

Game Theoretic Models for Energy Production

Michael Ludkovski and Ronnie Sircar

Abstract We give a selective survey of oligopoly models for energy production which capture to varying degrees issues such as exhaustibility of fossil fuels, development of renewable sources, exploration and new technologies, and changing costs of production. Our main focus is on dynamic Cournot competition with exhaustible resources. We trace the resulting theory of competitive equilibria and discuss some of the major emerging strands, including competition between renewable and exhaustible producers, endogenous market phase transitions, stochastic differential games with controlled jumps, and mean field games.

Keywords: dynamic games, energy production, Cournot and Bertrand markets, exhaustible resources

1 Introduction

The recent decline in oil prices, from around \$100 per barrel in June 2014 to less than \$50 in January 2015 is a dramatic illustration of the evolution of energy production as a result of competition between different sources. Indeed, the price drop was prompted in large part by OPEC's strategic decision not to decrease its oil output in the face of increased production of *shale oil* in the US, itself arising from new technologies that were spurred by investment in exploration and research in times of higher oil prices. These complex interactions are in addition to long-running concerns about dwindling fos-

Michael Ludkovski

Department of Statistics and Applied Probability, University of California Santa Barbara, South Hall, Santa Barbara, CA 93106-3110; e-mail: ludkovski@pstat.ucsb.edu

Ronnie Sircar

ORFE Department, Princeton University ORFE Department, Princeton University, Sherrerd Hall, Princeton, NJ, 08544 e-mail: sircar@princeton.edu

oil fuel reserves ('peak oil'), as well as climate change and the transition to sustainable energy sources.

We survey a (necessarily) selective line of work that builds models successively incorporating various of these features starting from a competitive oligopolistic view of an idealized global energy market, in which *game theory* describes the outcome of competition. In particular, the oligopoly will be taken to be in a *Cournot framework*, in which players choose *quantities* of production, and then prices are determined by aggregate supply. This seems reasonable for energy production in which major players determine their output relative to their production costs, as in the expected scenario that OPEC will cut production in order to increase the market price of oil. The complementary framework of Bertrand markets, in which player set *prices*, is more typically suitable for consumer goods markets, among other examples.

We begin with static, or one-period games, as an introduction to some of the effects that can arise, for instance the non-competitiveness of producing a relatively expensive renewable source, such as wind, against a cheap fossil fuel in plentiful supply. However, the very nature of the complexities calls for a dynamic model in which there are (to use a much over-employed cliché) *game changers* over time. Changes in the competitive environment may come from:

- dwindling reserves of oil or coal, ramping up their scarcity value;
- discoveries of new oil reserves (there were over 30 major finds in 2009, for instance);
- technological innovation such as *fracking*, which has led to extraction of shale oil and gas;
- government subsidies for renewable energy sources such as solar and wind power;
- varying costs of energy production, for instance cheaper solar power due to falling silicon prices and improved solar cell efficiency.

Many if not all of these phenomena are unpredictable and dramatic, and motivate the development of *stochastic models*, particularly with potentially significant 'jumps' (for instance in costs or reserves). Moreover, dealing with stochastic dynamic (nonzero-sum) games involving many influential interacting energy producers creates computational challenges, and some approximation methods include numerical discretization, parameter asymptotics, and continuum (mean field) games.

2 Static Cournot Games

The classic work of Cournot (1838) gives perhaps the earliest example of a Nash equilibrium for describing the outcome of a game. Cournot was concerned with competition between producers of an inexhaustible resource

(mineral water): their effect on sales was such that the more each bottled and brought to the market, the lower the price for mineral water that they would receive.

A Cournot market is described by $N \geq 1$ profit-maximizing producers (or players) that compete in a non-cooperative way. In a market with homogenous goods, the players compete based on production quantity (producing identical goods). The market is specified by an inverse demand curve $P(\cdot)$, which maps aggregate production to market price, and as such is a *decreasing* function. The players choose their production levels q_i . Given total production level $Q = q_1 + \dots + q_N$, the market clearing price is $P(Q)$. A simple illustrative example is linear inverse demand, $P(Q) = \eta - Q$, where η is the saturation level beyond which prices collapse to zero (and may become negative, meaning a producer would have to pay to have his good taken away). Such linear demand can be derived from the behavior of a representative consumer with a quadratic utility function (see, for instance, Vives (2001)), and allows to present explicit equilibrium calculations.

The players produce at per-unit (constant) cost of production $s_i \geq 0$, which in general will be different, reflecting the costs of producing from heterogeneous energy sources. The profit of player i is the quantity he produces multiplied by price minus cost:

$$\pi(q_i, Q_{-i}, s_i) = \begin{cases} q_i (P(Q_{-i} + q_i) - s_i) & \text{if } q_i > 0, \\ 0 & \text{if } q_i = 0, \end{cases} \quad (1)$$

where $Q_{-i} = \sum_{j \neq i} q_j$ is total production by the players other than i . The last line of (1) allows for the possibility that $P(0^+) = +\infty$, but if a player does not produce anything, then he makes zero profit.

