

DEPARTMENT OF ECONOMICS
OxCarre (Oxford Centre for the Analysis of
Resource Rich Economies)

Manor Road Building, Manor Road, Oxford OX1 3UQ
Tel: +44(0)1865 281281 Fax: +44(0)1865 281163
reception@economics.ox.ac.uk www.economics.ox.ac.uk



OxCarre Research Paper 92

Optimal Oil Production and the World Supply of Oil

Nikolay Aleksandrov
Nomura International

Raphael Espinoza*
IMF

&

Lajos Gyurko
Mathematical Institute, University of Oxford

***OxCarre External Research Associate**

Optimal Oil Production and the World Supply of Oil

Nikolay Aleksandrov*, Raphael Espinoza[§] and Lajos Gyurkó[¶]

June 2012

Abstract

We study the optimal oil extraction strategy and the value of an oil field using a multiple real option approach. The numerical method is flexible enough to solve a model with several state variables, to discuss the effect of risk aversion, and to take into account uncertainty in the size of reserves. Optimal extraction in the baseline model is found to be volatile. If the oil producer is risk averse, production is more stable, but spare capacity is much higher than what is typically observed. We show that decisions are very sensitive to expectations on the equilibrium oil price using a mean reverting model of the oil price where the equilibrium price is also a random variable. Oil production was cut during the 2008-2009 crisis, and we find that the cut in production was larger for OPEC, for countries facing a lower discount rate, as predicted by the model, and for countries with government finances less dependent on oil revenues. However, the net present value of a country's oil reserves would be increased significantly (by 100 percent, in the most extreme case) if production was cut completely when prices fall below the country's threshold price. If several producers were to adopt such strategies, world oil prices would be higher but more stable.

Keywords: Oil production ; Real Options ; Capacity Expansion ; Equilibrium Price of Oil; OPEC

JEL Classification: C61 ; Q30 ; Q43

*Nomura International plc., 1 Angel Lane, London EC4R 3AB, UK
email: nikolay.aleksandrov@chch.oxon.org

[§]Corresponding author
International Monetary Fund, Research Department, 700 19th Street NW, Washington DC 20431, and
University of Oxford, Oxford Centre for the Analysis of Resource-Rich Economies
email: respinoza@imf.org

[¶]University of Oxford, Mathematical Institute, 24-29 St Giles, Oxford OX1 3LB, UK, and
Oxford-Man Institute of Quantitative Finance, University of Oxford, Eagle House, Walton Well Road,
Oxford, OX2 6ED, UK
email: gyurko@maths.ox.ac.uk

1 Introduction

In this paper we investigate the optimal oil extraction strategy of a small oil producer facing uncertain oil prices.¹ We use a multiple real option approach. Extracting a barrel of oil is similar to exercising a call option, *i.e.* oil production can be modeled as the right to produce a barrel of oil with the payoff of the strategy depending on uncertain oil prices. Production is optimal if the payoff of extracting oil exceeds the value of leaving oil under the ground for later extraction (the continuation value). For an oil producer, the optimal extraction path corresponds to the optimal strategy of an investor holding a multiple real option with finite number of exercises (finite reserves of oil). At any single point in time, the oil producer is also limited in the number of options he can exercise, because of capacity constraints.

Our first contribution is to present the solution to the stochastic optimization problem as an exercise rule for a multiple real option and to solve the problem numerically using the Monte Carlo methods developed by Longstaff and Schwartz (2001), Rogers (2002), and extended by Aleksandrov and Hambly (2010), Bender (2011), and Gyurkó, Hambly and Witte (2011). The Monte Carlo regression method is flexible and it remains accurate even for high-dimensionality problems, *i.e.* when there are several state variables, for instance when the oil price process is driven by two state variables, when extraction costs are stochastic, or when the size of reserves is a random variable.

We solve the real option problem for two countries with similar capacity profiles but vastly different reserves (Brazil and the U.A.E.) and compute the threshold below which it is optimal to defer production. In our baseline model, we find that Brazil, which has small reserves, should only produce when prices are high (higher than US\$ 73 per barrel at 2000 constant prices), whereas for the U.A.E., full production is optimal as soon as prices exceed US\$ 39. Optimal production is found to be volatile given the stochastic process of oil prices. As a result, we show that the net present value of oil reserves would be substantially higher if countries were willing to vary production when oil prices change. This result has important implications for oil production policy and for the design of macroeconomic policies that depend on inter-temporal and inter-generational equity considerations. It also implies that the world supply curve would be very elastic to prices if all countries were optimizing production as in the baseline model - and as a result, prices would tend to be higher but much less volatile.

We investigate why observed production is not as volatile as what is predicted by the baseline calibration of the model. One possible explanation is that producers are risk averse. Under this assumption, production is accelerated and is more stable, but a risk averse producer should also maintain large spare capacity, a result at odds with the evidence that oil producers almost always produce at full capacity. A second potential explanation is that producers are uncertain about the actual size of their oil reserves. Using panel data on recoverable reserves, we show however that, historically, this uncertainty has been diminishing with time and therefore this explanation is incomplete, since even mature oil exporters maintain low spare capacity. A third explanation may be that the oil price process, and in particular the

¹We are grateful for comments and suggestions from two anonymous referees, Christian Bender, Ben Hambly, Vicky Henderson, Sam Howison, Robert Kaufmann, Rick van der Ploeg, Oral Williams, and seminar and conference participants at Oxford, Princeton, the IMF, the EEA Annual Meeting and the ECB Workshop on Coping with Volatile Oil and Commodity Prices. The views expressed in this paper are those of the author and do not necessarily represent those of the IMF or IMF policy.

equilibrium oil price, is unknown to the decision makers. Indeed, the optimal reaction to an increase in oil prices depends on whether the price increase is perceived to be temporary or to reflect a permanent shift in prices. If shocks are known to be primarily temporary, production should increase in the face of oil price increases. But if shocks are thought to be accompanied by movements in the equilibrium price, the continuation value jumps at the same time as the immediate payoff from extracting oil. In that case an increase in price may not result in an increase in production. Faced with uncertain views on the optimal strategy, the safe decision might well be to remain prudent with changes in production.

In practice, world oil production is partially cut in the face of negative demand shocks. The last section of the paper investigates whether the reduction in oil production during the 2008-2009 crisis can be explained by the determinants predicted by the model. We find that the cut in production was larger for OPEC, for countries facing a lower discount rate, as predicted by the model, and for countries with government finances less dependent on oil revenues.

Section 2 provides a survey of the related literature on optimal production and real options, while Sections 3, 4 and 5 cover the model formulation and calibration. Section 6 describes briefly the oil sector in the two countries used as applications. Section 7 presents the basic set of results and Section 8 discusses some limitations of the model. Section 9 investigates the determinants of production strategies during the 2008-2009 crisis and Section 10 concludes on the price-elasticity of the world supply of oil.

2 Related literature

2.1 Optimal Oil Production

The study of the economy of non-renewable resource extraction started with Hotelling (1931), who showed in a deterministic general equilibrium model that the price of the resource would grow at the rate of interest in competitive markets with constant extraction costs. General equilibrium models later included the effect of uncertainty in technology, the size of the resources, or the availability of substitutes. Partial equilibrium models in which the prices are given, but the decision to extract is a function of the stochastic price process, have a shorter history in the non-renewable resource literature, starting with Tourinho (1979a). Tourinho (1979a, 1979b) analyzed for the first time the valuation of resources in the context of a real ‘call’ option to exploit a field, using the Black and Scholes framework.

Paddock, Siegel and Smith (1988) later developed a model that became a popular approach for decisions on upstream oil investments, in which a company has the option to explore an area and in case oil is discovered, to commit to an immediate development investment before a given date (the time to expiration). If the firm does not exercise the option to develop the field until this date, the firm must return the concession rights back to a national authority.² The model took into consideration resource depletion when estimating the value of the oil field. However, the issue of when to extract the resource after the field is developed is completely absent since the only decision to take is the optimal timing to develop the field.

²Real option models are surveyed in Dixit and Pindyck (1994) and, for applications to oil investments, in Dias (2004).

Cherian *et al.* (1998) studied optimal production of a nonrenewable resource as a control problem in continuous time. The authors solved the Bellman nonlinear partial differential equation (PDE) numerically using the Markov chain approximation technique of Kushner (1977) and Kushner and Dupuis (1992). The cost of extraction in Cherian *et al.* (1998) is a function of time, the extraction rate and the total extracted amount, a model that fits well extraction costs for ore mining. However, a drawback of these numerical solutions to the Bellman PDEs is that they are feasible only for problems with low dimensionality, and it is therefore difficult to solve the problem with two-factor oil price processes, stochastic extraction costs, or stochastic reserves.

Work close to ours in terms of modeling assumptions is Caldentey *et al.* (2006), who study the optimal operation of a copper mining project when the copper spot price follows a mean reverting stochastic process. The project is modeled as a collection of blocks (minimal extraction units) each with its own mineral composition and extraction costs. The authors are interested in maximizing the economic value of the project by controlling the sequence and rate of extraction as well as investing on costly capacity expansions. Our model is more general since we allow for multiple exercise (*i.e.* the company can choose to extract different volumes every period) and because in our model the firm is able to scale down production.

