

A Probabilistic Graphical Approach to Computing Electricity Price Duration Curves under Price and Quantity Competition.

Pascal Michailat and Shmuel Oren

Industrial Engineering & Operations Research Department

University of California at Berkeley

Berkeley, California 94720-1777

Email: {pascal,oren}@ieor.berkeley.edu

Abstract—The electricity price duration curve (EPDC) represents the probability distribution function of the electricity price considered as a random variable. The price uncertainty comes both from the demand side and the supply side, since the load varies continuously, and not all generators may be available at all times. The production costs of electricity also fluctuate with the price of fuel. EPDCs have many application including the valuation of incremental generation assets or forward contracts on the energy produced by such assets, estimating capacity cost recovery and valuation of energy call options. Traditional approaches for calculating EPDCs were based on approximation methods such as the method of cumulants using Edgeworth expansions of multivariate probability distributions. This paper presents a new approach to compute numerically the EPDC under price and quantity competition models. This numerical method provide both exact numerical results and modeling flexibility. It is based on inference algorithms in probabilistic graphical models (PGMs) which exploit conditional independence relationships among the random variables.

I. INTRODUCTION

In regulated electricity industries, regulators have always been concerned with securing the supply of electricity and ensuring adequate and fair remuneration of the regulated monopolies entrusted with the supply of electricity. Since the electricity industry was viewed as a domain of natural monopolies it was the role of public authorities to control whether these monopolies were making adequate investment in generation and transmission facilities so as to ensure a high level of reliability at just and reasonable rates.

But the liberalization of electricity markets and the increased commercial use of interconnections between regional distribution networks have led to important interdependences of regional networks and to a retreat of regulatory authorities. However, electricity crises and massive blackouts that occurred around the world over the last few years have raised public concern regarding the supply-demand equilibrium and the means to secure enough investment in production and transportation of electricity in a competitive market. Attracting investors requires, among other thing, accurate means to predict electricity prices and cost recovery capability in the medium and long term. Predicting electricity prices is a complex task as the market is subject to physical uncertainty

(e.g., fuel prices, availability of the generators, availability of the transmission lines), economic uncertainty (e.g., number of participants in the market, strategy employed by the players, price-elasticity of the load), and finally regulatory uncertainty (e.g., market rules set up by the regulator).

The electricity price duration curve (EPDC) is a tool that captures price uncertainty in the form of a probability distribution describing the probability or the fraction of time that the market clearing price will exceed any particular level. Such information can be used, for instance, to predict the fraction of time that a resource offered at a fix marginal price or an option contract on energy at a given strike price will be "in the money", i.e. will be competitive relative to the prevailing whole sale price characterized by the EPDC. The EPDC can also be used to price energy call options and generation capacity and to evaluate the inframarginal profit of a generation plant with known marginal production cost. Such valuation is necessary in order to determine whether the plant will be able, on average, to recover its amortized fixed cost. The pricing of energy call options with specified strike prices is an important potential application of EPDCs since such instruments are gaining support as a mechanism for assuring generation adequacy. DeVries [1] provides strong support for physically covered call options as the capacity mechanism of choice, in terms of stabilization of the investment level, robustness against regional shortage, effectiveness in securing generation in an open market, and robustness against market power in the electricity market. A detailed description of a capacity mechanism involving physically covered call options is given in [2]. EPDCs are also useful in pricing bilateral forward supply contract, which are the prevailing form of electricity transaction even in electricity systems with highly developed spot markets. Allaz and Vila [3] provide theoretical support to the importance of forward contracting showing that forward contracts should not only be considered as hedging instruments, but can also increase consumer surplus by increasing competition in oligopoly markets.

Traditional methods for computing EPDC were based on the numerical method of cumulants used to approximate the convolution of independent random variables characterizing

the availability of generation plants. Originally this approach was applied to compute the distribution of marginal cost in a regulated electricity system dispatched centrally in merit order (e.g. [4]). This approach has been recently employed in [5], [6], to approximate the EPDC in a Bertrand, Cournot and supply function equilibrium (SFE) setting. In this paper, we propose a new exact method that can be used to forecast the electricity price distribution. This method is very flexible, and could accommodate many different models. It is based on probabilistic graphical models (PGMs). These probabilistic models provide a general framework for dealing with problems involving a very large number of random variables. Each node of a graph is naturally associated with a random variable. The edges of the graph reflect dependencies between the random variables: Graphical models take advantage of the conditional independence relations between the nodes of the graph to provide a compact representation of a problem. This compactness will allow for more accuracy and reliability in the inference and parameter estimation. We illustrate the applicability of this approach by predicting electricity prices in price competition and in quantity competition under uncertainty.

We will consider two sources of physical uncertainty: uncertainty of the load created by a random demand shock, and randomness of the supply caused by the possible outages of the generators. We examine the effect of such uncertainty on market prices of electricity. In our static price competition model, we assume that the generators behave as price takers and offer their output at marginal cost while an independent system operator (ISO) dispatches them optimally in merit order. In our two-player dynamic quantity competition model, each player learns at the beginning of each period the availability of her generators, but ignores the load level and the state of her opponent's generators. We consider two alternative information structures concerning what players know about the previous period's game. In our first model, each player learns the realization of the demand and the availability of her opponent's generator a posteriori, once the period is over. They use this information and their knowledge of the underlying random processes to form beliefs regarding the state of the system in the current period. In our second model, only the electricity market price is public. However, the players will use this limited information to infer last period's residual demand, form beliefs regarding the residual demand they are likely to face in the following period. In both cases, the players engage in a non-cooperative strategic interaction.

The structure of the paper is as follows. In Section II, we briefly introduce probabilistic graphical models and define the class of PGMs we use in this research. In Section III, we describe our stochastic model for the electricity market, and in Section IV and V, we define more precisely the price and quantity competition frameworks that we consider, present the graphical models used to predict the electricity price, and provide numerical examples. Section VI concludes.

II. PROBABILISTIC GRAPHICAL MODELS

A. General Introduction to PGMs

Intuitively, a graphical model can be thought as a structure carrying probabilistic relations between a set of random variables corresponding to the nodes of the graph. This probabilistic structure appears in the potential functions defined over subsets of nodes of the graph. Probability theory ensures the consistency of the aggregate construction. The model is used to answer queries about the random variables and their probabilistic relations. We may want to find the marginal probability distribution of a random variable or learn the most likely value of some parameters of the statistical model.

The advantage of graphical models comes from the economical representation of the joint probability distribution that they allow, based on certain conditional independence assumptions on the underlying probabilistic model. The relations between random variables are therefore only local, and the different algorithms defined on graphical models extensively use these local relations to perform rapidly various inference operations.