2.1 Nash Equilibrium

Definition 1. A Nash equilibrium is a vector $\mathbf{q}^* = (q_1^*, q_2^*, \dots, q_N^*) \in [0, \infty)^N$ such that, for all i ,

$$\pi(q_i^*, Q_{-i}^*, s_i) = \max_{q_i \in [0, \infty)} \pi(q_i, Q_{-i}^*, s_i), \quad (2)$$

where $Q_{-i}^* = \sum_{j \neq i} q_j^*$. That is, each player's equilibrium production q_i^* maximizes his own profit $\pi(\cdot, Q_{-i}^*, s_i)$ when the other $N - 1$ players produce their equilibrium quantities. If, in addition, $q_i^* > 0$ for all i , then \mathbf{q}^* is an interior Nash equilibrium.

Under suitable conditions on the price function P and the cost vector $\mathbf{s} = (s_1, s_2, \dots, s_N)$, a Nash equilibrium exists and is unique. An important issue arising from this is that it may be too costly for some players to participate.

We refer to (Vives, 2001, Chapter 4) for a discussion and references on general existence results for static Cournot games.

Assumption 1. *The price function P is twice continuously differentiable, with $P' < 0$ everywhere on $(0, \infty)$; and there exists $\eta \in (0, \infty)$ such that $P(\eta) = 0$.*

We order the firms by their costs and assume the latter are strictly less than the *choke price* $P(0^+)$:

$$0 \leq s_1 \leq s_2 \leq \dots \leq s_N < P(0^+). \quad (3)$$

When some firms have equal costs, the ordering is arbitrary and does not affect the result that follows. The behavior of P is best characterized in terms of the *relative prudence* of P , namely

$$\rho(Q) = -\frac{Q P''(Q)}{P'(Q)}. \quad (4)$$

We also define

$$\bar{\rho} = \sup_{Q \in (0, \infty)} \rho(Q). \quad (5)$$

The following is taken from Harris et al (2010).

Theorem 1. *Suppose that $\bar{\rho} < 2$. Then there is a unique Nash equilibrium which can be constructed as follows. Let $\bar{Q}^* = \max \{Q_n^* \mid 1 \leq n \leq N\}$, where Q_n^* is the unique non-negative solution to the scalar equation*

$$Q P'(Q) + n P(Q) = \sum_{j=1}^n s_j.$$

The unique Nash equilibrium production quantities are given by

$$q_i^*(\mathbf{s}) = \max \left\{ \frac{P(\bar{Q}^*) - s_i}{-P'(\bar{Q}^*)}, 0 \right\}, \quad 1 \leq i \leq N,$$

and the corresponding profits are

$$G_i(\mathbf{s}) = q_i^*(\mathbf{s})(P(\bar{Q}^*) - s_i), \quad 1 \leq i \leq N.$$

In particular, q_i^ and G_i are Lipschitz continuous, and the number of active players (that is, players with $q_i^* > 0$), in the unique equilibrium is $m = \min \{n \mid Q_n^* = \bar{Q}^*\}$.*

Moreover, in the case of a price function with a constant prudence, $\rho(Q) \equiv \rho$

$$P(Q) = \begin{cases} \frac{\eta}{1-\rho} \left(1 - \left(\frac{Q}{\eta}\right)^{1-\rho}\right) & \rho \neq 1 \\ \eta(\log \eta - \log Q) & \rho = 1, \end{cases} \quad (6)$$

one can weaken the requirement $\bar{\rho} < 2$ to simply $\rho < N + 1$.

The cost profile \mathbf{s} is the main parameter of a Cournot game and the respective sensitivity analysis is crucial, especially in dynamic models. Intuitively, we expect that if player- i costs s_i decrease, her production and profits will rise, and the production and profits of the other players will fall. However, the precise impact depends on the properties of the price function $Q \mapsto P(Q)$; for example there are well known examples where higher costs *increase* production for all players (Vives, 2001). Analysis of the precise dependence of equilibria on \mathbf{s} , including explicit formulas for the sensitivity of q_i^* to \mathbf{s} under constant prudence price functions of (6), is given in Harris et al (2010).

2.2 Blockading

The non-negativity constraint on production endogenizes the market structure in terms of the cost profile \mathbf{s} . Oligopolies with symmetric production costs generate a trivial market structure, namely either all firms active or all firms inactive. In contrast, in models where firms are asymmetric, some firms may be inactive in equilibrium. Moreover, in dynamic models asymmetric costs induce different entry times into the market. This aspect is especially pertinent to energy markets, where producers using different fuels and technologies have widely different costs of production: for example, oil and coal sources are much cheaper than renewables, such as solar or wind.

