, Nordvik, and Rohrer (2017), Asker, Collard-Wexler, and De Loecker (2019), and Newell and Prest (2019). Relative to this literature, our main contribution is to construct a general equilibrium model of oil and output production.

The third strand is a literature that relies on structural vector autoregressions to estimate the importance of demand and supply shocks on oil prices and oil production. Examples include Kilian (2009), Kilian and Murphy (2014), Baumeister and Hamilton (2019), Herrera and Rangaraju (2020), and Zhou (2020). These analyses generally apply to stationary environments. Our contribution relative to this literature is to provide a structural model that can be used to study an important structural change in the oil industry: the advent of fracking.

The fourth strand consists of papers that study the impact of fracking on aspects of the U.S. economy, such as stock market valuations (Gilje, Ready, and Roussanov 2016) and real GDP (Melek, Plante, and Yücel 2017). Relative to this literature, our contribution is to study how the advent of fracking shapes the dynamics of oil prices and world economic activity.

This paper is organized as follows. We describe our model 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. In Section 4, we use our model to study the impact of the advent of fracking on the economy. Section 5 includes different robustness analyses. Section 6 concludes.

2 The oil market in a general equilibrium model

In this section, we describe a model of the oil industry embedded in a general equilibrium model of the world economy.

2.1 Representative Household

The economy is populated by a representative household who owns the oil companies, supplies labor (N_t), invests in physical capital (K_t), and consumes a final good (C_t). The household problem is

$$\begin{aligned} \max_{\{C_t, K_{t+1}, L_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{1}{1-\gamma} \left(C_t - \varphi \frac{N_t^{1+\nu}}{1+\nu} \right)^{1-\gamma}, \\ \text{s.t. } C_t + K_{t+1} \leq (1+r_t - \delta)K_t + w_t N_t + \Pi_t^O, \end{aligned} \quad (1)$$

where Π_t^O is the total profit from oil companies. The variables r_t and w_t denote the rental rate of capital and the real wage rate, respectively. The parameter δ denotes the rate of depreciation of physical capital.

The first-order conditions of the household problem are

$$N_t = \varphi^{-\frac{1}{\nu}} w_t^{\frac{1}{\nu}}, \quad (2)$$

$$1 = \mathbb{E}_t \left[\tilde{\beta}_{t+1} (1 + r_{t+1} - \delta) \right], \quad (3)$$

where

$$\tilde{\beta}_{t+1} \equiv \beta \frac{M_{t+1}}{M_t}.$$

The variable M_t denotes the marginal utility of consumption:

$$M_t \equiv \left(C_t - \varphi \frac{N_t^{1+\nu}}{1+\nu} \right)^{-\gamma}.$$

2.2 Final goods producers

There is a continuum with measure one of competitive final goods producers. Their output (Y_t) is produced according to a CES production function that combines oil (O_t) and value added. Value added is produced with a Cobb-Douglas production function that combines physical capital and labor. We normalize the price of the final good to one.

The problem of a final-good producer is

$$\begin{aligned} & \max_{\{L_t, K_t, O_t, Y_t\}} Y_t - w_t N_t - r_t K_t - p_t O_t, \\ \text{s.t. } & Y_t = \left[(1 - s_o) A_t (K_t^\alpha N_t^{1-\alpha})^{\frac{\epsilon-1}{\epsilon}} + s_o O_t^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{\epsilon}{\epsilon-1}}, \end{aligned} \quad (4)$$

where p_t is the price of oil. The variable A_t denotes the exogenous level of total factor productivity, which follows an AR(2) in logarithms:

$$\ln A_t = \rho_1^A \ln A_{t-1} + \rho_2^A \ln A_{t-2} + \epsilon_t^a.$$

The disturbance ϵ_t^a is normally distributed with mean zero and variance σ_a^2 . Throughout the paper, we refer to ϵ_t^a as shocks to oil demand. The AR(2) specification allows the model to fit the data much better than an AR(1) specification. Demand shocks need to be very persistent to induce a large response in oil investment. This response allows the model to be consistent with the positive correlation between oil prices and oil investment, as well as with the high volatility of oil investment.

The first-order conditions for the firm's problem are

$$p_t = s_o Y_t^{\frac{1}{\epsilon}} O_t^{-\frac{1}{\epsilon}}, \quad (5)$$

$$w_t = (1 - \alpha)(1 - s_o) A_t Y_t^{\frac{1}{\epsilon}} (K_t^\alpha N_t^{1-\alpha})^{-\frac{1}{\epsilon}} K_t^\alpha N_t^{-\alpha}, \quad (6)$$

$$r_t = \alpha(1 - s_o) A_t Y_t^{\frac{1}{\epsilon}} (K_t^\alpha N_t^{1-\alpha})^{-\frac{1}{\epsilon}} K_t^{\alpha-1} N_t^{1-\alpha}. \quad (7)$$

Equation (5) represents the demand for oil. The presence of Y_t makes this demand dynamic. The level of Y_t reflects past oil prices that influenced past capital accumulation decisions. In addition, Y_t responds to anticipated movements in future oil prices through changes in physical-capital investment.

2.3 The oil sector

The oil sector is composed of an OPEC cartel and a non-OPEC competitive fringe. OPEC firms are all conventional oil producers. The set of non-OPEC firms includes both conventional oil producers and fracking firms.

2.3.1 Non-OPEC conventional oil producers

There is a continuum with measure one of competitive non-OPEC conventional oil producers. These firms maximize their value (V^N), which is given by

$$V^N = E_0 \sum_{t=0}^{\infty} \beta^t M_t \left[p_t \theta_t^N R_t^N - I_t^N - \psi (\theta_t^N)^\eta R_t^N \right]. \quad (8)$$

Here, I_t^N denotes investment, θ_t^N the extraction rate (the ratio of production to reserves), and R_t^N oil reserves. To simplify, we abstract from taxes and royalties for both non-OPEC and other producers.³

The term $\psi (\theta_t^N)^\eta R_t^N$ 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 extraction costs are convex in the extraction rate. In addition, we assume that the time $t + 1$ extraction rate is chosen at time t .

³When taxes or royalties are a fraction of the cash flow, $p_t \theta_t^N R_t^N - I_t^N - \psi (\theta_t^N)^\eta R_t^N$, they do not affect a firm's incentives to produce or invest.

We adopt a parsimonious way of modeling lags in investment by introducing *exploration capital*, which we denote by X_t . The law of motion for X_t is as follows:

$$X_{t+1}^N = (1 - \lambda)X_t^N + (I_t^N)^\xi (L^N)^{1-\xi}. \quad (9)$$

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

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.

One interpretation of equation (9) 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 . This interpretation is consistent with the way in which we estimate λ using our micro data.

An alternative interpretation is that there is a continuum of identical firms investing in exploration capital, but only a fraction λ discovers oil in every period. Since the representative household holds a continuum of these firms, the idiosyncratic risk of oil discovery is diversified away. Aggregating the behavior of these individual firms would yield a structure that is equivalent to our representative-firm model. This interpretation is also consistent with our micro data analysis.

Oil reserves evolve as follows:

$$R_{t+1}^N = (1 - \theta_t^N)R_t^N + \lambda X_{t+1}^N. \quad (10)$$

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

The notion of exploration capital embodied in equations (9) and (10) 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 this lag.⁴ When $\lambda = 0$ and $\eta > 1$, the

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

non-OPEC firm problem is similar to the one in [Anderson, Kellogg, and Salant \(2017\)](#)'s model. Changes in the extraction rate in our model play a role similar to drilling new wells in their model.

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

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

$$\mathbb{E}_t \tilde{\beta}_{t+1} p_{t+1} = \mathbb{E}_t \tilde{\beta}_{t+1} \eta \psi (\theta_{t+1}^N)^{\eta-1} + \mathbb{E}_t \tilde{\beta}_{t+1} \mu_{t+1}^{1,N}, \quad (11)$$

where $\mu_t^{1,N}$ is the Lagrange multiplier corresponding to equation (10). The extraction rate at time $t + 1$ is chosen at time t so as to equate the expected value of a barrel of oil, $\mathbb{E}_t \tilde{\beta}_{t+1} p_{t+1}$, to the sum of the expected value of the marginal cost of extraction, $\mathbb{E}_t \tilde{\beta}_{t+1} \eta \psi (\theta_{t+1}^N)^{\eta-1}$ and the expected value of a barrel of oil reserves at the end of time $t + 1$, $\mathbb{E}_t \tilde{\beta}_{t+1} \mu_{t+1}^{1,N}$.

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

$$\mu_t^{1,N} = \mathbb{E}_t \left\{ \tilde{\beta}_{t+1} \left[p_{t+1} \theta_{t+1}^N + (1 - \theta_{t+1}^N) \mu_{t+1}^{1,N} - \psi (\theta_{t+1}^N)^\eta \right] \right\}. \quad (12)$$

For a given value of θ_{t+1}^N , each extra barrel of oil reserves results in additional revenue with a value $\tilde{\beta}_{t+1} p_{t+1} \theta_{t+1}^N$ and additional extraction costs with a value $\tilde{\beta}_{t+1} \psi (\theta_{t+1}^N)^\eta$. A fraction $1 - \theta_{t+1}^N$ of the barrel of oil reserves remains in the ground and has a value $\tilde{\beta}_{t+1} \mu_{t+1}^{1,N}$.

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

$$\mu_t^{2,N} = \lambda \mu_t^{1,N} + (1 - \lambda) \mathbb{E}_t \tilde{\beta}_{t+1} \mu_{t+1}^{2,N}, \quad (13)$$

where $\mu_t^{2,N}$ is the Lagrange multiplier corresponding to equation (9). The value of increasing exploration capital by one unit, $\mu_t^{2,N}$, has two components. A fraction λ of the exploration capital materializes into oil reserves and has a value $\mu_t^{1,N}$. A fraction $1 - \lambda$ remains as exploration capital and has an expected value $\mathbb{E}_t \tilde{\beta}_{t+1} \mu_{t+1}^{2,N}$.

The first-order condition for I_t^N is

$$1 = \xi (I_t^N)^{\xi-1} (L^N)^{1-\xi} \mu_t^{2,N}. \quad (14)$$

This condition equates the cost of investment (one unit of output) to the marginal product of investment in generating exploration capital, $\xi (I_t^N)^{\xi-1} (L^N)^{1-\xi}$, evaluated at the value of exploration capital, $\mu_t^{2,N}$.

The four optimality conditions, (11)-(14), together with the law of motion for oil reserves, (10), and exploration capital, (9), are a subset of the implementability constraints faced by the OPEC cartel.

2.3.2 Non-OPEC fracking producers

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

$$\max_{\{I_t^F, \theta_t^F, R_{t+1}^F, X_{t+1}^F\}} V^F = E_0 \sum_{t=0}^{\infty} \beta^t M_t \left[p_t \theta_t^F R_t^F - I_t^F - \psi^F (\theta_t^F)^{\eta^F} R_t^F \right], \quad (15)$$

subject to

$$X_{t+1}^F = (1 - \lambda^F) X_t^F + (I_t^F)^\xi (L^F)^{1-\xi}, \quad (16)$$

$$R_{t+1}^F = (1 - \theta_t^F) R_t^F + \lambda^F X_{t+1}^F. \quad (17)$$

Here, I_t^F denotes investment, θ_t^F the extraction rate, X_t^F exploration capital, R_t^F oil reserves, and L^F the land available to the representative fracking firm. Implicit in this formulation is the assumption that fracking firms can adjust their extraction rate within the period. We provide evidence in favor of this assumption in Section 3.

The first-order condition for θ_t^F is

$$p_t = \eta^F \psi^F (\theta_t^F)^{\eta^F - 1} + \mu_t^{1,F}, \quad (18)$$

where $\mu_t^{1,F}$ is the Lagrange multiplier corresponding to equation (17). The optimality conditions for X_{t+1}^F , R_{t+1}^F , I_t^F take the same form as those for non-OPEC firms (12-14).

2.3.3 The OPEC cartel

We assume that OPEC firms operate as a cartel with commitment, while non-OPEC firms are a competitive fringe. [Stiglitz \(1976\)](#) and [Hassler, Krusell, and Olovsson \(2010\)](#) solve for an equilibrium in which the oil market is controlled by a monopolist that faces a constant elasticity demand. Our model is much 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 of oil at time t is a function of all past investment decisions.

We assume that all OPEC firms are conventional oil producers. The OPEC cartel maximizes the discounted utility value of its profits. The cartel takes into account the impact of its decisions on the price of oil as well as on other aggregate variables such as world output and real interest rates. We assume that the OPEC cartel can commit to a policy and not deviate from it. Such commitment is not time consistent. For example, OPEC has an incentive to announce high future production

levels to deter investment by non-OPEC firms and, in the future, renege on this commitment to push up oil prices.

The cartel operates a production technology identical to non-OPEC but has imperfect control over the effective extraction rate. The cartel chooses an extraction rate θ_t^O , but the effective extraction rate is $e^{u_t}\theta_t^O$ where u_t is an i.i.d. shock drawn from a normal distribution with mean zero and variance σ_u^2 . This process can be thought of as a combination of departures from commitment and other disruptions, such as wars. A major departure from commitment is represented by a large positive shock and a significant war by a large negative shock.⁵

Two natural questions arise about this supply shock specification. The first is: why not model u_t as a persistent process? The answer is that such persistence does not change the properties of the model because it is optimal for OPEC to adjust its extraction rate to offset predictable movements in u_t . Only unpredictable movements in u_t have an impact on OPEC production. The second question is: why assume that supply shocks affect only OPEC producers? In the robustness section, we estimate a version of the model in which non-OPEC producers are also subject to supply shocks. The estimated variance of these shocks is close to zero. Including supply shocks to non-OPEC producers does not improve the model's empirical performance.

The cartel maximizes

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t M_t \left[p_t e^{u_t} \theta_t^O R_t^O - I_t^O - \psi (e^{u_t} \theta_t^O)^\eta R_t^O \right], \quad (19)$$

subject to the laws of motion for OPEC oil reserves and exploration capital (X_t^O),

$$R_{t+1}^O = (1 - e^{u_t} \theta_t^O) R_t^O + \lambda X_{t+1}^O, \quad (20)$$

$$X_{t+1}^O = (1 - \lambda) X_t^O + (I_t^O)^\xi L_O^{1-\xi}, \quad (21)$$

as well as the implementability constraints. These constraints can be divided into five sets. The first set includes the household budget constraint, (1), and the first-order conditions for the household, (2) and (3). The second set includes the production function for final goods (4) and the first-order conditions of final goods producers, (5–7). The third set consists of the optimality conditions for non-OPEC conventional oil producers. These conditions include the laws of motion for exploration capital, (9), and oil reserves, (10), as well as the first-order conditions (11–14). The fourth set consists of the optimality conditions for non-OPEC fracking producers. These conditions include the laws of motion for oil reserves, (16), and exploration capital, (17), as well as the

⁵We abstract from idiosyncratic supply shocks. Since there is a continuum of firms and these firms are owned by diversified households, these shocks have no impact on the aggregate economy.

first-order conditions for θ_t^F , (18), and the first-order conditions for X_{t+1}^F , R_{t+1}^F , I_t^F , which take the same form as those for non-OPEC conventional oil producers, (11–14). The fifth set includes the equilibrium condition for oil markets:

$$e^{u_t} \theta_t^O R_t^O + \theta_t^N R_t^N + \theta_t^F R_t^F = O_t.$$

We solve the model using the “timeless perspective” approach proposed by [Woodford \(1999\)](#) and [Woodford \(2011\)](#). This approach can be implemented in two steps. First, we solve for a non-stochastic steady state where the Lagrange multipliers of the OPEC problem on the implementability conditions are constant. This steady state is independent of the initial state of the economy. If the economy converges to a steady state in the long run, it converges to this steady state. Second, we linearize the model around this non-stochastic steady state in order to obtain a recursive, time-invariant system.

2.4 The Hotelling rule revisited

The classic [Hotelling \(1931\)](#) rule emerges as a particular case of our model in which there are no OPEC firms, no new oil reserves can be found ($\lambda = 0$), extraction costs are linear ($\eta = 1$), and households are risk neutral ($\gamma = 0$). When λ is equal to zero, investment does not result in more oil reserves, so oil is an exhaustible resource. Equation (13) implies that, in this case, the value of exploration capital is zero: $\mu_t^{2,N} = 0$. Combining equations (11) and (12), 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 $t + 2$.

For the case in which $\lambda \geq 0$ and $\eta \geq 1$ and $\gamma = 0$, 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

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

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

When $\lambda = 0$ (no more oil can be found), there exists no steady state in which θ_t and p_t are constant. 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.⁶

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

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

Our model's implication that p_t is constant in the long run is consistent with the fact that the average annual growth rate in the real price of oil is not statistically different from zero.⁷ For the period 1900-2010, this average is 0.012 with a standard error of 0.020. For the period 1970-2010, this average is 0.036 with a standard error of 0.041. The property that average growth rates of real prices estimated over long time periods are close to zero is shared by many other commodities (see, e.g., [Deaton and Laroque 1992](#), [Harvey, Kellard, Madsen, and Wohar 2010](#), [Chari and Christiano 2014](#), and [Stuermer and Schwerhoff 2020](#)).

3 Estimation and quantitative analysis

In this section, we estimate the structural parameters of our model using both micro and aggregate data and study the model's quantitative properties. We start by describing our micro data, which we use to obtain estimates for two key model parameters: the average lag between investment and production ($1/\lambda$) and the elasticity of extraction costs with respect to production (η). We use the generalized method of moments (GMM) to estimate the remaining model parameters, targeting a set of second moments for aggregate oil-related variables constructed using the micro data.