We provide a multiple real option solution to the stochastic optimization problem. Production capacity is exogenous but it is a function of time. We also assume there is a minimum extraction capacity that is non-zero. This assumption captures the fact that some minimal extraction is often needed to finance the functioning of the firm, the transfers to the government, or even the spending of the government in countries where oil proceeds are the major source of government revenues.

2.2 Numerical Solutions for Real Options

The optimal extraction problem faced by an oil producer is similar to many other stochastic optimization problems described in macroeconomics (e.g. investment under uncertainty, see for instance Sakar, 2000), in finance (e.g. the choice of bank capital given uncertain cash flows, see Milne and Whalley, 2001, and Peura and Keppo, 2006) and in the natural resources literature (e.g. the optimal exploitation of forest, see Alvarez and Koskela, 2006). Closed formulae for the price of early exercise options have not been derived yet, even for the simplest cases. In particular, there is no analytical method for solving multiple real option problems. The difficulty in any option model is to compute the *continuation value* (the expected value of delaying the extraction of a barrel). Following the analysis of Arrow, Blackwell and Girshick (1949), the problem was recognized and discussed as an abstract optimal stopping problem by Snell (1952), and the first application of the optimal stopping problem to finance appeared in Bensoussan (1984).

The literature has however suggested various analytical approximations and numerical methods. For most option pricing problems, three numerical methods are available: lattices, finite difference, and Monte Carlo methods.³ The first two approaches work best for simple options on a single underlying (a single state variable). When there are more state variables and the

³A new method based on generalized linear complementarity problems (GLCPs) has recently been proposed by Nagae and Akamatsu (2008).

dimension of the problem increases, however, the Monte Carlo approach is preferred as the performance of the lattice and finite difference schemes is poor (the computational effort with these two methods grows exponentially with the number of state variables).

The first attempt to apply Monte Carlo techniques to American option pricing is due to Tilley (1993), while Broadie and Glasserman (1997) developed the first algorithm in which the suggested lower and upper bound estimates are proved to converge to the true value. Their approach can also deal with high-dimensional American options, but the computational effort still grows exponentially with the number of possible exercise dates. There has been a renewal of interest in the recent years for these methods (Jaillet, Ronn and Tompaidis, 2004; Meinshausen and Hambly, 2004; Carmona and Touzi, 2008). The solution we use was developed in Aleksandrov and Hambly (2010), Bender (2011) and Gyurkó, Hambly and Witte (2011) as an extension to the Monte Carlo method proposed by Longstaff and Schwartz (2002) and Tsitsiklis and van Roy (2001).

The method relies on approximating the value function by linear regression on a suitable space of basis functions (see section 3 for more details). The fitted value from the regression gives an estimate for the continuation value. By construction this optimal stopping policy gives a *lower* bound for the option price - only the exact decision rule would give the maximum value and an approximation can only give a lower estimate. The method is comparatively easy to implement and for properly chosen regression functions gives a good estimate of the value function, even with models with several states variables (see Glasserman, 2003).

3 Model Formulation and Numerical Solution

3.1 Model Formulation

We present here the stochastic optimization problem. We consider an economy in discrete time defined up to a finite time horizon of T years at which we assume reserves will be depleted.⁴ The maximum yearly capacity for extraction (the production capacity) is assumed to be exogenous, although in practice it would be a function of past investments. Production capacity is noted k_t (in billion barrels per year), $t = 1, 2, \dots, T$.⁵ The optimal extraction strategy maximizes the discounted utility of the cash flow from oil sales $V_t^{*,m,\mathbf{k}}$ at time t , subject to the capacity constraints $\mathbf{k} = \{k_0, k_1, k_2, \dots, k_T\}$ and to the total oil reserves constraint m . If the oil producer decides to extract one barrel of oil at time t , the profit is S_t , where S_t is the price of oil (net of the extraction cost c_t). Profits are always positive when the producer decides to extract oil since she is not forced to produce making losses. We abstract from the costs of shutting down an oil production unit.

We assume that oil prices follow a discrete Markov chain process $(S_t)_{t=0,1,\dots,T} \in \mathbb{R}^d$ (we come back to the discussion on oil prices in section 4). We define an extraction policy $\pi_{\mathbf{k}}$ to be a set of ‘stopping’ times (*i.e.* times at which the real options are exercised) $\{\tau_i\}_{i=1}^m$, $\tau_1 \leq \tau_2 \leq \dots \leq \tau_m$, such that the number of exercises (*i.e.* annual production) is lower than production capacity each year: $\#\{j : \tau_j = s\} \leq k_s$.

⁴In several of our applications, we assume that countries extract every year a minimum amount of barrels, in which case extraction occurs indeed over a finite horizon.

⁵ k_t must be a multiple of the discretization unit that we use and that represents one real option to extract oil

The instantaneous utility of extracting h_t units of oil on top of the minimum level l of annual extraction at time t is denoted by $u((h_t + l)S_t)$ where $u(x) = \frac{x^{1-\gamma}-1}{1-\gamma}$ is the Constant Relative Risk Aversion utility function defined over \mathbb{R}^+ , and $\gamma \geq 0$, $\gamma \neq 1$ is the coefficient of risk aversion.⁶ $1/B_t$ is the discount factor, *i.e.* if the rate of impatience is equal to a constant risk-free rate r , $B_t = e^{rt}$. Then, the optimal consumption problem can be formulated as follows.

Definition 1

$$V_t^{*,m,\mathbf{k}} = \sup_{\pi_{\mathbf{k}}} V_t^{\pi_{\mathbf{k}},m} = \sup_{\pi_{\mathbf{k}}} \mathbb{E}_t \left[\sum_{i=1}^m \frac{u((h_{\tau_i} + l)S_{\tau_i})}{B(t, \tau_i)} \right].$$

The corresponding optimal policy is $\pi^* = \{\tau_1^*, \tau_2^*, \dots, \tau_m^*\}$.

3.2 Numerical Solution

We estimate the optimal extraction strategy using Monte-Carlo techniques. In particular, our numerical approach is based on the dynamic programming formulation:

$$V_t^{*,n,\mathbf{k}}(s) = \sup_{0 \leq h \leq \min\{k_t, n\}} \left\{ u((h + l)s) + \frac{1}{B_{t,t+1}} \mathbb{E} \left[V_{t+1}^{*,n-h,\mathbf{k}}(S_{t+1}) \mid S_t = s \right] \right\} \quad (1)$$

for $t = 0, \dots, T - 1$ and $V_T^{*,n,\mathbf{k}}(s) = u(\min\{k_T, n\}s)$.

The conditional expectation in (1), is approximated using least squares regression techniques. In the case where the level of reserves is assumed to change due to extraction only, we use the approximation:

$$\forall n, \quad \mathbb{E} \left[V_{t+1}^{*,n,\mathbf{k}}(S_{t+1}) \mid S_t = s \right] \approx \sum_{i=1}^k \beta_{t,i}^n \psi_i(s) \quad (2)$$

where the ψ_i are the basis functions that we use to approximate the value function. In all the specifications, we use local functions (*i.e.* functions defined over intervals or boxes, in the higher dimensional cases). In the baseline model, the basis functions are polynomials of the oil price and its logarithm, and when risk aversion is positive, different powers of the utility function were included as well.

The least squares regression approach implies that the vector of regression coefficients $\beta_t^n = (\beta_{t,1}^n, \dots, \beta_{t,k}^n)^T$ satisfies the linear equation

$$\begin{aligned} \mathbb{E} \left[(\phi_1(S_t), \dots, \phi_k(S_t))^T (\phi_1(S_t), \dots, \phi_k(S_t)) \right] \beta_t^n \\ = \mathbb{E} \left[(\phi_1(S_t), \dots, \phi_k(S_t))^T V_{t+1}^{*,n,\mathbf{k}}(S_{t+1}) \right]. \end{aligned} \quad (3)$$

In the case where the level of reserves follows a stochastic process, the dynamic programming formulation has to be refined for two reasons. First, the level of reserves might drop below the minimum extraction level l before the end of the time horizon. Second, we have to take

⁶When $\gamma = 1$, the utility function is defined as $u(x) = \log(x)$.

into account the autoregressive nature of the reserves process (see equation (9) in Section 5). These considerations lead to the following formulation:

$$V_t^{*\mathbf{k}}(s_t, r_t, r_{t-1}) = \sup_{\min\{r_t, l\} \leq h \leq \min\{l+k_t, r_t\}} \left\{ u((h+l)s) + \frac{1}{B_{t,t+1}} \mathbb{E} \left[V_{t+1}^{*\mathbf{k}}(S_{t+1}, R_{t+1}, R_t) \middle| S_t = s_t, R_t = r_t - h, R_{t-1} = r_{t-1} \right] \right\}. \quad (4)$$

In this case, we approximate the continuation value by the following multivariate regression:

$$\mathbb{E} \left[V_{t+1}^{*\mathbf{k}}(S_{t+1}, R_{t+1}, R_t) \middle| S_t = s, R_t = r, R_{t-1} = \hat{r} \right] \approx \sum_{i=1}^k \beta_{t,i}^n \psi_i(s, r, \hat{r}). \quad (5)$$

The set of basis functions included polynomials of the oil price, the log oil price and the reserve level and its lag. The vector of regression coefficients $\beta_t^n = (\beta_{t,1}^n, \dots, \beta_{t,k}^n)^T$ satisfies the linear equation

$$\begin{aligned} & \mathbb{E} \left[(\phi_1(S_t, R_t, R_{t-1}), \dots, \phi_k(S_t, R_t, R_{t-1}))^T (\phi_1(S_t, R_t, R_{t-1}), \dots, \phi_k(S_t, R_t, R_{t-1})) \right] \beta_t^n \\ &= \mathbb{E} \left[(\phi_1(S_t, R_t, R_{t-1}), \dots, \phi_k(S_t, R_t, R_{t-1}))^T V_{t+1}^{*\mathbf{k}}(S_{t+1}, R_{t+1}, R_t) \right]. \end{aligned} \quad (6)$$

The regression coefficients in (2) (respectively in (5)) are estimated by replacing the expectations in (3) (and in (6)) by their Monte-Carlo estimates based on random grids and backward recursion. This method is referred to as *a priori* and described in detail in Gyurkó, Hambly and Witte (2011).

