

Nakov, Anton; Nuño, Galo

Working Paper

Saudi Aramco and the oil market

ECB Working Paper, No. 1354

Provided in Cooperation with:

European Central Bank (ECB)

Suggested Citation: Nakov, Anton; Nuño, Galo (2011) : Saudi Aramco and the oil market, ECB Working Paper, No. 1354, European Central Bank (ECB), Frankfurt a. M.

This Version is available at:

<https://hdl.handle.net/10419/153788>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



EUROPEAN CENTRAL BANK

EUROSYSTEM

WORKING PAPER SERIES

NO 1354 / JUNE 2011

**SAUDI ARAMCO
AND THE
OIL MARKET**

by Anton Nakov
and Galo Nuño



EUROPEAN CENTRAL BANK

EUROSYSTEM



WORKING PAPER SERIES

NO 1354 / JUNE 2011

SAUDI ARAMCO AND THE OIL MARKET ¹

by Anton Nakov ²
and Galo Nuño ³



In 2011 all ECB publications feature a motif taken from the €100 banknote.



NOTE: This Working Paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

This paper can be downloaded without charge from <http://www.ecb.europa.eu> or from the Social Science Research Network electronic library at http://ssrn.com/abstract_id=1857398.

¹ We have benefited from comments and suggestions by Dirk Krueger, Alessia Campolmi, and conference participants at ESSIM, Computing in Economics and Finance, Dynare, and seminar participants at Milano-Bicocca, ECB and BBVA. The views expressed are those of the authors and do not necessarily coincide with the views of the Bank of Spain, ECB, or the Eurosystem.

² European Central Bank, DG Research, Neue Mainzer Strasse 66, D-60311 Frankfurt am Main, Germany; e-mail: Anton.Nakov@ecb.europa.eu

³ Banco de España, Economía and Asuntos Internacionales, Alcalá 48, 28014 Madrid, Spain.

© European Central Bank, 2011

Address

Kaiserstrasse 29
60311 Frankfurt am Main, Germany

Postal address

Postfach 16 03 19
60066 Frankfurt am Main, Germany

Telephone

+49 69 1344 0

Internet

<http://www.ecb.europa.eu>

Fax

+49 69 1344 6000

All rights reserved.

Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

Information on all of the papers published in the ECB Working Paper Series can be found on the ECB's website, <http://www.ecb.europa.eu/pub/scientific/wps/date/html/index.en.html>

ISSN 1725-2806 (online)

CONTENTS

| | |
|--|----|
| Abstract | 4 |
| Non technical summary | 5 |
| 1 Introduction | 7 |
| 2 Model | 9 |
| 2.1. Oil-importing region | 9 |
| 2.2. Competitive fringe of oil exporters | 10 |
| 2.3. Dominant oil exporter | 11 |
| 3 Secular features of the oil market | 12 |
| 3.1. Oil prices and efficiency | 12 |
| 3.2. Oil price markup and capacity utilization | 12 |
| 4 Quantitative analysis of oil supply and demand shocks | 13 |
| 4.1. Calibration | 13 |
| 4.2. Impulse-responses | 14 |
| 4.3. Event study: the Persian Gulf war of 1990-91 | 15 |
| 5 Oil taxes and subsidies in general equilibrium | 16 |
| 5.1. Fringe investment subsidy: production efficiency vs. competition | 16 |
| 5.2. Oil consumption tax | 17 |
| 6 Conclusions | 18 |
| References | 19 |
| Tables and figures | 21 |
| Appendix | 25 |

Abstract

We present a general equilibrium model of the global oil market, in which the oil price, oil production, and consumption, are jointly determined as outcomes of the optimizing decisions of oil importers and oil exporters. On the supply side the oil market is modelled as a dominant firm – Saudi Aramco – with competitive fringe. We establish that a dominant firm may exist as long as it enjoys a cost advantage over the fringe. We provide an expression for the optimal markup and compute the spare capacity maintained by such a firm. The model produces plausible dynamics in response to oil supply and oil demand shocks. In particular, it reproduces successfully the jump in oil output of Saudi Aramco following the output collapse of Iraq and Kuwait during the first Gulf War, explaining it as the profit-maximizing response of the dominant firm. Oil taxes and subsidies affect the oil price and welfare through their effect on the trade-off between oil production efficiency and oil market competition.

Keywords: oil price, oil production, dominant firm, Saudi Aramco, oil tax

JEL classification: E32, Q43

Non-Technical Summary

An extensive literature studies the transmission of oil price shocks assuming that the oil price is given exogenously. Yet oil price shocks do not happen in isolation. Rather, like other relative price changes, they are triggered by deeper shocks to preferences, technology, and policy. The international oil price is determined, jointly with the oil quantity produced and consumed, as an equilibrium outcome of the interaction of different market participants. Indeed, Barsky and Kilian (2002, 2004), and Kilian (2009), have found overwhelming evidence that the oil price is affected significantly by global economic conditions. This is important both from a positive point of view, as shown by Kilian (2009), and from a normative perspective, as demonstrated by Nakov and Pescatori (2010) in the context of monetary policy.

Despite all this evidence, few equilibrium models determine the oil price, and even fewer determine endogenously both the oil price and the oil quantity. This paper presents a general equilibrium framework designed to approximate the industrial structure of the oil market, while jointly determining the oil price, oil production and oil consumption, as equilibrium outcomes of the optimizing decisions of oil importers and oil exporters. On the supply side we model the oil market as a dominant firm with competitive fringe, in the spirit of Salant (1976). We build this industrial organization model into a general equilibrium framework, assuming that the dominant firm internalizes the behavioral responses of fringe producers and oil consumers. Thus, the dominant firm in our model internalizes the fact that a change in its oil output and its demand for inputs affects oil demand, the supply of fringe producers, and the oil price.

We calibrate our dominant firm to data on Aramco, the 100% state-owned oil company of Saudi Arabia. As Adelman (1995) and Alhajji and Huettner (2000) convincingly argue, the behavior of Aramco can hardly be described as competitive: it is by far the largest oil firm in the world, accounting for more than a tenth of global oil production and a fifth of total proven reserves. Even though it has the lowest production costs, it is the only firm which maintains ample spare capacity and restricts investment in developing available reserves (Smith, 2009). Moreover, its output is negatively correlated with that of other OPEC producers (*ibid.*). Indeed, its output has been highly volatile compared to other major producers, even though Aramco itself has experienced few shocks affecting its production directly.

We use our model to address three sets of questions. First, in Section 3 we derive a set of analytical results related to the secular features of the economy that are independent of the particular calibration. Namely, the technological conditions required for the existence of a dominant oil firm, its degree of capacity utilization, and the conditions for balanced growth. We establish that a dominant firm may exist profitably as long as it enjoys a permanent cost advantage with respect to the fringe, and we offer an expression for the optimal markup as a function of this technological advantage. We also find that such a firm would optimally maintain spare capacity relative to the fringe, and that the oil price must grow over time to sustain balanced growth.

Second, in Section 4 we study the dynamic properties of our calibrated economy. With only two productivity shocks (one affecting the output of the oil-importing region and the other that of fringe producers), the model is capable of reproducing the monthly volatility of the oil price, of the oil output of Aramco and that of the fringe, as well as of US industrial production. We show that the co-movement between the oil price and value added of the oil-importing region may be positive or negative depending on the origin of the shock. We then evaluate the model's performance relative to the data in one particular episode, the first Persian Gulf War. The model reproduces quite well the more than 50% jump in the output of Saudi Arabia in response

to the combined output collapse of Iraq and Kuwait, as well as the initial jump of the oil price from \$19 to \$34 per barrel. From the point of view of our model, this behavior of the dominant firm is entirely consistent with its own profit-maximizing objective, as opposed to alternative non-economic (e.g. foreign policy) considerations. It is optimal for the dominant firm to increase its output substantially, but not so much as to fully offset the output collapse of the fringe, so that the oil price rises in tandem with oil sales. Our model thus explains the high volatility of Aramco's oil output as the endogenous response to shocks affecting fringe producers and oil importers.