3.1 Data on oil markets

Our empirical work relies on new proprietary data compiled by Rystad Energy that contains annual data on production, reserves, operational costs, and investment for all oil fields. The data

⁶One 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.

⁷We obtain oil prices for the period 1900-1948 from [Harvey, Kellard, Madsen, and Wohar \(2010\)](#). 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.

Table 1: Descriptive statistics

Time period	1970–2010			2011–2019		
	OPEC conv.	Non-OPEC conv.	Non-OPEC fracking	OPEC conv.	Non-OPEC conv.	Non-OPEC fracking
Average market share	42%	58%	0%	42%	53%	5%
Average production (Gbbl)	9,227	12,548	41	11,330	14,132	1,454
Average n. of operating oil fields	799	5,653	139	1,220	9,211	1,086

Notes: This table presents descriptive statistics of the three groups of firms in the data: OPEC conventional producers, non-OPEC conventional producers, and non-OPEC fracking producers. We split the sample into two subperiods: 1970–2010 and 2011–2019. In 2011, fracking producers surpassed for the first time a 1 percent market share in world oil production. *Gbbl* stands for billion barrels of oil.

cover roughly 14,000 oil fields operated by 3,200 companies across 109 countries. It includes information on both conventional and fracking oil fields. The latter include both tight-liquid and shale fields.

Our sample covers the period 1970-2019. The Rystad sample includes data prior to 1970. We excluded these data from our sample because it pertains to a period in which U.S. regulatory agencies sought to keep U.S. oil prices stable by setting production targets (see, e.g., [Hamilton 1983](#), [Kilian 2014](#), and [Fernandez-Villaverde 2017](#)).

We construct investment expenditures as the sum of exploration and capital expenditures. Exploration expenditures include costs from acquiring acreage, doing seismic analysis, and drilling wildcats or appraisal wells to discover and delineate oil fields. Capital expenditures include the costs of building facilities and drilling wells.

Production operating expenditures comprise operational expenses directly related to production: the lease cost and the fixed and variable costs. The latter include electricity, machinery, salaries, and tariffs. All expenditures are deflated with the U.S. GDP deflator.

Table 1 reports some key statistics for the periods 1970-2010 and 2011-2019. These statistics include the average values of production, the market share of OPEC and non-OPEC producers, the share of production of fracking producers, and the number of oil fields in operation.

3.2 Estimating technological parameters of oil producers

To estimate λ , we compute the lag between the first year of investment and the first year of production (T_i) for every oil field in our data set. If the arrival of production occurs according to

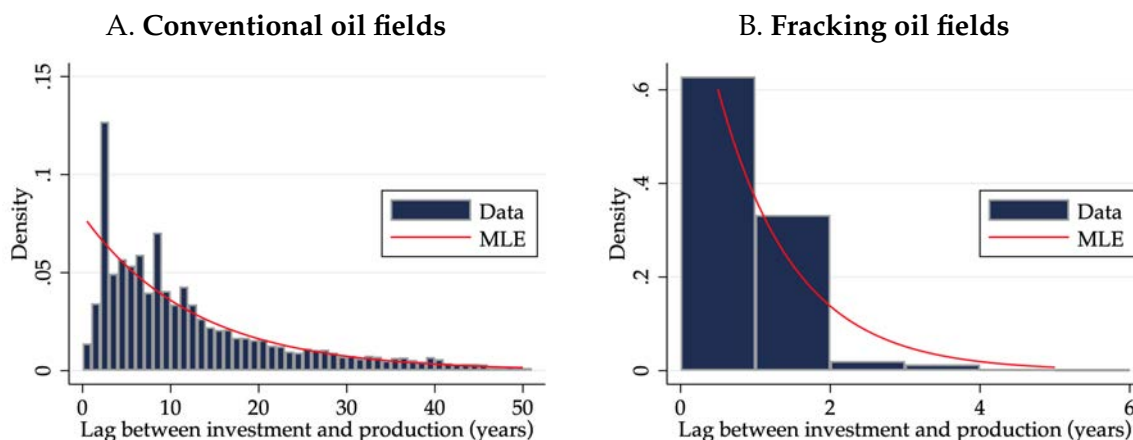
a Poisson process, the lag between investment and production follows a geometric distribution with mean λ . The maximum likelihood estimator of λ is the inverse of the average lag between investment and production

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

where N denotes the number of oil fields. Our estimates of λ are 0.08 for conventional oil fields and 0.98 for fracking fields. These estimates imply that the average lag between investment and production is roughly 12 years for conventional oil fields and one year for fracking fields. These estimates are consistent with textbook discussions of the lag between investment and oil production (see, e.g., chapter 1 of [Suicmez, Jing, Polikar, Allain, Pentland, and Dyson 2018](#)).

Figure 1 shows the empirical distribution of this lag, together with the implied geometric distribution for our estimate of $\hat{\lambda}$ for conventional and fracking fields.⁸

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 the first year of production across conventional (panel A) and fracking (panel B) oil fields. The MLE for λ , the Poisson arrival rate of production, is 0.08 for conventional oil fields and 0.98 for fracking fields. The red lines are the implied geometric distribution for the different estimated values of λ .

We also use our micro data to estimate η , the parameter that controls the convexity of extraction costs (see equation (8)). Since our data are collected at the oil field level, the elasticity of production costs with respect to the extraction rate is likely to mostly reflect changes in the pace at which

⁸Our estimate of the average production lag is higher than that reported in [Arezki, Ramey, and Sheng \(2016\)](#). 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, Ramey, and Sheng \(2016\)](#) estimate the lag between oil discovery and production, which is shorter.

additional wells are drilled.⁹ As [Anderson, Kellogg, and Salant \(2017\)](#) emphasize, drilling rigs and crews are a relatively fixed resource in the short run, so drilling faster to raise oil production increases costs of leasing equipment, hiring personnel, and so on. The parameter η represents these cost convexities.

Our estimate is based on the following regression:

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

where $C(\theta_{it}, R_{it})$ 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 detrended data. We chose the number of lags according to the Akaike information criterion. This procedure resulted in the selection of an AR(2) process.

Our data include all oil fields with positive extraction rates between 1970 and 2019. We exclude the last year of an oil field's operation because the cost data for this year include the costs of shutting down the field, which are not related to the rate of extraction.

Table 2 contains our slope estimates for conventional and fracking oil fields. All specifications include fixed effects for oil field and operation year (i.e., the age of the oil field). Specification 1 includes all oil fields in our sample. Specification 2 includes only conventional non-OPEC oil fields. Specification 3 includes only fracking non-OPEC oil fields. Specification 4 includes only OPEC oil fields.

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. For this reason, we use specification 2 as our benchmark and set η to the point estimate implied by this specification (6.1).¹⁰

⁹[Anderson, Kellogg, and Salant \(2017\)](#) show that production from individual oil wells does not respond to oil price changes. They also show that, in contrast, the number of wells that are drilled (rig activity) does respond to oil price changes.

¹⁰As a robustness check, we added royalties to our measure of operational costs and estimated η using this broader

Table 2: Extraction rate adjustment costs regression

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

Variable	(1)	(2)	(3)	(4)
ln(extraction)	6.04*** (0.34)	6.11*** (0.37)	2.23*** (0.75)	6.79*** (1.21)
Oil field FE	✓	✓	✓	✓
Operation year FE	✓	✓	✓	✓
Sample	All	Non-OPEC	Non-OPEC	OPEC
Technology	All	Conventional	Fracking	Conventional
IV	✓	✓	✓	✓
1 st stage F-stat	228	194	3	24
Clusters (oil fields)	13,506	11,103	939	1,464
Observations	303,457	260,552	8,819	36,086

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.

Specification 3 in Table 2 reports our estimates of η^F for fracking oil fields. This estimate is 2.23, so fracking oil fields have extraction costs that are significantly and substantially less convex than conventional oil producers.

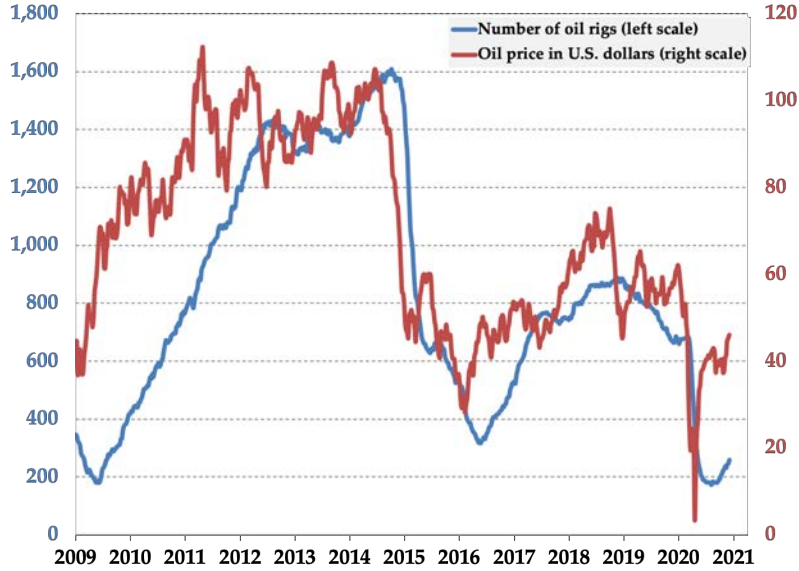
Our results are consistent with the findings of Bjornland, Nordvik, and Rohrer (2017) and Newell and Prest (2019). These authors use U.S. monthly well-level data to show that oil production from shale wells is much more flexible than conventional oil production.

Additional evidence 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. Figure 2 displays these data together with the nominal oil price. Between January 2009 and September 2014, oil prices rose from 42 dollars 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 dollars to 30 dollars per barrel. During this period, the number of oil rigs in operation fell from 1,600 to 400.

In sum, 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

cost measure. We obtain a slightly higher estimate: 7.5 with a standard error of 0.5.

Figure 2: 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.

conventional oil operations. Motivated by this evidence, we calibrate fracking firms so that the lag between investment and production is only one year and the extraction rate, θ_t^F , can be adjusted contemporaneously.

3.3 Calibrated parameters

We set α , the share of physical capital in value added, to $1/3$ and φ to 1.91, so that steady-state hours worked is equal to $1/3$. The parameter ψ matters only for the level of oil prices. We normalize it so that the price of oil, measured in units of output, equals one in the steady state. We set the household's discount factor, β , to 0.99, the coefficient of relative risk aversion, γ , to 2, and the Frisch elasticity of labor demand, ν , to 0.5. We set the annual rate of depreciation of physical capital, δ , to 0.1.

We first calibrate a version of our model with no fracking firms using data for the period 1970-2010. During this period, fracking represents less than 1 percent of the global oil supply. 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 (42 percent). The level of ξ and the total amount of land ($L^O + L^N$) are calibrated so that the steady-state extraction rates for OPEC and non-OPEC coincides with the average extraction rate in our data, 1.8 and 3.6 percent, respectively. We choose

the weight of oil in the production of final output, s_{or} , to match the average share of oil revenues in world GDP, 2.9 percent.

3.4 GMM estimation

We estimate a version of our model without fracking firms using data for the period 1970-2010. The parameters of the processes for productivity and supply shocks are estimated using GMM.¹¹ The first column of Table 4 presents the moments targeted in our estimation. Before we discuss the ability of our model to fit the empirical moments, 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 2010, 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 27 percent for nominal and real oil prices. In contrast, the annual standard deviation of nominal returns to the S&P 500 is roughly 16 percent, and the annual 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 statement “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 1970-2010 is 0.17. To put this number in perspective, this measure of volatility is 0.10 for U.S. manufacturing investment and 0.07 for U.S. aggregate investment.

The third fact is that investment in the oil industry is positively correlated with oil prices. The correlation between the growth rate of real oil prices and the growth rate of investment is 0.54. Table 3 reports this for each of the top 20 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 real 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 shocks to OPEC firms, which include deviations

¹¹Our weighting matrix is a diagonal matrix with diagonal elements equal to the inverse of the variance of the targeted moments.

Table 3: 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 between the growth rate of investment and the growth rate in oil prices for the top-20 oil producers, in descending order of production volume. In the table, x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$, and p_t and i_t represent the logarithm of the real price of oil and the logarithm of the firm's real investment, respectively.

Table 4: Data and model moments

	Moment	Data	(s.e.)	Model		Moment	Data	(s.e.)	Model
(1)	std(Δp_t)	0.27	(0.03)	0.22	(11)	corr($\Delta i_t^N, \Delta q_t^N$)	0.07	(0.12)	0.10
(2)	std(Δi_t^N)	0.18	(0.03)	0.16	(12)	corr($\Delta i_t^N, \Delta q_t^O$)	0.04	(0.13)	0.04
(3)	std(Δi_t^O)	0.17	(0.03)	0.21	(13)	corr($\Delta i_t^O, \Delta q_t^N$)	-0.11	(0.18)	0.08
(4)	std(Δq_t^N)	0.02	(0.003)	0.02	(14)	corr($\Delta i_t^O, \Delta q_t^O$)	-0.28	(0.13)	0.04
(5)	std(Δq_t^O)	0.07	(0.01)	0.08	(15)	corr($\Delta q_t^N, \Delta q_t^O$)	-0.07	(0.14)	0.40
(6)	corr($\Delta p_t, \Delta i_t^N$)	0.49	(0.14)	0.61	(16)	corr($\Delta p_t, \Delta p_{t-1}$)	-0.06	(0.09)	-0.31
(7)	corr($\Delta p_t, \Delta i_t^O$)	0.33	(0.11)	0.60	(17)	corr($\Delta i_t^N, \Delta i_{t-1}^N$)	0.13	(0.16)	0.05
(8)	corr($\Delta p_t, \Delta q_t^N$)	0.00	(0.09)	0.08	(18)	corr($\Delta i_t^O, \Delta i_{t-1}^O$)	0.29	(0.11)	0.04
(9)	corr($\Delta p_t, \Delta q_t^O$)	0.04	(0.14)	-0.69	(19)	corr($\Delta q_t^N, \Delta q_{t-1}^N$)	0.67	(0.10)	0.71
(10)	corr($\Delta i_t^N, \Delta i_t^O$)	0.67	(0.12)	1.00	(20)	corr($\Delta q_t^O, \Delta q_{t-1}^O$)	0.25	(0.20)	-0.31

Notes: This table presents the targeted moments from the data and the model-implied moments under the benchmark specification. Newey-West standard errors computed with five-year lags in parentheses. In the table, x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$, and the variables q_t^N and q_t^O represent the logarithm of oil production for non-OPEC and OPEC firms, respectively.

of the cartel from perfect commitment and disruptions such as the Iranian revolution and the Iran-Iraq war.

Table 4 compares the estimated moments targeted by our GMM procedure with the population moments implied by our model. We see that the model fit is relatively good given that we only have five parameters to target 20 moments.¹² As discussed in the introduction, the model is consistent with the high correlation between the real price of oil and real investment.

Table 5 reports our five parameter estimates. Our point estimate for ϵ is 0.15 with a standard error of 0.003. This estimate, which is similar to the ones reported in [Caldara, Cavallo, and Iacoviello \(2017\)](#) and [Hassler, Krusell, and Olovsson \(2019\)](#), implies that oil demand is very inelastic. A 1 percent increase in production reduces oil prices by 7 percent. To see why demand has to be inelastic to fit the data, it is useful to write the logarithm of the demand for oil by final good producers:

$$\ln p_t = \ln s_o + \frac{1}{\epsilon} \ln(Y_t) - \frac{1}{\epsilon} \ln(O_t).$$

Our model embeds two mechanisms that generate price volatility. The first mechanism is a low value of ϵ that makes low oil production volatility consistent with high price volatility. The second

¹²In addition, the model is consistent with the ratio of average investment to revenue in the data (10 percent), which is not targeted by the estimation procedure.

Table 5: Estimated parameters

Parameter	Estimate	(s.e.)
ϵ	0.15	(0.003)
ρ_1^a	1.72	(0.07)
ρ_2^a	-0.74	(0.08)
$std(\varepsilon_t^a)$	0.013	(0.002)
$std(u_t)$	0.05	(0.005)

Notes: This table presents the benchmark GMM estimates of the structural parameters.

mechanism is changes in the demand for oil associated with changes in final goods production induced by changes in A_t . The observed volatility and persistence of oil production and investment help determine the importance of the two channels that generate volatility in p_t and therefore identify ϵ . Volatile, persistent shocks to A_t 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.

The parameters of the process for A_t imply a hump-shaped impulse response function that peaks six years after the initial shock. This pattern is an important factor in driving investment volatility. It takes time for investment to increase oil production, so investment does not respond to short-lived shocks.

3.5 Identification

To discuss more formally our sources of identification, we use the method proposed by [Andrews, Gentzkow, and Shapiro \(2017\)](#). This method involves computing the elasticity of estimated parameters with respect to observed data moments. We report these elasticities in panel A of Table 6. Each row in this table shows how all the estimated model parameters would change if one data moment increased, holding all the other data moments constant. We also report the elasticities of model-implied moments with respect to the parameters in panel B of Table 6. Dark blue and red colors highlight large positive and negative elasticities, respectively.