To illustrate this effect, consider a very simple case of Theorem 1, namely a duopoly $N = 2$ with linear demand $P(Q) = 1 - Q$ (i.e., $\eta = 1, \rho = 0$). When there is one player with marginal cost of production $s_1 \in [0, 1)$, he chooses his optimal quantity $q_1 \geq 0$ to maximize his monopoly profit function

$$\Pi_1 = q_1(1 - q_1) - s_1 q_1.$$

The optimal quantity and profit are given by

$$q_1^*(s_1) = \frac{1}{2}(1 - s_1), \quad G_1(s_1) = \frac{1}{4}(1 - s_1)^2.$$

When there are two players with costs $(s_1, s_2) \in [0, 1]^2$ and non-negative production quantities (q_1, q_2) , the aggregate quantity is $Q = q_1 + q_2$ and each player's profit function is

$$\Pi_i = q_i(1 - q_i - q_j) - s_i q_i, \quad i = 1, 2; j \neq i.$$

In a Nash equilibrium $(q_1^*, q_2^*) \in [0, 1]^2$ for the duopoly, each player maximizes profit as a best response to the other player's equilibrium strategy:

$$G_i(s_1, s_2) = \max_{q_i \geq 0} q_i(1 - q_i - q_j^*) - s_i q_i, \quad i = 1, 2; j \neq i.$$

For costs $s_1, s_2 < \frac{1}{2}$, it is easy to see that both players have positive equilibrium productions

$$q_i^*(s_1, s_2) = \frac{1}{3}(1 - 2s_i + s_j), \quad G_i(s_1, s_2) = \frac{1}{9}(1 - 2s_i + s_j)^2, \quad (7)$$

where $i = 1, 2; j \neq i$. However, if player j 's cost is too high relative to player i 's, specifically $s_j > \frac{1}{2}(1 + s_i)$, then he is *blockaded* from production, meaning his equilibrium quantity is zero. In this case, player i has a monopoly and the Nash equilibrium is given by

$$q_i^* = \frac{1}{2}(1 - s_i), \quad q_j^* = 0, \quad G_i = \frac{1}{4}(1 - s_i)^2, \quad G_j^* = 0.$$

See Figure 1 for the resulting monopoly wedges.

A current example may be OPEC holding back on cuts in production to drive shale oil producers out of the market and into bankruptcy, which an Op-Ed in The New York Times on 27 January, 2015 described thus: “the plunge in oil prices offers a sobering reminder of the power of markets over policy”.




Fig. 1 Type of Game Equilibrium in a Cournot Duopoly with Linear Demand $P(Q) = 1 - Q$.

A full characterization of the static N -player game for a wide class of general inverse-demand functions is given in (Harris et al, 2010, Section 2), and for Bertrand games in (Ledvina and Sircar, 2011, Section 2); a comparison between Cournot and Bertrand in terms of the number of blockaded players is in Ledvina and Sircar (2012).

3 Exhaustible Resources & Dynamic Games

When a resource, for instance a fossil fuel, is in finite supply, the energy oligopolies are necessarily changing over time due to the increasing scarcity value of the *exhaustible resource*.

3.1 Monopoly & Hotelling's Rule

The seminal work of Hotelling (1931) introduced a mathematical model for management of an exhaustible resource stock. Hotelling considered a single producer (a monopolist), and set up a continuous-time calculus of variations problem for maximizing total discounted value of the resource between now and exhaustion point. A crucial insight of Hotelling is the fact that along the optimum path the marginal value of reserves must grow at the risk-free rate, precisely offsetting the time value of money. This spawned a large body of economic literature based on optimizing social planning in the context of resource management or on Ramsey-type growth models that aim to optimize investment across several economic sectors. The link to exhaustible resources has become especially relevant in the past decade in connection with sustainable production in the face of climate change. For example development of clean energy backstops to guard against exhaustibility of conventional fossil fuels is addressed in Tsur and Zemel (2003); Lafforgue (2008); Grimaud et al (2011) among others.

Consider a single oil producer who has reserves $x(t)$ at time t , with the dynamics

$$\frac{dx}{dt} = -q(x(t))\mathbf{1}_{\{x(t)>0\}}, \quad (8)$$

where $q(x(t))$ is his production (or extraction) rate. When his reserves run out, he no longer participates in the market. The producer extracts to maximize lifetime discounted profit, and his value function v is defined by

$$v(x) = \sup_q \int_0^{\tau_x} e^{-rt} q(x(t)) P(q(x(t))) dt.$$

Here the maximization is over control strategies $q \geq 0$, $r > 0$ is the discount rate, P is the inverse demand (or price) function satisfying Assumption 1, and τ_x is the exhaustion time

$$\tau_x = \inf\{t > 0 \mid x(t) = 0\}.$$

Moreover x here stands for the initial resource level: $x(0) = x$, and we have assumed zero extraction costs.

By standard dynamic programming arguments, $v(x)$ solves the Hamilton-Jacobi (ordinary) differential equation

$$rv = \sup_{q \geq 0} qP(q) - qv', \quad x > 0, \quad (9)$$

where $v' = dv/dx$, and the boundary condition describing exhaustibility is $v(0) = 0$. Denote by $G(s)$ the solution to the static monopoly problem with cost s :

$$G(s) = \max_{q \geq 0} q(P(q) - s),$$

and the corresponding optimizer $q^*(s)$, where the parameter again refers to the production costs. Then the ODE (9) is simply $rv = G(v')$, so that v' plays the role of a shadow cost, or scarcity value in the dynamic exhaustible resources monopoly problem, and the optimal extraction policy is $q^*(v'(x(t)))$.