Jordan [7] provides a complete tutorial demonstrating the potential of probabilistic graphical models and Jordan [8] is a collection of works presenting learning methods in graphical models. Wainwright [9] gives a brief introduction to the message-passing algorithms that we use in this paper to infer the price distribution. Finally, Murphy [10] provides more details on dynamic graphical models, their representation and associated inference algorithms.

B. Construction of a PGM to Describe Competition in the Electricity Market

A PGM is constituted of two different layers: the graph structure and the potential functions defined over subsets of nodes of the graph.

1) *Graph structure*: We will work with a *directed graphical model*, also known as *Bayes net*. Directed graphical models are more adapted here since we model a physical system where the causal relations can be postulated a priori. Given a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, each node $i \in \mathcal{V}$ has an associated random variable X_i . A clique $C \subseteq \mathcal{V}$ is a subset of fully connected nodes: $\forall i, j \in C, \exists e \in \mathcal{E}, e = (i, j)$. We define $\forall A \subseteq \mathcal{V}, X_A := \{X_i, i \in A\}$.

A graph imposes restrictions on the joint probability distribution p over the support defined by the vector $X_{\mathcal{V}}$. Particularly, p must satisfy the *conditional independence property* (CI). Let $\pi(s)$ be the set of parents of node s in \mathcal{G} . A topological ordering I of \mathcal{V} is such that all the nodes in $\pi(s)$ appear before s in I . $\forall i \in I, \nu(s) = \{1, \dots, s-1\} / \pi(s)$. $X_{\mathcal{V}}$ satisfies the CI property if $\forall s \in \mathcal{V}, X_s$ is independent of $X_{\nu(s)}$ given $X_{\pi(s)}$. We build a graph \mathcal{G} so that the known statistical properties of the random vector $X_{\mathcal{V}}$ are reflected onto the graph.

2) *Potential functions*: $\forall i, \mathcal{X}_i$ is the state space associated with the random variable X_i . We need to be able to obtain a concise representation of the joint distribution $p : \mathcal{X}_1 \times \dots \times \mathcal{X}_n \rightarrow [0, 1]$. The two possible options are either to impose

a particular functional form on p , such as a multinomial Gaussian distribution, or to discretize the state spaces \mathcal{X}_i and then be completely free to specify the joint distribution on this discretized space. We choose the second approach.

Because the distribution p on $X_{\mathcal{V}}$ satisfies the CI property, the Hammersley-Clifford theorem [7] tells us that p can also be factorized in the form $p(X_{\mathcal{V}}) = \prod_{s \in \mathcal{V}} p(X_s | X_{\pi(s)})$. Therefore we only need to specify the conditional probability distributions of each node of the graph given his parents: *the potential functions*, to fully describe the joint distribution p . Moreover, these potentials can be represented on a tabular form because the state spaces \mathcal{X}_i are discrete. One key parameter of the PGM representation is the maximal size of the conditional probability tables: $\max_{s \in \mathcal{V}} |\mathcal{X}_s| \times \prod_{t \in \pi(s)} |\mathcal{X}_t|$. We will have to limit this parameter to a reasonable size to avoid memory problems when working on the graph.

C. Inference on PGMs

To compute marginal probabilities at various nodes of the graph, we use the *junction tree algorithm*. It is a message-passing algorithm, but not directly on the original graph associated with the problem: The graph needs to be first transformed to an equivalent *clique-tree* where the cliques of the original graph are gathered to form “meta-nodes” so that the cycles disappear and we obtain a tree. Some conditions on the graph and the transformation need to be satisfied so that the inference gives consistent results, because some nodes may appear in several cliques. The junction tree algorithm allowed us to compute efficiently exact marginal and conditional probabilities.

In the case where we work with a *dynamic Bayesian net*, i.e., a graphical model whose structure is repeated to account for a time dimension, we do not want to use the junction tree algorithm which is too cumbersome. To reduce the computational time, we resort to an approximate inference algorithm: the *BK algorithm* described in [11]. The general intuition behind the BK algorithm is that some random variables may only be weakly correlated together, and it may thus not be necessary to keep track of all the correlations to obtain a good approximation of the belief state on the system. Consequently, the BK algorithm will represent the belief state over the system as a set of localized beliefs on subsets of the system. These subsets can range from the whole system as a subset, in which case we obtain an exact result, to a fully factorized version where every node constitutes a subset. The error induced by this approximation is bounded since the errors on the belief state contract exponentially as the system evolves in time.

III. THE STOCHASTIC ELECTRICITY MARKET MODEL

We try to compute the EPDC, which is the curve defined at period t by $y = \mathbb{P}\{p(t) > x\}$, with $p(t)$ the spot price of electricity in period t . Our model focuses on the uncertainty of the load and availability of the generating units. Particularly, we ignore variations of the production costs (that could result from random fuel costs), unit commitment and transmission constraints.

A. Supply Side Uncertainty

The supply side uncertainty comes from the stochastic availability of the generators: We associate a random variable $Y_i \in \{0, 1\}$ to each production plant i . $Y_i = 1$ if the generator provides electricity to the network, $Y_i = 0$ if it is shut down. The availability of the generators will be described by a continuous time Markov process with failure rate λ_i and repair rate μ_i .

B. Demand Side Uncertainty

The actual demand function at hour t is $L(t) = K(t) - \zeta p(t)$. The nominal load $K(t)$ is a random variable which corresponds to a demand shock. ζ reflects the price-elasticity of the demand and is constant in time. As explained earlier, the state space of $K(t)$ will have to be discretized. We form a demand state space \mathcal{D} and the distribution of $K|t$ is in $\Delta(\mathcal{D})$.

C. Choice of a Competition Model

When modeling competition between firms, one needs to choose the strategic variable: price or quantity, which will be used. The role of the strategic variables in presence of uncertainty of the residual demand is studied theoretically in [12].

The most natural model of competition is the Bertrand model of price competition [13]. The Nash equilibrium (NE) of the Bertrand game is to bid marginal cost and the players should not make any profit. These results constitute the *Bertrand paradox* since “It is hard to believe that firms in an industry with few firms never succeed in manipulating the market price to make profits.” [14].

The immediate alternative to this model is the Cournot model of quantity competition [15] which yields results that conform better with the real world: The players bid above marginal cost in equilibrium and make a positive profit [14]. Different variations of the Bertrand model have tried to reconcile it with real-world outcomes. Particularly, several dynamic models of price competition [16], [17], [18] have been developed that lead to possible price cycles¹, potential stabilization of the price above the marginal costs, etc.