To see how these tables can help us understand the sources of identification, consider the following question: why can't we increase the volatility of the shocks to supply and demand (σ_u^2 and σ_A^2) to bring the model-implied volatility of oil prices closer to the data? As we can see in the last two columns of panel B, increasing σ_u^2 and σ_A^2 at the same time by the same amount

Table 6: Sources of identification

A. Elasticity of parameters to moments						B. Elasticity of moments to parameters					
$\sigma^2(p)$	0.14	-0.17	-0.40	3.80	2.64	$\sigma^2(p)$	-0.19	-0.04	0.03	0.07	0.04
$\sigma^2(i_n)$	0.73	-0.44	-1.18	1.98	5.64	$\sigma^2(i_n)$	-0.00	6.93	-2.78	0.00	0.08
$\sigma^2(i_o)$	-0.58	0.09	0.19	-0.55	-0.94	$\sigma^2(i_o)$	-1.63	11.87	-4.90	0.00	0.11
$\sigma^2(q_n)$	-2.24	4.11	11.16	5.91	20.80	$\sigma^2(q_n)$	-0.01	-0.38	0.19	0.00	0.01
$\sigma^2(q_o)$	0.32	-0.05	-0.13	10.95	3.51	$\sigma^2(q_o)$	-0.01	-0.11	0.06	0.03	0.01
$\rho(p, i_n)$	0.03	-0.04	-0.10	-0.34	0.43	$\rho(p, i_n)$	0.03	0.86	-0.30	-0.19	0.19
$\rho(p, i_o)$	0.05	-0.07	-0.17	-0.58	0.74	$\rho(p, i_o)$	0.02	1.09	-0.41	-0.19	0.19
$\rho(p, q_n)$	0.03	0.14	0.31	0.01	-0.65	$\rho(p, q_n)$	0.61	13.81	-4.98	-0.02	0.02
$\rho(p, q_o)$	0.01	0.00	0.00	-0.41	0.21	$\rho(p, q_o)$	0.08	4.33	-1.28	-0.22	0.22
$\rho(i_n, q_n)$	-0.00	0.00	0.00	-0.00	-0.00	$\rho(i_n, q_n)$	-0.02	0.02	-0.01	-0.00	0.00
$\rho(i_o, q_n)$	-0.02	0.03	0.07	-0.02	-0.17	$\rho(i_o, q_n)$	0.02	-0.54	0.42	-0.00	0.00
$\rho(i_n, i_o)$	-0.01	0.01	0.03	-0.05	-0.05	$\rho(i_n, i_o)$	-0.04	-0.50	0.32	-0.02	0.02
$\rho(q_n, q_o)$	-0.01	0.02	0.04	-0.02	-0.10	$\rho(q_n, q_o)$	-0.27	-0.81	0.50	-0.00	0.00
$\rho(i_n, q_o)$	-0.01	0.01	0.04	-0.04	-0.09	$\rho(i_n, q_o)$	-0.17	-1.04	0.53	-0.01	0.01
$\rho(i_o, q_n)$	-0.02	0.02	0.05	-0.38	0.03	$\rho(i_o, q_n)$	-0.34	-3.13	1.75	-0.17	0.17
$\rho(p, p)$	0.06	0.00	-0.01	-0.36	0.29	$\rho(p, p)$	0.42	9.60	-3.59	-0.11	0.11
$\rho(i_n^o, i_n)$	-0.00	0.00	0.01	-0.01	-0.03	$\rho(i_n^o, i_n)$	-0.02	-1.61	0.70	-0.00	0.00
$\rho(i_o^o, i_o)$	-0.02	0.01	0.04	-0.04	-0.12	$\rho(i_o^o, i_o)$	-0.23	-1.46	0.63	-0.00	0.00
$\rho(q_n^o, q_n)$	-0.01	0.18	0.41	0.05	-0.98	$\rho(q_n^o, q_n)$	0.57	16.17	-5.65	-0.00	0.00
$\rho(q_o^o, q_o)$	-0.01	0.02	0.04	-0.16	-0.03	$\rho(q_o^o, q_o)$	-0.28	-0.45	0.77	-0.16	0.16
	ϵ	ρ_1^l	ρ_2^l	σ_u^2	σ_A^2		ϵ	ρ_1^l	ρ_2^l	σ_u^2	σ_A^2

Notes: The left panel presents the elasticity of estimated parameters to the targeted moments, following the approach of Andrews, Gentzkow, and Shapiro (2017). The right panel presents the elasticity of model-implied moments to the structural parameters of the model.

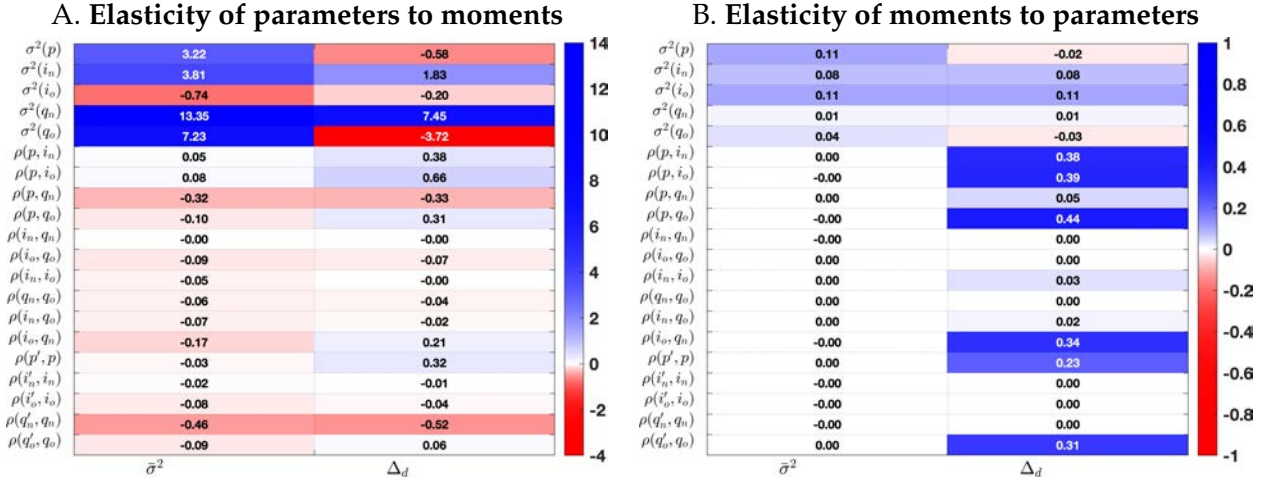
raises the volatility of oil prices and the volatility of investment and production in OPEC and non-OPEC while leaving the other moments basically unchanged. However, raising the volatility of production in both OPEC and non-OPEC worsens the model fit. Rows 4 and 5 of panel A show that if production of non-OPEC and OPEC were more volatile in the data, we would be able to raise σ_u^2 and σ_a^2 to increase the model-implied price volatility, bringing it closer to the data.

Another reason not to increase σ_u^2 is that, as element $(\rho(q_o^l, q_o), \sigma_u^2)$ of panel B shows, it would further reduce the first-order serial correlation of OPEC production, a moment that the model undershoots. Panel A shows that the serial correlation of OPEC production is indeed important in identifying σ_u^2 .

A similar argument applies to the identification of the elasticity of oil demand, ϵ . Panel B shows that lowering ϵ raises the volatility of oil prices, improving model performance along this dimension. However, lowering ϵ would also raise the volatility of OPEC investment and non-OPEC production, two moments that the model overshoots relative to the data. Rows 3 and 4 in the first column of panel A show that the volatility of OPEC investment and non-OPEC production are indeed important in identifying ϵ .

To further examine the forces that identify demand versus supply shocks, we construct a version of the last two columns of Table 6 written in terms of two transformed variables, $\bar{\sigma}^2$ and Δ_d .

Table 7: Sources of identification - demand vs. supply shocks



Notes: The left panel presents the elasticity of the estimated values of $\bar{\sigma}^2$ and Δ_d to the targeted moments, following the approach of [Andrews, Gentzkow, and Shapiro \(2017\)](#). The right panel presents the elasticity of model-implied moments with respect to these two structural parameters.

These variables are defined as follows:

$$\sigma_d^2 = \bar{\sigma}^2 \times \Delta_d,$$

$$\sigma_u^2 = \frac{\bar{\sigma}^2}{\Delta_d}.$$

A rise in $\bar{\sigma}^2$ increases the volatility of both demand and supply. A rise in Δ_d increases the importance of demand shocks relative to supply shocks.

Panel B of Table 7 displays the elasticity of model-implied moments to a change in either $\bar{\sigma}^2$ or Δ_d in isolation. This change affects multiple moments simultaneously. We see that increasing the importance of demand versus supply shocks (Δ_d) increases the volatility of non-OPEC production and reduces the volatility of OPEC production. These effects result from the fact that OPEC production volatility is primarily driven by supply shocks.

Panel A of Table 7 displays the elasticity of the parameters with respect to a change in a given data moment in isolation. This change affects multiple parameters simultaneously. We see that an increase in the volatility of OPEC production reduces the relative importance of demand shocks. In contrast, an increase in the volatility of non-OPEC production increases the relative importance of demand shocks. In sum, the relative volatility of OPEC and non-OPEC production is key in identifying the importance of supply versus demand shocks.

3.6 Other performance diagnostics

In this section, we evaluate whether our model is consistent with the estimates of the short-run elasticity of supply obtained from the micro data. This set of moments was not targeted by our estimation.

Estimating the short-run elasticity of oil supply. Oil producers can respond to an increase in the price of oil in two ways. The first is to increase the extraction rate in order to produce more oil from fields already in operation. The second is to increase the number of oil fields in operation. We show that the short-run elasticity of the extraction rate with respect to an exogenous change in the price of oil is positive but small. We also show that the elasticity 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.¹³ These estimates suggest that a rise in oil prices generates only a slight increase in the supply of oil from a given oil field.

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

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

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} represents other controls. 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 dynamics of oil fields discussed in [Arezki, Ramey, and Sheng \(2016\)](#) and [Anderson, Kellogg, and Salant \(2017\)](#).

Specification 1 in Table 8 is a simple OLS regression. The resulting slope coefficient estimate can be biased downward 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 2-4 use this instrument. Specifications 2, 3, and 4 include all oil fields, non-OPEC oil fields, and OPEC oil fields, respectively. We obtain estimates of β that range from 0.2 to 0.21. These estimates are broadly comparable to those obtained by [Caldara, Cavallo, and Iacoviello \(2017\)](#), by combining a narrative analysis of episodes of large drops in oil production with country-level instrumental variable regressions.

¹³Our data pertain to 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.

Table 8: Price elasticity of extraction rates

Dep. variable: ln(extraction rate)

Variable	(1)	(2)	(3)	(4)
ln(price)	0.05*** (0.005)	0.20*** (0.009)	0.21*** (0.010)	0.21*** (0.027)
IV	✗	✓	✓	✓
Sample	All	All	Non-OPEC	OPEC
Clusters (oil fields)	11,553	10,958	9,471	1,217
Observations	229,985	229,390	201,386	28,004

Notes: This table presents estimates of the elasticity of extraction rates with respect to the price of oil. The data used include all oil fields with positive extraction rates in 1970-2010, excluding the last year of operation. The regression specifications include oil field and operation year fixed effects as well as a year trend. Standard errors, clustered at the oil field level, are reported in parentheses. The instrument for price is the one-year-ahead forecast of detrended world real GDP. (***) - significant at a 1 percent level.

The following calculation is useful for evaluating the magnitude of this elasticity. The average extraction rate in our sample is 2.8 percent. A one standard deviation (27 percent) increase in the price of oil raises the extraction rate from 2.8 percent to 2.9 percent, resulting only in a 5 percent increase in production.

An oil field generally contains many oil rigs. Production increases can come from the intensive margin (higher production from existing rigs) or the extensive margin (drilling new rigs). Using a sample of Texas oil rigs, [Anderson, Kellogg, and Salant \(2017\)](#) show that the intensive-margin price elasticity is close to zero, so production increases come from the extensive margin. We conduct our analysis at the level of the oil field, so our supply elasticity encompasses both the intensive and the extensive margins.

Table 9 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 worldwide and the independent variable is the logarithm of real oil prices. Specification 2 uses our forecast of the cyclical component of world real GDP as an instrument for the logarithm of real oil prices. Specifications 3 and 4 use as dependent variables the number of non-OPEC fields and OPEC fields, respectively. All four specifications yield elasticity estimates that are statistically insignificant.

In Table 12, Appendix A, we report estimates of the elasticity with respect to oil prices of the

Table 9: Price elasticity of oil fields in operation

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

Variable	(1)	(2)	(3)	(4)
ln(price)	-0.12 (0.06)	-0.50 (0.34)	-0.28 (0.21)	-0.21 (0.14)
Year trend	✓	✓	✓	✓
IV	✗	✓	✓	✓
Dep. variable	All fields	All fields	Non-OPEC fields	OPEC fields
Observations	49	49	49	49

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 five-year lags are reported in parentheses. The instrument for price is the one-year-ahead forecast of detrended world real GDP.

number of new oil fields that are producing and the number of oil fields that are retiring. Both of these elasticities are statistically insignificant. 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.

We now discuss the implications of our model for the regressions reported in Table 8. We simulate data from our model and run regression 8 using the one-year-ahead forecast of world real GDP as an instrument for the price of oil. In our model, the elasticity of response of extraction rates to changes in oil prices is 0.24, which is close to our empirical estimates reported in Table 8.

In estimating the structural parameters of our model, we target the unconditional correlation between oil prices and oil production. This correlation, which is affected by both demand and supply shocks, is close to zero in the data. The instrumental variable regression reported above isolates the impact of demand shocks. For this reason, this regression focuses on an aspect of the empirical performance of the model that we did not target in our estimation procedure.

4 The impact of fracking

In this section, we study the long-run impact of fracking on the level and volatility of oil prices and world output. We then discuss the transition dynamics induced by the advent of fracking.

4.1 Calibrating the model with fracking

In the previous section, we estimate the values of λ^F and η^F using micro data. Rystad Energy forecasts that fracking fields will represent 20 percent of oil production by 2050. To construct the steady state with fracking, we calibrate L^F so that the production of fracking firms represents 20 percent of global oil production. We calibrate ψ^F so that in the steady state with fracking, the average extraction cost per barrel of oil is the same for fracking and conventional non-OPEC firms.

4.2 Long-run effects of fracking

We now study the properties of the model with fracking by linearizing the model's equilibrium equations around the steady state with fracking. Table 10 compares the implications of versions of the model with and without fracking for some key moments.

Table 10: Implications of fracking for key aggregate moments

Moment	No fracking	Fracking
Average p_t		-46%
Average y_t		+5%
Volatility of Δp_t	0.22	0.12
Volatility of Δy_t	0.035	0.036
OPEC's market share	42.3%	40.1%

Notes: This table presents the aggregate long-run impact of the advent of fracking. In the table, x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$.

Fracking has a large impact on oil prices. The price of oil is 46 percent lower in the steady state with fracking than in the no-fracking steady state. These lower oil prices raise world output by 5 percent. At the same time, the share of OPEC in oil production falls from 42.3 to 40.1 percent.

The oil price decline associated with fracking results from two forces. The first is that fracking weakens OPEC's cartel power. The second is that fracking expands oil supply. To understand the relative importance of these two forces, we study the impact of fracking in a version of our model calibrated with the same parameters where OPEC firms are competitive. In this competitive model, oil prices are 17 percent lower and world output 0.5 percent higher in the steady state with fracking when compared with the no-fracking steady state. These results suggest that more than one-third of the decline in oil prices can be attributed to the decline in the cartel's market power

and the other half to the expansion in oil supply. The importance of OPEC in the determination of oil prices is consistent with the results in [Asker, Collard-Wexler, and De Loecker \(2019\)](#).

In response to the advent of fracking, world output increases much more when OPEC is a cartel (5 percent) than when all firms are competitive (0.5 percent). The intuition for this result is as follows. In the no-fracking steady state, the supply of oil is much lower in a model in which OPEC is a cartel. As a result, the marginal product of oil is higher in the cartel model. This higher marginal product, combined with a higher increase in oil supply, produces a larger expansion in world output in response to fracking.

In the model, fracking reduces the long-run volatility of oil prices and world oil production by 43 and 33 percent, respectively. At the same time, the volatility of world real GDP *rises* by 2.9 percent. This rise occurs because fracking allows the economy to be more responsive to productivity shocks. Without fracking, a positive productivity shock raises oil prices, which dampens the effect of the productivity shock. With fracking, oil supply is more elastic, which amplifies the effect of productivity shocks.

The correlation between oil prices and oil production, as well as between oil investment and oil production, is higher in the model with fracking, reflecting the response of fracking firms to high-frequency movements in prices. Oil production is less volatile in the version of the model with fracking. This result reflects two opposing effects. The aggregate output response to *demand* shocks is higher in the model with fracking, as fracking firms are more nimble. The aggregate output response to *supply* shocks is lower in the model with fracking. Fracking firms can respond within the period to supply shocks, smoothing the response of oil prices to these shocks.

Production by fracking firms exceeded, for the first time, 1 percent of global production in 2011. Using this date to mark the beginning of the fracking period leaves us with too few time-series observations to draw firm empirical conclusions about the impact of fracking. But even though we are still in the beginning of the fracking expansion, we can find some echoes of the model's implications in the data. The volatility of real oil prices, measured by the annual standard deviation of percentage changes in real oil prices, is 5 percent lower in the 2011–2019 period when compared to the 1970–2010 period. The volatility of world oil production, measured by the standard deviation of annual percentage change in oil production, is 60 percent lower in the period 2011–2019 when compared to the period 1970–2010.