The numerical solution can be affected by three types of errors: (i) discretization errors that come from the transformation of a continuous time process into a discrete one (this can be reduced simply by reducing the length of the period); (ii) the projection error due to the fact that the conditional expectations are approximated by linear combinations of basis functions (this error can be reduced by increasing the number of basis functions); (iii) finally, the statistical error arising when we estimate the regression coefficients using Monte-Carlo techniques (this can be controlled by choosing a sufficiently large Monte-Carlo sample).⁷

4 Oil price and extraction cost models

The multiple real option approach is flexible enough to allow numerical solutions even when the dimensionality of the problem (the number of state variables) increases. This allows us to use complex oil price processes and to take into account stochastic extraction costs. The first oil price model used is the mean-reverting process, but we also simulate the two-factor oil price model of Schwartz and Smith (2000). In addition, a process for stochastic extraction costs is added to the model.

⁷In an earlier version of the paper (Aleksandrov and Espinoza, 2011), we used the lower bound of the value function provided by the Monte-Carlo Least Squares method as well as the upper bound of the value function provided by the dual method of Aleksandrov and Hambly (2010) to construct a confidence interval for the value function when risk aversion is set to 0. We compared the upper and lower bound approaches for our baseline model and found that the relative difference between the upper and the lower bound does not exceed 3%, which confirms the accuracy of the numerical approximations used.

A mean reverting price process

The mean-reverting (one factor) model has been the model of choice in the literature, as in Schwartz (1997). The logarithm of the oil price satisfies the stochastic differential equation

$$dX_t = \kappa(\alpha - X_t)dt + \sigma dW_t,$$

The process is estimated using annual ‘real’ oil prices for the period 1957 - 2008 (see Table 1) as the discretized process⁸ $X_t - X_{t-1} = a + bX_{t-1} + \sigma Z$ where Z is a standard normal random variable independent of X_{t-1} . The estimated long-term mean of the oil price is US\$49 (at constant prices), and the half-life of oil price shocks is 14 years.

Table 1: Parameters of the oil price process (OLS on yearly data for the period 1957 - 2008)

Price process	a	b	σ
Parameter	0.183	-0.047	0.26
Standard error	(0.180)	(0.057)	

Table 2: Parameters of the Schwartz-Smith (2000) price process estimated on futures data

Parameters	κ	σ_χ	μ_{ξ^*}	σ_ξ	$\rho_{\xi\chi}$
Estimate	1.49	0.286	0.0115	0.145	0.3
Standard error	(0.03)	(0.01)	(0.0013)	(0.05)	(0.044)

Table 3: Annualized parameters of the extraction cost process (OLS on quarterly data for the period 1999-Q1 to 2009-Q1)

Cost process	c	σ_c
Parameter	0.08	0.39
Standard error	(0.06)	

A model with stochastic equilibrium price

The mean reversion model assumes that the long-term mean (or the equilibrium level) of the oil price is known. Schwartz and Smith (2000) proposed a two-factor model for the oil price that captures uncertainty in the equilibrium oil price. More precisely, the logarithm of the price X_t has two components: short-term deviations Y_t , which follow a mean-reverting process (with mean 0), and the equilibrium level M_t , which is assumed to follow a Brownian motion.

$$\begin{cases} X_t = Y_t + M_t \\ dY_t = -\kappa Y_t dt + \sigma_\chi dW_t \\ dM_t = \mu_{\xi^*} dt + \sigma_\xi d\tilde{W}_t \end{cases} \quad (7)$$

In addition, the shocks in the short term are correlated with the shocks to the equilibrium price, *i.e.* $\text{corr}(W_t, \tilde{W}_t) = \rho_{\xi\chi}$. The model is implemented using the following discretization.⁹

⁸The oil price is deflated by the US CPI, with index 100 in 2000. The standard errors in Table 1 are the OLS standard errors, which are valid if the series is stationary. However, the augmented Dickey-Fuller test and the Phillips-Perron test were not able to reject the null hypothesis of non-stationarity.

⁹The two factors are not observable, but Schwartz and Smith (2000) estimate them using spot and future prices over the period 1990-1996.

$$\begin{cases} Y_t = Y_{t-1}(1 - \kappa\Delta t) + \sigma_\chi\sqrt{\Delta t}Z \\ M_t = M_{t-1} + \mu_{\xi^*}\Delta t + \sigma_\xi\sqrt{\Delta t}(\rho_{\xi\chi}Z + \sqrt{1 - \rho_{\xi\chi}^2}\tilde{Z}) \end{cases} \quad (8)$$

where Z and \tilde{Z} are independent and normally distributed, and Δt is set to 1. The parameters we use are those estimated by Schwartz and Smith (2000) and presented in Table 2. The half-life of short-term shocks is less than a year but the equilibrium price is also variable, with a volatility of 14.5 percent. In order to account for the two factor model in the Least Squares Monte-Carlo algorithm, the two state variables are added to the set of basis functions ψ .

Oil prices correlated with extraction costs

Extraction costs can reach high levels. For instance, the cost of extraction for *Petróleo Brasileiro S.A.* (Petrobras) exceeded US\$30 per barrel in 2008. These costs can be added to the model assuming they follow a geometric Brownian motion:

$$\log C_t = \log C_{t-1} + c + \sigma_c(\rho Z + \sqrt{1 - \rho^2}\hat{Z}),$$

where \hat{Z} is a standard normal random variable independent of Z , and ρ is the correlation between oil prices and extraction costs. The correlation between oil prices and the extraction costs faced by Petrobras has been low in Brazil because the government contributed to the cost, but when taking into account this cost to the government, the correlation between quarterly percentage changes is high; it exceeded 70 percent over the period 1999-2009. We first estimate a discretized model $\log(C_t) - \log(C_{t-1}) = c + \sigma_c\hat{Z}$, with parameter estimates shown in Table 3 and then investigate how positive correlations affect the optimal strategy. Again, the new state variable is added to the set of basis functions in the Least Square Monte-Carlo algorithm.

5 Learning the Size of Oil Reserves

Farzin (2001) argued that proven oil reserves are the output of a production process and depend on the payoff (the oil price), past extraction (to capture depletion) and a level of technology. In particular, he estimated on U.S. data that a 10 percent increase in oil prices triggers discoveries of 1 percent of additional reserves. Lund (2000) modeled a Bayesian learning process in the amount of oil reserves. The prior distribution on oil reserves is updated by information on the amount of oil extracted and the pressure of the well. In Lund (2000), when pressure declines, it indicates that the remaining volume of oil is near depletion. If the pressure is constant, it indicates that the reserves are at least higher than production and therefore the range of uncertainty in reserves declines with extraction.

We model reserves taking into account these findings. We estimate a model where increases in reserves depend on the oil price, past extraction, time (which captures technological innovations), and where the conditional variance of reserves is also a function of time. This later component proxies for learning in the size of reserves in the spirit of Lund (2000). Lund (2000)'s formulation for reserves is appropriate to model reserves of a specific field, but not for a country as the model would not be able to capture potential discoveries in reserves due to higher oil prices or better technology. A cursory look at the EIA data however confirms that uncertainty in reserves has been declining with time. Figure 1 shows that for the median

oil producer, the standard deviation of annual changes in proven reserves has decreased by about 50 percent in 19 years. This discussion suggests the following model for reserves:

$$\Delta \log(R_t) = a_0 + a_1 \text{time} + a_2 \log(P_{t-1}) + a_3 \Delta \log(R_{t-1}) + a_4 CUM_{t-1} + \epsilon_t \quad (9)$$


where the conditional variance of the normally-distributed ϵ_t (given available information) is a function of time:¹⁰

$$\epsilon_t \sim \mathcal{N}(0, \sigma_t^2) \quad (10)$$

$$\sigma_t^2 = \exp(b_0 + b_1 \text{time}) \quad (11)$$

The model is estimated on the proven reserves of crude oil from the EIA, for the panel of 98 countries for which reserves data was available, over the period 1980-2009. After removing extreme changes in reserves that would not be realistic for relatively mature oil producers such as the U.A.E. and Brazil (higher than 50 percent and lower than 50 percent), the data still shows excess kurtosis of 6.5. This suggests that the model described by equation 9, with normally distributed errors, may be improved by choosing a different specification for ϵ_t . We therefore estimated the model 9 by quasi-maximum likelihood (Bollerslev and Wooldridge, 1992) assuming that ϵ_t is distributed as a Student's t distribution, with degrees of freedom $v = 3$, $v = 5$ (our preferred specification as it implies an excess kurtosis of 6 for ϵ_t), and $v = 10$ (closer to normal distribution).