Finally, since the economy is distorted by the market power of the dominant firm, in Section 5 we assess the desirability, from the point of view of the oil-importing region, of two policy interventions: a tax on oil consumption, and a subsidy to oil production by the fringe. Unlike partial equilibrium settings assuming an exogenous oil price, our framework is well suited for tracing the general equilibrium effects of these measures, working through the endogenously determined price of oil. We find that an oil consumption tax is not necessarily welfare improving, and that its effects depend on whether there are constant or decreasing returns to scale in oil production. In particular, in the presence of constant returns, an oil production subsidy to the fringe does a better job at mitigating the market distortion than an oil consumption tax. The optimal subsidy to fringe production strikes a balance between reducing the market power of the dominant firm and allocating resources to the more efficient oil production technology.

1. Introduction

An extensive literature studies the transmission of oil price shocks assuming these are given exogenously, typically as an AR(1) driving process (e.g. Kim and Loungani, 1992; Leduc and Sill, 2004; Carlstrom and Fuerst, 2005). Yet oil price shocks do not happen in isolation. Rather, like other relative price changes, they are triggered by deeper shocks to preferences, technology, and policy. The international oil price is determined, jointly with the oil quantity produced and consumed, as an equilibrium outcome of the interaction of different market participants. Indeed, Barsky and Kilian (2002, 2004), and Kilian (2009), have found overwhelming evidence that the oil price is affected significantly by global economic conditions. This is important both from a positive point of view, as shown by Kilian (2009), and from a normative perspective, as demonstrated by Nakov and Pescatori (2010) in the context of monetary policy.

Despite all this evidence, few equilibrium models treat the oil price as endogenous, and even fewer determine endogenously both the oil price and the oil quantity. Recent progress includes papers by Bodenstein et. al. (2011), and Leduc and Sill (2007), in which the oil price is determined endogenously while oil supply is given as an exogenous endowment. Backus and Crucini (2000) partially endogenize also oil output by modelling it as the sum of two terms: an exogenous shock meant to represent unpredictable OPEC supply changes; and a term related to economic activity, which represents competitive oil supply.

This paper presents a general equilibrium framework designed to approximate the industrial structure of the oil market, while jointly determining the oil price, oil production and consumption, as equilibrium outcomes of the optimizing decisions of oil importers and oil exporters. On the supply side we model the oil market as a dominant firm with competitive fringe, in the spirit of Salant (1976).¹ He used a similar setup to study the Nash-Cournot equilibrium in which the dominant firm takes as given the output path of the competitive fringe. Instead, we build this industrial organization model into a general equilibrium framework, assuming that the dominant firm internalizes *à la* Stackelberg the behavioral responses of fringe producers and oil consumers. Thus, the dominant firm in our model understands and exploits the fact that its oil output and demand for inputs affects the supply of fringe producers, oil demand, and the oil price.

The focus of our paper is thus on explaining the behavior of the dominant firm and on the distortionary effects it has on the global oil market. It differs from that of Backus and Crucini, who emphasize the terms of trade effects of exogenous oil endowment shocks, and from Bodenstein et. al. who focus on the real exchange rate and the trade balance. Hassler et. al. (2010) solve the problem of a single monopolist, analyzing the implications for climate change. In contrast, we study the problem of a dominant firm with a large competitive fringe, focusing on the implications for business cycle dynamics. In particular, although any individual fringe firm is small and has no impact on the oil price, collectively fringe producers limit substantially the market power of the dominant firm.

We calibrate our dominant firm to data on Saudi Aramco, the 100% state-owned oil company of Saudi Arabia. As Adelman (1995) and Alhajji and Huettner (2000) convincingly argue, the behavior of Aramco can hardly be described as competitive: it is by far the largest oil firm in the world (see Table 1), accounting for more than a tenth of global oil production and a fifth of total proven reserves. Even though it has the lowest production costs, it is the only firm which maintains ample spare capacity and restricts investment in developing available reserves

¹In the energy literature the “dominant firm with competitive fringe” view of the oil industry is quite common. Examples include Mabro (1975), Adelman (1986, 1995), Alhajji and Huettner (2000a,b), Almuoguer and Herrera (2007).

(Smith, 2009).² Moreover, its output is negatively correlated with that of other OPEC producers (ibid.). Indeed, its output has been highly volatile compared to other major producers, even though Aramco itself has experienced few shocks affecting its production directly (Fig.1).³

We use our model to address three sets of questions. First, in Section 3 we derive a set of analytical results related to the secular features of the economy that are independent of the particular calibration. Namely, the technological conditions required for the existence of a dominant oil firm, its degree of capacity utilization, and the conditions for balanced growth. We establish that a dominant firm may exist profitably as long as it enjoys a permanent cost advantage with respect to the fringe, and we offer an expression for the optimal markup as a function of this technological advantage. We also find that such a firm would optimally maintain spare capacity relative to the fringe.

Second, in Section 4 we study the dynamic properties of our calibrated economy. With only two productivity shocks, affecting the output of the oil-importing region and of fringe producers, the model is capable of reproducing the monthly volatility of the oil price, of the oil output of Aramco and that of the fringe, as well as of US industrial production. We show that the comovement between the oil price and GDP of the oil-importing region may be positive or negative depending on the origin of the shock. We then evaluate the model's performance relative to the data in one particular episode, the first Persian Gulf War. The model reproduces quite well the more than 50% jump in the output of Saudi Arabia in response to the combined output collapse of Iraq and Kuwait, as well as the initial jump of the oil price from \$19 to \$34 per barrel. From the point of view of our model, this behavior of the dominant firm is entirely consistent with its own profit-maximizing objective, as opposed to alternative non-economic (e.g. foreign policy) considerations. It is optimal for the dominant firm to increase its output substantially, but not so much as to fully offset the output collapse of the fringe, so that the oil price rises in tandem with oil sales. Our model thus explains the high volatility of Aramco's oil output as the endogenous response to shocks affecting fringe producers and oil importers.

Finally, since the economy is distorted by the market power of the dominant firm, in Section 5 we assess the desirability, from the point of view of the oil-importing region, of two policy interventions: a tax on oil consumption, and a subsidy to oil production by the fringe. Unlike partial equilibrium settings assuming an exogenous oil price, our framework is well suited for tracing the general equilibrium effects of these measures, working through the endogenously determined price of oil. We find that an oil consumption tax is not necessarily welfare improving, and that its effects depend on whether there are constant or decreasing returns to scale in oil production. In particular, in the presence of constant returns, an oil production subsidy to the fringe does a better job at mitigating the market distortion than an oil consumption tax. The optimal subsidy to fringe production strikes a balance between reducing the market power of the dominant firm and allocating resources to the more efficient oil production technology.

²According to the International Energy Agency, Aramco's spare capacity in 2009 stood at 25%. To put this in perspective, it is equivalent to the entire production of oil-exporting countries such as Kuwait or Venezuela.

³According to the U.S. Energy Information Administration's "Official oil market chronology", the only instances when Saudi oil production was directly affected by exogenous events were a fire at the Abqaiq facilities which halved production in 1977, the "tanker war" in 1984, when several Saudi tankers were destroyed, and the attacks in 1991 by Iraqi missiles during the first Gulf War. Apart from these singular episodes, changes in Saudi oil production were the result of business decisions, and not the consequence of exogenous disruptions to their production capabilities.

2. Model

Our model of the global oil market comprises three regions: one oil-importing and two oil-exporting. The oil-importing region employs labor in the production of final goods, part of which are consumed domestically, and the rest are exported to the two oil-producing regions. The oil-importing region imports oil for use in consumption. The inclusion of oil as a consumption good rather than as a factor of production is in part motivated by Kilian (2008). His estimates suggest that the primary channel of transmission of oil price shocks to the economy is through their effect on households' demand for goods and services, rather than as an input in production. In addition, this is a simple way of postulating inelastic oil demand, which is key for matching the high oil price volatility (relative to oil consumption volatility) observed in the data.