4.2.1 Welfare gains from fracking

In this subsection, we discuss the long-run welfare gains from fracking. Fracking affects economic welfare in two ways. First, it lowers average oil prices, which results in higher average world output. Second, fracking affects the volatility of the economy. As discussed above, oil prices are more stable in an economy with fracking. This stability results in a larger response of world output to productivity shocks and a larger cost of business cycles.

The welfare gain from the decline in oil prices associated with fracking is large: it is equivalent to a permanent increase in consumption of 0.9 percent. The welfare cost of business cycle fluctuations is higher by 0.05 percent of equilibrium consumption in an economy with fracking. So, the net welfare gain from fracking is equivalent to a permanent 0.85 percent increase in consumption. An important caveat to these calculations is that they abstract from the environmental costs of both conventional oil production and fracking.¹⁴

4.3 Impulse response functions

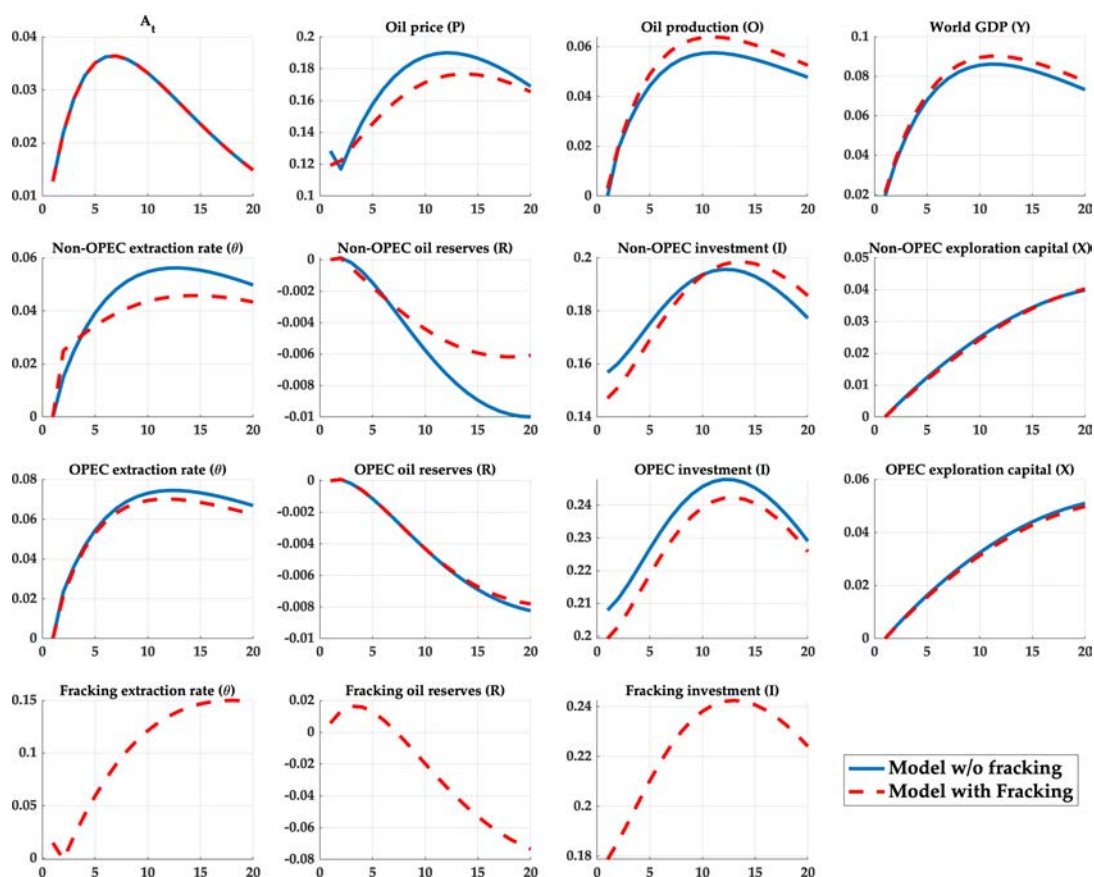
We can use impulse response functions to better understand the impact of fracking on the response of the economy to shocks. Figures 3 and 4 show the impulse responses for productivity and supply shocks, respectively. The dashed and solid lines pertain to the model with and without fracking, respectively.

The solid blue line in Figure 3 depicts the impulse response function to a one standard deviation shock to A_t for the model with no fracking firms. The shock follows a hump-shaped pattern with a peak around year 7. On impact, firms cannot change their extraction rates, so the price increases. In the following periods, the price of oil continues to rise, but this rise is moderated by an increase in extraction rates by both OPEC and non-OPEC firms. Production rises and reserves are depleted. Since the shock is very persistent, investment rises to increase future production and take advantage of the extended period of high oil prices.

Consider now the economy with fracking. The overall response is similar with and without fracking. The impulse response with fracking is the dashed red line in Figure 3. Recall that the costs of increasing the extraction rate and the lag between investment and oil production are smaller for fracking firms than for non-fracking firms. For this reason, fracking firms increase their extraction rate by more than non-fracking firms. The short lag between investment and oil productions allows fracking firms to increase their reserves in the first five years. World oil production

¹⁴Barnett, Brock, and Hansen (2020) propose a framework for studying this environmental impact.

Figure 3: Impulse response to a demand shock

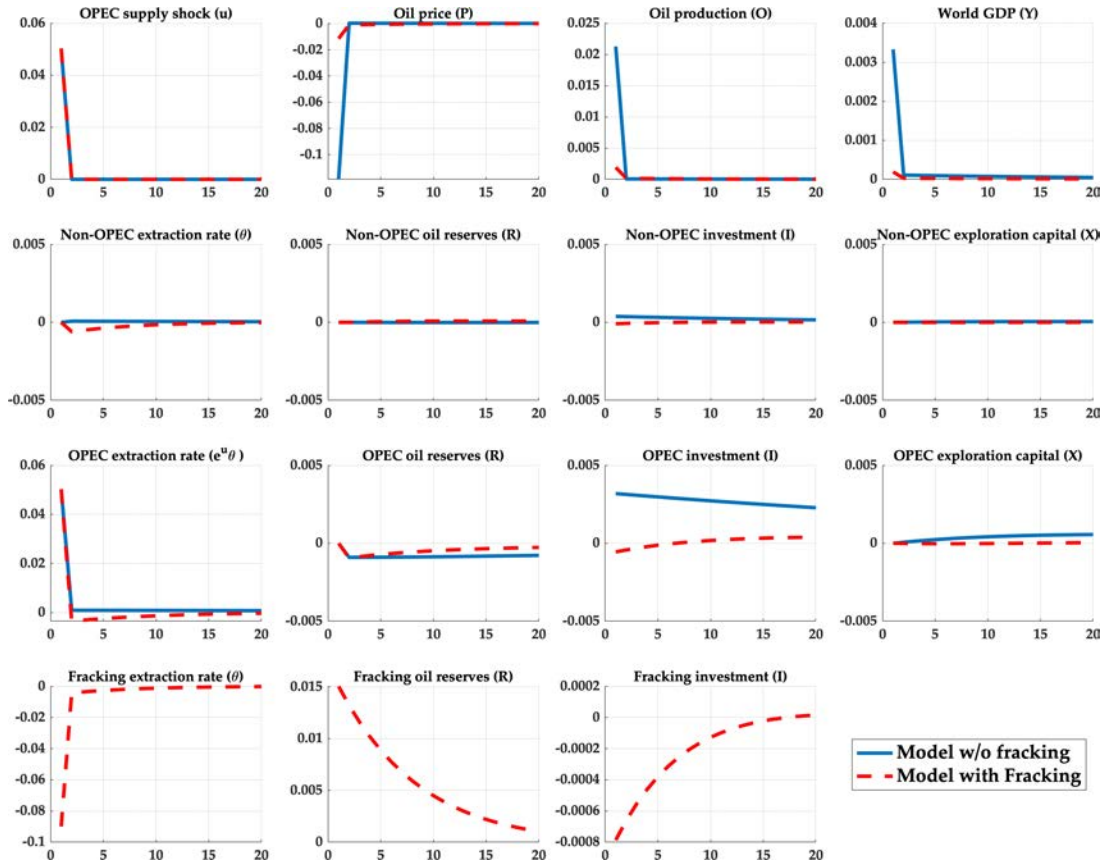


Notes: This figure presents the impulse response functions to a one standard deviation demand shock (A_t) in the model with and without fracking. The red and blue lines correspond to the response with and without fracking, respectively.

risers by more and oil prices by less in the economy with fracking. The smaller rise in oil prices implies a stronger positive effect of the shock to A_t on world output in the economy with fracking.

Consider the response of the economy without fracking to a positive one standard deviation supply shock. This response is represented by the solid blue line in Figure 4. Recall that OPEC has imperfect control over its extraction rate. This rate can be affected by positive shocks (e.g. deviations from commitment) or negative shocks (e.g. wars or other disturbances). Since non-fracking firms can only adjust their extraction rate with a one-period lag, they cannot cut production in response to this shock. As a result, there is a temporary fall in oil prices and an expansion in world output.

Figure 4: Impulse response to a supply shock



Notes: This figure presents the impulse response functions to a one standard deviation OPEC supply shock in the model with and without fracking. The red and blue lines correspond to the response with and without fracking, respectively.

Consider now the economy with fracking. Fracking firms cut production in response to the OPEC supply shock, reducing the fall in oil prices by tenfold and thus reducing oil price volatility.

4.4 Demand and supply shocks

We now use our model to answer a classic question: what is the role of demand and supply shocks in driving oil industry fluctuations? Table 11 reports the model-implied variance decomposition for oil prices, production, and investment with respect to supply and demand shocks.

Consider first the model without fracking. Both demand and supply shocks are quantitatively important. Table 11 shows that eliminating demand shocks lowers the volatility of oil prices by 39 percent. This fall in price volatility is consistent with the importance of macroeconomic perform-

ance in driving oil prices emphasized in Barsky and Kilian (2001) and Barsky and Kilian (2004). But in our model, supply shocks are more important than demand shocks as drivers of oil price volatility. Eliminating supply shocks reduces the volatility of oil prices by 61 percent.

Table 11 also shows that the volatility of investment for both OPEC and non-OPEC is predominantly driven by demand shocks. These shocks are long-lived, and so they elicit a large response of investment (see Figure 3). In contrast, supply shocks are temporary, so they have little impact on investment (see Figure 4).

Table 11: Variance decomposition

Moment	No fracking		Fracking	
	Demand	Supply	Demand	Supply
Δp_t	39.0%	61.0%	98.4%	1.6%
Δi_t	99.9%	0.01%	99.99%	0.01%
Δq_t	42.8%	57.2%	99.1%	0.9%

Notes: This table presents the variance decomposition of five key variables to demand and supply shocks in the model without fracking and the one with fracking. In the table, x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$.

The volatility of production by OPEC firms is dominated by supply shocks. In contrast, demand shocks drive the volatility of production by non-OPEC firms. This result reflects the temporary nature of supply shocks. Since extraction rates are chosen one period in advance, non-OPEC cannot respond to OPEC supply shocks. Overall, the relative importance of demand and supply in production is similar to that in prices: supply and demand shocks account for roughly 60 and 40 percent of the variance, respectively.

Consider now fluctuations around the steady state where fracking accounts for 20 percent of the oil supply. Fracking firms can respond to supply shocks, neutralizing their effects. As a result, supply shocks are not important drivers of the volatility of prices, aggregate oil production, and investment.

4.5 Transitional dynamics to steady state with fracking

In this section, we study the transition from a steady state without fracking to one where fracking accounts for 20 percent of oil production. The values of the state variables are considerably different in these two steady states. For this reason, characterizing transition dynamics by linearizing

the model around one of the two steady states produces a poor approximation to the dynamics of the true nonlinear system.

To study the transition, we solve the full nonlinear system. Because we are interested in the trend behavior of different variables, we abstract from shocks. We assume, as an approximation, that the economy converges to the new steady state in 300 years. Under this assumption, the transition dynamics are the solution to a large nonlinear system of equations ($48 \times 300 = 14,400$ equations and unknowns). Working with these equations, we managed to reduce the dimension of the system. The resulting nonlinear system is still large but much more manageable ($5 \times 300 = 1,500$ equations and unknowns). We describe the model derivations and the algorithm used to solve the nonlinear system in Appendix B.

We assume that land available for fracking, L_t^F , expands throughout the transition period in a way that makes the model broadly consistent with the average market share of fracking in the period 2011-2019 and the evolution of Rystad Energy's prediction for the average decadal market share of fracking in the period 2020-2050. In order to match these market share dynamics, we assume that L_t^F follows a deterministic AR(1) process ($(L_{t+1}^F - \bar{L}^F) = \rho(L_t^F - \bar{L}^F)$) with an initial value of zero. The terminal value, \bar{L}^F , is the land available for fracking in the terminal steady state. We set ρ equal to 0.88. This value of ρ minimizes the sum of the square differences between the market shares implied by the model and the realized and forecasted average decadal market shares for the period 2011-2050. In order to keep the numerical error associated with solving the nonlinear system of equations low, we start the transition at a point where fracking reserves are equal to 0.005 instead of zero. Figure 5 displays the data and model-implied market shares of fracking.¹⁵

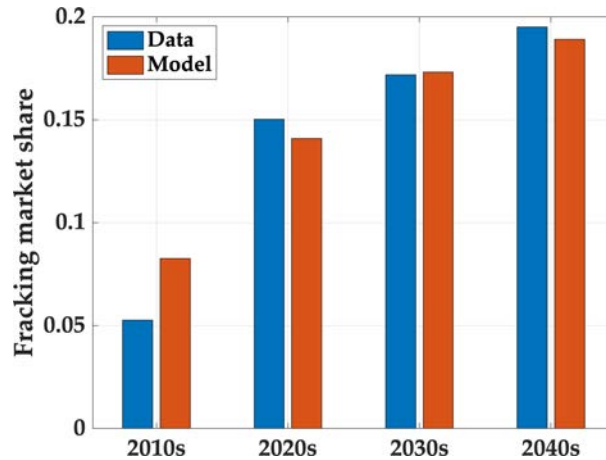
Solving for the transition requires taking a stand on the initial values of the Lagrange multipliers associated with OPEC's implementability constraints. We assume these initial Lagrange multipliers are equal to zero. This assumption implies that OPEC has a one-time opportunity to ignore previous commitments and reoptimize its plan going forward.¹⁶

Figure 6 plots the transition dynamics for oil production, oil prices, world output, and market

¹⁵We implemented an alternative specification in which ψ follows a deterministic AR(1) process. We interpret variations in ψ as reflecting technical progress in fracking. This specification does not do as well at matching the data displayed in Figure 5 as the specification described in the main text.

¹⁶We also computed an alternative solution in which the initial Lagrange multipliers on OPEC's implementability conditions are equal to the value of the Lagrange multipliers in the first period of the transition. Under this assumption, if the initial state variables were equal to their values in the fracking steady state, the economy would remain in this steady state. The transition dynamics are broadly similar to the ones displayed in Figure 6.

Figure 5: Transition dynamics: fracking production share



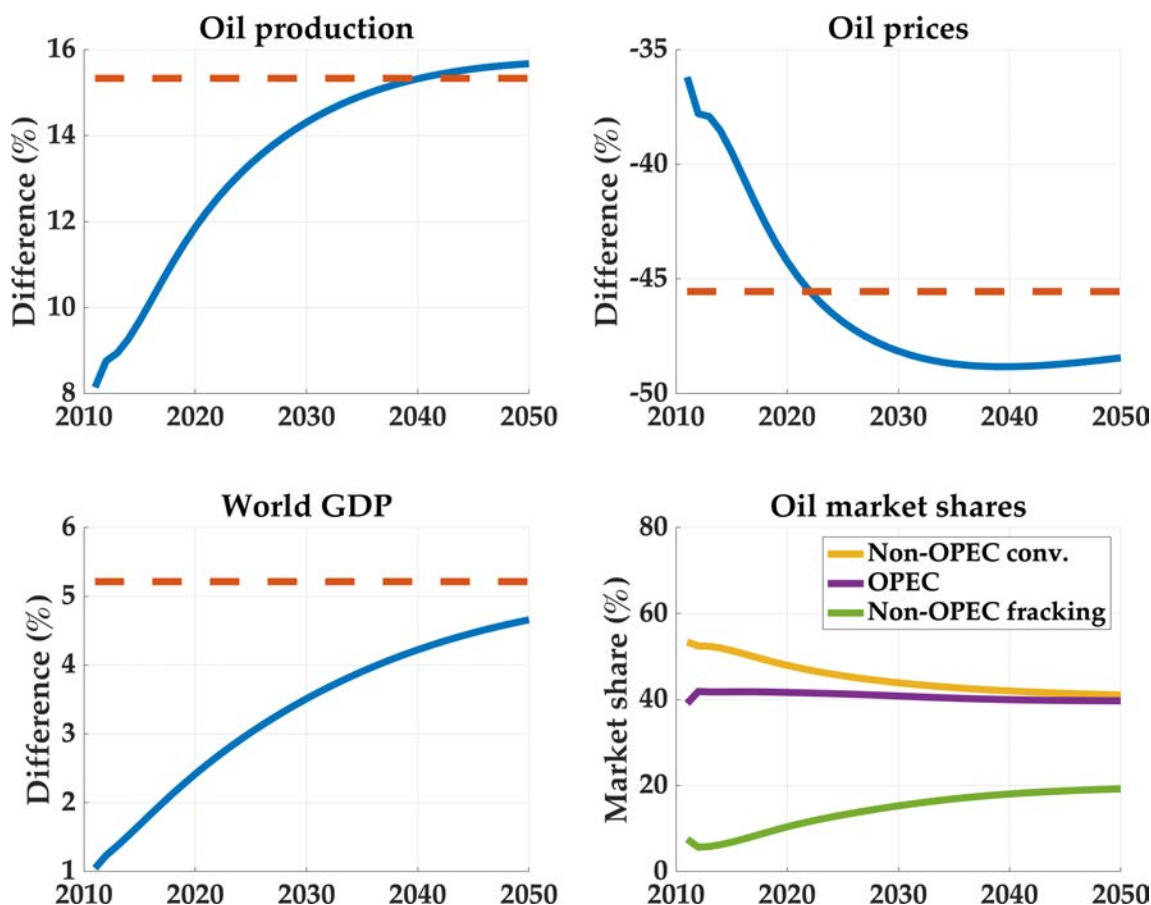
Notes: This figure presents the average decadal production share of fracking oil fields in the model and the data (realized and forecasted by Rystad Energy).

shares of OPEC, non-OPEC conventional oil producers, and non-OPEC fracking firms. The upper left panel shows that oil production rises in the short run by 8 percent and continues to increase over time. The upper right panel shows that there is a sizable short-run decline in oil prices (35 percent). This decline results from the expansion in production combined with the low elasticity of oil demand. The initial decline in oil prices boosts world output. Some of this increase in output is invested, increasing future demand for oil. However, this demand increase is outstripped by the increase in oil supply triggered by the expansion of fracking. As a result, oil prices fall throughout the transition, undershooting their level in the long-run steady state. This undershooting reflects the fact that OPEC has a one-time opportunity to ignore previous commitments in response to the unanticipated advent of fracking.¹⁷ The lower left panel shows that world output expands by 1 percent in the short run and continues to rise as fracking expands.