The question of the most relevant strategic variable has also been raised in modeling electricity markets. Because the exercise of market power has been proved to be a reality in electricity markets such as California [1] or the U.K. [19], it has often been argued that a quantity competition model would be more relevant [20].

However, we will study both competition frameworks. This offers two different perspectives on the forecasting of electricity prices which could be adequate under different conditions. It also shows the flexibility of the PGM approach which can tackle these two different problems efficiently.

IV. PRICE COMPETITION MODEL

This first model is a *probabilistic production costing model*: First introduced by Baleriaux et al. [21], these models consider

¹The Edgeworth cycles.

load uncertainty and take the availability of the producing units into account. Additionally, they assume that the generating units are dispatched by a central planner according to economic merit order, i.e., in increasing order of marginal costs. The rationing rule is as follows: if several generators have the same marginal cost and are marginal, they produce an equal amount of energy to satisfy the load, and thus share equally the profit. In the context of price competition, assuming a production costing model is tantamount to assuming that perfect competition reigns, and that each generator bids at marginal cost.

Our model can therefore be adapted to describe a very competitive market, but is an inadequate characterization of a market where players can exercise market power. We deal with the case of an oligopoly later in Section V. This model is close to the one used in [5]. However, we are fully aware that bidding at marginal cost when the generators are capacity-constrained is not a Nash equilibrium of the game. Indeed, the generators could always be better-off by behaving like a monopoly on the residual demand when they happen to be marginal.

Finally, we want our market price to take only a finite number of values. Yet, with every generator bidding its whole capacity at a fixed cost, the aggregate supply is a stair case function. If the demand is such that it intercepts the supply function on a vertical segment, then the market clearing price could take a continuum of values. Therefore, for the purpose of this analysis we will assume that the market price is the marginal cost of the marginal generator, or the marginal cost of the next generator in the merit order if the intercept of the supply and the demand is on a vertical segment. Importantly, this means that our EPDC is an upper bound on the real EPDC, and the profits obtained for the generators also constitutes upper bounds on the profit obtained in reality. The above assumption is commonly found in the production costing modeling literature [6], but its implications are rarely pointed out.

A. Notations

$i = 1, \dots, N$: generator
$j = 1, \dots, M$: generator type
Ext	: external supply of power
$k = 1, \dots, F$: firm
\mathcal{M}_j	: $\{i \text{ of type } j\}$
$ \mathcal{M}_j = m_j$	
\mathcal{F}_k	: $\{i \text{ belonging to firm } k\}$
$ \mathcal{F}_k = f_k$	
q_j (MWh)	: capacity of type j
c_j (\$/MWh)	: marginal production cost of type j

B. Market Structure

The market include F competing firms which are asymmetric: each firm k has f_k generators and each generator belongs to one of the M generator types with their own characteristics. Very importantly, if a firm cannot supply enough energy because its generators may not be working, the firm can always resort to an external purchase of electricity, with an unlimited

capacity and a high marginal cost: c_{ext} . Alternatively, c_{ext} can be interpreted as the value of loss load (VOLL) when the shortage is covered by load curtailment.

C. Price Competition

We saw that each production unit will offer its energy at marginal cost. We assume $c_1 < \dots < c_M$. The ISO dispatches the cheapest types first until the total energy supplied can meet the demand. We also make the additional simplification about the market clearing price: $\exists j, p(t) = c_j$.

Under our assumptions: $p(t) = c_{J(t)}$ where $J(t)$ is the *marginal generator type* dispatched. Determining the EPDC boils down to determining the distribution of $J(t)$.

We define $\forall j, W_j = \sum_{i \in \mathcal{M}_j} Y_i$. For a given realization of $K(t)$ and of the W_j , $J(t)$ is uniquely determined by

$$J(t) = \min \left\{ h \left| \sum_{j=1}^h q_j W_j - (K(t) - \zeta c_h) > 0 \right. \right\}.$$

Consequently we obtain the following conditional probabilities:

$$\begin{aligned} \forall j, \mathbb{P} \{J(t) = j | K(t), W_1, \dots, W_M\} &= 1 \\ \Leftrightarrow J(t) &= \min \left\{ h \left| \sum_{j=1}^h q_j W_j - (K(t) - \zeta c_h) > 0 \right. \right\} \quad (1) \end{aligned}$$

D. Graphical Model Representation

1) *Graph structure*: The structure of the graphical model is presented in Figure 1.

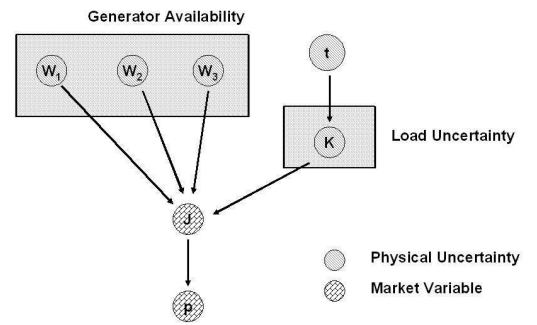


Fig. 1. Graphical model describing the price competition.

2) *Potential functions*: Now that the structure of the graph is defined, we need to endow this graph with potentials functions at each node. Because we focus on medium-term forecasts, we reach the steady state of the Markov process followed by the availability of the generators Y_i . Thus, Y_i is a Bernoulli random variable with $\mathbb{P}(Y_i = 1) = \frac{\mu_j}{\lambda_j + \mu_j}$ if $i \in \mathcal{M}_j$. Consequently, W_j is binomial with parameters m_j and $\frac{\mu_j}{\lambda_j + \mu_j}$.

We assume that the time t is uniformly distributed on $\{1, \dots, 24\}$. The EPDC obtained is therefore an average over a day of the hourly EPDCs, which are straightforward to obtain.

TABLE I
 DESCRIPTION OF THE GENERATOR TYPES

Type	1	2	3	4	5	6	7	8	9
q_j	400	50	20	350	155	76	12	100	197
c_j	0.4	1.0	8.0	14.0	14.0	14.0	14.8	14.8	14.8
$P\{Y_j = 1\}$.92	.99	.9	.92	.96	.98	.98	.96	.95

 TABLE II
 MARKET COMPOSITION

Unit Type	Firm 1	Firm 2	Firm 3	Total
1	0	0	1	1
2	0	2	2	4
3	1	1	1	3
4	2	0	0	2
5	2	0	0	2
6	2	1	1	4
7	2	2	0	4
8	0	1	2	3
9	2	0	2	4
Total	11	7	9	27

The nominal load is uniformly distributed between a_t and b_t in each period t and from that we build the discrete distribution of $K|t$.