The first-order condition defining q^* is

$$P(q^*) + q^*P'(q^*) = s, \quad \Rightarrow \quad G(s) = -(q^*)^2P'(q^*). \quad (10)$$

Then, differentiating the ODE (9) with respect to x , we have

$$rv' = -q^*(v') \cdot v'' (2q^*(v') \cdot P'(q^*(v')) + (q^*(v'))^2 \cdot P''(q^*(v'))),$$

and differentiating the equation (10) for $q^*(s)$ with respect to s gives

$$q^{*'}(2P'(q^*) + q^*P''(q^*)) = 1.$$

Therefore, we have $rv' = -q^*(v')v''$, which evaluated along the optimal trajectory defined by $\frac{dx}{dt} = -q^*(v'(x(t)))$ (up until the exhaustion time τ_x) leads to

$$rv'(x(t)) = -q^*(v'(x(t)))v''(x(t)) = \frac{dx}{dt}v''(x(t)) = \frac{d}{dt}v'(x(t)).$$

This is known as Hotelling's rule (for a monopolist with exhaustible resources):

$$\frac{d}{dt}v'(x(t)) = rv'(x(t)), \quad (11)$$

or $v'(x(t)) = v'(x(0))e^{rt}$.

3.2 Multiple Players with Exhaustible Resources

Incorporating other players into a genuine game framework elevates the single player control/ODE problem to the setting of nonzero-sum differential games and partial differential equations (PDEs). Here existence and regu-

larity theory is scarce (outside of the case of linear-quadratic (LQ) games). General dynamic programming equations are studied in Başar and Olsder (1999), and some applications are presented in Dockner et al (2000). The solution approach generally goes through the feedback strategy representation, which allows to express optimal policies in terms of local properties of the game-value functions. For stationary models, one can utilize Euler-Lagrange methods, while in non-stationary or stochastic contexts, more involved analysis is necessary using Hamilton-Jacobi-Bellman-Isaacs tools. For a nonzero-sum dynamic game between N players, each with their own resources, the computation of a solution generally requires dealing with coupled systems of N fully nonlinear PDEs, with one value function per player. This quickly becomes very challenging as N grows and explains the focus on the duopoly $N = 2$ case.

To illustrate the complexity in the duopoly case, we let $x_i(t)$ be the reserves of each player at time t , which are depleted at their extraction rates q_i :

$$\frac{dx_i}{dt} = -q_i(\mathbf{x}(t)), \quad i = 1, 2 \quad \text{where} \quad \mathbf{x}(t) = (x_1(t), x_2(t)).$$

With assumed zero extraction costs and the same discount rate $r > 0$, each player maximizes lifetime discounted profit in (best) response to the extraction policy of the other. A Nash (or Markov perfect) equilibrium (q_1^*, q_2^*) , if it exists, describes the value functions

$$v_i(\mathbf{x}) = \sup_{q_i \geq 0} \int_0^{\tau_{x_i}} e^{-rt} q_i(\mathbf{x}(t)) P(q_i(\mathbf{x}(t)) + q_j^*(\mathbf{x}(t))) dt, \quad i = 1, 2; j \neq i, \quad (12)$$

where $\tau_{x_i} = \inf\{t > 0 \mid x_i(t) = 0\}$ are the exhaustion times starting at $x_i = x_i(0)$. The state-space approach in (12) restricts attention to policies specified in closed-loop feedback form $q_i(\mathbf{x})$, linking to the single-agent control frameworks and removing technical challenges related to equilibrium existence. Moreover, it naturally generalizes to the stochastic extensions discussed below. Dynamic programming arguments lead to the following equations for the value functions in $x_1, x_2 > 0$:

$$rv_i = \sup_{q_i \geq 0} \left\{ q_i P(q_i + q_j^*) - q_i \frac{\partial v_i}{\partial x_i} \right\} - q_j^* \frac{\partial v_i}{\partial x_j}, \quad j \neq i,$$

which, using the notation introduced in (1), we can write as

$$rv_i = \sup_{q_i \geq 0} \pi \left(q_i, Q_{-i}^*, \frac{\partial v_i}{\partial x_i} \right) - q_j^* \frac{\partial v_i}{\partial x_j}.$$

This identifies the infinitesimal problem in the dynamic programming equation as the static Nash equilibrium problem with scarcity costs $s_i = \frac{\partial v_i}{\partial x_i}$. The interaction of exhaustibility and blockading allows to endogenize market

structure. As players deplete their reserves, their marginal costs may rise sufficiently to make further production uneconomical and causing them to drop out of competition.

Using the notation of Theorem 1 for the solution of the static game, we can write the dynamic game PDEs as

$$rv_i = G_i(\mathcal{D}v) - q_j^*(\mathcal{D}v) \frac{\partial v_i}{\partial x_j}, \quad i = 1, 2; j \neq i, \quad \mathcal{D}v = \left(\frac{\partial v_i}{\partial x_i}, \frac{\partial v_j}{\partial x_j} \right).$$

For the linear pricing function example, the functions G_i and q_i^* were given in (7). In a model of *only* exhaustible resources, when player i runs out, player j has a monopoly until he also exhausts, which lead to boundary conditions on $x_i = 0$ and at $(0, 0)$ respectively. A more nuanced (and optimistic) view of future energy production allows that when an oil producer exits, he is replaced by an inexhaustible producer (such as from solar) with infinite (or sustainable) supply. This type of model is analyzed by asymptotic and numerical methods in Harris et al (2010).