Oil is a homogeneous commodity supplied by a dominant firm and a fringe of competitive oil producers. The fringe take the oil price as given when choosing their production level. The dominant firm faces a downward sloping residual demand curve and picks the profit-maximizing points on that curve at each point in time and in each state of nature. Oil exporters produce oil only and their revenue is recirculated to the oil-importing region in the form of demand for final consumption and investment goods.

Our model of the oil industry is in the spirit of Salant (1976) who used a similar structure to study the Nash-Cournot equilibrium in which the dominant firm takes as given the output path of the fringe. Instead, we build this industrial structure into a general equilibrium framework, assuming that the dominant firm is aware that it can manipulate *à la* Stackelberg the choices both of the competitive fringe, and of the oil importer. For example, the dominant firm understands that a change in its own oil supply and demand for goods from the oil-importing region will have an impact on oil demand, oil supply by fringe producers, and the oil price.

Except for the difference in market power, which is founded on a technological gap between the two types of oil producers, the competitive fringe and the dominant firm are modelled symmetrically. This implies that, by appropriate choice of parameters of the oil production technology, our model spans the space from the extreme case of perfect competition to the opposite extreme of single monopoly. Our preferred calibration is one in which the dominant oil producer maintains an average market share of 12%, which corresponds roughly to the market share of Saudi Arabia since 1973.

2.1. Oil-importing region

A representative household has a period utility function which depends on consumption, C_t , oil O_t , and labor L_t , and takes the form

$$U(C, O, L) = \log(C) + \nu_t O^{1-\eta}/(1-\eta) - L^{1+\psi}/(1+\psi),$$

where η and ψ are oil demand and labor supply elasticities. We allow for a trend in ν_t to capture some secular features of the data, while ensuring the existence of a balanced growth path.

The household faces the period budget constraint

$$C_t + s_t O_t = w_t L_t + D_t, \tag{1}$$

which equates income from labor, $w_t L_t$, and dividends from the ownership of firms, D_t , to outlays on consumption, C_t , and oil, $s_t O_t$; s_t denotes the real price of oil, while w_t denotes the real wage rate. We assume that all oil must be consumed within the period of production.⁴

⁴As Litzenger and Rabinowitz (1995) point out, oil storage "above ground" is limited because of the high physical storage cost. Most of the oil "stored" above ground is oil in transit in pipelines or tankers. Relaxing this assumption to allow for delay between production and consumption does not affect significantly our main results.

The household chooses C_t , O_t , and L_t , to maximize expected present discounted utility

$$\max_{C_t, O_t, L_t} E_o \sum_{t=0}^{\infty} \beta^t U(C_t, O_t, L_t)$$

subject to the budget constraint (1), where $0 < \beta < 1$ is a time preference parameter.

Final goods are produced with labor only by a representative price-taking firm, according to

$$Y_t = Z_t L_t^\alpha \quad (2)$$

with $\alpha \leq 1$. Aggregate total factor productivity Z_t follows an AR(1) process in log differences: $\hat{g}_t^z = \rho_z \hat{g}_{t-1}^z + \varepsilon_t^z$, where $\hat{g}_t^z \equiv g_t^z - g^z$, with $g_t^z \equiv \log(Z_t/Z_{t-1})$, $g^z = E(g_t^z)$ and ε_t^z is a Gaussian innovation with mean zero and variance σ_z^2 .

Utility maximization by households implies the following oil demand curve

$$\nu_t C_t = s_t O_t^\eta \quad (3)$$

and labor supply curve

$$C_t L_t^\psi = w_t. \quad (4)$$

Profit maximization by price-taking firms yields the following labor demand curve

$$w_t L_t = \alpha Y_t. \quad (5)$$

Barring borrowing across regions, final goods consumption equals gross domestic product net of the value of oil imports

$$C_t = Y_t - s_t O_t. \quad (6)$$

2.2. Competitive fringe of oil exporters

A representative household maximizes the present discounted flow of utility from consumption,⁵

$$\max_{\tilde{C}_t, \tilde{I}_t} E_o \sum_{t=0}^{\infty} \beta^t \log(\tilde{C}_t)$$

subject to the period budget constraint

$$\tilde{C}_t + \tilde{I}_t = \tilde{r}_t \tilde{K}_{t-1} + \tilde{D}_t$$

where consumption, \tilde{C}_t , and investment, \tilde{I}_t , are both purchased from the oil-importing region, \tilde{r}_t is the rental price of capital \tilde{K}_t rented out by the household to oil firms, and \tilde{D}_t are oil firm dividends rebated lump sum to the household.

The household invests in capital according to

$$\tilde{K}_t = (1 - \delta) \tilde{K}_{t-1} + \tilde{I}_t,$$

where δ is the depreciation rate of installed capital.⁶

⁵We decorate variables belonging to the competitive fringe by tildes (e.g. \tilde{X}_t), and variables belonging to the dominant firm by hats (e.g. \hat{X}_t).

⁶We abstract from the existence of significant technological lags in the buildup of oil production capital, leaving the analysis of this feature for future work.

A representative fringe firm, owned by the household, maximizes period profits

$$\tilde{D}_t = \max_{\tilde{X}_t, \tilde{K}_{t-1}} \left(s_t \tilde{O}_t - \tilde{X}_t - \tilde{r}_t \tilde{K}_{t-1} \right)$$

subject to the production technology

$$\tilde{O}_t = \tilde{Z}_t \tilde{X}_t^{\tilde{\gamma}} \tilde{K}_{t-1}^{1-\tilde{\gamma}}, \quad (7)$$

taking the oil price as given. The firm buys the intermediate good \tilde{X}_t from the oil-importing region, and rents the capital \tilde{K}_t from the household. Total factor productivity $\tilde{Z}_t = \tilde{A}_t \tilde{Z}_0 \exp(g^{\tilde{z}}t)$ follows an exogenous process with a secular trend component $\tilde{Z}_0 \exp(g^{\tilde{z}}t)$, and a stationary AR(1) component in logs, $\log(\tilde{A}_t) = \rho_{\tilde{A}} \log(\tilde{A}_{t-1}) + \varepsilon_t^{\tilde{A}}$ with persistence $\rho_{\tilde{A}}$ and a Gaussian innovation $\varepsilon_t^{\tilde{A}}$ with mean zero and variance $\sigma_{\tilde{A}}^2$.

Optimal oil supply by competitive fringe producers implies

$$\tilde{\gamma} s_t = \tilde{X}_t / \tilde{O}_t, \quad (8)$$

while their optimal capital accumulation is given by

$$1 = \beta E_t \left[\frac{\tilde{C}_t}{\tilde{C}_{t+1}} \left((1 - \tilde{\gamma}) s_{t+1} \frac{\tilde{O}_{t+1}}{\tilde{K}_t} + (1 - \delta) \right) \right].$$

2.3. Dominant oil exporter

The dominant producer's economy has a structure symmetric to that of fringe producers, except that there is a single oil firm. The firm produces oil, \hat{O}_t , according to

$$\hat{O}_t = \hat{Z}_t \hat{X}_t^{\hat{\gamma}} \hat{K}_{t-1}^{1-\hat{\gamma}}, \quad (9)$$

using an imported intermediate good \hat{X}_t and capital \hat{K}_{t-1} . We assume that the dominant firm's technology evolves deterministically according to $\hat{Z}_t = \hat{Z}_0 \exp(g^{\hat{z}}t)$, that is, unlike the fringe, the dominant firm's output is not directly affected by productivity shocks. Capital is accumulated by purchasing \hat{I}_t units of the investment good from the oil-importing region and the representative household receives a stream of log utility from consumption \hat{C}_t .

The substantial difference with the competitive fringe is that the dominant oil firm has market power: it is aware of the dependence of fringe oil supply, of oil demand, and of the equilibrium oil price on its supply decision. We assume that the dominant firm chooses a state-contingent plan which maximizes the expected present discounted utility of its owners, subject to the optimal behavior of households and firms in the oil-importing region, of competitive fringe producers, and clearing in the oil market.