The lower right panel shows the market shares of OPEC, conventional non-OPEC producers, and fracking firms. There are some short-run movements in market shares mostly driven by the assumption that the extraction rates of conventional oil producers are chosen one period in advance. The key result in this panel is that, throughout the transition, most of the rise in the market share of fracking is compensated by a decline in the market share of conventional non-OPEC oil producers. OPEC manages to hold on to its market share by increasing its level of oil production.

¹⁷When the initial OPEC multipliers on implementability conditions are positive, this undershooting is much weaker.

Figure 6: Transition dynamics following the advent of fracking



Notes: This figure displays the deterministic transition dynamics following the introduction of fracking. The first three panels display percentage deviations from the no-fracking steady state. The bottom right panel displays the market shares of non-OPEC conventional producers, OPEC producers, and non-OPEC fracking producers.

Interestingly, we see this pattern emerging in the data. Consider first the evolution of OPEC's market share. Between 1970 and 2010, the average share of OPEC is 42 percent. By 2019, this share had declined by only 2 percent, to 40 percent. Consider next the evolution of the market share of conventional non-OPEC producers. Between 1970 and 2010, this share is on average 58 percent. By 2019, this share had declined by 7 percent, to 51 percent.¹⁸

The relative stability of OPEC's market share cannot be accounted for by a model in which

¹⁸We obtain similar results if we HP-filter the data to eliminate cyclical fluctuations. By 2019, shares of OPEC and non-OPEC conventional producers drop by 1 and 8 percent, respectively.

all firms are competitive. In such a model, both OPEC market shares and the market share of non-OPEC conventional oil producers decline in response to the advent of fracking.

5 Robustness

In subsection 5.1, we discuss two robustness exercises. First, we consider a competitive version of the model. In this model, all oil producers, including those part of OPEC, act as perfect competitors. Second, we consider a version of the cartel model in which a fraction of OPEC firms deviate from the cartel's decisions in order to maximize their individual profit.

In subsection 5.2, we consider a version of the model with supply shocks to non-OPEC firms. In addition, we discuss a version of the model with shocks to the cost of investment rather than shocks that directly affect oil supply.

5.1 Market structure

In this subsection, we discuss two alternative versions of the market structure in the oil industry. In the first version, OPEC firms act competitively. In the second version, a fraction of cartel firms deviate from the cartel's prescribed policy and act to maximize their individual profit.

5.1.1 Competitive model

In the competitive version of the benchmark model, each OPEC firm maximizes its objective (19), subject to the law of motion for reserves (20) and exploration capital (21). We follow the same calibration strategy as in the benchmark model with one exception. Since in the competitive version of the model, OPEC and non-OPEC firms choose the same extraction rate, we cannot match the individual rates of each individual group. We choose the total amount of land to match the average extraction rate across OPEC and non-OPEC. We are left with one degree of freedom, so we choose the value of ξ to be the same as in the benchmark model.

We reestimate the structural parameters of this model targeting the same set of moments used for the benchmark estimation. Table 13, Appendix A, presents the resulting parameter estimates. We see that the parameter estimates are relatively similar to those of the benchmark model. The most salient difference is that ε is 0.135 instead of 0.15, so the demand for oil is more inelastic. The fit of the competitive model is fairly similar to that of the benchmark model. Table 14, Appendix A, reports these moments.

We use the estimated competitive model to study the long-run impact of fracking. The calibration of the fracking sector follows the same approach used in the benchmark model. In the steady state with fracking, the price of oil is 35 percent lower and the oil price volatility is 47 percent lower than in the steady state without fracking. Recall that the analogous results for the benchmark model are 46 and 43 percent, respectively. In sum, both the cartel and the competitive model imply that fracking results in a large decline in the level and volatility of oil prices. However, the cartel model predicts a larger decline in the level of oil prices than the competitive model.

Under the competitive specification, the rise in fracking market share is compensated by a proportional decline in non-OPEC and OPEC conventional producers. The market share of non-OPEC conventional producers falls from 58 to 46 percent, and that of OPEC falls from 42 to 34 percent. These market share declines are starkly different from those in our benchmark model. When OPEC acts as a cartel, the advent of fracking reduces non-OPEC's market share from 58 to 40 percent and OPEC's market share from 42 to 40 percent. The modest decline in OPEC's market share observed during the 2011–2019 period is consistent with our benchmark specification but not with the competitive version of the model.

5.1.2 Model in which a fraction of cartel firms deviates

In this subsection, we consider a model in which a fraction of OPEC firms abandon the policy prescribed by the cartel and, instead, maximize their individual profits. This extended model encompasses the two market structures we have considered so far: (i) OPEC acts as a cartel with full commitment, and (ii) all firms act competitively.

We follow the same estimation procedure used for the benchmark model. The extended model has an additional parameter to be estimated: the fraction of land owned by firms that deviate. Tables 15 and 16, Appendix A, display the estimated parameters as well as data and model-implied moments. We estimate the fraction of OPEC members that deviate to be 36 percent and estimate their production share to be 40 percent of total OPEC production. However, both of these shares are imprecisely estimated. This imprecision is not surprising given that the empirical performance of the cartel and the competitive model is relatively similar empirical performance.

The performance of the version of the model with firms that deviate is slightly better than the performance of the benchmark model. In particular, this model generates higher oil price volatility.

5.2 Supply shocks

In this subsection, we discuss versions of our model with supply shocks to non-OPEC firms as well as a specification with investment cost shocks.

5.2.1 Supply shocks to non-OPEC producers

Our benchmark model assumes that only OPEC firms are subject to supply shocks. In this subsection, we consider a specification with supply shocks to non-OPEC producers. The only difference between this model and the benchmark model is in the maximization problem for non-OPEC producers. This problem is to maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t M_t \left[p_t e^{u_t^N} \theta_t^N R_t^N - I_t^N - \psi \left(e^{u_t^N} \theta_t^N \right)^\eta R_t^N \right],$$

subject to the law of motion for oil reserves

$$R_{t+1}^N = \left(1 - e^{u_t^N} \theta_t^N \right) R_t^N + \lambda X_{t+1}^N,$$

and exploration capital, (9). We model u_t^N as an i.i.d. process. Recall that producers can offset any forecastable movements in u_t^N , so introducing persistence in the process for u_t^N does not improve the performance of the model.

The estimation algorithm sets the volatility of u_t^N to a value close to zero (0.007). The performance of the model is very similar to our benchmark specification (see Tables 17 and 18 in Appendix A). Recall that OPEC production is more volatile than non-OPEC production. Introducing supply shocks to non-OPEC makes, other things equal, non-OPEC production more volatile, worsening the performance of the model.

5.2.2 Investment shocks

We now consider a version of the model in which both OPEC and non-OPEC producers are subject to investment shocks instead of supply shocks. The objective functions for OPEC and non-OPEC are

$$\begin{aligned} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t M_t \left[\theta_t^O R_t^O p_t - \zeta_t^O I_t^O - \psi \left(\theta_t^O \right)^\eta R_t^O \right], \\ \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t M_t \left[\theta_t^N R_t^N p_t - \zeta_t^N I_t^N - \psi \left(\theta_t^N \right)^\eta R_t^N \right], \end{aligned}$$

where ζ_t^O and ζ_t^N are investment shocks to OPEC and non-OPEC, respectively. We assume that these shocks follow orthogonal AR(2) processes.

In Tables 19 and 20 in Appendix A, we display our parameter estimates as well as the data and model-implied moments. Despite having five additional parameters, this model's performance is worse than that of the benchmark model. In particular, the volatility of prices and OPEC production generated by the model with investment shocks is roughly 50 percent lower than in the data.

6 Conclusion

Using micro data on oil fields, we show that fracking firms are much more nimble than conventional oil producers. The costs of fracking firms are less convex in the extraction rate than the costs of conventional firms. In addition, the lag between investment and production is much shorter for fracking firms than for conventional oil producers.

We estimate a model of the world economy in which OPEC acts as a cartel. We use the model to study the impact of the advent of fracking. Consistent with the data, our model implies that the rise in the market share of fracking firms is compensated almost exclusively by a decline in the market share of non-OPEC conventional oil producers. OPEC maintains its market share by increasing production.

Fracking affects the world economy in three ways. First, the volatility of oil prices falls because the supply of oil becomes more elastic. Second, the volatility of world output rises because the economy becomes more responsive to oil demand shocks. Without fracking, a positive demand shock implies a larger rise in oil prices. This rise dampens the effect of the demand shock on the economy. Third, the average level of oil prices falls because the oil production of both fracking firms and OPEC increases.

Overall, the impact of fracking on welfare is equivalent to increasing consumption by 0.85 percent in every period. An important caveat to this welfare calculation is that it abstracts from the environmental impact of both conventional oil production and fracking. Evaluating these environmental welfare costs in a general equilibrium model of oil production is an interesting topic for future research.

References

Adao, Bernardino, Borghan Narajabad, and Ted P Temzelides, "Renewable technology adoption and the macroeconomy," 2017.

- Anderson, Soren T, Ryan Kellogg, and Stephen W Salant**, “Hotelling Under Pressure,” *Journal of Political Economy*, 2017.
- Andrews, Isaiah, Matthew Gentzkow, and Jesse M Shapiro**, “Measuring the sensitivity of parameter estimates to estimation moments,” *The Quarterly Journal of Economics*, 2017, 132 (4), 1553–1592.
- Arezki, Rabah, Valerie A Ramey, and Liugang Sheng**, “News Shocks in Open Economies: Evidence from Giant Oil Discoveries,” *The Quarterly Journal of Economics*, 2016.
- Asker, John, Allan Collard-Wexler, and Jan De Loecker**, “(Mis) allocation, market power, and global oil extraction,” *American Economic Review*, 2019, 109 (4), 1568–1615.
- Backus, David K and Mario J Crucini**, “Oil prices and the terms of trade,” *Journal of international Economics*, 2000, 50 (1), 185–213.
- Barnett, Michael, William Brock, and Lars Peter Hansen**, “Pricing uncertainty induced by climate change,” *The Review of Financial Studies*, 2020, 33 (3), 1024–1066.
- Barsky, Robert B and Lutz Kilian**, “Do we really know that oil caused the great stagflation? A monetary alternative,” *NBER Macroeconomics annual*, 2001, 16, 137–183.
- and —, “Oil and the Macroeconomy since the 1970s,” *The Journal of Economic Perspectives*, 2004, 18 (4), 115–134.
- Baumeister, Christiane and James D Hamilton**, “Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks,” *American Economic Review*, 2019, 109 (5), 1873–1910.
- Bjornland, Hilde C., Frode Martin Nordvik, and Maximilian Rohrer**, “Supply Flexibility in the Shale Patch: Evidence from North Dakota,” *manuscript, Norges Bank*, 2017.
- Blanchard, Olivier J and Jordi Gali**, “The Macroeconomic Effects of Oil Shocks: Why are the 2000s so different from the 1970s?,” Technical Report, National Bureau of Economic Research 2007.
- Bodenstein, Martin, Christopher J Erceg, and Luca Guerrieri**, “Oil shocks and external adjustment,” *Journal of International Economics*, 2011, 83 (2), 168–184.

- Caldara, Dario, Michele Cavallo, and Matteo M Iacoviello**, “Oil price elasticities and oil price fluctuations,” 2017.
- Chari, V.V. and Lawrence Christiano**, “The Optimal Extraction of Exhaustible Resources,” Technical Report, Federal Reserve Bank of Minneapolis 2014.
- Deaton, Angus**, “Commodity prices and growth in Africa,” *The Journal of Economic Perspectives*, 1999, 13 (3), 23–40.
- **and Guy Laroque**, “On the behaviour of commodity prices,” *The Review of Economic Studies*, 1992, 59 (1), 1–23.
- Ezekiel, Mordecai**, “The Cobweb Theorem,” *The Quarterly Journal of Economics*, 1938, 52 (2), 255–280.
- Fernandez-Villaverde, Jesus**, “Energy: The Mover of Output in Global Economic History,” Technical Report, University of Pennsylvania 2017.
- Finn, Mary G**, “Perfect competition and the effects of energy price increases on economic activity,” *Journal of Money, Credit and banking*, 2000, pp. 400–416.
- Gilje, Erik, Robert Ready, and Nikolai Roussanov**, “Fracking, Drilling, and Asset Pricing: Estimating the Economic Benefits of the Shale Revolution,” Technical Report, National Bureau of Economic Research 2016.
- Hamilton, James D**, “Oil and the Macroeconomy since World War II,” *Journal of Political Economy*, 1983, 91 (2), 228–248.
- Harvey, David I, Neil M Kellard, Jakob B Madsen, and Mark E Wohar**, “The Prebisch-Singer hypothesis: four centuries of evidence,” *The Review of Economics and Statistics*, 2010, 92 (2), 367–377.
- Hassler, John, Per Krusell, and Conny Olovsson**, “Oil monopoly and the climate,” *The American Economic Review*, 2010, 100 (2), 460–464.
- , – , **and –** , “Directed technical change as a response to natural-resource scarcity,” Technical Report, Sveriges Riksbank Working Paper Series 2019.
- Herrera, Ana Maria and Sandeep Kumar Rangaraju**, “The effect of oil supply shocks on US economic activity: What have we learned?,” *Journal of Applied Econometrics*, 2020, 35 (2), 141–159.

- Hotelling, Harold**, "The economics of exhaustible resources," *Journal of Political Economy*, 1931, 39 (2), 137–175.
- Kellogg, Ryan**, "The effect of uncertainty on investment: evidence from Texas oil drilling," *The American Economic Review*, 2014, 104 (6), 1698–1734.
- Kilian, Lutz**, "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *The American Economic Review*, 2009, 99 (3), 1053–1069.
- , "Oil price shocks: causes and consequences," *Annu. Rev. Resour. Econ.*, 2014, 6 (1), 133–154.
- **and Daniel P Murphy**, "The role of inventories and speculative trading in the global market for crude oil," *Journal of Applied Econometrics*, 2014, 29 (3), 454–478.
- Kim, In-Moo and Prakash Loungani**, "The role of energy in real business cycle models," *Journal of Monetary Economics*, 1992, 29 (2), 173–189.
- Leduc, Sylvain and Keith Sill**, "A quantitative analysis of oil-price shocks, systematic monetary policy, and economic downturns," *Journal of Monetary Economics*, 2004, 51 (4), 781–808.
- Lippi, Francesco and Andrea Nobili**, "Oil and the macroeconomy: a quantitative structural analysis," *Journal of the European Economic Association*, 2012, 10 (5), 1059–1083.
- Melek, Nida Çakir, Michael Plante, and Mine K Yücel**, "The US Shale Oil Boom, the Oil Export Ban, and the Economy: A General Equilibrium Analysis," Technical Report, National Bureau of Economic Research 2017.
- Nerlove, Marc**, "Adaptive expectations and cobweb phenomena," *The Quarterly Journal of Economics*, 1958, 72 (2), 227–240.
- Newell, Richard G and Brian C Prest**, "The unconventional oil supply boom: Aggregate price response from microdata," *The Energy Journal*, 2019, 40 (3).
- Rotemberg, Julio J and Michael Woodford**, "Imperfect Competition and the Effects of Energy Price Increases on Economic Activity," *Journal of Money, Credit and Banking*, 1996, 28 (4 Part 1), 549–577.
- Rouwenhorst, K Geert**, "Time to build and aggregate fluctuations: A reconsideration," *Journal of Monetary Economics*, 1991, 27 (2), 241–254.

Stiglitz, Joseph E, “Monopoly and the rate of extraction of exhaustible resources,” *The American Economic Review*, 1976, 66 (4), 655–661.