Figure 1: Standard deviation of change in reserves (10-year rolling window, median country)



The results are shown in Table 4. Columns 1 to 3 show the most general model for different degrees of freedom of the Student's t distribution. The coefficients for the variables that significant are robust, and we decide to stick to a model with $v = 5$. We use a General-To-Specific approach and drop the least significant variables (columns 4 to 6). Reserves depletion (CUM)¹¹ and oil prices were not found to be significant. However, the past change in oil reserves was significant and the time trend in the variance equation was highly significant, for all our specifications. Column 7 shows that there are large differences in the variance equation for countries with reserves larger than 5 billion barrels. We therefore estimate a

¹⁰ $time$ is a linear trend that takes the value 0 in 1980 and 32 in 2012, which is the first year of the simulations presented in section 6.

¹¹Farzin (2001) uses both CUM_{t-1} and CUM_{t-2} but the two variables are almost identical in our sample so we dropped CUM_{t-2} .

Table 4: Model for proven reserves (Standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$)

Level Eq.	(1) $v=3$	(2) $v=5$	(3) $v=10$	(4) $v=5$	(5) $v=5$	(6) $v=5$	(7) $v=5$	(8) $v=5$ (res > 5 bn b)
$\Delta \log R_{t-1}$	0.0142** (0.00622)	0.0239** (0.0118)	0.0226 (0.0192)	0.0239** (0.0118)	0.0197* (0.0112)	0.0195* (0.0110)	0.0144 (0.0101)	0.0473* (0.0248)
CUM_{t-1}	-1.65e-05 (1.50e-05)	-2.07e-05 (3.25e-05)	-4.99e-05 (9.48e-05)					
$\log P_{t-1}$	-0.000578 (0.00177)	0.00262 (0.00371)	0.00889 (0.00622)	0.00265 (0.00371)	0.000356 (0.00186)			
$\log P_{t-2}$	-1.03e-05 (0.00160)	-0.00402 (0.00361)	-0.0134** (0.00634)	-0.00402 (0.00361)				
time	0.000154 (0.000120)	0.000100 (0.000171)	-7.18e-05 (0.000235)	9.65e-05 (0.000170)	0.000121 (0.000155)	0.000158 (0.000125)		
Constant	-0.00100 (0.00246)	0.00272 (0.00598)	0.0159 (0.0110)	0.00263 (0.00597)	-0.00376 (0.00522)	-0.00342 (0.00287)		
Variance Eq.								
time	-0.162*** (0.0210)	-0.109*** (0.0156)	-0.0745*** (0.0134)	-0.109*** (0.0155)	-0.113*** (0.0143)	-0.114*** (0.0142)	-0.147*** (0.0207)	-0.0453** (0.0204)
Const.	-3.494*** (0.256)	-4.000*** (0.218)	-4.128*** (0.197)	-4.000*** (0.218)	-3.944*** (0.192)	-3.950*** (0.190)	-3.321*** (0.239)	-5.427*** (0.351)
time I _[res>5]							-2.120*** (0.423)	
Const. I _[res>5]							0.103*** (0.0292)	
Obs.	1,611	1,611	1,611	1,611	1,777	1,802	1,802	565

separate model for these countries, shown in column 8, and use these estimates to calibrate our model of reserves. The coefficient of the time trend in the variance equation implies that every year the conditional variance in the volume of proven reserves is reduced by around 4.5 percent. This corresponds to a 50 percent decline of the standard deviation in 15 years - in line with what can be deduced by a reading of Figure 1.

6 Application to Brazil and the U.A.E.

We apply the model to two countries: Brazil and the U.A.E. The two countries were chosen to be representative of a non-OPEC country, with limited oil reserves and very high needs of oil revenues, and an OPEC country, with large reserves, and an intergenerational approach to managing its resources. Nonetheless, Brazil and the U.A.E. have roughly similar production capacities and ambitious expansion plans. The two countries are large producers and will allow us to pin down a rough approximate of the global supply curve (see last section).

In 2009, Brazil had 12.6 billion barrels of proven oil reserves, a number that leaves Brazil with the second largest oil reserves in South America after Venezuela. The production of oil reached 2.7 million barrels per day (bbl/d) in 2010 (U.S. Energy Information Agency, EIA) and Brazil has been one of the fastest non-OPEC growing oil industry in the recent years, with numerous projects brought onstream. Petrobras' expansion capacity investment plans amount to US\$ 224.7 billion over the next 5 years (Petrobras, 2011), and estimates of the costs for exploration and production reach US\$ 117.7 billion. With 95 percent of crude oil produced by the state-controlled company Petr leo Brasileiro S.A. (Petrobras), and production licenses issued by the Petroleum Agency (ANP), the choice of extracting oil or leaving it under the ground in order to generate future income is an important policy choice. Discussions on the exploitation of the oil resources gained prominence recently after new oil reserves were discovered.

The United Arab Emirates (U.A.E.) holds the sixth largest proven oil reserves in the world (estimated at 97.8 billion barrels, around 8.5 percent of total world reserves) and is an OPEC member. In 2010, the U.A.E. produced 2.8 million barrels per day of crude oil but the govern-

ment plans to expand capacity so as to make the U.A.E. a ‘swing’ producer, which would give the country a clearer negotiating power within OPEC. Projects that are under development and the plans to scale up the oil sector in the U.A.E. were estimated to cost around US\$ 20 billion in investment projects (IMF, 2007).

We assume in all calculations that the extraction decisions are made on a yearly basis. The smallest unit that is used to change production is 0.2 billion barrels per year, equivalent to 0.55 million barrels per day. We also make the assumption that each country extracts at least 0.5 billion barrels per year to finance operations and ensure a minimum of revenues. For Brazil the time horizon chosen is 16 years in the model with known reserves, which leaves us with 25 units of 0.2 billion barrels per year (after the minimal extraction rate of 0.5 billion barrels per year are subtracted). In the model with unknown reserves, the horizon is extended to 25 years. For the U.A.E. we consider a time horizon of 100 years and 240 extractable units. The EIA provides us with projections of production capacity for the next twenty years (see Table 5), and we assume capacity is constant after 2030. The calibration used for the numerical results is reported in the same table.

Table 5: Oil production capacity outlook 2010-2030 (Billion Barrels per Year). Source: EIA, System for the Analysis of Global Energy Markets (2006);

Extraction capacity	2010	2015	2020	2025	2030
United Arab Emirates (in billion barrels per year)	1.0	1.3	1.4	1.6	1.7
Brazil (in billion barrels per year)	1.0	1.3	1.4	1.5	1.6
Calibration, in units (1 unit = 0.2 billion barrels)	5	6	7	8	8

7 Results

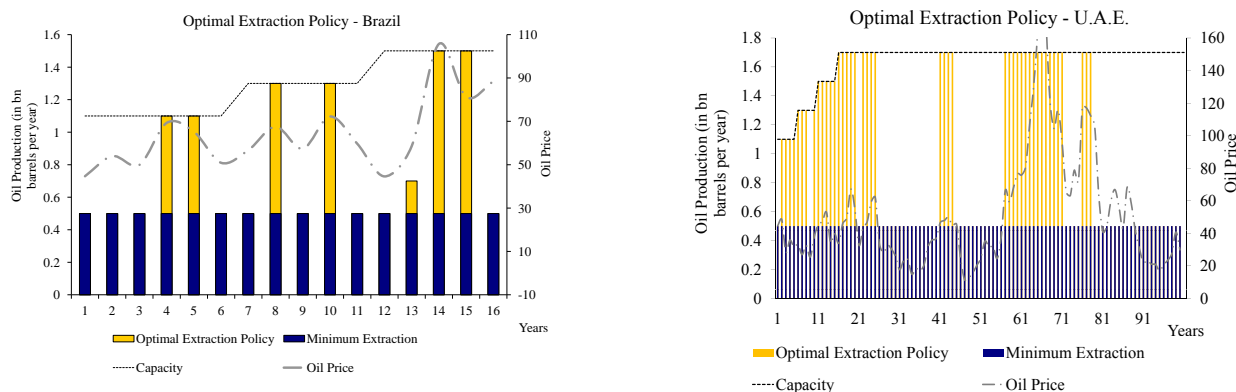
The model parameters used to determine the optimal policy are assumptions of the model. We discuss the changes in optimal policy when these parameters change. The baseline model is a model with risk aversion set to 0, mean reverting oil prices, no extraction costs, and known reserves. We investigate the dependence of optimal policy on: (i) the expected volatility of oil prices ; (ii) interest rates; (iii) the size of reserves; (iv) production capacity; (v) risk aversion; (vi) extraction costs and (vii) uncertainty in the size of reserves.

Baseline. We show in the left hand side chart of Figure 2 the optimal extraction path for Brazil for a hypothetical oil price simulated using the mean-reverting oil price model described earlier. Extraction at the minimum level remains optimal for three years, until the oil price exceeds US \$70 in the fourth year of the simulation. A striking feature of the optimal extraction strategy for both Brazil and the U.A.E., and that was found with most simulations of the oil price under the baseline model, is that production is quite volatile (although we did not impose a so-called ‘bang-bang’ solution, as will be clearer when looking at the impact of risk aversion). This result is at odds with the actual behavior of oil producers, who most often extract at full capacity until oil fields are exhausted. We discuss in the remainder of the paper whether different specifications of the model yield more realistic results.