Thus, the decision problem of the dominant oil producer is to

$$\max E_o \sum_{t=0}^{\infty} \beta^t \log(\hat{C}_t)$$

subject to demand equations (2)-(6), fringe production (7), (8), its production technology (9) and the oil market clearing condition⁷

$$O_t = \hat{O}_t + \tilde{O}_t.$$

Equilibrium is defined as the solution to the above problem under commitment. The set of first-order conditions is given in appendix B.

⁷We assume that the dominant firm has limited information and does not internalize the capital accumulation condition of competitive fringe producers.



3. Secular features of the oil market

This section describes the balanced-growth properties of the model. In particular, it focuses on the long-run trends of the oil price and oil consumption efficiency consistent with balanced growth, as well as on the optimal oil price markup and capacity utilization of the dominant oil firm.

3.1. Oil prices and efficiency

The model incorporates secular trends in the growth rate of final goods technology (Z_t), oil production technology (\tilde{Z}_t and \hat{Z}_t), and oil efficiency (ν_t). In a steady-state with balanced growth, the ratio $s_t O_t / Y_t$ must remain stable over time. Since O_t grows at rate $g^z + g^{\tilde{z}}$, while Y_t grows at rate g^z , for the ratio $s_t O_t / Y_t$ to remain stable, s_t should grow at rate $-g^{\tilde{z}}$. This is summarized in the following balanced-growth-path

Condition 1 *The real price of oil s_t grows at rate $-g^{\tilde{z}}$ over time.*

The variable ν_t scales the utility of consumption of oil in terms of the utility of consumption of final goods. One of the first-order optimality conditions is the oil demand curve, $\nu_t C_t = s_t O_t^\eta$. Along the balanced growth path, C_t grows at rate g^z while $s_t O_t^\eta$ grows at rate $\eta (g^z + g^{\tilde{z}}) - g^{\tilde{z}}$. Hence, along the balanced growth path ν_t must grow at rate $(\eta - 1) (g^z + g^{\tilde{z}})$. Assuming balanced growth, therefore, supposes the following

Condition 2 *Oil efficiency ν_t grows at rate $(\eta - 1) (g^z + g^{\tilde{z}})$.*

Since in the data the trend in oil production is positive, $(g^z + g^{\tilde{z}}) > 0$, for values of $\eta > 1$, energy efficiency must increase over time.

3.2. Oil price markup and capacity utilization

The dominant oil firm extracts a pure rent by picking the profit-maximizing point on the residual demand curve, where marginal revenue equals marginal cost. Thus, one can derive

Proposition 3 *The average price mark-up of the dominant oil producer is given by*

$$\mu = \Upsilon r^{\hat{\gamma} - \tilde{\gamma}}, \quad (10)$$

where $r = g_z / \beta + \delta - 1$ is the return on capital used in oil production and $\Upsilon = \hat{Z}_0 \hat{\gamma}^{\hat{\gamma}} (1 - \hat{\gamma})^{1 - \hat{\gamma}} / [\tilde{Z}_0 \tilde{\gamma}^{\tilde{\gamma}} (1 - \tilde{\gamma})^{1 - \tilde{\gamma}}]$ is the technological advantage of the dominant firm over the fringe.

Proof. See Appendix A ■

Corollary 4 *In the case in which $\tilde{\gamma} = \hat{\gamma}$ (the two technologies are symmetric), a dominant producer can exist profitably as long as it enjoys an average cost advantage, $\hat{Z}_0 > \tilde{Z}_0$.*

Proof. See Appendix A ■

Corollary 5 *When $\tilde{\gamma} = \hat{\gamma}$, the degree of competition is related to the size of the technological gap $\tilde{Z}_0 / \hat{Z}_0 \in [0, 1]$. As $\tilde{Z}_0 / \hat{Z}_0 \rightarrow 1$, the model approaches perfect competition in oil production.*

Writing the production function for oil as $O = Zu(X, K)K$ we propose the following

Definition 6 *The capacity utilization rate of installed capital is given by $u(X, K)$, with $u(0, K) = 0$, $\partial u/\partial X > 0$, and $\partial^2 u/\partial X^2 < 0$. In the Cobb-Douglas case, $u(X, K) = (X/K)^\gamma$.*

Based on the above definition, we can derive an expression for the capacity utilization of the dominant oil producer *relative* to that of the competitive fringe,

$$u(\hat{X}, \hat{K})/u(\tilde{X}, \tilde{K}) = \left(\hat{X}/\hat{K}\right)^{\hat{\gamma}} / \left(\tilde{X}/\tilde{K}\right)^{\tilde{\gamma}} = (\hat{\gamma}\widehat{MC})^{\frac{\hat{\gamma}}{1-\hat{\gamma}}} / (\tilde{\gamma}\widetilde{MC})^{\frac{\tilde{\gamma}}{1-\tilde{\gamma}}}, \quad (11)$$

where \widehat{MC} and \widetilde{MC} are defined in Appendix A (in equations (17) and (18), respectively). Capacity utilization in the data is computed as the ratio of actual output to “sustainable production capacity”, defined as the production level that can be reached within 30 days and sustained for 90 days. In the model, capacity utilization is computed from equation (11), assuming that the competitive fringe operates at 100 percent of capacity. The dominant producer utilizes less capacity than the fringe as long as the parameters of the model satisfy

$$(\hat{\gamma}\widehat{MC})^{\frac{\hat{\gamma}}{1-\hat{\gamma}}} < (\tilde{\gamma}\widetilde{MC})^{\frac{\tilde{\gamma}}{1-\tilde{\gamma}}}.$$

4. Quantitative analysis of oil supply and demand shocks

4.1. Calibration

The working frequency of our model is monthly. We calibrate the trend and time preference parameters as follows. The secular growth rate of technology of the oil-importing region is set to $\exp(g^z) = 1.03^{1/12}$, consistent with an average world output growth rate of 3 percent per year for the period from 1973 to 2009. Given this, we set the time preference parameter to $\beta = 1.01^{-1/12}$ equivalent to an average real interest rate of 4 percent per year.

Based on the stationary market share of Saudi Arabia in the data, we impose equality between the growth rates of the dominant oil producer and the fringe, $g^{\hat{z}} = g^{\tilde{z}}$. The average growth rate of total oil production is 0.8 percent per year in the data. In the model this must equal the sum of the growth rate of oil production inputs, g^z , and the growth rate of oil production technology itself, $g^{\tilde{z}}$, implying a value for $\exp(g^{\tilde{z}}) = 0.9982$ on a monthly basis. Given Proposition 1, the latter implies that the real oil price must grow at an annual rate of 2.2 percent, which is consistent with the average growth rate actually observed in the data.

Second, we set three parameters governing the oil-importing region’s labor disutility, technology, and capital depreciation, to their typical values in the RBC literature: the labor disutility parameter is set to $\psi = 1$, the labor share in the production of final goods is set to $\alpha = 0.67$; and the depreciation rate is set to $\delta = 1.10^{1/12} - 1$, consistent with 10% annual depreciation of installed capital.

Third, we set $\eta = 21$, consistent with Smith’s (2009) estimates of the short-run price elasticity of oil demand of about 0.05.

Fourth, we calibrate jointly ν and the parameters of the two oil production technologies $(\tilde{Z}_0, \hat{Z}_0, \tilde{\gamma}, \hat{\gamma}, \nu)$ as follows. We normalize the initial level of productivity of the fringe to $\tilde{Z}_0 = 1$. We then search for a vector $(\hat{Z}_0, \tilde{\gamma}, \hat{\gamma}, \nu)$ which allows the model to match an average “oil share” of 5 percent of GDP, as well as the following three averages for Saudi Arabia: a global market share of 12.3%, capacity utilization of 75%, and a price mark-up of 25%. We should stress that our measure of marginal costs includes not only the intermediate input costs but also capital rents. The targets for Saudi Arabia are calculated in the following way. The market share average is computed directly from data on oil output provided by the Energy Information Administration for the period from January 1973 to April 2009. The capacity utilization rate is estimated

around 75 percent based on data from IEA (2009). Data on the average markup is hard to obtain, but we assume it to be around 25 percent based on estimates of long-run marginal costs provided by IEA. This number is also in line with estimates of the marginal revenue of OPEC members cited by Smith (2009), who considers it a lower bound for the marginal revenue of Saudi Arabia. We are able to approximate well these four targets by setting $\tilde{Z}_0 = 2$, $\tilde{\gamma} = 0.4$, $\hat{\gamma} = 0.5$, and $\nu = 10^{-10}$.