Alternatively, one can consider models with a single exhaustible producer, and hence a single state variable, along with $N - 1$ renewable producers that do not need to worry about reserves. This maintains game effects but minimizes mathematical complexity (see Harris et al (2010); Ledvina and Sircar (2012)), and allows to study the effect of blockading: how low must oil reserves go before it becomes profitable to start producing from more expensive but sustainable sources? As (levelized) costs of setting up and maintaining energy production from different sources are different, the entry points for solar and wind, for instance, may likewise be very different. This generates phase transitions in the dynamic game characterized by the (endogenously determined) number of active producers $n(t)$. The respective blockading points can be computed explicitly for the linear price function model, see Ledvina and Sircar (2012) where it is also shown that a modified, piecewise, version of Hotelling's rule holds in the presence of competition. Namely, there exist blockading times $\tau_1^b \leq \tau_2^b \leq \dots$, such that for $t \in [\tau_{n-1}^b, \tau_n^b)$ there are n energy producers (one exhaustible oil producer and $n - 1$ active renewables), and the marginal value function for the oil producer along the equilibrium extraction path grows according to

$$\frac{d}{dt} v'(x(t)) = \left(\frac{1}{2} + \frac{1}{2n} \right) rv'(x(t)), \quad t \in [\tau_{n-1}^b, \tau_n^b). \quad (13)$$

The above relationship recovers the classical Hotelling rule when $n = 1$ (oil monopoly) and blunts the sharp price increases associated with 'peak oil' (the growth rate of v' declines as oil runs out and renewables enter).

4 Renewability & Exploration

Exhaustibility is manifested through consideration of the resource reserves which must remain non-negative. Exhaustible resource stocks can be divided into three types: non-renewable, renewable, and replenishable. Non-renewable resources are only available one time, and once used-up are gone forever. Thus, the level of remaining reserves $x(t)$ is non-increasing. The extreme case of $x(t) = 0$ represents complete exhaustion and requires a boundary condition to specify the resulting utility for the producers. A common example are fossil fuels in a global physical context. Once fossil fuels are removed, one can imagine that production is permanently suspended $v_i(0) = 0$; alternatively one could switch to a more expensive (green) backstop, as described above.

Renewable resources, like fisheries or forests, can grow back on their own if left unexploited. A common setup is a logistic growth model with a finite capacity, such that $\frac{dx}{dt} = F(x(t))$, an ordinary differential equation (see e.g., Benckroun et al (2009)). With renewable resources the main concern is over-exploitation: if production is too high, stocks can be damaged in the long-term (or completely exhausted). However, sustainable extraction is possible all on its own, and only requires enforceable discipline. Mathematically, sustainability/renewability leads to a stationary model where a local-in-time equilibrium between production and extraction generates a global solution (as a long-term steady state); in contrast non-renewable game equilibria are inherently time-dependent and in particular strongly affected by the “terminal condition” of running out of reserves. While steady-state models are not suited for most energy sources, they are common for describing pollution stock dynamics.

Replenishable resources capture the middle ground—reserves can grow, but this requires separate effort/costs. This is meant to represent costly search for, say, new mines or oil fields and matches the industrial exploration-and-production (E&P) cycle. Indeed, with economic incentives the reserve base is not fixed and can be increased. For example, while oil is exhaustible, it is also replenishable since there is a difference between total abstract reserves on Earth, and what is actually commercially “proven” and drives production decisions. Under replenishable reserves, both the upward and downward dynamics in $x(t)$ are controlled. Mathematically, exploration is modeled via a separate control a_t (which may be coupled to production level q_t). Under exploration, the boundary case $x(t) = 0$ requires separate consideration in terms of whether players can “resurrect” themselves. Assuming that future discoveries can finance present exploration (Pindyck, 1978; Ludkovski and Sircar, 2011) leads to an implicit boundary condition for $v_i(0)$.

In a pure resource model, each player has their own, independent, reserves $x_i(t)$, which she has complete control over. In such models, players interact solely through the price mechanism; reserves then add a separate individual *marginal cost of production*, introducing a new source of asymmetry between the players. Alternatively, especially for renewable resources, one may add

further game effects by tying together reserves. This can be done by postulating a single, common reserve stock $x(t)$ (akin to the classical tragedy of the commons (Benckroun, 2003)), or by considering individual reserves levels that have coupled dynamics. For example Colombo and Labrecciosa (2013) used

$$\frac{dx_i}{dt} = \delta x_i(t) 1_{\{X(t) < \bar{X}/2\}} + \delta(\bar{X}_i - x_i(t)) 1_{\{X(t) > \bar{X}/2\}},$$

where $X(t) = \sum_i x_i(t)$ is the total resource stock, and δ is the resource growth rate. Thus, up to the aggregate sustainable level $\bar{X}/2$, resource stocks grow exponentially; beyond $\bar{X}/2$ growth rates lessen linearly, turning negative above the individual carrying capacity \bar{X}_i .