Sensitivity to the volatility of oil prices. We show in the top left chart of Figure 3 the sensitivity of the solution to different assumptions on expected oil price volatility, for a given *realization* of the oil price path. As was emphasized by the ‘waiting to invest’ literature,

high volatility makes it optimal to defer the exercise of options. During the first four years of the simulation, as oil prices are low, it is only optimal to extract oil for a producer that would believe that oil price volatility is low ($\sigma = 0.13$). A producer who thinks prices are stable would always find it optimal to extract at full capacity and would quickly exhaust its oil field, missing the opportunity to extract when prices are higher. On the other hand, a producer who thinks volatility is 0.52 would delay extraction until it is forced to extract (when the contract to exploit the field expires).

Figure 2: Optimal Extraction - Brazil and the U.A.E.



Sensitivity to the real interest rate. Dependence on the real interest rate is shown in the top right chart of Figure 3. In these simulations, exercise occurs earlier at higher interest rates, a result that differs from the one found for simple European call options. In the Black-Scholes formula, using the risk neutral form of the stock price, the drift for the oil price is higher the higher the interest rate, which implies that the underlying price goes above the strike with larger (risk-neutral) probability when the interest rate is higher. As a result, the continuation value is higher when the interest rate is higher. In our model, the interest rate only appears in the discount factor and has no impact on the drift of the oil price, which is why the continuation value can be lower with higher interest rates.

Sensitivity to the size of reserves. The threshold price at which production at full capacity is optimal is of course a function of reserves. The right hand side chart of Figure 2 show that the U.A.E. should extract at full capacity for fairly low oil prices. The threshold price for Brazil was found to be US \$73 (in 2000 terms, *i.e.* around US \$95 of 2011).¹² For the U.A.E., the threshold would be US \$39 (around US \$50 of 2011). However, we found that smaller differences in the amount of reserves (one to two billion barrels) do not change significantly the shape of the optimal production strategy (see third chart in Figure 3).

Impact of increasing capacity. We compute the increase in net present value of the oil fields (the value function when risk aversion is set to 0) that is obtained when increasing

¹²US CPI inflation between 2000 and 2011 was 30 percent.

capacity, for a given amount of total reserves. The increase in NPV is our estimate of the shadow price of the yearly production capacity constraint. This number can be used to assess the profitability of a project for which investment costs are known. We consider two countries with oil reserves estimated at 5 and 10 billion barrels (25 and 50 units respectively).¹³ The results are presented in Table 6. A capacity expansion of 200 million barrels per year in 2025 (1 additional unit in Year 15) improves the value of the oil reserves by US\$ 2.1 billion, for a country whose reserves are 5 billion barrels.¹⁴ The increase would reach US\$ 6.27 billion for a country that holds 10 billion barrels. Additional capacity has more value for countries with larger reserves, a result that, if it already existed, was not echoed in the oil production literature. The intuition is that the additional option provided by higher capacity has more value if it is available for many years and if production capacity was a tighter constraint at times of high prices. Our results therefore suggest that project evaluation needs to be performed in the context of the overall oil strategy of the country, since the viability and profitability (*i.e.* net present value) of projects cannot be assessed project-by-project. However, the value of ‘optionality’ created by higher capacity is small compared with the value of reserves that are normally made available when new wells or platforms are open. We show in Table 7 the value of such a project if the additional reserves are worth 1 billion barrels (5 units of 200 million barrels). The value of the project is now higher for the country with lower reserves because the value of optionality is dwarfed by the value of additional oil reserves, and an additional one billion barrels of oil reserves is worth more for a country that has low reserves than for a country that has high reserves. Indeed, for a given finite horizon and a given production capacity, a country with low reserves can ‘distribute’ the extraction of newly-found reserves more optimally than a country that has high reserves.

Table 6: Added value by expanded capacity.

Year (unit = 0.2 bn bbl/year)	1-5	5 -10	10-15	15-	25 units	50 units
Current capacity	3	4	5	5	US\$ 325.9bn	US\$ 622.5bn
Increased capacity	3	4	5	6	US\$ 328bn	US\$ 628.8bn
Added value					US\$ 2.1bn	US\$ 6.27bn

Table 7: Added value by expanded capacity and increased access to reserves.

Year (unit = 0.2 bn bbl/year)	1-5	5 -10	10-15	15-	25/30	50/55
Current capacity	3	4	5	5	US\$ 325.9bn	US\$ 622.5bn
Increased capacity	3	4	5	6	US\$ 391.4bn	US\$ 685.2bn
Added value					US\$ 65.6bn	US\$ 62.7bn

Risk Aversion. The fourth chart in Figure 3 shows how increasing risk aversion makes earlier production optimal. In the first periods of the simulation, production at full capacity becomes optimal with a small level of risk aversion because the utility costs of lowering production when income is low are high. Production is accelerated when risk aversion is increased, a result similar to that of Alvarez and Koskela (2006). However, as prices increase in period 3, production is maintained at a lower but stable level, above minimum extraction, when risk aversion is higher. This allows the country to preserve its reserves and thereby

¹³The oil price process parameters are as before. We start with oil price $S = 54.6$. The time horizon is 100 years.

¹⁴We assume that production capacity is maintained at the highest level after expansion, although production capacity is usually declining with time because of depletion.

to ‘insure’ itself against future drops in the oil price. Although risk aversion might explain why production is less volatile, risk-averse producers should maintain large spare capacity, a result at odds with the evidence that only Saudi Arabia, of all producers, maintains sizable spare capacity.

Extraction Costs. Extraction costs, even when highly correlated with oil prices, do not affect significantly the optimal strategy because the net payoff of extracting a barrel of oil remains volatile.

Uncertainty in the size of oil reserves. We compute the optimal extraction strategy for different calibrations of the reserves process. As is done with the oil price, the path of discovery of reserves is kept identical across simulations, but the decision rules have been estimated assuming different processes for reserves. The differences across simulations can therefore be interpreted as due to the formation of expectations by the oil producer. First, we compare a simulation where reserves are known (but moving with time) with a simulation where reserves are thought to follow the stochastic process described in equation 9, with shocks to reserves parameterized with a constant variance of 5 percent. The bottom right chart in Figure 3 shows that additional uncertainty in the size of reserves triggers earlier extraction. The intuition behind this result is that since the life of the contract is finite, there is an opportunity cost of delaying extraction if additional oil resources are found later.

Learning the size of oil reserves. Second, we show the optimal extraction policy for different learning speeds (of the size of reserves). The first simulation assumes that the variance of shocks to reserves is 5 percent in the first year and declines to reach 3.2 percent in year 10 (‘slow learning’: the model for reserves was calibrated assuming $b_0 = -1.5$ and $b_1 = -0.046$ in equation 11). The second simulation assumes that the variance of shocks to reserves is also 5 percent in the first year but declines to 1.9 percent in year 10 (fast learning: the model for reserves was calibrating assuming $b_0 = 0.23$; $b_1 = -0.1$). We show two sets of simulations: one in which reserves are discovered (example (a); top left chart in Figure 4) and one in which reserves are depleted faster than in the simulation with known reserves (example (b); top right chart of Figure 4). The charts show that in the first period of the model, fast learning tends to defer extraction, for the same reason that uncertainty tended to accelerate extraction: the opportunity cost of delaying extraction if reserves are found later is lower since the producer will learn quickly about the size of reserves. Production is also more stable after some years because reserves uncertainty is reduced considerably in 10 years.

8 Limitations of the Model

The previous section showed that the supply elasticity of the model is often much higher than what can be inferred from the observed supply response to prices. Krichene (2005) indeed argued that short-term supply is insensitive to price. There are several potential explanations to this finding, and we discuss below the three that we think are most relevant before investigating the determinants of oil production volatility empirically.

- There are technical limitations and fixed costs that reduce the benefits of varying production. For instance, there is a minimum turn-down that is necessary to have the plant on, to avoid rusting, losing staff, etc. Similarly, if extraction costs are a function of oil production, the solution of the model will differ somewhat and optimal oil

production will be less volatile. However, these costs do not seem to have prevented several oil producers (with varied types of plants) to cut dramatically production during the recent crisis. According to the Joint-Oil Data Initiative (JODI) monthly database, the peak-to-through cut in oil production during the last crisis reached or exceeded 30 percent in Azerbaijan, Brunei, Denmark, Malaysia, Norway and the UK.