Finally, the conditional distribution of $J|K, W_1, \dots, W_M$ comes directly from equation (1).

E. Numerical Example

In order to compare our method with approximate numerical techniques, we use the exact same data as in [5], [6]. The generators are grouped into $M = 9$ different types whose characteristics are summarized in Table I. Table II presents the market composition of our model. There are a total of $N = 27$ generators, which belong to $F = 3$ different firms. The load data are presented in Table III. These data come from the hourly load pattern for the PJM East Region for weekdays of Fall 2003 and have been scaled down to fit the market composition of our power system [6]. We construct a demand space of 16 points². We choose $\zeta = 38.5$. The cost of the external supply of electricity is $c_{ext} = 100\$/MWh$.

Figure 2 presents the EPDC obtained with the PGM method for the price competition model. The stair shape come from the initial assumption that $\exists j, p(t) = c_j$, and from the fact that the PGM method offers exact results. The corresponding marginal probabilities for the marginal generator $J(t)$ are presented on Table IV and have been computed exactly with the junction tree algorithm. One important advantage of this exact inference method is that we can obtain the exact probability of external supply (or load curtailment): $\mathbb{P}\{J(t) = Ext\} = 0.0018$. This probability is usually approximated to 0 by the traditional numerical approximation methods. Our method shows that the clearing price will reach c_{ext} with a non-zero probability. This

²We tried various demand space dimensions and the variations between 10 and 16 demand points are negligible.

 TABLE III
 LOAD PATTERN (MWh)

Hour	a_t	b_t	Hour	a_t	b_t
1	1650	2300	13	2100	2850
2	16000	2200	14	2000	2800
3	1550	2200	15	1950	2800
4	1550	2200	16	1900	2800
5	1550	2250	17	1950	2900
6	1650	2450	18	2050	3100
7	1800	2750	19	2050	3100
8	1900	2850	20	2000	3100
9	2050	2850	21	2000	3050
10	2100	2850	22	2000	2950
11	2150	2850	23	1950	2700
12	2150	2850	24	1750	2450

is important because that is the only time when peaking units can make a profit, which is essential for covering their fixed costs.

Table V displays the expected profit for the different generator types and compares them with the results from [6]. We notice that generator types 7, 8, and 9 have zero expected profit with the method of cumulants, which approximates the small probabilities to 0 by taking only a finite order expansion of the distribution functions. With the PGM method, we observe a positive expected profit for these more expansive peak-load units. The method of cumulants may be misleading, since investing in peak-load units can indeed be profitable. In the context of resource adequacy, this is a valuable insight. The PGM method makes the study of the profitability of peak-load generators possible.




Fig. 2. EPDC from the PGM method with the price competition model.

V. QUANTITY COMPETITION MODELS

Our price competition model presented in Section IV essentially assumes perfect competition and does not consider any strategic behavior. In this section, we model the strategic behaviors of two players producing electricity under uncertain

TABLE IV
PROBABILITY MASS FUNCTION OF $J(t)$ FROM PGM.

j	1	2	3	4	5
$P\{J(t) = j\}$	0	0	0	0.1080	0.2300
j	6	7	8	9	Ext
$P\{J(t) = j\}$	0.2575	0.0719	0.2044	0.1264	0.0018

TABLE V
EXPECTED DAILY PROFIT PER GENERATOR WITH PERFECT PRICE
COMPETITION.

Generator Type	Cumulants	PGM
1	131,669	122,780
2	15,839	16,013
3	2,930	2,798
4	2,034	2,906
5	904	1,654
6	448	852
7	0	44
8	0	346
9	0	605

conditions. Each player owns a unique generating unit. To better model the strategic interactions, we move to a quantity competition framework. We follow a game theoretic approach that was first inspired by [22]. This approach is particularly adapted here since the PGM method requires the strategy space of the players to be discrete.

A. Additional Notations

$z \in \{a, b\}$: player
$-z \in \{a, b\}$: opponent of player z
u (resp. v)	: type of player z (resp. $-z$)

B. Timing and Information Structure

In this model we account for strategic interactions in a game theoretic setting with discrete strategies. Since the dynamic game is more realistic and brings a deeper understanding of the interaction between the two players, we work with a multi-period setting. The information structure is hence a keystone of the analysis. We elaborate two scenarios regarding the information from the previous period available to the players.

In the first scenario (Subsection V-D), each player z learns the realization of the nominal load in period t : K_t and the availability of the generator of her opponent in period t : $Y_{-z,t}$ at the end of period t . The information is delayed, but will prove useful since the transition probabilities from one state to the other are known by the players.

In the second scenario (Subsection V-E), the only information available to the players at the end of each period is the market clearing price, and the realized demand or availability of the opponent's generator are not observed even after the market clears. However, we will see that knowing her own bid and the market price in the previous period is sufficient for each player to learn the residual demand she was actually facing.

In both scenarios, each player z learns at the beginning of the current period t the state of her generator: $Y_{z,t}$. She also uses the information available from period $t-1$ to form beliefs about the state of the system in the current period.

C. One-Stage Quantity Quantity Game

1) *Setting of the game*: Each player z possesses a single generator whose availability is described by the random variable $Y_z \in \{0, 1\}$. We say that player z is of type 0 if $Y_z = 0$ and of type 1 otherwise. Again, each player can resort to an external supply of electricity (which can come from curtailable load). Therefore the uncertainty on the availability of the load translates to an uncertainty of the marginal cost of player z . If she has a generator of type j we define: $c_z^1 = c_j$ and $c_z^0 = c_{ext}$.

A strategy of a player in an extensive form game is a complete plan of actions: It specifies an action in each information set where the player moves. Therefore, a strategy of player z is $x_z = (x_z^0, x_z^1)$ where $x_z^u \in \mathcal{S}_z^u$ is the strategy played by player z if she is of type u . Because we are in a quantity competition model, the strategy space \mathcal{S}_z^u is basically the set of quantities that player z can offer to the market if she is of type u . Particularly, $0 \leq \max \mathcal{S}_z^1 \leq q_j$: The player cannot bid more than the capacity of her unique generator when it is working.