Turning our attention to the shocks affecting reserves, the classical formulation provides the deterministic dynamics

$$\frac{dx_i}{dt} = -q_i(t) + F(X_i(t), a_i(t)),$$

where the first term represents lower reserves due to extraction, and the second term represents reserves growth (thanks to either exogenous or endogenous factors). Thus, the future level of reserves is completely determined by the players' strategies and can be extrapolated to any future date t . While mathematically convenient, this is not very realistic, since practical forecasts of future stocks clearly involve a lot of uncertainty (consider for example forecasting of fishery stocks in 2020, or the fossil fuel exhaustion point somewhere in the next few hundred years). This uncertainty permeates even central planner growth models, so is not solely a feature of uncertainty about future equilibria.

4.1 Shocks to Reserves

Taking a stochastic tack, some models have therefore incorporated stochastic dynamics for reserves, now denoted by a *stochastic process* $X_i(t)$:

$$dX_i(t) = -q_i(t) dt + \sigma_i dW_t^i, \quad (14)$$

arguing that reserve levels are uncertain and subject to ongoing up/down revisions described by the Brownian motions W^i . Within a Bertrand competition, Ledvina and Sircar (2011) justified similar Brownian shocks through small fluctuations in the respective demand levels.

Stochastic shocks become especially pertinent for replenishable stocks, where the attendant exploration efforts yield intrinsically stochastic outcomes. Thus, starting with the seminal work of Kamien and Schwartz (1978), there has been a long literature on stochastic *exploration*. In particular, discrete upward jumps in reserves, modeled as a Poisson process, have been

advocated, leading to $dX_i(t) = -q_i(t) dt + \delta dN_t^i$ where δ are (random) increments and N^i is a controlled point process. A common setup is to specify controlled intensity of N^i , i.e., $\lambda_t^i = G(a_i(t))$ where λ^i is the hazard rate of arrivals of N^i . This leads to HJB-I system of equations for the game value functions, see Ludkovski and Sircar (2011).

4.2 More Stochasticity

Beyond the aforementioned stochastic shocks to reserves, one can imagine other factors that generate random environment for the Cournot producers. This is especially so over the medium- and long-run contexts that are often used to motivate the models. Clearly on a longer scale essentially every aspect of the market, including demand, costs, reserves, etc., is subject to unpredictable changes.

To capture macroeconomic cycles, Ludkovski and Yang (2014) considered stochastic demand, so that $P_t = P(D_t, \mathbf{q}_t)$ has exogenous shocks from the stochastic factor D_t . For example, taking D_t to be a 2-state independent Markov chain allows to maintain tractability, while representing the low- and high-demand regimes that can be associated with commodity booms and recessions. In combination with non-renewable resources, stochastic demand generates the phenomenon of *strategic mothballing*, whereby producers may temporarily shut-down production during low demand periods. Another regime-switching model with exhaustibility but a single agent is in Dasgupta and Stiglitz (1981).

From a different angle, Dasarathy and Sircar (2014) considered non-constant production costs to mimic the non-stationary economics of extracting more and more difficult to access reserves. Indeed, as well-documented empirically, extraction costs of say crude oil steadily rise as conventional, cheap sources are depleted and replaced by non-conventional oil sands, deep off-shore and shale fields: see Figure 2.

Accordingly, Dasarathy and Sircar (2014) take costs $s_i(x)$ to depend on reserves, such that $s_i(x)$ increases as x decreases for exhaustible players, and decreases as x decreases for renewable players (due to government subsidies as conventional energy sources are depleted). The resulting dynamic game can force the exhaustible player to leave early, i.e. X_t never reaches zero.

4.3 Other Types of Strategic Interactions

Beyond production/exploration, the literature has also considered other player controls. One major idea from industrial organization (IO) concerns Research and Development (R&D) which generates additional benefits (such

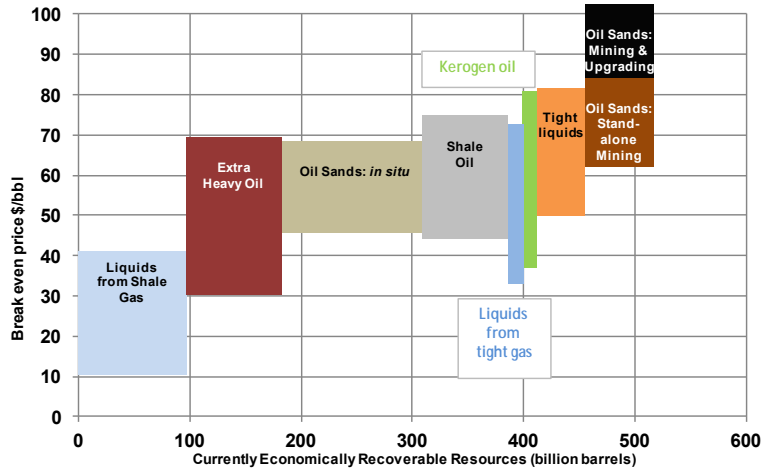


Fig. 2 Estimated oil extraction costs from varying sources. Source: IEA (2012).

as lower production costs or first-mover advantages) to the innovator. In the context of R&D, efforts to innovate may lead to spillovers (Cellini and Lambertini, 2009; Dawid et al, 2013), introducing a different source of coupling between players. Spillovers are well-documented empirically and tend to lower R&D investments and therefore reduce productivity growth. Game-theoretically, spillovers can be viewed as either raising the innovation rate of competitors in a static set-up, or removing first-mover advantages after innovation success. A notable reference is Fölster and Trofimov (1997) who consider an oligopoly where each of the N firms maximizes R&D effort. The random first innovator is determined stochastically and temporarily collects extra profits. Cellini and Lambertini (2009) studied a deterministic R&D game with spillovers.