Stuermer, Martin and Gregor Schwerhoff, “Non-renewable Resources, Extraction Technology, and Endogenous Growth,” 2020.

Suicmez, Vural Sander, Xudong Jing, Marcel Polikar, M Olivier Allain, Christopher Pentland, and Michael Dyson, *Imperial College Lectures In Petroleum Engineering, The-Volume 4: Drilling And Reservoir Appraisal*, World Scientific, 2018.

Woodford, Michael, “Commentary: How should monetary policy be conducted in an era of price stability?,” *New challenges for monetary policy*, 1999, 277316.

– , *Interest and prices: Foundations of a theory of monetary policy*, princeton university press, 2011.

Zhou, Xiaoqing, “Refining the workhorse oil market model,” *Journal of Applied Econometrics*, 2020, 35 (1), 130–140.

Appendix

A Additional tables

Table 12: Price elasticity of oil fields in operation

<i>Dep. variable: ln(number of new oil fields)</i>				<i>Dep. variable: ln(number of retired oil fields)</i>			
Variable	(1)	(2)	(3)	Variable	(1)	(2)	(3)
ln(price)	-0.01 (0.53)	-0.40 (0.60)	-0.39 (0.53)	ln(price)	0.81 (1.09)	0.23 (0.44)	-0.14 (0.29)
Year trend	✓	✓	✓	Year trend	✓	✓	✓
IV	✓	✓	✓	IV	✓	✓	✓
Dep. variable	All fields	Non-OPEC fields	OPEC fields	Dep. variable	All fields	Non-OPEC fields	OPEC fields
Observations	49	49	49	Observations	48	48	48

Notes: This table presents estimates of the elasticity of the number of *new* and *retiring* oil fields with respect to the price of oil. Newey-West standard errors computed with five-year lags are reported in parentheses. The instrument for oil prices is the one-year-ahead forecast of detrended world real GDP.

Table 13: Estimated parameters: OPEC behaves competitively

Parameter	Estimate	(s.e.)	Benchm. estimate
ϵ	0.135	(0.03)	0.15
ρ_1^a	1.69	(0.07)	1.72
ρ_2^a	-0.71	(0.07)	-0.74
$std(\varepsilon_t^a)$	0.012	(0.002)	0.013
$std(u_t)$	0.048	(0.007)	0.05

Notes: This table presents the GMM estimates of the structural parameters when OPEC firms behave competitively. The final column presents our benchmark estimates for comparison.

Table 14: Data and model moments: OPEC behaves competitively

Moment	Data	(s.e.)	Model		Moment	Data	(s.e.)	Model	
			Competitive	Benchmark				Competitive	Benchmark
(1) $\text{std}(\Delta p_t)$	0.27	(0.03)	0.23	0.22	(11) $\text{corr}(\Delta i_t^N, \Delta q_t^N)$	0.07	(0.12)	0.08	0.10
(2) $\text{std}(\Delta i_t^N)$	0.18	(0.03)	0.19	0.16	(12) $\text{corr}(\Delta i_t^N, \Delta q_t^O)$	0.04	(0.13)	0.03	0.04
(3) $\text{std}(\Delta i_t^O)$	0.17	(0.03)	0.19	0.21	(13) $\text{corr}(\Delta i_t^O, \Delta q_t^N)$	-0.11	(0.18)	0.08	0.08
(4) $\text{std}(\Delta q_t^N)$	0.02	(0.003)	0.02	0.02	(14) $\text{corr}(\Delta i_t^O, \Delta q_t^O)$	-0.28	(0.13)	0.03	0.04
(5) $\text{std}(\Delta q_t^O)$	0.07	(0.01)	0.07	0.08	(15) $\text{corr}(\Delta q_t^N, \Delta q_t^O)$	-0.07	(0.14)	0.29	0.40
(6) $\text{corr}(\Delta p_t, \Delta i_t^N)$	0.49	(0.14)	0.61	0.61	(16) $\text{corr}(\Delta p_t, \Delta p_{t-1})$	-0.06	(0.09)	-0.29	-0.31
(7) $\text{corr}(\Delta p_t, \Delta i_t^O)$	0.33	(0.11)	0.61	0.60	(17) $\text{corr}(\Delta i_t^N, \Delta i_{t-1}^N)$	0.13	(0.16)	0.04	0.05
(8) $\text{corr}(\Delta p_t, \Delta q_t^N)$	0.00	(0.09)	0.11	0.08	(18) $\text{corr}(\Delta i_t^O, \Delta i_{t-1}^O)$	0.29	(0.11)	0.04	0.04
(9) $\text{corr}(\Delta p_t, \Delta q_t^O)$	0.04	(0.14)	-0.71	-0.69	(19) $\text{corr}(\Delta q_t^N, \Delta q_{t-1}^N)$	0.67	(0.10)	0.67	0.71
(10) $\text{corr}(\Delta i_t^N, \Delta i_t^O)$	0.67	(0.12)	1.00	1.00	(20) $\text{corr}(\Delta q_t^O, \Delta q_{t-1}^O)$	0.25	(0.20)	-0.39	-0.31

Notes: This table presents the targeted data moments and the model-implied moments for a specification in which OPEC firms behave competitively. Newey-West standard errors computed with five-year lags are reported in parentheses. In the table, x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$.

Table 15: Estimated parameters: fraction of OPEC firms deviate

Parameter	Estimate	(s.e.)	Benchm. estimate
ϵ	0.148	(0.03)	0.15
ρ_1^a	1.66	(0.07)	1.72
ρ_2^a	-0.69	(0.08)	-0.74
$\text{std}(\varepsilon_t^a)$	0.016	(0.003)	0.013
$\text{std}(u_t)$	0.051	(0.007)	0.05
ζ_d	0.4	(0.35)	-

Notes: This table presents the GMM estimates of the structural parameters when a fraction of OPEC oil members do not follow the cartel's prescribed policy. The final column presents our benchmark estimates for comparison. The parameter ζ_d is the estimate for the share of OPEC production produced by OPEC members who deviate.

Table 16: Data and model moments: fraction of OPEC firms deviate

Moment	Data	(s.e.)	Model		Moment	Data	(s.e.)	Model	
			Partial enforcement	Benchmark				Partial enforcement	Benchmark
(1) $\text{std}(\Delta p_t)$	0.27	(0.03)	0.24	0.22	(11) $\text{corr}(\Delta i_t^N, \Delta q_t^N)$	0.07	(0.12)	0.09	0.10
(2) $\text{std}(\Delta i_t^N)$	0.18	(0.03)	0.19	0.16	(12) $\text{corr}(\Delta i_t^N, \Delta q_t^O)$	0.04	(0.13)	0.04	0.04
(3) $\text{std}(\Delta i_t^O)$	0.17	(0.03)	0.19	0.21	(13) $\text{corr}(\Delta i_t^O, \Delta q_t^N)$	-0.11	(0.18)	0.09	0.08
(4) $\text{std}(\Delta q_t^N)$	0.02	(0.003)	0.02	0.02	(14) $\text{corr}(\Delta i_t^O, \Delta q_t^O)$	-0.28	(0.13)	0.04	0.04
(5) $\text{std}(\Delta q_t^O)$	0.07	(0.01)	0.08	0.08	(15) $\text{corr}(\Delta q_t^N, \Delta q_t^O)$	-0.07	(0.14)	0.41	0.40
(6) $\text{corr}(\Delta p_t, \Delta i_t^N)$	0.49	(0.14)	0.68	0.61	(16) $\text{corr}(\Delta p_t, \Delta p_{t-1})$	-0.06	(0.09)	-0.24	-0.31
(7) $\text{corr}(\Delta p_t, \Delta i_t^O)$	0.33	(0.11)	0.68	0.60	(17) $\text{corr}(\Delta i_t^N, \Delta i_{t-1}^N)$	0.13	(0.16)	0.04	0.05
(8) $\text{corr}(\Delta p_t, \Delta q_t^N)$	0.00	(0.09)	0.11	0.08	(18) $\text{corr}(\Delta i_t^O, \Delta i_{t-1}^O)$	0.29	(0.11)	0.04	0.04
(9) $\text{corr}(\Delta p_t, \Delta q_t^O)$	0.04	(0.14)	-0.62	-0.69	(19) $\text{corr}(\Delta q_t^N, \Delta q_{t-1}^N)$	0.67	(0.10)	0.75	0.71
(10) $\text{corr}(\Delta i_t^N, \Delta i_t^O)$	0.67	(0.12)	1.00	1.00	(20) $\text{corr}(\Delta q_t^O, \Delta q_{t-1}^O)$	0.25	(0.20)	-0.29	-0.31

Notes: This table presents the targeted moments from the data and the model-implied moments under the specification in which a fraction of OPEC oil members do not follow the cartel's prescribed policy. Newey-West standard errors computed with five-year lags are in parentheses. In the table, x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$.

Table 17: Estimated parameters: supply shocks to non-OPEC

Parameter	Estimate	(s.e.)	Benchm. estimate
ϵ	0.16	(0.03)	0.15
ρ_1^a	1.75	(0.12)	1.72
ρ_2^a	-0.77	(0.12)	-0.74
$\text{std}(\varepsilon_t^a)$	0.012	(0.004)	0.013
$\text{std}(u_t)$	0.047	(0.006)	0.05
$\text{std}(u_t^N)$	0.007	(0.002)	-

Notes: This table presents the GMM estimates of the structural parameters for a model in which non-OPEC members are also subject to supply shocks. The final column presents our benchmark estimates for comparison. The parameter u_t^N is the non-OPEC specific supply shock.

Table 18: Data and model moments: supply shocks to non-OPEC

Moment	Data	(s.e.)	Model		Moment	Data	(s.e.)	Model	
			Non-OPEC shocks	Benchmark				Non-OPEC shocks	Benchmark
(1) $\text{std}(\Delta p_t)$	0.27	(0.03)	0.19	0.22	(11) $\text{corr}(\Delta i_t^N, \Delta q_t^N)$	0.07	(0.12)	0.10	0.10
(2) $\text{std}(\Delta i_t^N)$	0.18	(0.03)	0.17	0.16	(12) $\text{corr}(\Delta i_t^N, \Delta q_t^O)$	0.04	(0.13)	0.05	0.04
(3) $\text{std}(\Delta i_t^O)$	0.17	(0.03)	0.23	0.21	(13) $\text{corr}(\Delta i_t^O, \Delta q_t^N)$	-0.11	(0.18)	0.08	0.08
(4) $\text{std}(\Delta q_t^N)$	0.02	(0.003)	0.02	0.02	(14) $\text{corr}(\Delta i_t^O, \Delta q_t^O)$	-0.28	(0.13)	0.04	0.04
(5) $\text{std}(\Delta q_t^O)$	0.07	(0.01)	0.07	0.08	(15) $\text{corr}(\Delta q_t^N, \Delta q_t^O)$	-0.07	(0.14)	0.35	0.40
(6) $\text{corr}(\Delta p_t, \Delta i_t^N)$	0.49	(0.14)	0.60	0.61	(16) $\text{corr}(\Delta p_t, \Delta p_{t-1})$	-0.06	(0.09)	-0.26	-0.31
(7) $\text{corr}(\Delta p_t, \Delta i_t^O)$	0.33	(0.11)	0.59	0.60	(17) $\text{corr}(\Delta i_t^N, \Delta i_{t-1}^N)$	0.13	(0.16)	0.05	0.05
(8) $\text{corr}(\Delta p_t, \Delta q_t^N)$	0.00	(0.09)	0.09	0.08	(18) $\text{corr}(\Delta i_t^O, \Delta i_{t-1}^O)$	0.29	(0.11)	0.04	0.04
(9) $\text{corr}(\Delta p_t, \Delta q_t^O)$	0.04	(0.14)	-0.65	-0.69	(19) $\text{corr}(\Delta q_t^N, \Delta q_{t-1}^N)$	0.67	(0.10)	0.58	0.71
(10) $\text{corr}(\Delta i_t^N, \Delta i_t^O)$	0.67	(0.12)	1.00	1.00	(20) $\text{corr}(\Delta q_t^O, \Delta q_{t-1}^O)$	0.25	(0.20)	-0.31	-0.31

Notes: This table presents the targeted data moments and the model-implied moments under the specification in which non-OPEC members are also subject to supply shocks. Newey-West standard errors computed with five year lags are reported in parentheses. In the table, x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$.

Table 19: Estimated parameters: investment cost shocks

Parameter	Estimate	(s.e.)	Benchm. estimate
ϵ	0.18	(0.001)	0.15
ρ_1^a	1.64	(0.08)	1.72
ρ_2^a	-0.67	(0.09)	-0.74
$\text{std}(\varepsilon_t^a)$	0.014	(0.003)	0.013
$\rho_1^{\zeta_N}$	1.14	(0.16)	-
$\rho_2^{\zeta_N}$	-0.15	(0.16)	-
$\text{std}(\varepsilon_t^{\zeta_N})$	0.16	(0.04)	-
$\rho_1^{\zeta_O}$	1.31	(2.8)	-
$\rho_2^{\zeta_O}$	-0.32	(3.2)	-
$\text{std}(\varepsilon_t^{\zeta_O})$	0.19	(2.3)	-

Notes: This table presents the GMM estimates of the structural parameters for a model in which non-OPEC and OPEC members are subject to investment cost shocks. The final column presents our benchmark estimates for comparison. The parameters ζ_N and ζ_O are the non-OPEC and OPEC specific investment cost shocks, respectively. Both processes follow an AR(2) process.

Table 20: Data and model moments: investment cost shocks

Moment	Data	(s.e.)	Model	
			Inv. cost shocks	Benchmark
(1) $\text{std}(\Delta p_t)$	0.27	(0.03)	0.13	0.22
(2) $\text{std}(\Delta i_t^N)$	0.18	(0.03)	0.19	0.16
(3) $\text{std}(\Delta i_t^O)$	0.17	(0.03)	0.22	0.21
(4) $\text{std}(\Delta q_t^N)$	0.02	(0.003)	0.02	0.02
(5) $\text{std}(\Delta q_t^O)$	0.07	(0.01)	0.03	0.08
(6) $\text{corr}(\Delta p_t, \Delta i_t^N)$	0.49	(0.14)	0.58	0.61
(7) $\text{corr}(\Delta p_t, \Delta i_t^O)$	0.33	(0.11)	0.63	0.60
(8) $\text{corr}(\Delta p_t, \Delta q_t^N)$	0.00	(0.09)	-0.08	0.08
(9) $\text{corr}(\Delta p_t, \Delta q_t^O)$	0.04	(0.14)	-0.07	-0.69
(10) $\text{corr}(\Delta i_t^N, \Delta i_t^O)$	0.67	(0.12)	0.99	1.00

Moment	Data	(s.e.)	Model	
			Inv. cost shocks	Benchmark
(11) $\text{corr}(\Delta i_t^N, \Delta q_t^N)$	0.07	(0.12)	0.10	0.10
(12) $\text{corr}(\Delta i_t^N, \Delta q_t^O)$	0.04	(0.13)	0.08	0.04
(13) $\text{corr}(\Delta i_t^O, \Delta q_t^N)$	-0.11	(0.18)	0.05	0.08
(14) $\text{corr}(\Delta i_t^O, \Delta q_t^O)$	-0.28	(0.13)	0.06	0.04
(15) $\text{corr}(\Delta q_t^N, \Delta q_t^O)$	-0.07	(0.14)	0.99	0.40
(16) $\text{corr}(\Delta p_t, \Delta p_{t-1})$	-0.06	(0.09)	-0.14	-0.31
(17) $\text{corr}(\Delta i_t^N, \Delta i_{t-1}^N)$	0.13	(0.16)	0.15	0.05
(18) $\text{corr}(\Delta i_t^O, \Delta i_{t-1}^O)$	0.29	(0.11)	0.29	0.04
(19) $\text{corr}(\Delta q_t^N, \Delta q_{t-1}^N)$	0.67	(0.10)	0.51	0.71
(20) $\text{corr}(\Delta q_t^O, \Delta q_{t-1}^O)$	0.25	(0.20)	0.48	-0.31

Notes: This table presents the targeted data moments and the model-implied moments for a specification in which non-OPEC and OPEC members are subject to investment cost shocks. Newey-West standard errors computed with five year lags are in parentheses. In the table, x_t represents the logarithm of X_t , Δx_t is equal to $x_t - x_{t-1}$.