We also need to define the inverse demand from the linear demand function defined above: if player a of type u and player b of type v bid x_a^u and x_b^v respectively, then:

$$p(t) = \frac{K(t)}{\zeta} - \frac{1}{\zeta} (x_a^u + x_b^v).$$

The universe describing the physical state of the market is $\Omega = \mathcal{D} \times \{0, 1\} \times \{0, 1\}$ and we define the joint probability matrix Γ with $\Gamma(k, u, v) := \mathbb{P}\{K = k, Y_a = u, Y_b = v\}$. We also define a set of conditional probability matrices. For example, the conditional probability for a type- u player a to face a type- v player b and a nominal load $K = k$ is:

$$\Theta_a^u(k, v) = \frac{\Gamma(k, u, v)}{\sum_{(k,v) \in \mathcal{D} \times \{0,1\}} \Gamma(k, u, v)}.$$

Eventually, we define a set of conditional payoff matrices. For instance $\pi_z^u(k, v)$ is the matrix describing the payoffs obtained by a type- u player z facing a demand $K = k$ and a type- v player $-z$. It is defined for each strategy $x_z^u \in \mathcal{S}_z^u$ and $x_{-z}^v \in \mathcal{S}_{-z}^v$ by:

$$\pi_z^u(k, v)(x_z^u, x_{-z}^v) = \left(\frac{k}{\zeta} - \frac{1}{\zeta} (x_z^u + x_{-z}^v) - c_z^u \right) x_z^u.$$

2) *Structure and payoffs of the game*: Each player z only knows Y_z at the beginning of the period. Thus, we have a game of imperfect information since the player does not know the residual demand she is facing. We transform the game of imperfect information to a game of incomplete information. Nature chooses first the type of player a , player b and the level of the nominal load K . Then both players choose their strategy x_a and x_b in order to maximize their utility.

The expected payoffs of player z of type u are defined for every strategy $x_z^u \in \mathcal{S}_z^u$ and $x_{-z} \in \mathcal{S}_{-z}$ by:

$$\begin{aligned} G_z^u(x_z^u, x_{-z}) &= \mathbb{E}_{K \times Y_{-z}} [\pi_z^u(k, v)(x_z^u, x_{-z}^v)] \\ &= \sum_{k \in \mathcal{D}, v \in \{0,1\}} \Theta_z^u(k, v) \pi_z^u(k, v)(x_z^u, x_{-z}^v) \end{aligned}$$

3) *Equilibrium concept*: If there are multiple Nash equilibria (NE), or if there are none, it is hard to imagine how the players can “coordinate” to play a given equilibrium point. And with the structure of the game, the existence of a unique pure strategy NE is very hypothetical. We resolve this potential ambiguity by assuming that the players play the *maximin strategy* [23], which always exists. Doing so, the players maximize the minimum expected payoff that they can obtain in the game, given that the opponent can play any possible strategy. Therefore, each player will obtain at least the equilibrium expected payoff in the game.

This maximin equilibrium may be a NE as in the PoolCo example of [22]. The main advantage of this solution concept is that the players do not assume any sort of rationality from their opponent. Moreover, in a multi-period setting, it is very hard to predict the reaction of the opponent because it requires keeping track of the beliefs of the opponent from the first period on. Thus, coordinating on a NE seems illusory.

D. Dynamic Quantity Game with Delayed Information

We keep the same setting as in the one-stage game. However, we introduce an inter-temporal stochastic process. We assume a Markov process [24] with transition matrices M_K, M_a, M_b for K, Y_a, Y_b . We also assume that even if they are not observed at the beginning of the period, K and Y_{-z} are revealed to each player z at the end of each period.

The players will update their beliefs at the beginning of each period after observing last period’s game. Because they know the transition matrix of K and Y_{-z} , and they observed K_{t-1} and $Y_{z,t-1}$, they can form subjective probabilities over K_t and $Y_{z,t}$. These probabilities are then used to compute the expected profit G_z^u from the conditional profits $\pi_z^u(k, v)$ through the conditional probabilities $\Theta_z^u(k, v)$. More precisely, at the beginning of period t when player z is of type u :

$$\Theta_z^u(k, v) = [K_{t-1} M_K e_k] [Y_{-z,t-1} M_{-z} e_v]$$

where e_k, e_v are the canonical basis vector and $(K_{t-1}, Y_{-z,t-1}) \in \{0, 1\}^{|\mathcal{D}|} \times \{0, 1\}^2$ are observed random variables. For instance, we represent the observation of K as a vector $K = [1, 0, \dots, 0]$ if the realization of K is the first element of \mathcal{D} . The “beliefs” will automatically be built from rows of the transition matrices M_K, M_a, M_b .

1) *Graphical Model Representation*: The structure of the graphical model is presented in Figure 3. X_a, X_b are the random variables representing the bids of player a and b .

The unique *observed variable* is the electricity price P . All the other variables are *hidden*. The potentials on K, Y_a, Y_b come from the Markov process defined above. The potential of P is immediately derived from the inverse demand function. Finally, the potentials on X_a, X_b are derived from two

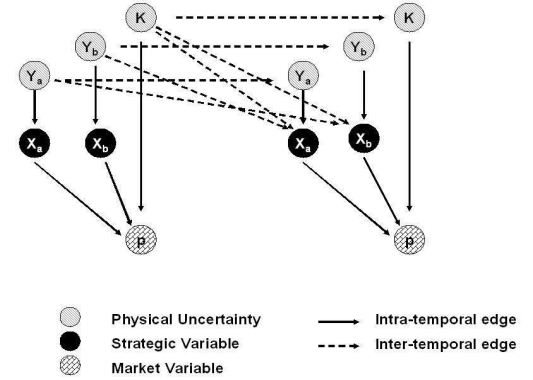


Fig. 3. Graphical model describing the dynamic quantity game of delayed information.

components: the beliefs of a and b which are formed with the variables K, Y_a, Y_b of the last period, and the maximin functions derived from these beliefs and the realization of the type of the generator.

The PGM formalism can perform two different tasks. First, the observed price variable allow to infer the most likely set of priors on all the other hidden variables (“post obs.”), instead of simply imputing some arbitrary priors into the model. This allows starting with a prior on the different variables that reflects the real state of the market, and makes the study of the transient properties of the system (“1 lag”, “5 lag”, etc.) relevant.

Second, we use our model to predict the future prices of electricity. Our dynamic model respects the Markov property [24], so we can transform it to a hidden Markov model (HMM) characterized by a prior joint distribution on all the hidden variables \mathbb{P}_0 and a transition matrix T from any state t to the state $t + 1$. The joint distribution in k periods in the future is then obtained with $\mathbb{P}_k = \mathbb{P}_0 T^k$. And if $T^k \xrightarrow{+\infty} T^\infty$, then the stationary behavior of the system is given by $\mathbb{P}_\infty = \mathbb{P}_0 T^\infty$. Since T is a finite stochastic matrix, it admits 1 for eigenvalue and all the other eigenvalues λ (real or complex) are such that $|\lambda| < 1$. After diagonalizing T , it is straightforward to obtain T^∞ .