Another link is to the theory of real options, by modeling the strategic opportunities available to producers as one-shot events to be optimally timed. For example, a classical setting concerns producers competing to initiate a new project (such as development of a renewable energy backstop to an exhaustible resource, see Hung and Quyen (1993)) that carries first-mover advantage and leads to a so-called preemption game. Such timing games partition the global model into distinct phases, providing a different mechanism to endogenize market structure. They can be seen as intermediate ground between a static and differential game.

5 Mean Field Games

In dynamic oligopoly problems with a finite number of players, the HJB system of PDEs does not admit an explicit solution, except possibly in the monopoly case. As a result, one needs numerical means for computing the value functions, as well as the equilibrium strategies, which of course quickly becomes infeasible as the number of players goes beyond three. Moreover, even in the two-player case, these equations are hard to handle. To overcome this problem, one may study the market dynamics when the number of firms tends to infinity by using the concept of a mean field game (MFG). MFGs were proposed by Lasry and Lions (2006, 2007) and independently by Huang et al (2006) to handle certain types of competition in the continuum limit of an infinity of small players.

The interaction is modeled by assuming that each player only sees and reacts to the statistical distribution of the states of other players. Optimization against the distribution of other players leads to a backward (in time) Hamilton-Jacobi-Bellman (HJB) equation; and in turn their actions determine the evolution of the state distribution, encoded by a forward Kolmogorov equation. We refer to the survey article by Guéant et al (2011) and the recent monograph of Bensoussan et al (2013) for further background. In our context, the mean-field interaction is captured by making the market price be a function of equilibrium global supply, which in turn is affected by the distribution of reserves $m_t(\cdot)$, which is a measure on \mathbb{R}_+ . Thus, inverse demand curve translates into a functional relationship between reserve distribution $m_t(\cdot)$ and resulting price P_t , according to $D(t, P_t) = \int_{\mathbb{R}_+} q_t^*(x) m_t(dx)$. Numerical resolution of MFG equations is an active area of research; simultaneously dealing with the forward-backward system of PDEs typically requires a fixed-point iteration scheme. At the same time, removing the awkward dependence on the number of players simplifies the equations and provides a unified theory in terms of measure-valued processes.

The analysis of the relationship between Markov perfect equilibria in finite- N Cournot games and the MFG limit was carried in Chan and Sircar (2014) considering the general situation of differentiated goods and asymmetric players. Chan and Sircar (2014) also show that in the continuum MFG limit, the linear Bertrand and Cournot models are equivalent, removing the usual distinction observed with a finite number of players.

In a related work, Guéant et al (2010) analyzed a MFG formulation of oil oligopoly for both the deterministic case (reserves endogenously determined by production rates) and the stochastic case (reserves include Brownian noise, generating a parabolic forward Kolmogorov equation for m_t). Guéant et al (2010) proposed an iterative numerical scheme and presented some numerical examples for both linear and constant elasticity of substitution (CES) demand curves. One particular focus was on the marginal cost of exhaustibility (Hotelling rent) and also on substitution effects in a 2-energy model. See

also Bauso et al (2012) who treated a *robust* version of above which adds another first-order quadratic term to the MFG equations.

6 Summary of Game Models for Exhaustible Resources

In the table below g refers to a green producer, so that $1 + g$ is a duopoly with one exhaustible and one renewable player, see Section 3.2. Infinity of players corresponds to mean-field models. Demand refers to the shape of the price function $P(Q)$.

	# Players	Type	Demand	Randomness	Replenish
Hotelling (1931)	1	–	linear	Determ.	No
Dasgupta and Stiglitz (1981)	N	Cournot	constant	single-shock	No
Deshmukh and Pliska (1983)	1	–	regimes	Poisson	Yes
Benchekroun (2008)	N	Cournot	linear	Determ.	Yes
Benchekroun et al (2009)	N	Cournot	linear	Determ.	Yes
Harris et al (2010)	$1+g$	Cournot	linear	Brownian	No
Ludkovski and Sircar (2011)	$1+g$	Cournot	linear	Poisson	Yes
Ledvina and Sircar (2012)	$1+N$	Bertrand	linear	Determ.	No
Ludkovski and Yang (2014)	$1+g$	Cournot	linear	Poisson	Yes
Colombo and Labrecciosa (2013)	N	Cournot	linear	Determ.	Yes
Dasarathy and Sircar (2014)	$1+N$	Cournot	linear	Poisson	Yes
Guéant et al (2010)	∞	Cournot	CES	Determ.	No
Chan and Sircar (2014)	∞	Bertrand	linear	Brownian	No