With these parameter values, we can compute the permanent consumption loss for the oil-importing region as a result of the strategic behavior of the dominant firm. From the point of view of the oil-importing region the optimal allocation is for all oil to be supplied by the lowest-cost producer at marginal cost. In this hypothetical scenario the superior technology of the dominant firm is operated under free competition, there is no spare capacity, and oil profits are competed away. In that case the oil price would be 20% lower on average, oil production 1.1% higher, and consumption of the oil-importing region would be permanently higher by 0.7%. Over time, this amounts to a substantial loss of resources wasted on producing oil using the inferior technology of the fringe.

Finally, in order to analyze the dynamics, we calibrate the two shock processes as follows.⁸ First, we fix the two monthly persistence parameters to $\rho_Z = 0.944$ and $\rho_{\bar{A}} = 0.944$, which is equivalent to 0.5 on an annual basis.⁹ We then select the standard deviations of the two innovations $\sigma_Z = 0.004$ and $\sigma_{\bar{A}} = 0.04$ so as to best match the monthly standard deviations of the following five series: the log-differences of total oil output, of Aramco's oil output, and of the fringe oil output; the log-difference of the real oil price; and the log-difference of US industrial production (which we use as a proxy for the output volatility of the oil-importing region as a whole).

Tables 2 and 3 summarize the fit of our model to relevant first (Table 2) and second (Table 3) moments of the data.¹⁰ The model is quite successful at reproducing, with only two shocks, the historical volatility of the five series of interest. Namely, while it slightly underpredicts the volatility of total oil output, and somewhat overpredicts the volatility of total final goods output, it matches quite well the historical volatility of the real oil price, Saudi Arabia's oil production, and fringe output. Key parameters for matching the two most volatile series (the real oil price and Saudi output) are the low elasticity of oil demand and the higher variable input share of the dominant oil firm. In particular, the low elasticity of oil demand with respect to the price produces a highly volatile oil price, while the higher variable-input share of the dominant firm (relative to the fringe) explains the more volatile output of Saudi Arabia as the result of maintaining higher average spare capacity.

4.2. Impulse-responses

Figure 2 shows the general equilibrium responses of several variables of interest to each of the two shocks.

The solid line shows the responses to a one-standard deviation drop in the productivity of

⁸Ideally, one would like to estimate these processes from the data. Computing them as Solow residuals is subject to well-known difficulties related to the correct measurement of inputs. On the other hand, a full-blown Bayesian estimation of the model is beyond the scope of this paper.

⁹We simply pick values for the autoregressive coefficients which are consistent with a half-life of one year. This may overstate somewhat the persistence of the *growth rate* of US TFP, but we really mean the demand shock to represent a persistent TFP increase in developing countries such as China, not the US.

¹⁰Data on oil supply are taken from the Energy Information Administration's *Monthly Energy Review*. The real oil price is the nominal spot West Texas Intermediate price, deflated by U.S. CPI. Data on the nominal oil price, U.S. CPI, and U.S. industrial production are taken from the FRED II online database.

the fringe. This is an example of a transitory “negative oil supply shock” with a half-life of one year. As a result, the oil output of the fringe falls by 1.4% on impact, while the output of the dominant oil producer increases by as much as 6%, raising its market share by around 0.8 percentage points. The increased production of the dominant oil firm is not sufficient to fully offset the output decline of the fringe; hence, total oil supply falls by around 0.4%, while the oil price rises by 8%. The latter induces a contraction of domestic consumption in the oil-importing region by around 0.3 percentage points as oil’s share of final goods output increases. These impulse-responses are in line with the observed behavior in episodes of “negative oil supply shocks”, for instance during the Iranian revolution and during the first Gulf war (see figure 1). In both cases the fall in fringe production – of Iran in the first episode and Iraq and Kuwait in the second – was accompanied by a surge in the oil price and a sharp increase in Saudi Arabia’s oil output.

The dashed lines in figure 2 show the responses to a type of “positive oil demand shock”, namely an unexpected one-standard deviation rise in the *growth rate* of TFP of the oil-importing region, resulting in a permanent rise of the *level* of TFP. As a consequence of this shock, the output of the oil-importing region rises gradually towards its new, higher, balanced growth path, as does consumption. The oil price at first rises gradually, peaks at 5% above its balanced growth path around two years after the initial impulse, and then decays back towards its steady-state path. Both the fringe oil producers and the dominant oil supplier increase their output, although the dominant producer does so a bit faster so that its market share slightly increases. These responses are consistent with episodes of (negative) oil demand shocks such as the dotcom bust and the sub-prime crisis. In both cases Saudi Arabia reduced significantly its oil output in response to dwindling global demand and a falling oil price. In contrast to the supply shock, in this case the oil price rise is associated with a rise in the GDP of the oil-importing region following the cumulative rise in its TFP. As pointed out by Kilian (2009), and Nakov and Pescatori (2010), it is natural that different fundamental shocks should induce different co-movement between the oil price, oil supply, and the GDP of the oil-importing region.

The above exercises build confidence that our simple model can explain some of the patterns observed in the data. In particular, the model matches the historical volatilities of oil prices and quantities, and is able to capture the main co-movements among the series. Profit maximization on behalf of the dominant producer with spare capacity goes a long way in explaining why Saudi Arabia increases strongly its output in response to supply disruptions elsewhere, while it reduces its production aggressively when oil importers are hit by a recession.

4.3. *Event study: the Persian Gulf war of 1990-91*

In August 1990 Iraq invaded Kuwait on the pretext that Kuwait was stealing oil across the border by “slant-drilling” into Iraq’s Rumaila oil field. Because of the war and the trade embargo imposed by the UN, the combined oil supply of the two countries fell by more than 5.3 million barrels per day (mbd) (see fig.1), while the oil price surged from \$19 per barrel in July to \$34 per barrel in September 1990. The reaction of Saudi Arabia was quite powerful: it raised its own oil production from 5.4 mbd in July 1990 to 7.6 mbd in September, and to 8.4 mbd by the end of 1990, a 56% increase in just a few months! Other major oil producers did not increase their production significantly: Iran and the US increased their oil output by only 0.2 mbd, while the Soviet Union actually reduced its production.¹¹

The prevailing hypothesis for the decision of Saudi Arabia to drastically increase its own oil

¹¹A similar pattern of collapsing Iraqi oil supply accompanied by a sizable increase in Saudi Arabia’s oil output was observed also in the second Persian Gulf war of 2003.

supply is that it was based on non-economic considerations, such as the desire to support the already flagging US economy facing a war with Iraq (Yergin, 2008). Yet, as we saw in figure 2, this same behavior of increasing its own supply (but not enough to fully offset the impact of the initial shock on the oil price), is also consistent with the purely economic objective of profit maximization by the dominant firm. Indeed, given a negative supply shock experienced by the fringe, it is optimal for the dominant firm to let profits rise through a combined increase of both the oil price and the quantity of oil sold.

In figure 3 we assess the ability of our model to account quantitatively for this episode. In particular, we estimate a sequence of shocks to fringe productivity such that the fringe oil output predicted by the model matches exactly its actual amount in the data (see the top left panel). We then contrast the actual with the model-implied evolution of Saudi Arabia's oil output, total oil production, and the real price of oil.

The figure illustrates that the model is quite successful at predicting the actual surge in Saudi Arabia's oil production, and the fact that this surge was less than sufficient to prevent a decline in total oil output. As for the oil price, the model predicts well the initial jump from around \$20 to around \$35 per barrel; after that, however, the model predicts that the oil price would remain around \$35 while the actual oil price gradually fell back to \$20 in the course of the year. Part of the explanation for the subsequent decline of the oil price may be the US recession from July 1990 to March 1991. This is consistent, from the point of view of our model, with an autonomous drop in US TFP, a channel which we have deliberately shut down in this exercise.