2) *Numerical Examples*: The EPDC obtained when two players with different repair rates for their generator face each other in a physical system where the nominal load K is independent of the past is displayed in Figure 4(a). The transition matrices of the Markov process for K is such that $\mathbb{P}\{K_t | K_{t-1}\} = \mathbb{P}\{K_t\} = \frac{1}{|\mathcal{D}|}$. We set $\mu_a = .9$ and $\mu_b = .1$. We also compare the expected profit in that case. Player a repairs her generator more rapidly, and performs better than player b , as illustrated in Figure V-D.2. The system converges rapidly to the long-term behavior (after two periods) since the dependence to the past is very weak.

We then consider a system with *sticky random variables*. Particularly, M_K is closer to the identity matrix and $\mathbb{P}\{K_t = K_{t-1} | K_{t-1}\} = .8$. The two players are similar,

except that player 1 have marginal costs that are more concentrated: We set $c_a = [15, 10]$ and $c_b = [50, 5]$ such that $\mathbb{E}[c_a] = \mathbb{E}[c_b]$ in the long-run. The EPDC obtained in that case is in Figure 4(c), and the profits are compared in Figure 4(d). Player b performs better, but we do not have a strong intuition of why it is the case. Interestingly, the EPDC converges more slowly to the steady state.

E. The Dynamic Quantity Competition with Heuristics

We now relax the assumptions we made in the previous model about the information available at the end of every period to the players. We only made the arguably more realistic assumption that at the end of period t , each player z observes the price p_t and knows her offer $X_{z,t}$. Since $p = \frac{1}{\zeta}(K - X_a - X_b)$, each player can infer $K - X_{-z}$ that was realized in the previous period. From this, each player z can form a set of beliefs on the possible realizations of the residual demand $R_z = K - X_{-z}$ in the current period: $\hat{p}(R_z = r)$.

1) *Belief formation*: Since the residual demand in the last period $R_{z,t-1}$ is observed: $\hat{p}(R_{z,t}) = \hat{p}\{R_{z,t}|R_{z,t-1}\}$. To compute \hat{p} , each player needs $\mathbb{P}\{K_t, X_{-z,t}|K_{t-1}, X_{-z,t-1}\}$. The first simplifying assumption made by both players is that $X_{-z,t}$ is conditionally independent of K_{t-1} given $X_{-z,t-1}$. Thus, the players only need to find $\hat{p}\{X_{-z,t}|X_{-z,t-1}\}$. Because the game is too complex to devise what the opponent's strategy will be, each player resorts to a heuristic.

We distinguish between three kinds of players:

- A *myopic player* who believes that:

$$\hat{p}\{X_{-z,t} = X_{-z,t-1}|X_{-z,t-1}\} = 1.$$

Intuitively, the myopic player conjectures that her opponent always offers the exact same quantity.

- An *agnostic player* who believes that:

$$\hat{p}\{X_{-z,t}|X_{-z,t-1}\} = \hat{p}\{X_{-z,t}\} = \frac{1}{|\mathcal{S}_{-z}|}.$$

Intuitively, the agnostic player does not presume any possible move from her opponent, i.e., anything can happen.

- A *quasi-myopic player* with parameter α who more generally believes:

$$\hat{p}\{X_{-z,t} = X_{-z,t-1}|X_{-z,t-1}\} = \alpha$$

and for $X \in \mathcal{S}_{-z}/X_{-z,t-1}$:

$$\hat{p}\{X_{-z,t} = X|X_{-z,t-1}\} = \frac{1 - \alpha}{|\mathcal{S}_{-z}| - 1}.$$

$\alpha = 0$ correspond to the agnostic player and $\alpha = 1$ correspond to the myopic player.

2) *Expected profit maximization*: Since player z formed beliefs on the residual demand she is going to face in the current period, she can compute and then maximize her expected profit. She simply solves if she is of type u :

$$\max_{x_z^u} \mathbb{E}_{\hat{p}} \left[\left\{ \frac{1}{\zeta}(R - x_z^u) - c_z^u \right\} x_z^u \right].$$

3) *Graphical Model Representation*: The structure of the graphical model is presented in Figure 5. As explained, $X_{z,t}$ depends on $K_{t-1} - X_{-z,t-1}$ and this dependence appears on the graph since $\{K_{t-1}, X_{-z,t-1}\}$ are parents of X_z for both players z .

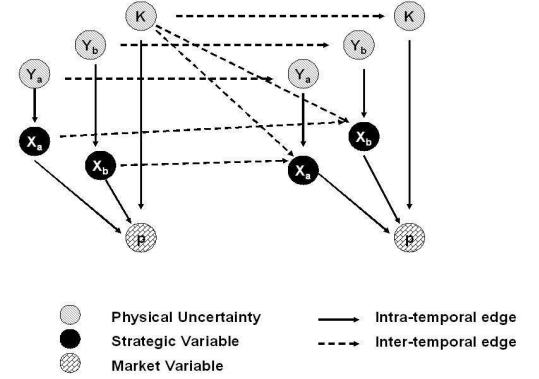


Fig. 5. Graphical model describing the quantity, multi-stage heuristic competition.

4) *Numerical Examples*: We provide here two numerical examples. In the first one, we model a highly variable physical system in the sense that for any random variable $Z \in \{K, Y_a, Y_b\}$: $\mathbb{P}\{Z_t \neq Z_{t-1}|Z_{t-1}\} \gg \mathbb{P}\{Z_t = Z_{t-1}|Z_{t-1}\}$. In the second model, on the other hand, $\mathbb{P}\{Z_t \neq Z_{t-1}|Z_{t-1}\} \ll \mathbb{P}\{Z_t = Z_{t-1}|Z_{t-1}\}$. We try and determine how the different heuristics perform in both cases.

The EPDC obtained when a myopic player faces an agnostic player in a quantity competition setting with a highly variable physical system is displayed in Figure 6(a). As expected, the EPDC converges quickly to the steady state. We then compare the expected profit generated by both heuristics. Since the physical system is highly variable, we expected the myopic bidder to perform worse than the agnostic bidder. This is verified in Figure 6(b). In Figure 6(c), we study the same markets but with a quasi-myopic player with $\alpha = .5$ facing an agnostic player. The reduction in the myopic bias lowers the difference in expected profits between the two players, the agnostic player still performing better.