- We have abstracted from re-investment costs to re-start production and from fixed operational costs. Fixed operational costs do not affect the optimal strategy when the horizon of exploitation is known, but they would matter if extending the life of a field was possible, but costly for a variety of reasons. For instance, in countries such as Brazil or Norway, contracts on the exploitation of a field expire after 25-30 years, but contracts might be extended (at some economically significant cost) after negotiating with the government.¹⁵
- The optimization model is not appropriate if a country's discount rate is contingent on oil production. This will be the case for oil producers whose budget and external accounts are highly dependent on oil revenues. For such countries, cutting production at the time prices are low will push those accounts in deficit and trigger borrowing at increased rates, which governments and oil companies will be unwilling to do. However, given the significant losses in NPV (130 percent in the extreme case, see section 10), strategies that lead to more volatile production should be considered, especially for producers with larger reserves and for those who already accumulated large financial reserves.
- Decisions are very sensitive to expectations of the equilibrium oil price, as can be seen from the bottom left chart in Figure 4. When the decision rule is computed based on an equilibrium oil price of US\$ 30, production at full capacity is optimal for 8 years in a row (given the realized price shown in the figure), even though the half-life is 14 years (and therefore even though temporary shocks have a large effect of the continuation value). Decision making becomes even more complex when oil prices follow a mean-reverting model with stochastic equilibrium price (the Schwartz and Smith (2000) model). The bottom right chart of Figure 4 shows a simulated path of an *observed* oil price together with the underlying *unobserved* equilibrium price, where the simulation parameters are those shown in Table 2. The series labeled $\rho_{\xi\chi} = 0.3$ shows optimal extraction when the decision rules have been calculated using the Monte Carlo simulations based on these same parameters. In that sense, this series shows the optimal extraction when expectations are based on the 'correct' model of oil prices. Roughly speaking, extraction is optimal when spot prices exceed equilibrium prices. The series labeled $\rho_{\xi\chi} = 0.8$ shows optimal extraction when the decision rules have been calculated using the Schwartz-Smith model, but with the correlation between short-term shock and equilibrium shocks set to 0.8. The series therefore shows the optimal extraction when the decision maker expects observed short-term shocks to be accompanied by shocks to the equilibrium price of oil. As a result, in the first 10 years of the production horizon, production is increased slowly even as prices go up because the continuation value increases when shocks are thought to be permanent (since there is a value to wait). Faced with uncertainty on the right model of oil prices, the prudent decision may be to maintain a relatively stable level of production.

¹⁵We are grateful to an anonymous referee for pointing this to us.

9 Determinants of Production Policies during 2008-2009 Crisis

We finally investigate statistically the determinants of oil production volatility. It seems difficult to base such an investigation on estimates of production volatility over a large horizon because this volatility is affected by changes in extraction capacity, new investments, field exhaustion as well as other unknown factors (geopolitical, etc.). In addition, oil production is typically not volatile, which is the ‘puzzle’ this paper investigates. However, the recent crisis provides a simple way to test the determinants of production cuts over a short horizon. Indeed, the fall in prices between 2008 and 2009 occurred over a very short period and was the most important driver of the fall in production. The drawback of the method is that it relies on a single episode and therefore the dimension of the dataset is reduced to one cross-section. For all the oil producers with reserves higher than 1 billion barrels, we compute the change in oil production after the Lehman Brothers’ collapse, comparing the average production in the 12 months leading to October 2008 with the average production in the following 12 months.¹⁶

Table 8: Determinants of cuts in production during the crisis of 2008-2009

OLS	(1)	(2)	(3)	(4)	(5)
Log(GDP per capital)	3.609* (1.803)	3.586* (1.756)		3.836** (1.838)	
Interest rate	-2.254* (1.142)	-2.172* (1.084)	-2.162* (1.171)		
Log(Reserves)	-0.508 (1.597)				
Government oil revenues / Total revenues	-21.77** (8.695)	-21.89** (8.466)	-19.30** (9.041)	-17.56* (8.742)	-15.11* (8.800)
OPEC membership	23.49*** (6.224)	22.42*** (5.103)	21.04*** (5.463)	18.71*** (5.116)	16.98*** (5.299)
Constant	-18.15 (18.75)	-19.24 (17.97)	14.36* (7.792)	-34.16* (17.67)	2.333 (3.325)
Observations	23	23	23	24	25
R-squared	0.572	0.569	0.470	0.460	0.341
Standard errors in parentheses					
*** $p < 0.01$, ** $p < 0.05$					

Following our results in section 7, we would want to test whether the drop in oil production can be explained by:

- **The size of oil reserves.** We expect that countries with smaller reserves cut production further when the oil price dropped. The data on oil reserves is taken from the EIA.
- **The discount rate** We expect that countries that are more impatient will not cut as much production in the face of a collapse of the oil price. The discount factor in a utility function is not observable, but we use two proxies. The first one is the interest rate on external public debt, taken from JPMorgan’s EMBI database. For the countries that do not have sovereign debt instruments, because they already accumulated large external assets, we use publicly available rates of return on their oil funds (for Algeria, Brunei and Norway)¹⁷ or an estimated rate of return for a diversified sovereign wealth fund

¹⁶Production data is a monthly average of production, in million barrels per day. The source is Joint Oil Data Initiative.

¹⁷See individual IMF Country Reports for 2009-2010.

(for Libya, Iran and the GCC).¹⁸ The second proxy is the level of GDP per capita (in US\$ PPP)¹⁹: poorer countries that do not have access to international capital markets should give higher value to current oil revenues (see van der Ploeg and Venables, 2011, who show that, for capital-constrained resource-rich economies, domestic interest rates are high and the optimal path of consumption is tilted towards the near future).

- **Risk aversion.** We expect that countries with higher risk aversion will be less likely to change production as the oil price falls. Risk aversion is not observable, and there is no simple proxy that can be used. However, we think that the share of oil revenues in total government revenues can capture a similar link between risk and cuts in production. Indeed, *ceteris paribus*, a government that relies heavily on oil revenues for its finances is unlikely to vary production frequently, because the other elements in the budget are not flexible.²⁰
- **Membership in OPEC.** OPEC has stated explicitly an objective of price stabilization, which would suggest that OPEC members should react more aggressively than other exporters to price changes.

A look at the data (Figure 5) reveals that cuts in production after the crisis were indeed larger for countries that faced lower discount rates. No relationship appears with the share of oil revenues in government revenues if one does not distinguish between OPEC and non-OPEC members. However, for non-OPEC members, there is indeed a negative relationship between the importance of oil revenues for the government and the production cuts that were undertaken in end-2008 and 2009. The OLS regressions are presented in Table 8, and they confirm what can be seen from the data. Controlling for OPEC membership, countries that are more ‘patient’ (richer countries and countries facing lower interest rates/lower opportunity costs) cut production further (see column 1). In addition, the share of government revenues due to oil is significant and with the expected sign: governments that depend on oil are less likely to cut production when prices fall. This could either be because the share of government revenues is a good proxy for the willingness to stabilize revenues or because interest rates/discount rates are more sensitive to production policies for governments whose finances depend strongly on oil revenues. The size of reserves is not significant when controlling for OPEC membership and was dropped in the regression shown in column 2 (the other coefficients were not affected). Dropping income per capita and/or the interest rate does not affect the results either (columns 3-5).

10 Concluding remarks and the World Supply of Oil

We proposed a Monte Carlo real option approach as a solution to the optimization problem of a price-taker oil producer. This approach allows us to replicate some results of the ‘waiting to invest’ literature in a multiple-period setting and to discuss the effect of interest rates, size

¹⁸Based on a portfolio of half US 10-year yields (3.7 percent) and half the long-run performance of the Dow Jones (8 percent).

¹⁹The data is taken from IMF, 2010.

²⁰The data on the share of oil revenues is for 2007 and comes from individual country IMF reports. For Brazil, the data is deduced from Gobetti (2010). For Canada, the data is deduced from Ahmad and Mottu (2002). For Denmark, the data is taken from Danish Energy Agency (2008). The data for India is deduced from Table 7 in Segal and Sen (2011). For the UK, the data is from HM Revenues and Customs, 2011. The data for Venezuela comes from an old IMF country report (1999). Data was unavailable for Argentina, Australia, China, Egypt and the U.S.

of reserves, risk aversion, expectations on the price of oil, and learning about the size of oil reserves. The Monte Carlo numerical solution is accurate and flexible enough to solve the problem with several state variables.

The baseline model shows that optimal extraction should be very volatile. Indeed, a 10 billion barrels oil field is worth US\$ 357.9 billion under a constant extraction policy, but is worth US\$ 835.4 billion under the optimal policy when there is no minimum extraction required for technical or financing reasons. The benefits from following an optimal strategy would therefore reach 133% in terms of net present value, a dramatic improvement. This proportion is approximately constant for different values of the oil reserves. Assuming that a minimum production of 0.5 billion barrels per year is required every year (40 percent of production is constrained), the gains reach 25.6 percent. If 60 percent of annual capacity must be extracted every year, optimizing over the remaining 40 percent still increases the PV of oil fields by 22 percent.

The model gives us two points in the theoretical world market supply curve (Brazil and the U.A.E.), and allows us to derive an approximation of the rest of the supply curve. We ranked countries with capacity higher than 500 thousand bb/d by their reserves and cumulated their production²¹ to draw a rough approximation of the theoretical supply curve under the baseline specification of the model (mean reverting oil prices, known reserves, linear utility and no extraction costs). Figure 6 shows that in such a world, the supply elasticity would be very large and oil price volatility would never reach the levels attained in 2008-2009. Indeed, demand would need to fall as low as 30 million bb/d to see prices declining to US\$ 40. On average, production would be lower (as several countries would find the oil price to be below their threshold price). As a result, prices would be higher but less volatile.