B Transition dynamics derivations

The objective function of the cartel is

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t M_t \left[e^{u_t} \theta_t^O R_t^O P_t - I_t^O - \psi (e^{u_t} \theta_t^O)^\eta R_t^O \right].$$

Denote $\tilde{\beta}_{t+1} \equiv \beta \frac{M_{t+1}}{M_t}$. The cartel's set of constraints is given by

$$\begin{aligned} R_{t+1}^O &= (1 - e^{u_t} \theta_t^O) R_t^O + \lambda X_{t+1}^O, & [\mu_t^{O,1}] \\ X_{t+1}^O &= (1 - \lambda) X_t^O + (I_t^O)^\xi L_O^{1-\xi}, & [\mu_t^{O,2}] \\ (1 - \theta_t^N) R_t^N + \lambda X_{t+1}^N &= R_{t+1}^N, & [\phi_{\mu_1}] \\ (1 - \lambda) X_t^N + (I_t^N)^\xi L_N^{1-\xi} &= X_{t+1}^N, & [\phi_{\mu_2}] \\ \mathbb{E}_t \left\{ \tilde{\beta}_{t+1} \left[P_{t+1} \theta_{t+1}^N + (1 - \theta_{t+1}^N) \mu_{t+1}^1 - \psi (\theta_{t+1}^N)^\eta \right] \right\} &= \mu_t^1, & [\phi_R] \\ \lambda \mu_t^1 + (1 - \lambda) \mathbb{E}_t \tilde{\beta}_{t+1} \mu_{t+1}^2 &= \mu_t^2, & [\phi_X] \\ \xi (I_t^N)^{\xi-1} L_N^{1-\xi} \mu_t^2 &= 1, & [\phi_I] \\ \mathbb{E}_t \tilde{\beta}_{t+1} \left[P_{t+1} - \mu_{t+1}^{1,N} \right] &= \mathbb{E}_t \tilde{\beta}_{t+1} \eta \psi (\theta_{t+1}^N)^{\eta-1}, & [\phi_\theta] \\ (1 - \theta_t^F) R_t^F + \lambda X_{t+1}^F &= R_{t+1}^F, & [\phi_{\mu_1}^F] \\ (1 - \lambda) X_t^F + (I_t^F)^\xi (L^F)^{1-\xi} &= X_{t+1}^F, & [\phi_{\mu_2}^F] \\ \mathbb{E}_t \left\{ \tilde{\beta}_{t+1} \left[P_{t+1} \theta_{t+1}^F + (1 - \theta_{t+1}^F) \mu_{t+1}^{1,F} - \psi (\theta_{t+1}^F)^\eta \right] \right\} &= \mu_t^{1,F}, & [\phi_R^F] \\ \lambda \mu_t^{1,F} + (1 - \lambda) \mathbb{E}_t \tilde{\beta}_{t+1} \mu_{t+1}^{2,F} &= \mu_t^{2,F}, & [\phi_X^F] \\ \xi (I_t^F)^{\xi-1} (L^F)^{1-\xi} \mu_t^{2,F} &= 1, & [\phi_I^F] \\ P_t - \mu_t^{1,F} &= \eta \psi (\theta_t^F)^{\eta-1}, & [\phi_{\theta^F}] \\ P_t &= s_o Y_t^{\frac{1}{\epsilon}} O_t^{-\frac{1}{\epsilon}}, & [\phi_P] \\ e^{u_t} \theta_t^O R_t^O + \theta_t^N R_t^N + \theta_t^F R_t^F &= O_t, & [\phi_O] \\ w_t &= (1 - \alpha)(1 - s_o) A_t Y_t^{\frac{1}{\epsilon}} K_t^{\frac{\epsilon-1}{\epsilon}} N_t^{\frac{\epsilon-1}{\epsilon}(1-\alpha)-1}, & [\phi_w] \\ r_t &= \alpha(1 - s_o) A_t Y_t^{\frac{1}{\epsilon}} K_t^{\frac{\epsilon-1}{\epsilon}\alpha-1} N_t^{\frac{\epsilon-1}{\epsilon}(1-\alpha)}, & [\phi_r] \\ Y_t &= \left[(1 - s_o) A_t (K_t \alpha N_t^{1-\alpha})^{\frac{\epsilon-1}{\epsilon}} + s_o O_t^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{\epsilon}{\epsilon-1}}, & [\phi_Y] \\ (1 - \delta) K_t + Y_t &= C_t + K_{t+1} + I_t^N + I_t^F + I_t^O + \psi (\theta_t^N)^\eta R_t^N + \psi (\theta_t^F)^\eta R_t^F + \psi (\theta_t^O)^\eta R_t^O, & [\phi_{RC}] \\ N_t &= \varphi^{-\frac{1}{\nu}} w_t^{\frac{1}{\nu}}, & [\phi_N] \\ 1 &= \mathbb{E}_t \left[\tilde{\beta}_{t+1} (1 + r_{t+1} - \delta) \right], & [\phi_K] \\ \frac{\tilde{\beta}_t}{\beta} \left(C_{t-1} - \varphi \frac{N_{t-1}^{1+\nu}}{1+\nu} \right)^{-\gamma} &= \left(C_t - \varphi \frac{N_t^{1+\nu}}{1+\nu} \right)^{-\gamma}, & [\phi_\beta] \end{aligned}$$

where the set of controls is given by $\Omega_t = \{I_t^O, R_{t+1}^O, X_{t+1}^O, \theta_{t+1}^O, R_{t+1}^N, X_{t+1}^N, I_t^N, \theta_{t+1}^N, \mu_t^{1,N}, \mu_t^{2,N}, R_{t+1}^F, X_{t+1}^F, I_t^F, \theta_t^F, \mu_t^{1,F}, \mu_t^{2,F}, P_t, O_t, w_t, r_t, K_{t+1}, Y_t, C_t, N_t, \tilde{\beta}\}$, and Ω_{-1} is given together with the level of productivity and supply shock A_0 and u_0 . The first two constraints are the laws of motion for OPEC, and the remaining constraints are the implementability constraints. The associated Lagrange multipliers are indicated to the right of each constraint. We scale each constraint by

$\beta^t \left(C_t - \varphi \frac{N_t^{1+\nu}}{1+\nu} \right)^{-\gamma}$ except for the last constraint, which is scaled by β^t only.

The 25 optimality conditions are as follows:

$$\begin{aligned}
[R_{t+1}^O] \quad \mu_t^{O,1} &= \mathbb{E}_t \left\{ \tilde{\beta}_{t+1} \left[e^{u_{t+1}} \theta_{t+1}^O P_{t+1} + (1 - e^{u_{t+1}} \theta_{t+1}^O) \mu_{t+1}^{O,1} - \psi (e^{u_t} \theta_{t+1}^O)^\eta \right. \right. \\
&\quad \left. \left. - e^{u_{t+1}} \theta_{t+1}^O \phi_{t+1}^O + \psi (\theta_{t+1}^O)^\eta \phi_{t+1}^{RC} \right] \right\}, \\
[X_{t+1}^O] \quad \mu_t^{O,2} &= \lambda \mu_t^{O,1} + \mathbb{E}_t \tilde{\beta}_{t+1} (1 - \lambda) \mu_{t+1}^{O,2}, \\
[I_t^O] \quad 1 - \phi_t^{RC} &= \xi (I_t^O)^{\xi-1} L_O^{1-\xi} \mu_t^{O,2}, \\
[\theta_{t+1}^O] \quad \mathbb{E}_t \tilde{\beta}_{t+1} \left[\eta \psi e^{\eta u_t} (\theta_{t+1}^O)^{\eta-1} - \eta \psi (\theta_{t+1}^O)^{\eta-1} \phi_{t+1}^{RC} \right] &= \\
&\quad \mathbb{E}_t \tilde{\beta}_{t+1} \left[e^{u_{t+1}} P_{t+1} - e^{u_{t+1}} \mu_{t+1}^{O,1} - e^{u_{t+1}} \phi_{t+1}^O \right], \\
[R_{t+1}^N] \quad \mathbb{E}_t \tilde{\beta}_{t+1} \left[(1 - \theta_{t+1}^N) \phi_{t+1}^{\mu,1} - \psi (\theta_{t+1}^N)^\eta \phi_{t+1}^{RC} \right] &= \phi_t^{\mu,1} - \mathbb{E}_t \tilde{\beta}_{t+1} \theta_{t+1}^N \phi_{t+1}^O, \\
[X_{t+1}^N] \quad \lambda \phi_t^{\mu,1} + \mathbb{E}_t \tilde{\beta}_{t+1} (1 - \lambda) \phi_{t+1}^{\mu,2} &= \phi_t^{\mu,2}, \\
[I_t^N] \quad \xi (I_t^N)^{\xi-1} L_N^{1-\xi} \phi_t^{\mu,2} - \phi_{RC} &= (1 - \xi) \xi (I_t^N)^{\xi-2} L_N^{1-\xi} \mu_t^I \phi_t^I, \\
[\theta_{t+1}^N] \quad \mathbb{E}_t \tilde{\beta}_{t+1} \left\{ [P_{t+1} - \mu_{t+1}^{1,N} - \psi \eta (\theta_{t+1}^N)^{\eta-1}] \phi_{t+1}^R - \eta \psi (\theta_{t+1}^N)^{\eta-1} R_{t+1}^N \phi_{t+1}^{RC} \right\} &= \\
&\quad \mathbb{E}_t \tilde{\beta}_{t+1} \left\{ R_{t+1}^N \phi_{t+1}^{\mu,1} + \eta (\eta - 1) \psi (\theta_{t+1}^N)^{\eta-2} \phi_{t+1} \theta - R_{t+1}^N \phi_{t+1}^O \right\}, \\
[\mu_t^{1,N}] \quad \lambda \phi_t^X + (1 - \theta_t^N) f(s_t | s^{t-1}) \phi_{t-1}^F &= \phi_t^R + f(s_t | s^{t-1}) \phi_{t-1} \theta, \\
[\mu_t^{2,N}] \quad (1 - \lambda) f(s_t | s^{t-1}) \phi_{t-1}^X + \xi (I_t^N)^{\xi-1} L_N^{1-\xi} \phi_t^I &= \phi_t^X, \\
[R_{t+1}^F] \quad \mathbb{E}_t \tilde{\beta}_{t+1} \left[(1 - \theta_{t+1}^F) \phi_{t+1}^{\mu,1,F} - \psi (\theta_{t+1}^F)^\eta \phi_{t+1}^{RC} \right] &= \phi_t^{\mu,1,F} - \mathbb{E}_t \tilde{\beta}_{t+1} \theta_{t+1}^F \phi_{t+1}^O, \\
[X_{t+1}^F] \quad \lambda \phi_t^{\mu,1,F} + \mathbb{E}_t \tilde{\beta}_{t+1} (1 - \lambda) \phi_{t+1}^{\mu,2,F} &= \phi_t^{\mu,2,F}, \\
[I_t^F] \quad \xi (I_t^F)^{\xi-1} (L_t^F)^{1-\xi} \phi_t^{\mu,2,F} - \phi_{RC} &= (1 - \xi) \xi (I_t^F)^{\xi-2} (L_t^F)^{1-\xi} \mu_t^{2,F} \phi_t^{I,F}, \\
[\theta_t^F] \quad [P_t - \mu_t^{1,F} - \psi \eta (\theta_t^F)^{\eta-1}] \phi_t^{R,F} - \eta \psi (\theta_t^F)^{\eta-1} R_t^F \phi_t^{RC} &= \\
&\quad R_t^F \phi_t^{\mu,1,F} + \eta (\eta - 1) \psi (\theta_t^F)^{\eta-2} \phi_t^{\theta,F} - R_t^F \phi_t^O, \\
[\mu_t^{1,F}] \quad \lambda \phi_t^{X,F} + (1 - \theta_t^F) f(s_t | s^{t-1}) \phi_{t-1}^{R,F} &= \phi_t^{R,F} + \phi_t^{\theta,F}, \\
[\mu_t^{2,F}] \quad (1 - \lambda) f(s_t | s^{t-1}) \phi_{t-1}^{X,F} + \xi (I_t^F)^{\xi-1} (L_t^F)^{1-\xi} \phi_t^{I,F} &= \phi_t^{X,F}, \\
[P_t] \quad \phi_t^P + f(s_t | s^{t-1}) (\theta_t^N \phi_{t-1}^R + \phi_{t-1} \theta + \theta_t^F \phi_{t-1}^{R,F}) + \phi_t^{\theta,F} &= e^{u_t} \theta_t^O R_t^O, \\
[O_t] \quad \frac{1}{\epsilon} s_o Y_t^{\frac{1}{\epsilon}} O_t^{-\frac{1+\epsilon}{\epsilon}} \phi_t^P &= \phi_t^O + Y_t^{\frac{1}{\epsilon}} s_o O_t^{-\frac{1}{\epsilon}} \phi_t^Y, \\
[w_t] \quad \phi_t^w &= \frac{1}{\nu} N_t w_t^{-1} \phi_t^L, \\
[r_t] \quad \phi_t^r &= f(s_t | s^{t-1}) \phi_{t-1}^K, \\
[K_{t+1}] \quad \mathbb{E}_t \tilde{\beta}_{t+1} \left[\left(1 - \frac{\epsilon-1}{\epsilon} \alpha \right) r_{t+1} \phi_{t+1}^r - \frac{\epsilon-1}{\epsilon} \alpha w_{t+1} \phi_{t+1}^w \right] &K_{t+1}^{-1} - \phi_t^{RC} = \\
&\quad \mathbb{E}_t \tilde{\beta}_{t+1} (r_{t+1} \phi_{t+1}^Y - (1 - \delta) \phi_{t+1}^{RC}), \\
[Y_t] \quad \phi_t^Y + \phi_t^{RC} &= \frac{1}{\epsilon Y_t} (P_t \phi_t^P + w_t \phi_t^w + r_t \phi_t^r), \\
[C_t] \quad -\phi_t^{RC} + \gamma \left(C_t - \varphi \frac{N_t^{1+\nu}}{1+\nu} \right)^{-1} &\left[e^{u_t} \theta_t^O R_t^O P_t - I_t^O - \psi (\theta_t^O)^\eta R_t^O \right] = \\
&\quad -\gamma \beta \left(C_t - \varphi \frac{N_t^{1+\nu}}{1+\nu} \right)^{-1} \left[\phi_t^\beta - \mathbb{E}_t \tilde{\beta}_{t+1} \phi_{t+1}^\beta \right], \\
[N_t] \quad -\gamma \beta \varphi N_t^\nu \left(C_t - \varphi \frac{N_t^{1+\nu}}{1+\nu} \right)^{-1} &\left[\phi_t^\beta - \mathbb{E}_t \tilde{\beta}_{t+1} \phi_{t+1}^\beta \right] + \left(1 - \frac{\epsilon-1}{\epsilon} (1 - \alpha) \right) w_t N_t^{-1} \phi_t^w + \phi_t^L = \\
&\quad \gamma \varphi N_t^\nu \left(C_t - \varphi \frac{N_t^{1+\nu}}{1+\nu} \right)^{-1} \left[e^{u_t} \theta_t^O R_t^O P_t - I_t^O - \psi (\theta_t^O)^\eta R_t^O \right] + \frac{\epsilon-1}{\epsilon} (1 - \alpha) r_t N_t^{-1} \phi_t^r + w_t \phi_t^Y \\
[\tilde{\beta}_t] \quad f(s_t | s^{t-1}) \left\{ [P_t \theta_t^N + (1 - \theta_t^N) \mu_t^{1,N} - \psi (\theta_t^N)^\eta] \phi_{t-1}^R + (1 - \lambda) \mu_t^{2,N} \phi_{t-1}^X \right\} & \\
&+ f(s_t | s^{t-1}) \left\{ [P_t \theta_t^F + (1 - \theta_t^F) \mu_t^{1,F} - \psi (\theta_t^F)^\eta] \phi_{t-1}^{R,F} + (1 - \lambda) \mu_t^{2,F} \phi_{t-1}^{X,F} \right\} = \\
&f(s_t | s^{t-1}) \left\{ [\eta \psi (\theta_t^N)^{\eta-1} + \mu_t^{1,N} - P_t] \phi_{t-1}^\theta + (1 + r_t - \delta) \phi_{t-1}^K \right\} - \beta \phi_t^\beta.
\end{aligned}$$

The equilibrium conditions are all of the following form: $\mathbb{E} [\Omega_{t-1}, \Omega_t, \Omega_{t+1}, z_t, z_{t+1}] = 0$, where Ω_t is the set of control variables chosen at time t , and z_t are the set of exogenous states at time t . The extraction rate at time $t + 1$ is, for example, included in Ω_t . We assume that $\Omega_{T+1} = \Omega^*$, where Ω^* is the new steady state with fracking. We want to find a vector $\vec{\Omega}_t$ from $t = 0$ to T , so that all optimality conditions hold with equality.

B.1 Algorithm

We start with a guess for $\{\vec{\theta}_F, \vec{\theta}_N, \vec{\theta}_O, \vec{P}_t, \vec{\phi}_{RC}\}$. This vector has $5 \times T$ endogenous variables. The following steps describe how to obtain the remaining endogenous parameters as well as the $5 \times T$ vector of residual equations.

1. Obtaining μ_1^N . We have that $\mu_t^{1,N} = P_t - \eta\psi(\theta_t^N)^{\eta-1}$.

2. Obtaining $\tilde{\beta}$. We have that

$$\tilde{\beta}_{t+1} = \frac{\mu_t^{1,N}}{P_{t+1}\theta_{t+1} + (1 - \theta_{t+1}^N)\mu_{t+1}^{1,N} - \psi(\theta_{t+1}^N)^\eta}.$$

3. Obtaining μ_2^N . We have that $\mu_t^{2,N} = \lambda\mu_t^{1,N} + (1 - \lambda)\tilde{\beta}_{t+1}\mu_{t+1}^{2,N}$.

4. Obtaining I^N . We have that $I_t^N = \left(\xi L_N^{1-\xi} \mu_t^{2,N}\right)^{\frac{1}{1-\xi}}$.

5. Obtaining X_N . We have the initial state variable X_0^N . Using the investment and exploration capital laws of motion, we get $X_{t+1}^N = (1 - \lambda)X_t^N + (I_t^N)^\xi L_N^{1-\xi}$.

6. Obtaining R_N . We have the initial state variable R_0^N . Using the exploration capital and reserves laws of motion, we have $R_{t+1}^N = (1 - \theta_t^N)R_t^N + \lambda X_{t+1}^N$.