We now switch to a sticky physical system. We keep all the parameters constant in expected value, but there is less variability in the system. Figure 6(d) presents the EPDC obtained when a myopic player faces an agnostic player. Clearly, the EPDC converges far more slowly to the steady state. Finally, Figure 6(e) and Figure 6(f) present the same comparison as in the previous example, with a myopic and a quasi-myopic player facing an agnostic opponent. Interestingly, the myopic player outperforms the agnostic player in that case, and the quasi-myopic player with $\alpha = .5$ also performs better than the agnostic player in the transient state (less than 30 lags), but seemed to be outperformed in the long-run. The superiority of the myopic heuristic in that case was expected since the

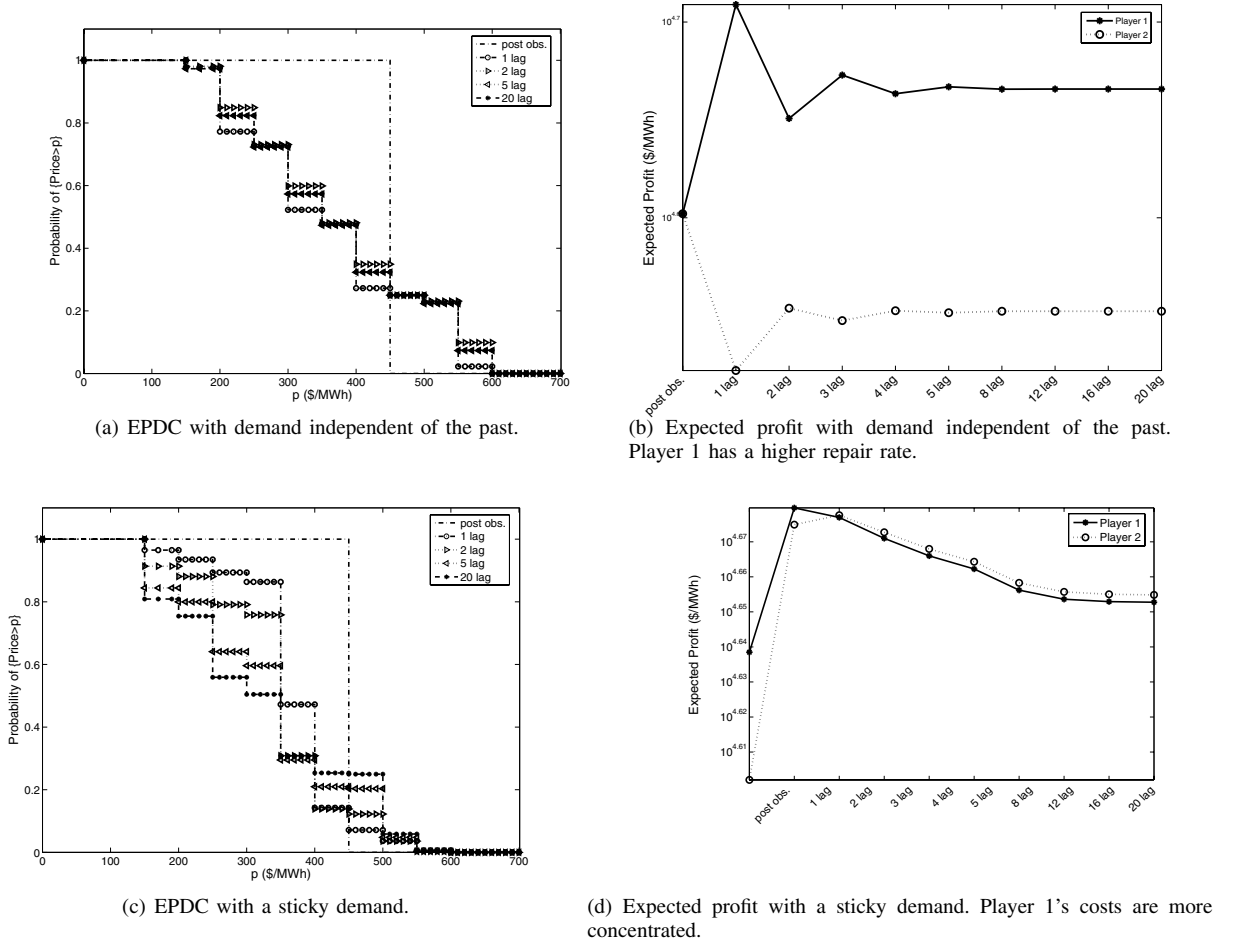


Fig. 4. Quantity game with different Markov processes for K .

system is moving very slowly: Predicting a constant strategy for the opponent seems to make more sense than inferring anything from the past period.

VI. SUMMARY AND DISCUSSION

The contributions of this paper are threefold. First, we show how to use PGMs to study static stochastic competition models. Here, the electricity price is considered as a random variable because of physical uncertainties such as generators' availability and load level. The distribution of the price is determined exactly, and represented by the electricity price duration curve.

Second, we develop a method to study in depth the dynamic behavior of strategic games with uncertainty. We focus on a quantity competition game in an oligopoly market. We present two versions of this game which differ in the amount of information made available to the players. First, we assume that the past physical state of the system is entirely revealed at the end of every period, and we then explore another version which does not require any particular assumption on the knowledge of the players. This latter approach requires behavioral hypotheses regarding the way players form their conjectures on various events where very limited information

is available and rational expectations and subjective probabilities are too complex to compute.

Third, we study heuristics for dynamic decision-making with very limited information, and provide numerical illustrations of various phenomena that could occur. We consider two extreme heuristics: a myopic case where the player assumes that her opponent always plays the same strategy, and an agnostic case where the player assumes that her opponent can play anything with equal probabilities, and we eventually introduce a behavioral parameter to describe the degree of myopia of the player: α .

Two ways to extend this research seem particularly promising. The first one is to provide theoretical foundations for the observations made in simulating the quantity competition games: Can we predict theoretically the behavior of biased players based on the stochastic process driving the physical system, can we bind the differences in expected profit between the heuristics? Second, it would be interesting to fit our model to real-world data, in order to estimate both the parameters of the physical system and the behavioral parameters describing the heuristics used by the players.