The benefits from varying production might be reduced by extraction costs, but costs tend to depend on capacity rather than on production itself. Technical costs may also limit the scope for varying production, but even optimizing over half of capacity would yield substantial benefits to oil producers. We suggested risk aversion could explain why production is stable in practice but risk-averse producers should also maintain spare capacity, a result at odds with the evidence. We also showed that uncertainty in the size of reserves could explain why full extraction is optimal but for mature producers, this uncertainty has been shrinking with time. A third potential explanation is that uncertainties on the price process (for instance on whether a shock is temporary or permanent) could explain why production is less volatile. Finally, it is possible that because countries and companies' borrowing conditions are worsened by pro-cyclical extraction policies, volatile extraction policies are not optimal. Indeed, the econometric finding that governments that are highly dependent on oil revenues cut less production during the 2008-2009 crisis is compatible with this interpretation.

²¹Kuwait expected production profile and reserves are very similar to those of the U.A.E. and therefore its decision rule should be the same. Algeria, Mexico and Brazil should also have comparable policies (threshold around US\$90) while Venezuela's threshold price would lie in between. Norway and Egypt policies should be to extract at prices higher than US\$90 only.

References

- Ahmad, E., and E. Mottu, 2002, "Oil Revenue Assignments: Country Experiences and Issues," IMF Working Paper 02/203, Washington DC: International Monetary Fund.
- Aleksandrov, N., and R. Espinoza, 2011, "Optimal Oil Extraction as a Multiple Real Option," OxCarre Research Paper 64, Department of Economics, University of Oxford.
- Aleksandrov, N., and B. Hambly, 2010, "A Dual Approach to Multiple Exercise Option Problems under Constraints," *Mathematical Methods of Operations Research*, Vol. 71, pp. 503-533.
- Alvarez, L., and E. Koskela, 2006, "Does Risk Aversion Accelerate Optimal Forest Rotation Under Uncertainty?," *Journal of Forest Economics*, Elsevier, Vol. 12, pp. 171-184.
- Arrow, K. J., D. Blackwell, and M.A. Girshick, 1949, "Bayes and Minimax Solutions of Sequential Decision Problems," *Econometrica*, Vol. 17, pp. 213-244.
- Bender, C., 2011, "Dual Pricing of Multi-exercise Options under Volume Constraints," *Finance and Stochastics*, Vol.15, pp.1-26.
- Bensoussan, A., 1984, "On the Theory of Option Pricing," *Acta Applicandae Mathematicae*, Vol. 2, pp. 139-158.
- Bollerslev, T., and J. M. Wooldridge, 1992, "Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances," *Econometric Reviews*, Vol. 11, pp. 143-172.
- Broadie, M., and P. Glasserman, 1997, "Pricing American-style securities using simulation," *Journal of Economic Dynamics and Control* Vol. 21(89), pp. 1323-1352.
- Caldentey, R., R. Epstein and D. Saure, 2006, "Optimal Exploitation of a Nonrenewable Resource," mimeo, New York University.
- Carmona, R., and N. Touzi, 2008, "Optimal Multiple Stopping and Valuation of Swing Options," *Mathematical Finance*, Vol. 18, pp. 239-268.
- Cherian, J., J. Patel and I. Khripko, 1998, "Optimal Extraction of Nonrenewable Resources when Prices Are Uncertain and Costs Cumulate," NUS Business School Working Paper (Singapore: NUS).
- Danish Energy Agency, 2008, Denmark's Oil and Gas Production 2008, Chapter 7.
- Dias, M.A., 2004, "Valuation of Exploration and Production Assets: An Overview of Real Option Models," *Journal of Petroleum Science and Engineering*, Vol. 44, pp. 93-114.
- Dixit, A.K., and R.S. Pindyck, 1994, *Investment under Uncertainty*, Princeton NJ: Princeton University Press.
- Farzin, Y.H., 2001, "The impact of oil prices on additions to US proven reserves," *Resource and Energy Economics*, Vol. 23, pp. 271-291.
- Glasserman, P., 2003, *Monte Carlo Methods in Financial Engineering*, New York NY: Springer-Verlag.
- Gobetti, S. W., 2010, "Federalism and Royalties in Brazil: What Does Change in Pre-salt Context?", Presentation at the Conference on Oil and Gas in Federal Systems.
- Gyurkó, L., B. Hambly and J. Witte, 2011, "Monte Carlo methods via a dual approach for some discrete time stochastic control problems", mimeo available at <http://arxiv.org/abs/1112.4351>.
- Hotelling, H., 1931, "The Economics of Exhaustible Resources," *Journal of Political Economy*, Vol. 39, pp. 137-175.
- HM Revenues and Customs, 2011, *Statistics on Corporate Tax*, Table 11.11.
- IMF, 2007, *Article IV Consultations for the U.A.E., 2007*. Washington DC: International Monetary Fund.
- IMF, 2010, *World Economic Outlook*. Washington DC: International Monetary Fund.
- Jaillet, P., E.I. Ronn and S. Tompaidis, 2004, "Valuation of Commodity-based Swing Op-

- tions,” *Management Science*, Vol. 50, pp. 909-921.
- Krichene, N., 2005, “A Simultaneous Equation Model for World Crude Oil and Natural Gas Markets,” IMF Working Paper, N. 05/32, Washington DC: International Monetary Fund.
- Kushner, H.J., 1977, *Probability Methods for Approximations in Stochastic Control and for Elliptic Equations*, Academic Press, New York.
- Kushner, H.J., and P.G. Dupuis, 1992, *Numerical Methods for Stochastic Control Problems in Continuous Time*, Springer, New York.
- Longstaff, F.A., and E.S. Schwartz, 2001, “Valuing American Options by Simulation: A Least-Square Approach,” *The Review of Financial Studies*, Vol. 14(1), pp. 113-147.
- Lund, M.W., 2000, “Valuing Flexibility in Offshore Petroleum Projects,” *Annals of Operations Research*, Vol. 99, pp.325-349.
- Meinshausen, N., and B.M. Hambly, 2004, “Monte Carlo Methods for the Valuation of Multiple-Exercise Options,” *Mathematical Finance*, Vol. 14, pp. 557-583.
- Milne, A., and E. Whalley, 2001, “Bank Capital Regulation and Incentives for Risk-Taking,” Working Paper, City University Business School, London.
- Nagae, T., and T. Akamatsu, 2008, “A Generalized Complementarity Approach to Solving Real Option Problems,” *Journal of Economic Dynamics and Control*, Vol. 32, pp. 1754-1779.
- Paddock, J.L., D.R. Siegel and J.L. Smith, 1988, “Option Valuation of Claims on Real Assets: The case of Offshore Petroleum Leases,” *Quarterly Journal of Economics*, Vol. 103, pp. 479-508.
- Peura, S., J. Keppo, 2006, “Optimal Bank Capital with Costly Recapitalization,” *Journal of Business*, Vol. 79, pp. 2163-2201.
- Petrobras, 2011, *Petrobras Business Plan Presentation, 2011-2015*.
- Rogers, L.C., 2002, “Monte Carlo Valuation of American Options,” *Mathematical Finance*, Vol. 12, pp. 271-286.
- Sakar, S., 2000, “On the Investment Uncertainty Relationship in a Real Options Model,” *Journal of Economic Dynamics and Control*, Vol. 24, pp. 219-225.
- Schwartz, E., 1997, “The Stochastic Behavior of Commodity Prices: Implications For Valuation and Hedging,” *Journal of Finance*, Vol. 52, pp. 923-973.
- Schwartz, E., and J.E. Smith, 2000, “Short-Term Variations and Long-Term Dynamics in Commodity Prices”, *Management Science*, Vol. 46(7), pp. 893-911.
- Segal, P., and A. Sen, 2011, “Oil Revenues and Economic Development: The Case of Rajasthan, India,” Oxford Institute for Energy Studies Working Paper WPM 43.
- Snell, J.L., 1952, “Applications of Martingale System Theorems,” *Transactions of the American Mathematical Society*, Vol. 73, pp. 293-312.
- Tilley, J.A., 1993, “Valuing American Options in a Path Simulation Model,” *Transaction of Society Actuaries*, Vol. 45, pp. 83-104.
- Tourinho, O. A., 1979a, *The Option Value of Reserves of Natural Resources*, Research Program in Finance Working Paper Series No.94, Institute of Business and Economic Research, Berkeley CA: University of California, Berkeley.
- Tourinho, O. A., 1979b, *The Valuation of Reserves of Natural Resources: An. Option Pricing Approach*, Ph.D. dissertation, Berkeley CA: University of California, Berkeley.
- Tsitsiklis, J.N., and B. van Roy, 2001, “Regression Methods for Pricing Complex American-Style Options,” *IEEE Transactions on Neural Networks*, Vol. 12, pp. 694-703.
- Van der Ploeg, F., and A. Venables, 2011, “Harnessing Windfall Revenues: Optimal Policies for Resource-Rich Developing Economies,” *The Economic Journal*, Vol. 121, pp. 1-30.