5. Oil taxes and subsidies in general equilibrium

Finally, we turn to the question of policy. We saw in section 3 that the presence of a dominant oil firm forces a markup distortion on the global economy. A natural question is whether the oil-importing region can adopt any suitable fiscal policies to mitigate this distortion. One possible policy is raising a tax on oil consumption;¹² another is giving an investment subsidy to fringe oil producers. Here we analyze these possibilities from the perspective of our general equilibrium model. Since the equilibrium oil price, oil supply, and demand are all likely to be affected by any oil taxes or subsidies, our model, which endogenously determines these variables, is better suited to addressing the question of tax incidence than partial equilibrium frameworks which take the oil price or oil supply as given.

5.1. Fringe investment subsidy: production efficiency vs. competition

We first analyze the case of an oil production subsidy. The latter might work by offsetting what is effectively an oil production tax: the presence of a dominant firm supplying less oil than the competitive level despite lower costs. The way we implement it is as a subsidy to the investment good purchased by the competitive fringe to build up oil production capacity. This implies that the household's budget constraint of the competitive fringe becomes

$$\tilde{C}_t + (1 + \phi)\tilde{I}_t = \tilde{r}_t\tilde{K}_{t-1} + \tilde{D}_t \quad (12)$$

with $\phi < 0$ denoting an investment subsidy (and $\phi > 0$ denoting a tax). We assume that the resources needed to finance this subsidy are raised in a lump-sum manner from the oil-importing

¹²See Bergstrom (1982). See also "Raise the Gas Tax" by Gregory Mankiw in *The Wall Street Journal*, October 20, 2006; as well as the "Pigou Club Manifesto". Mankiw explicitly makes the point that a consumption tax in the U.S. might be useful for capturing some of the oil rents of Saudi Arabia.

region,

$$C_t + s_t O_t = Y_t + T_t \quad (13)$$

with $T_t = \phi \tilde{I}_t$.

Notice that, in principle, such a subsidy can eliminate completely the oil profits of the dominant firm:

Proposition 7 *The oil price markup can be eliminated completely by subsidizing the capital investment of competitive fringe producers. The subsidy which achieves this is*

$$\phi = \Upsilon^{-1} [g_z/\beta - (1 - \delta)]^{\frac{\tilde{\gamma}-\tilde{\gamma}}{1-\tilde{\gamma}}} - 1. \quad (14)$$

Proof. See Appendix A. ■

Yet, there is an important tension in the model between reducing the pure rents of the dominant firm and making use of the more efficient oil production technology. In particular, the production subsidy which eliminates all oil rents is clearly suboptimal, as it involves a complete shift of resources from the more efficient to the less efficient oil production technology.

Figure 4 shows that the maximum welfare is achieved for an oil production subsidy of about 17% of the price of investment ($\phi = -0.17$), equivalent to 0.5% of the GDP of oil importers. The subsidy works by lowering the cost of oil production by the competitive fringe, increasing the global oil supply, and lowering the market share of the dominant firm. Beyond the optimal level, however, any additional subsidy lowers welfare because of the transfer of resources to the fringe and the loss of efficiency involved in using the inferior oil production technology. Importantly, the net gains for the oil-importing region of this policy are quite limited: around 0.03% in terms of domestic consumption, or a net welfare gain equivalent to a permanent increase of consumption by 0.05%, when accounting for the changes in oil consumption and hours worked.

5.2. Oil consumption tax

To study the effects of a proportional tax on oil consumption, we modify the budget constraint of the households of the oil-importing region as follows

$$C_t + (1 + \tau) s_t O_t = w_t L_t + T_t, \quad (15)$$

where $(1 + \tau) s_t$ is the price of oil paid by the consumer and $T_t = \tau s_t O_t$ is a lump-sum rebate.

In general, it is not clear *a priori* how the burden of the tax is shared between oil consumers and oil producers. It is possible, at least in principle, that a higher oil consumption tax, by discouraging oil consumption, reduces the price of oil so that some of the tax is effectively paid by oil producers.

Indeed, we find that oil consumption is discouraged by a positive oil consumption tax (see Figure 5). Resources previously used for oil consumption are now redirected to more final goods consumption and leisure. However, the fact that final goods consumption increases while labor hours are reduced does not necessarily imply that welfare rises. Indeed, the utility gain from increased goods consumption and leisure is more than offset by the loss of utility from less oil consumption. At the same time, an oil consumption subsidy ($\tau < 0$), while increasing oil consumption and the utility derived thereof, reduces final goods consumption and leisure in a way that total utility again is reduced. Thus, in our baseline model with constant returns to scale in oil production, the optimal oil consumption tax is zero from a welfare point of view,

despite the fact that final goods consumption increases while hours worked fall with a positive tax.

The above result, however, hinges on the assumption of constant returns to scale in oil production. As equation (17) shows, constant returns to scale imply that the oil price is determined by technological and preference parameters. A production subsidy in the previous case worked because it lowered the effective rental rate of capital; but with CRS a consumption tax does not work since it has no general equilibrium effect on the oil price and does not affect the market power of the dominant firm. That is, with constant returns to scale the consumption tax introduces an additional distortion without at all addressing the one already in place.

The above result can be overturned if oil is produced under decreasing returns to scale, as would be the case with a fixed factor of oil production. Suppose that capital in oil production is fixed, so that the oil production function, both of the dominant firm and of the competitive fringe, is

$$O_t = Z_t X_t^\gamma, \quad (16)$$

where $\gamma < 1$ and Z_t grows so as to ensure the balanced growth path property.

In this case, the welfare of the oil importer is maximized when the oil consumption tax is as high as 115% ($\tau = 1.15$). As before, a positive oil consumption tax reduces oil consumption and induces a shift towards more final goods consumption and leisure. Differently from the case with constant returns to scale, however, this time part of the tax ends up being paid effectively by the dominant oil firm, whose profits are squeezed by the negative effect of the tax on the price. This result is reminiscent of Bergstrom (1982)'s analysis of the possibility of capturing oil rents with a national excise tax in the context of a model of fixed oil endowment extracted at zero cost.

6. Conclusions

A closer look at the international oil market reveals an industrial structure which is quite different from the perfect competition paradigm. A single firm accounts for more than 10% of global oil supply. This firm restricts investment, maintains spare capacity, and reacts strongly to supply shocks affecting other oil producers. Much of the informal discussion about the role of Saudi Aramco as a "swing producer" focuses on its alleged "stabilization objectives" based on non-economic, e.g. foreign-policy, considerations.

We show that the same facts can be accounted for quantitatively by a fairly standard model of a dominant firm with competitive fringe. In our model, it is in the dominant firms' own interest to increase output when fringe production is hit by a negative shock. The offsetting is incomplete by design, as it is optimal to let the oil price rise in tandem with the firm's oil output.

We use our model to quantify the distortion from the presence of a dominant oil firm. While the loss for oil importers is non-negligible, we find no easy solution through taxes in the oil-importing region, or subsidies to fringe producers. Indeed, the only way such measures could work is by capturing some of the rents of the dominant oil firm by reducing its market power. But it is the dominant firm which is endowed with the low-cost oil production technology, and hence shifting resources away from it comes at the cost of lower oil production efficiency. Importantly, our analysis of oil taxation ignores the environmental impact of oil usage as well as the financial linkages among oil exporters and oil importers, two key issues which we leave for future research.