References

- Başar T, Olsder GJ (1999) Dynamic Noncooperative Game Theory. SIAM, Philadelphia
- Bauso D, Tembine H, Başar T (2012) Robust mean field games with application to production of an exhaustible resource. In: Proceedings of 7th IFAC symposium on robust control design, Aalborg, Denmark, pp 454–459
- Benchekroun H (2003) Unilateral production restrictions in a dynamic duopoly. *Journal of Economic Theory* 111(2):214–239
- Benchekroun H (2008) Comparative dynamics in a productive asset oligopoly. *Journal of Economic Theory* 138(1):237–261
- Benchekroun H, Halsema A, Withagen C (2009) On nonrenewable resource oligopolies: The asymmetric case. *Journal of Economic Dynamics and Control* 33(11):1867–1879
- Bensoussan A, Frehse J, Yam P (2013) Mean Field Games and Mean Field Type Control Theory. Springer
- Cellini R, Lambertini L (2009) Dynamic R&D with spillovers: Competition vs cooperation. *Journal of Economic Dynamics and Control* 33(3):568–582
- Chan P, Sircar R (2014) Bertrand and Cournot mean field games. *Applied Mathematics & Optimization* To appear, DOI: 10.1007/s00245-014-9269-x
- Colombo L, Labrecciosa P (2013) Oligopoly exploitation of a private property productive asset. *Journal of Economic Dynamics and Control* 37(4):838–853

- Cournot A (1838) *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. Hachette, Paris, english translation by N. T. Bacon published in *Economic Classics*, Macmillan, 1897, and reprinted in 1960 by Augustus M. Kelly
- Dasarathy A, Sircar R (2014) Variable costs in dynamic Cournot energy markets. In: Aid R, Ludkovski M, Sircar R (eds) *Energy, Commodities and Environmental Finance*, Fields Institute Communications, Fields Institute
- Dasgupta P, Stiglitz J (1981) Resource depletion under technological uncertainty. *Econometrica: Journal of the Econometric Society* pp 85–104
- Dawid H, Kopel M, Kort P (2013) R&D competition versus R&D cooperation in oligopolistic markets with evolving structure. *International Journal of Industrial Organization* 31(5):527–537
- Deshmukh SD, Pliska SR (1983) Optimal consumption of a nonrenewable resource with stochastic discoveries and a random environment. *The Review of Economic Studies* 50(3):543–554
- Dockner EJ, Jørgensen S, Long NV, Sorger G (2000) *Differential games in economics and management science*. Cambridge University Press, Cambridge
- Fölster S, Trofimov G (1997) Industry evolution and R&D externalities. *Journal of Economic Dynamics and Control* 21(10):1727–1746
- Grimaud A, Lafforgue G, Magné B (2011) Climate change mitigation options and directed technical change: A decentralized equilibrium analysis. *Resource and Energy Economics* 33(4):938–962
- Guéant O, Lasry JM, Lions PL (2010) Mean field games and oil production. Tech. rep., College de France, URL <http://mfglabs.com/wp-content/uploads/2012/12/cfe.pdf>
- Guéant O, Lasry JM, Lions PL (2011) Mean field games and applications. In: *Paris-Princeton Lectures on Mathematical Finance 2010*, Springer, pp 205–266
- Harris C, Howison S, Sircar R (2010) Games with exhaustible resources. *SIAM J Applied Mathematics* 70:2556–2581
- Hotelling H (1931) The economics of exhaustible resources. *Journal of Political Economy* 39 (2):137–175
- Huang M, Malhamé RP, Caines PE, et al (2006) Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Communications in Information & Systems* 6(3):221–252
- Hung N, Quyen N (1993) On R&D timing under uncertainty: the case of exhaustible resource substitution. *Journal of Economic Dynamics and Control* 17(5):971–991
- IEA (2012) *Medium-term oil and gas markets outlook*. Tech. rep., International Energy Agency, Paris, France
- Kamien MI, Schwartz NL (1978) Optimal exhaustible resource depletion with endogenous technical change. *Review of Economic Studies* 45(1)
- Lafforgue G (2008) Stochastic technical change, non-renewable resource and optimal sustainable growth. *Resource and Energy Economics* 30(4):540–554
- Lasry JM, Lions PL (2006) Jeux à champ moyen. i–le cas stationnaire. *Comptes Rendus Mathématique* 343(9):619–625
- Lasry JM, Lions PL (2007) Mean field games. *Japanese Journal of Mathematics* 2(1):229–260
- Ledvina A, Sircar R (2011) Dynamic Bertrand oligopoly. *Applied Mathematics & Optimization* 63(1):11–44
- Ledvina A, Sircar R (2012) Oligopoly games under asymmetric costs and an application to energy production. *Mathematics and Financial Economics* 6(4):261–293
- Ludkovski M, Sircar R (2011) Exploration and exhaustibility in dynamic Cournot games. *European Journal of Applied Mathematics* 23(3):343–372
- Ludkovski M, Yang X (2014) Dynamic Cournot models for production of exhaustible commodities under stochastic demand. In: Aid R, Ludkovski M, Sircar R (eds) *Energy, Commodities and Environmental Finance*, Fields Institute Communications, Fields Institute

- Pindyck R (1978) The optimal exploration and production of nonrenewable resources. *Journal of Political Economy* 86:841–862
- Tsur Y, Zemel A (2003) Optimal transition to backstop substitutes for nonrenewable resources. *Journal of Economic Dynamics and Control* 27(4):551–572
- Vives X (2001) *Oligopoly pricing: old ideas and new tools*. The MIT press