Figure 3: Sensitivity to key parameters

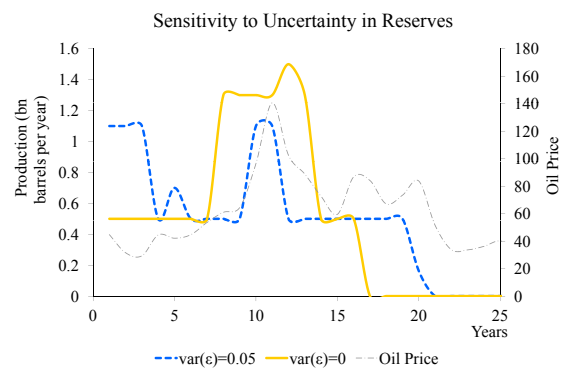
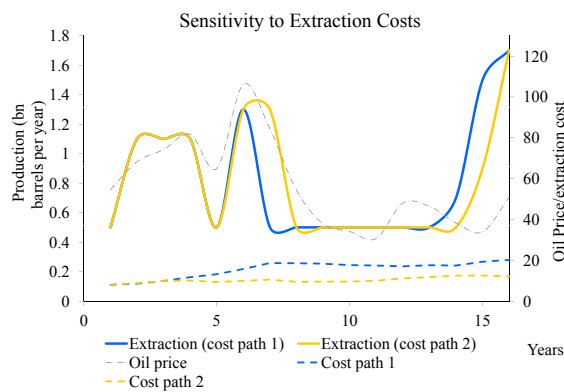
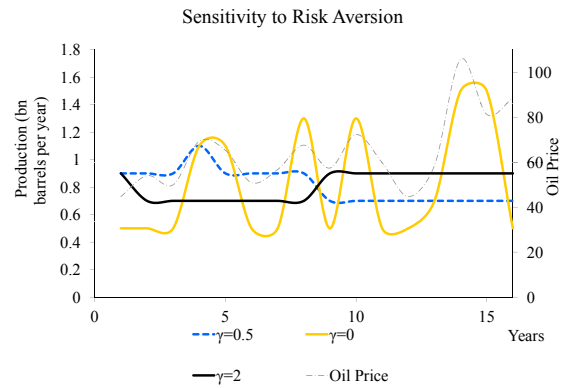
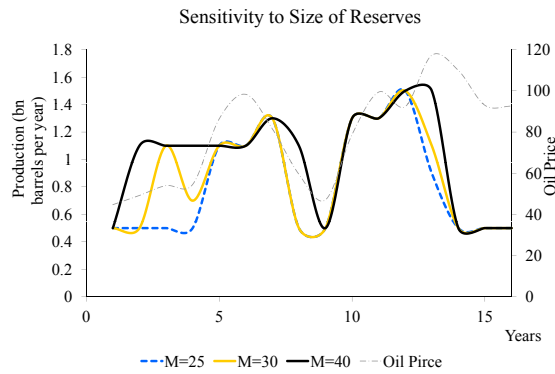
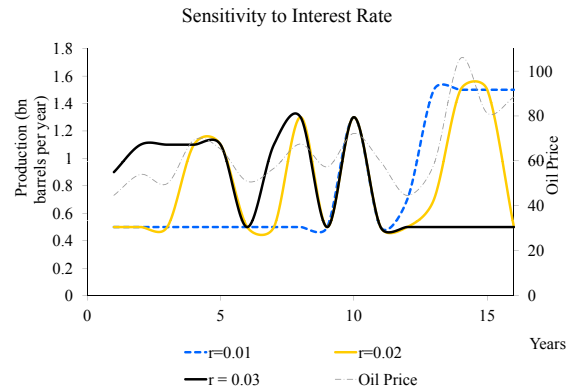
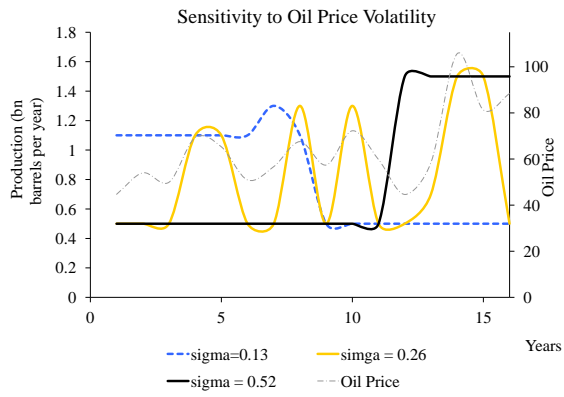


Figure 4: Sensitivity to key parameters (ctd.)

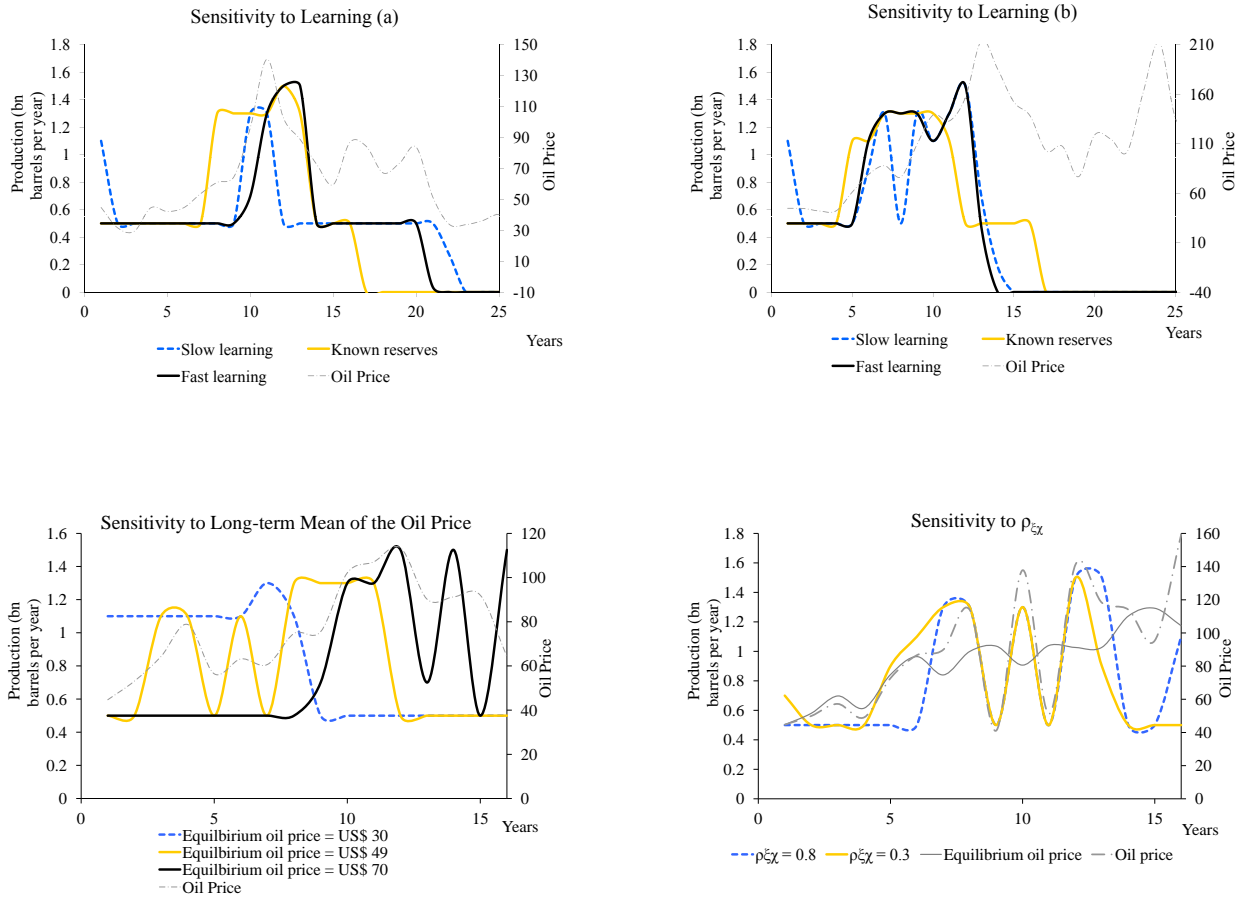


Figure 5: Determinants of the cut in oil production during the crisis of 2008-2009

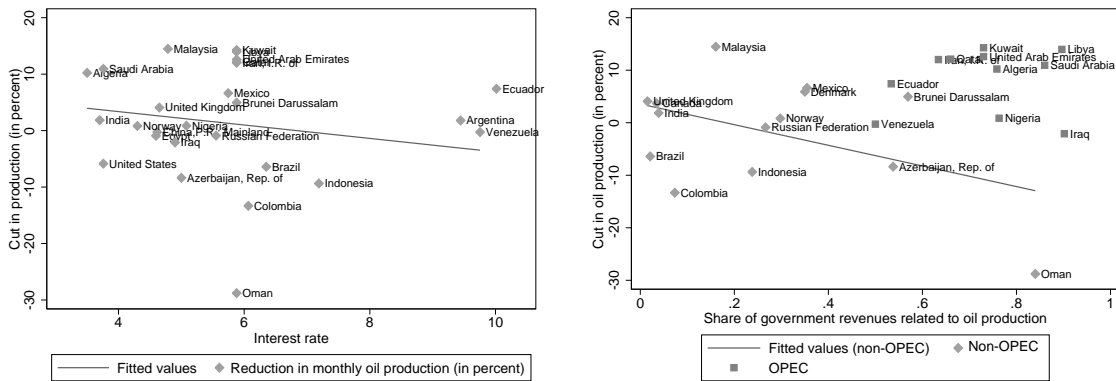


Figure 6: World oil market