References

- Adelman, M.: 1986, The Competitive Floor to World Oil Prices, *Energy Journal*, 7(4): 9–31.
- Adelman, M.: 1995, *The genie out of the bottle: World oil since 1970*, The MIT Press.
- Alhajji, A. and Huettner, D.: 2000a, OPEC and Other Commodity Cartels: a Comparison, *Energy Policy*, 28, 1151–1164.
- Alhajji, A. and Huettner, D.: 2000b, OPEC and world crude oil markets from 1973 to 1994: Cartel, oligopoly or competitive?, *Energy Journal* 21(3), 31–60.
- Almoguera, P. and Herrera, A.: 2007, Testing for the cartel in OPEC: Noncooperative collusion or just noncooperation?, *Mimeo*, MSU.
- Backus, D. and Crucini, M.: 2000, Oil prices and the terms of trade, *Journal of International Economics* 50(1), 185–213.
- Barsky, R. B. and Kilian, L.: 2002, *Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative*, National Bureau of Economic Research, Inc, 137–183.
- Barsky, R. B. and Kilian, L.: 2004, Oil and the macroeconomy since the 1970s, *Journal of Economic Perspectives* 18(4), 115–134.
- Bergstrom, T.C.: 1982, On capturing oil rents with a national excise tax, *The American Economic Review*, 72(1), 194–201.
- Bodenstein, M., Erceg, C. J. and Guerrieri, L.: 2011, Oil shocks and external adjustment, *Journal of International Economics*, 83(2), 168–184.
- Carlstrom, C. T. and Fuerst, T. S.: 2005, Oil prices, monetary policy, and counterfactual experiments, *Working Paper 0510*, Federal Reserve Bank of Cleveland.
- Hassler, J. and Krusell, P. and Olovsson, C.: 2010, Oil Monopoly and the Climate, *The American Economic Review, Papers and Proceedings*, 100(2), 460–464.
- International Energy Agency: 2008, *World energy outlook*, OECD, Paris.
- International Energy Agency: 2009, *World energy outlook*, OECD, Paris.
- Kilian, L.: 2008, The economic effects of energy price shocks, *Journal of Economic Literature* 46(4), 871–909.
- Kilian, L.: 2009, Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market, *American Economic Review* 99(3).
- Kim, I.-M. and Loungani, P.: 1992, The role of energy in real business cycle models, *The Journal of Monetary Economics* 29(2), 173–189.
- Leduc, S. and Sill, K.: 2004, A quantitative analysis of oil-price shocks, systematic monetary policy, and economic downturns, *The Journal of Monetary Economics* 51(4), 781–808.
- Leduc, S. and Sill, K.: 2007, Monetary policy, oil shocks, and TFP: Accounting for the decline in US volatility, *Review of Economic Dynamics* 10(4), 595–614.

- Litzenberger, R.H. and Rabinowitz, N.: 1995, Backwardation in oil futures markets: Theory and empirical evidence, *Journal of Finance*, 50(5), 1517–1545.
- Mabro, R.: 1975, Can OPEC hold the line?, *OPEC and the World Oil Market: The Genesis of the 1986 Price Crisis* .
- Nakov, A. and Pescatori, A.: 2010, Monetary policy tradeoffs with a dominant oil producer, *Journal of Money, Credit and Banking*, 42(1), 1-32.
- Salant, S. W.: 1976, Exhaustible resources and industrial structure: A Nash-Cournot approach to the world oil market, *The Journal of Political Economy* 84(5), 1079–1094.
- Smith, J.: 2009, World oil: Market or mayhem?, *Journal of Economic Perspectives* 23(3), 145–164.
- Strauch, R. Meyler A., Beck R., Consolo A., Costantini R., Fidora M., Gattini L., Landau B., Lima A., Lodge D., Lombardi M., Mestre R., Matthia: 2010. Energy markets and the euro area macroeconomy, *Occasional Paper Series* 113, European Central Bank.
- Yergin, D.: 1992, *The Prize: The Epic Quest for Oil, Money & Power*, Free Press.

Tables and figures

Table 1. Output, reserves, and spare capacity of the top 10 oil companies

| Company | Country | State-owned (%) | Output (mbd) | Reserves (bil. barrels) | Spare capacity (% of global supply) |
|-------------|--------------|--------------------|-----------------|----------------------------|--|
| Aramco | Saudi Arabia | 100 | 10.4 | 264.2 | 3.9 |
| NIOC | Iran | 100 | 4.4 | 138.4 | 0.3 |
| Pemex | Mexico | 100 | 3.5 | 12.2 | - |
| CNPC | China | 100 | 2.8 | 22.4 | - |
| Exxon Mobil | US | | 2.6 | 11.1 | - |
| KPC | Kuwait | 100 | 2.6 | 101.5 | 0.5 |
| PdVSA | Venezuela | 100 | 2.6 | 99.4 | 0.3 |
| BP | UK | | 2.4 | 10.1 | - |
| INOC | Iraq | 100 | 2.1 | 115.0 | 0.1 |
| Rosneft | Russia | 75 | 2.0 | 17.5 | - |

Sources: Smith (2009), *Petroleum Intelligence Weekly* (December 4, 2008), and *World Energy Outlook*, International Energy Agency (2009)

Table 2. Data and model averages

| % | Saudi share | Oil/GDP share | Oil price growth | Oil output growth | Final output growth |
|-------|----------------|------------------|---------------------|----------------------|------------------------|
| Data | 12.3 | 5.0 | 2.21 | 0.78 | 2.98 |
| Model | 12.7 | 5.4 | 2.21 | 0.77 | 3.00 |

Table 3. Data and model standard deviations*

| | Oil price | Oil output | Fringe output | Saudi output | Final output |
|-------|--------------|---------------|------------------|-----------------|-----------------|
| Data | 8.5 | 1.7 | 1.5 | 6.9 | 0.7 |
| Model | 8.3 | 1.2 | 1.8 | 6.4 | 1.2 |

*Standard deviations, in percentage points, of first log differences.