7. Obtaining μ_1^F . We have that $\mu_t^{1,F} = P_t - \eta_F\psi_F(\theta_t^F)^{\eta_F-1}$.

8. First set of residual equations. The fracking optimality with respect to R_t^F yields the first set of residual equations. We get an implied guess for the fracking extraction rate sequence as follows. First, we get an implied μ_1^F by backward induction:

$$\bar{\mu}_t^{1,F} = \tilde{\beta}_{t+1} \left[P_{t+1}\theta_{t+1}^F + (1 - \theta_{t+1}^F)\mu_{t+1}^F - \psi_F(\theta_{t+1}^F)^{\eta_F} \right].$$

Then we get an implied guess for θ_t^F :

$$\theta_t^F = \left[\frac{1}{\eta_F\psi_F} \left(P_t - \bar{\mu}_t^{1,F} \right) \right]^{\frac{1}{\eta_F-1}}.$$

9. Obtaining μ_2^F . We have that $\mu_t^{2,F} = \lambda_F\mu_t^{1,F} + (1 - \lambda_F)\tilde{\beta}_{t+1}\mu_{t+1}^{2,F}$.

10. Obtaining I^F . We have that $I_t^F = \left(\xi (L_t^F)^{1-\xi} \mu_t^{2,F}\right)^{\frac{1}{1-\xi}}$.

11. Obtaining X_F . In case $\lambda_F = 1$ we simply have $X_{t+1} = (I_t^F)^\xi (L_t^F)^{1-\xi}$.

12. Obtaining R_F . We have the initial state variable R_0^F . Using the exploration capital and reserves laws of motion, we have $R_{t+1}^F = (1 - \theta_t^F)R_t^F + \lambda_F X_{t+1}^F$.

13. Obtaining μ_1^O . Combining OPEC optimality for R_O and θ_O , we have

$$\mu_t^{1,O} = \tilde{\beta}_{t+1} \left[\mu_{t+1}^{1,O} + (\eta - 1)\psi(\theta_{t+1}^O)^\eta - (\eta - 1)\psi(\theta_{t+1}^O)^\eta \phi_{t+1}^{RC} \right].$$

14. Obtaining ϕ_t^O . Using the cartel's optimal extraction rate, we have

$$\phi_{t+1}^O = P_{t+1} - \mu_{t+1}^{O,1} + \eta\psi(\theta_{t+1}^O)^{\eta-1} (1 - \phi_{t+1}^{RC}).$$

Note that we don't have a value for ϕ_0^O . We guess the initial value of ϕ_O as well. We have an explicit value for θ_t^O , so that overall we have to guess $5 \times T$ endogenous variables. Finally, this equation provides an additional residual equation for ϕ_{T+1}^O , so that balances with the initial guess of ϕ_O .

15. Obtaining $\mu_t^{2,O}$. We use backward induction to obtain

$$\mu_t^{O,2} = \lambda\mu_t^{O,1} + \mathbb{E}_t \tilde{\beta}_{t+1} (1 - \lambda)\mu_{t+1}^{2,O}.$$

16. Obtaining I_t^O . OPEC optimality yields

$$I_t^O = \left(\frac{\xi L_O^{1-\xi} \mu_t^{O,2}}{1 - \phi_t^{RC}} \right)^{\frac{1}{1-\xi}}.$$

17. Obtaining X_t^O . We have the initial state variable X_0^O . Using the investment and the exploration law of motion, we get $X_{t+1}^O = (1 - \lambda)X_t^O + (I_t^O)^\xi L_O^{1-\xi}$.

18. Obtaining R_t^O . We have the initial state variable R_0^O . Using the exploration capital and reserves laws of motion, we have $R_{t+1}^O = (1 - \theta_t^O)R_t^O + \lambda X_{t+1}^O$.

19. Obtaining O_t . With all extraction rates and oil fields, we have

$$O_t = R_t^O \theta_t^O + R_t^N \theta_t^N + R_t^F \theta_t^F.$$

20. Obtaining Y_t . The demand equation pins down

$$Y_t = \left(\frac{P_t}{s_o} \right)^\epsilon O_t.$$

21. Obtaining $VA_t^{\frac{\epsilon-1}{\epsilon}}$. The aggregate production pins down value added:

$$VA_t^{\frac{\epsilon-1}{\epsilon}} = \frac{1}{(1-s_o)A_t} \left(Y_t^{\frac{\epsilon-1}{\epsilon}} - s_o O_t^{\frac{\epsilon-1}{\epsilon}} \right).$$

22. Obtaining N_t . We have two implementability constraints that involve N and w . Combining the two, we have

$$N_t = \left[\frac{1}{\varphi} (1-\alpha)(1-s_o)A_t Y_t^{\frac{1}{\epsilon}} VA_t^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{1}{1+\nu}}.$$

23. Obtaining w_t . Now we can use labor supply to obtain

$$w_t = \varphi N_t^\nu.$$

24. Obtaining K_t . Value added and labor immediately pin down K_t :

$$K_t = (VA_t N_t^{\alpha-1})^{\frac{1}{\alpha}}.$$

Since the initial level of K_0 is given, this is a residual equation.

25. Obtaining r_t . Capital demand yields the rate of return r_t . Capital demand implies

$$r_t = \alpha(1-s_o)A_t Y_t^{\frac{1}{\epsilon}} VA_t^{\frac{\epsilon-1}{\epsilon}} K_t^{-1}.$$

26. Second set of residual equations. Given the value of $\tilde{\beta}_t$ we have the following residual equation:

$$1 + r_{t+1} - \delta = \frac{1}{\tilde{\beta}_{t+1}}.$$

27. Obtaining C_t . The resource constraint provides an implied value for consumption at times $t = 0$ until $T - 1$.

$$C_t = (1-\delta)K_t + Y_t - K_{t+1} - I_t^N - I_t^F - I_t^O - \psi(\theta_t^N)^\eta R_t^N - \psi(\theta_t^O)^\eta R_t^O - \psi_F(\theta_t^F)^{\eta_F} R_t^F.$$

The equation at time T is missing K_{t+1} and doesn't uniquely pin down C_T . For this reason, we use the following condition.

28. Third set of residual equations. The definition of $\tilde{\beta}_t$ provides the following residual equation:

$$C_t = \left(C_{t-1} - \varphi \frac{N_{t-1}^{1+\nu}}{1+\nu} \right) \left(\frac{\tilde{\beta}_t}{\beta} \right)^{-\frac{1}{\gamma}} + \varphi \frac{N_t^{1+\nu}}{1+\nu}.$$

At time T , since we don't have a guess for C_T , this equation just pins down C_T , and the equation above pins down K_{T+1} .

29. Obtaining $\phi_t^{\mu_1}$. We get $\phi_t^{\mu_1}$ by backward induction:

$$\phi_t^{\mu_1} = \tilde{\beta}_{t+1} \left[\theta_{t+1}^N \phi_{t+1}^O + (1 - \theta_{t+1}^N) \phi_{t+1}^{\mu_1} - \psi (\theta_{t+1}^N)^\eta \phi_{t+1}^{RC} \right].$$

30. Obtaining $\phi_t^{\mu_2}$. We use backward induction to get

$$\phi_t^{\mu_2} = \lambda \phi_t^{\mu_1} + \tilde{\beta}_{t+1} (1 - \lambda) \phi_{t+1}^{\mu_2}.$$

31. Obtaining ϕ_t^I . We have that

$$\phi_t^I = \frac{\xi (I_t^N)^{\xi-1} L_N^{1-\xi} \phi_t^{\mu_2} - \phi_t^{RC}}{(1 - \xi) \xi (I_t^N)^{\xi-2} L_N^{1-\xi} \mu_t^2}.$$

32. Obtaining ϕ_t^X . We get ϕ_t^X from the following condition (which includes the promise keeping)

$$\phi_t^X = (1 - \lambda) \phi_{t-1}^X + \xi (I_t^N)^{\xi-1} L_N^{1-\xi} \phi_t^I.$$

33. Obtaining ϕ_t^R and ϕ_t^θ . We get ϕ_t^R using the following backward-looking equation:

$$\phi_t^R = \lambda \phi_t^X + (1 - \theta_t^N) \phi_{t-1}^R - \phi_{t-1}^\theta.$$

Then we obtain

$$\phi_t^\theta = \frac{\left(P_{t+1} - \mu_{t+1}^{1,N} - \psi \eta (\theta_{t+1}^N)^{\eta-1} \right) \phi_t^R - \eta \psi (\theta_{t+1}^N)^{\eta-1} R_{t+1}^N \phi_{t+1}^{RC} + R_{t+1}^N \phi_{t+1}^O - R_{t+1}^N \phi_{t+1}^{\mu_1}}{\eta (\eta - 1) \psi (\theta_{t+1}^N)^{\eta-2}}.$$

34. Obtaining $\phi_t^{\mu_{1,F}}$. We use backward induction to get

$$\phi_t^{\mu_{1,F}} = \tilde{\beta}_{t+1} \left[(1 - \theta_{t+1}^F) \phi_{t+1}^{\mu_{1,F}} - \psi_F (\theta_{t+1}^F)^{\eta_F} \phi_{t+1}^{RC} + \theta_{t+1}^F \phi_{t+1}^O \right].$$

35. Obtaining $\phi_t^{\mu_{2,F}}$. We use backward induction to obtain

$$\phi_t^{\mu_{2,F}} = \lambda \phi_t^{\mu_{1,F}} + \tilde{\beta}_{t+1} (1 - \lambda) \phi_{t+1}^{\mu_{2,F}}.$$

36. Obtaining $\phi_t^{I,F}$. We have

$$\phi_t^{I,F} = \frac{\xi (I_t^F)^{\xi-1} (L_t^F)^{1-\xi} \phi_t^{\mu_{2,F}} - \phi_{RC}}{(1 - \xi) \xi (I_t^F)^{\xi-2} (L_t^F)^{1-\xi} \mu_t^{2,F}}.$$

37. Obtaining $\phi_t^{X,F}$. We obtain $\phi_t^{X,F}$ as follows:

$$\phi_t^{X,F} = (1 - \lambda) \phi_{t-1}^{X,F} + \xi (I_t^F)^{\xi-1} (L_t^F)^{1-\xi} \phi_t^{I,F}.$$

38. Obtaining $\phi_t^{\theta,F}$. From the optimality condition for θ_F , we have

$$\phi_t^{\theta,F} = \frac{\left(P_t - \mu_t^{1,F} - \psi_F \eta_F (\theta_t^F)^{\eta_F - 1}\right)}{\eta_F (\eta_F - 1) \psi_F (\theta_t^F)^{\eta_F - 2}} \theta_t^{F,F} - \frac{\eta_F \psi_F (\theta_t^F)^{\eta_F - 1} R_t^F \phi_t^{RC} - R_t^F \phi_t^O + R_t^F \phi_t^{\mu_1,F}}{\eta_F (\eta_F - 1) \psi_F (\theta_t^F)^{\eta_F - 2}}.$$

In equilibrium, we have that the first term is equal to zero, so that

$$\phi_t^{\theta,F} = \frac{R_t^F \phi_t^{\mu_1,F} - \eta_F \psi_F (\theta_t^F)^{\eta_F - 1} R_t^F \phi_t^{RC} - R_t^F \phi_t^O}{\eta_F (\eta_F - 1) \psi_F (\theta_t^F)^{\eta_F - 2}}.$$

39. Obtaining $\phi_t^{R,F}$. We now have that

$$\phi_t^{R,F} = \lambda_F \phi_t^{X,F} + (1 - \theta_t^F) \phi_{t-1}^{R,F} - \phi_t^{\theta,F}.$$

40. Obtaining ϕ_t^P . We have

$$\phi_t^P = e^{u_t} \theta_t^O R_t^O - \phi_t^{\theta,F} - \left(\theta_t^N \phi_{t-1}^R + \phi_{t-1} \theta + \theta_t^F \phi_{t-1}^{R,F}\right).$$

41. Obtaining ϕ_t^Y . We have

$$\phi_t^Y = \frac{1}{Y_t^{\frac{1}{\epsilon}} s_o O_t^{-\frac{1}{\epsilon}}} \left(\frac{1}{\epsilon} s_o Y_t^{\frac{1}{\epsilon}} O_t^{-\frac{1+\epsilon}{\epsilon}} \phi_t^P - \phi_t^O \right).$$

42. Obtaining ϕ_t^r and ϕ_t^w . We have the following two optimality conditions:

$$w_t \phi_t^w + r_t \phi_t^r = \epsilon Y_t (\phi_t^Y + \phi_t^{RC}) - P_t \phi_t^P,$$

and

$$w_{t+1} \phi_{t+1}^w + r_{t+1} \phi_{t+1}^r = \frac{\epsilon}{\epsilon - 1} \frac{1}{\alpha} \left[r_{t+1} \phi_{t+1}^r - K_{t+1} \left(r_{t+1} \phi_{t+1}^Y + (1 - \delta) \phi_{t+1}^{RC} + \frac{1}{\beta_{t+1}} \phi_{t+1}^{RC} \right) \right].$$

This equation does not pin down ϕ_0^r .

Together we have

$$r_{t+1} \phi_{t+1}^r = \alpha (\epsilon - 1) Y_{t+1} (\phi_{t+1}^Y + \phi_{t+1}^{RC}) - \frac{\epsilon - 1}{\epsilon} \alpha P_t \phi_t^P + K_{t+1} \left(r_{t+1} \phi_{t+1}^Y + (1 - \delta) \phi_{t+1}^{RC} + \frac{1}{\beta_{t+1}} \phi_{t+1}^{RC} \right).$$

From this equation, we obtain

$$\phi_t^w = \frac{1}{w_t} (\epsilon Y_t (\phi_t^Y + \phi_t^{RC}) - P_t \phi_t^P - r_t \phi_t^r).$$

43. Obtaining ϕ_t^N . We have

$$\phi_t^L = \nu w_t \frac{1}{N_t} \phi_t^w.$$

44. Obtaining ϕ_t^K . We have

$$\phi_{t-1}^K = \phi_t^r.$$

Note that ϕ_{-1}^K yields ϕ_0^r . This expression is a form of promise keeping by the cartel.

45. Obtaining ϕ_t^β . Finally, we have

$$\begin{aligned} \beta\phi_t\beta &= \left[\eta\psi(\theta_t^N)^{\eta-1} + \mu_t^{1,N} - P_t \right] \phi_{t-1}^\theta + (1 + r_t - \delta)\phi_{t-1}^K \\ &\quad - \left[P_t\theta_t^N + (1 - \theta_t^N)\mu_t^{1,N} - \psi(\theta_t^N)^\eta \right] \phi_{t-1}^R - (1 - \lambda)\mu_t^{2,N}\phi_{t-1}^X \\ &\quad - \left[P_t\theta_t^F + (1 - \theta_t^F)\mu_t^{1,F} - \psi_F(\theta_t^F)^{\eta_F} \right] \phi_{t-1}^{R,F} - (1 - \lambda_F)\mu_t^{2,F}\phi_{t-1}^{X,F}. \end{aligned}$$

46. Fourth set of residual equations.

$$\begin{aligned} -\phi_t^{RC} + \gamma \left(C_t - \varphi \frac{N_t^{1+\nu}}{1+\nu} \right)^{-1} \left[e^{u_t}\theta_t^O R_t^O P_t - I_t^O - \psi(\theta_t^O)^\eta R_t^O \right] = \\ -\gamma\beta \left(C_t - \varphi \frac{N_t^{1+\nu}}{1+\nu} \right)^{-1} \left[\phi_t^\beta - \mathbb{E}_t \tilde{\beta}_{t+1} \phi_{t+1} \beta \right]. \end{aligned}$$

47. Fifth set of residual equations.

$$\begin{aligned} -\gamma\beta\varphi N_t^\nu \left(C_t - \varphi \frac{N_t^{1+\nu}}{1+\nu} \right)^{-1} \left[\phi_t^\beta - \mathbb{E}_t \tilde{\beta}_{t+1} \phi_{t+1} \beta \right] + \left(1 - \frac{\epsilon-1}{\epsilon}(1-\alpha) \right) w_t N_t^{-1} \phi_t^w + \phi_t^N = \\ \gamma\varphi N_t^\nu \left(C_t - \varphi \frac{N_t^{1+\nu}}{1+\nu} \right)^{-1} \left[e^{u_t}\theta_t^O R_t^O P_t - I_t^O - \psi(\theta_t^O)^\eta R_t^O \right] + \frac{\epsilon-1}{\epsilon}(1-\alpha)r_t N_t^{-1} \phi_t^r + w_t \phi_t^Y, \end{aligned}$$

which can be simplified with the equation above to

$$\varphi N_t^\nu \phi_t^{RC} + \frac{\epsilon-1}{\epsilon}(1-\alpha)r_t N_t^{-1} \phi_t^r + w_t \phi_t^Y = \left(1 - \frac{\epsilon-1}{\epsilon}(1-\alpha) \right) w_t N_t^{-1} \phi_t^w + \phi_t^N.$$

We now have a set of $5 \times T$ residual equations. For the algorithm to converge, we need to start with a good initial guess. To do so, we proceed as follows. We start by setting the initial conditions Ω_0 and z_0 , including L_0^F , to their terminal steady-state values. For these initial conditions, we know the sequence that solves the transition path, as the system doesn't deviate from the terminal steady state. We then gradually move the initial condition and the value for past Lagrange multipliers on implementability constraints toward the no-fracking steady state. In each iteration, we set the guess to be the previous solution for the transition

path. By doing so, the system always starts with a decent guess for the transition, and we manage to solve the system of equations. Once the system is solved for the desired initial conditions, we search for ρ , the convergence rate of the land controlled by fracking firms, to match the share of fracking in oil production over time.