Figure 1: Oil output (tbpd) by six large OPEC producers, 1973-2009

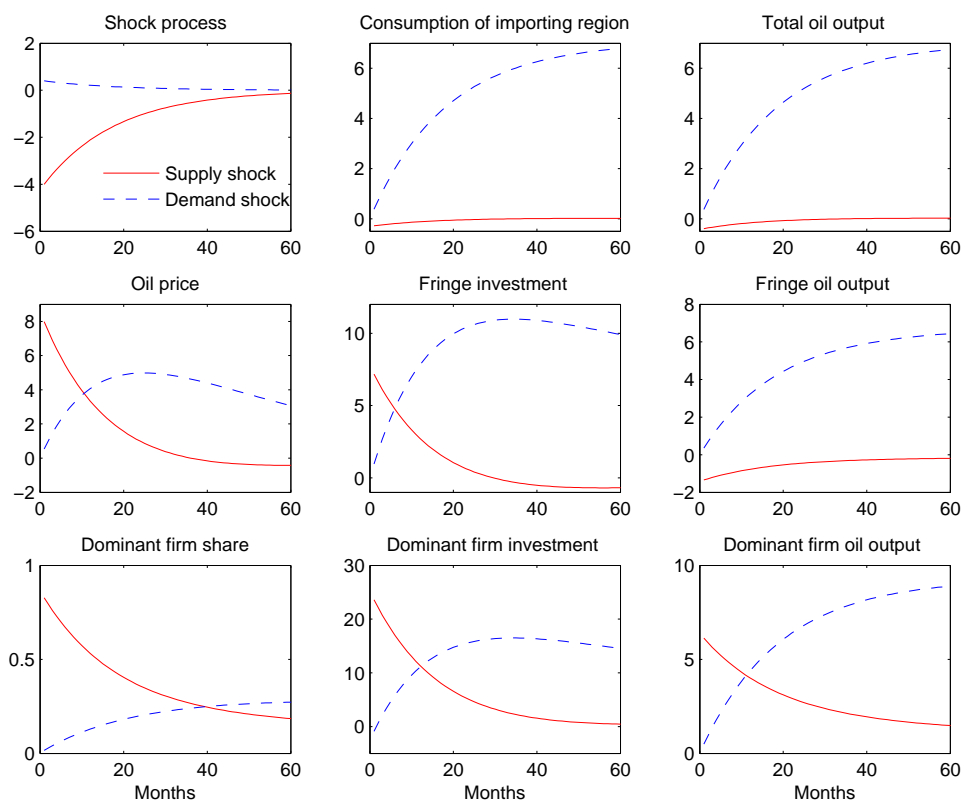


Figure 2: Responses to an oil supply and an oil demand shock

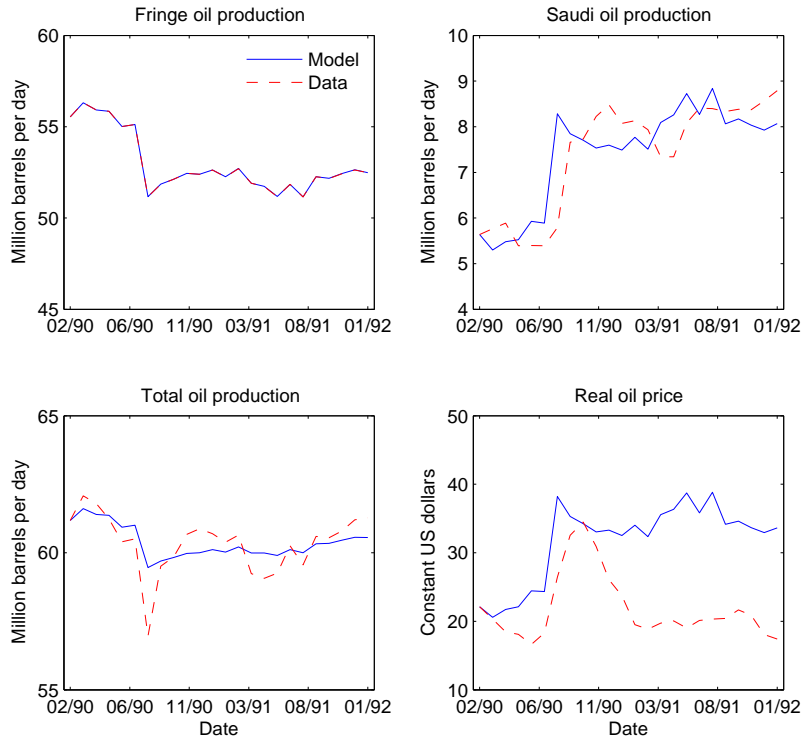


Figure 3: Oil output and prices during the 1991 Gulf war




Figure 4: The effects of a subsidy to fringe investment

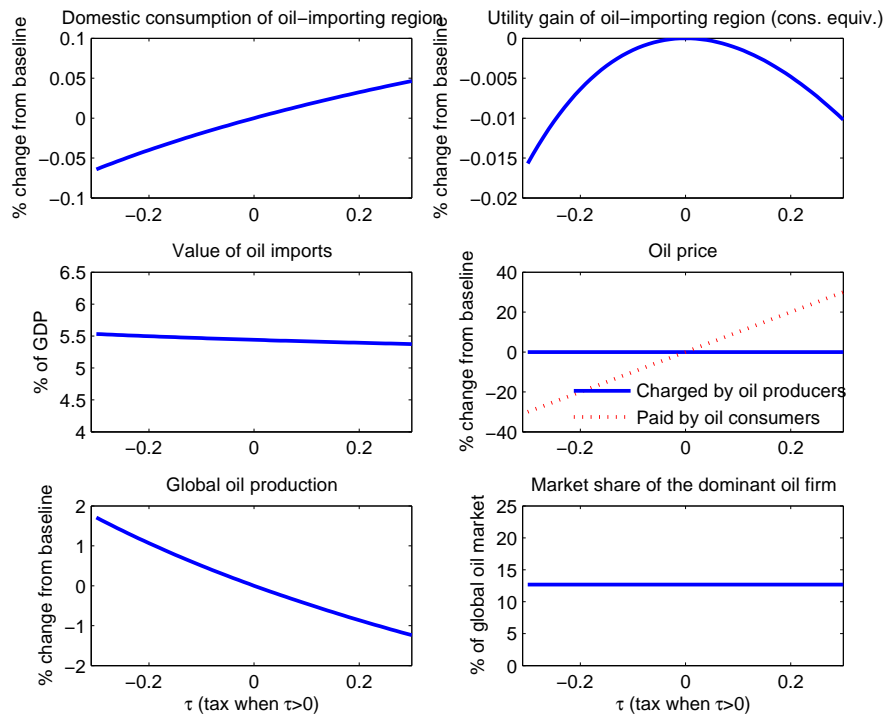


Figure 5: The effects of an oil consumption tax

Appendix A: Proofs

Proposition 3

Proof. Since there are no barriers to entry in the competitive fringe, fringe producers must earn zero profits. Thus, the real price of oil must equal the marginal cost of the competitive fringe,

$$s_t = \widehat{MC}_t = \widehat{r}_t^{1-\tilde{\gamma}} / \left[\widehat{Z}_t \tilde{\gamma}^{\tilde{\gamma}} (1 - \tilde{\gamma})^{1-\tilde{\gamma}} \right]. \quad (17)$$

The same formula replacing hats with tildes gives the marginal cost of the dominant producer,

$$\widetilde{MC}_t = \widehat{r}_t^{1-\hat{\gamma}} / \left[\widehat{Z}_t \hat{\gamma}^{\hat{\gamma}} (1 - \hat{\gamma})^{1-\hat{\gamma}} \right]. \quad (18)$$

The period t oil price markup for the dominant firm is then

$$\mu_t = \frac{s_t}{\widetilde{MC}_t}. \quad (19)$$

In steady-state the above expression reduces to (10), noticing that the steady-state rental price of capital, pinned down by preferences and technology, is the same across oil producers. ■

Corollary 4

Proof. In the symmetric case ($\tilde{\gamma} = \hat{\gamma}$), the average oil price markup (10) reduces to $\mu = \widehat{Z}_0 / \widetilde{Z}_0$. The dominant producer will thus be profitable if and only if $\widehat{Z}_0 > \widetilde{Z}_0$. ■

Proposition 7

Proof. Taking the investment subsidy into account, the markup of the dominant firm is

$$\mu = \Upsilon (g_z / \beta - (1 - \delta))^{\hat{\gamma}-\tilde{\gamma}} (1 + \phi)^{1-\tilde{\gamma}}. \quad (20)$$

Solving this for ϕ while setting $\mu = 1$ yields expression (14). ■

Appendix B: Equilibrium conditions of the dominant firm

Production function, capital accumulation equation, and resource constraint:

$$\widehat{O}_t = \widehat{Z}_t \widehat{X}_t^{\hat{\gamma}} \widehat{K}_{t-1}^{1-\hat{\gamma}} \quad (21)$$

$$\widehat{K}_t = (1 - \delta) \widehat{K}_{t-1} + \widehat{I}_t \quad (22)$$

$$\widehat{C}_t = s_t \widehat{O}_t - \widehat{I}_t - \widehat{X}_t \quad (23)$$

First-order conditions of the maximization problem:

$$\lambda_{3t} = \lambda_{2t} s_t (1 - \tilde{\gamma}) \quad (24)$$

$$\widehat{O}_t = \lambda_{1t} O_t^\eta - \lambda_{2t} \tilde{\gamma} \widehat{O}_t + \lambda_{4t} (O_t^\eta + \nu_t O_t) \quad (25)$$

$$0 = \eta s_t O_t^{\eta-1} (\lambda_{1t} + \lambda_{4t}) + \lambda_{3t} + \lambda_{4t} \nu_t s_t \quad (26)$$

$$0 = \lambda_{1t} (\alpha - 1 - \psi) s_t O_t^\eta + \alpha \lambda_{4t} \nu_t Z_t L_t^\alpha \quad (27)$$

$$\widehat{X}_t = \hat{\gamma} (s_t + \lambda_{3t}) \widehat{O}_t \quad (28)$$

$$1 = \beta E_t \left[\frac{\widehat{C}_t}{\widehat{C}_{t+1}} \left((1 - \tilde{\gamma}) \frac{(s_{t+1} + \lambda_{3t+1}) \widehat{O}_{t+1}}{\widehat{K}_t} + (1 - \delta) \right) \right] \quad (29)$$