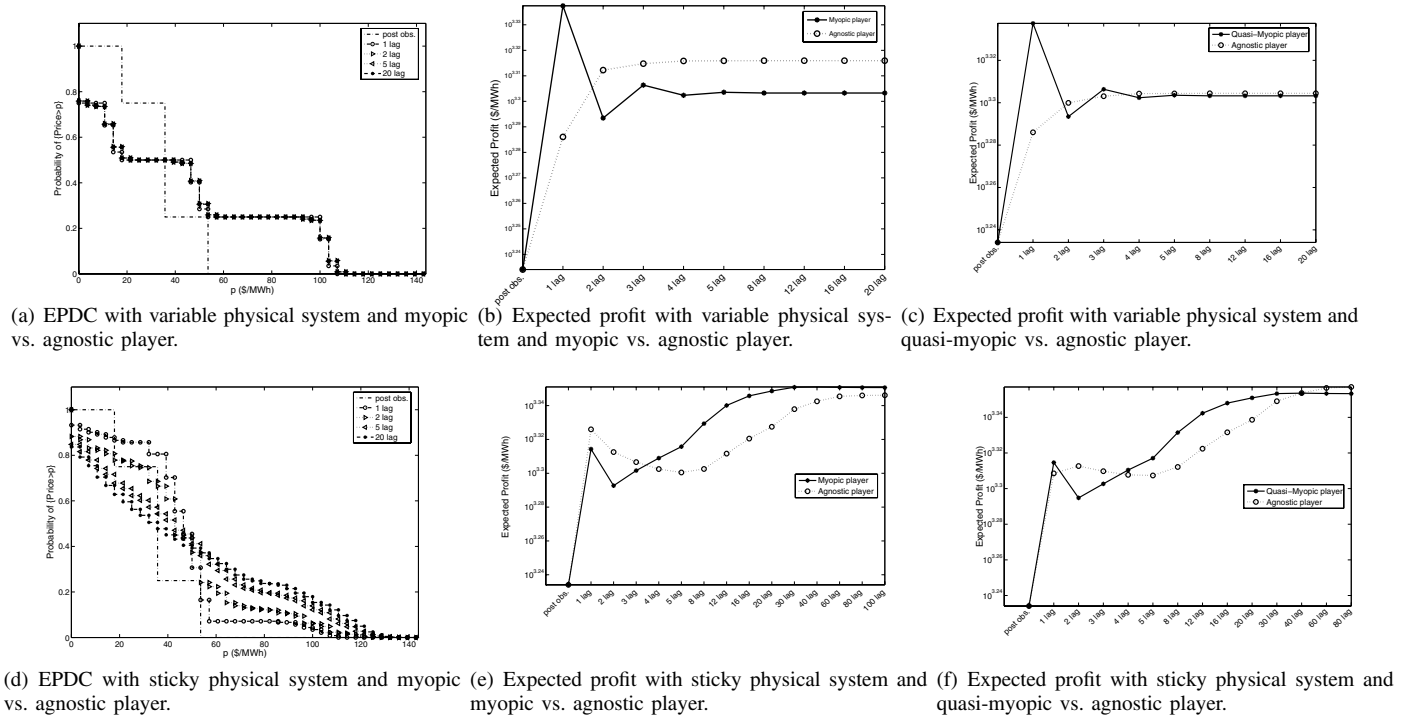


Fig. 6. Quantity competition with different heuristics and various Markov processes for K, Y_a, Y_b .

ACKNOWLEDGMENT

This work was supported by EPRI and EDF.

We used the Bayes Net Toolbox (BNT) for Matlab developed by K. Murphy and available at www.ai.mit.edu/~murphyk/Software/Bayes/bnsoft.html. Murphy [25] provides an overview of the possibilities offered by the BNT.

REFERENCES

- [1] L. de Vries, "Securing the public interest in electricity generation markets, the myths of the invisible hand and the copper plate," Ph.D. dissertation, T.U. Delft, 2004.
- [2] S. Oren, "Generation adequacy via call options obligations: Safe passage to the promised land," *The Electricity Journal*, 2005.
- [3] B. Allaz and J. L. Vila, "Cournot competition, forward markets and efficiency," *Journal of Economic Theory*, vol. 59, pp. 1–16, 1993.
- [4] J. Bloom, "Representing the production cost curve of a power system using the method of moments," *IEEE Transactions on Power Systems*, vol. 7, pp. 1370–1377, Aug. 1992.
- [5] J. Valenzuela and M. Mazumdar, "The electricity price duration curve under Bertrand and Cournot models," in *International Conference on Probabilistic Methods Applied to Power Systems, 2004*, Sept. 2004, pp. 28–43.
- [6] —, "Including generator availability in a stochastic Cournot model for electricity prices," Feb. 2005, unpublished.
- [7] M. Jordan, "An introduction to probabilistic graphical models," June 2003, in Preparation.
- [8] M. Jordan, Ed., *Learning in Graphical Models*, 1999.
- [9] M. Wainwright. (2003) Graphical model, message-passing algorithms and convex optimization. [Online]. Available: http://www.eecs.berkeley.edu/~wainwright/Talks/A_GraphModel_Tutorial.pdf
- [10] K. Murphy, "Dynamic Bayesian networks: Representation, inference and learning," Ph.D. dissertation, U.C. Berkeley, 2002. [Online]. Available: <http://www.cs.ubc.ca/~murphyk/Thesis/thesis.pdf>
- [11] X. Boyen and D. Koller, "Tractable inference for complex stochastic processes," in *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence—UAI 1998*. San Francisco: Morgan Kaufmann, 1998, pp. 33–42.

- [12] P. Klemperer and M. Meyer, "Price competition vs. quantity competition: The role of uncertainty," *The RAND Journal of Economics*, vol. 17, pp. 618–638, Winter 1986.
- [13] J. Bertrand, "Théorie mathématique de la richesse sociale," *Journal des Savants*, pp. 499–508, 1883.
- [14] J. Tirole, *The Theory of Industrial Organization*. Cambridge, MA: The MIT press, 1988.
- [15] A. Cournot, *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. Paris: L. Hachette, 1838.
- [16] K. Judd, "Cournot versus Bertrand: A dynamic resolution," Mimeo, April 1996.
- [17] E. Maskin and J. Tirole, "A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs," *Econometrica*, vol. 56, pp. 549–569, May 1988.
- [18] M. Dudev, "Dynamic Edgeworth-Bertrand competition," *The Quarterly Journal of Economics*, vol. 107, pp. 1461–1477, Nov. 1992.
- [19] R. Green and D. Newbery, "Competition in the british electricity spot market," *Utilities Policy*, vol. 1, pp. 245–254, 1991.
- [20] S. Borenstein and J. Bushnell, "Price competition vs. quantity competition: The role of uncertainty," *The Journal of Industrial Economics*, vol. 47, pp. 285–323, Sept. 1999.
- [21] H. Baleriaux, E. Jamouille, and F. L. de Guertechin, "Simulation de l'exploitation dun parc des machines thermiques de production d'électricité couplé à des stations de pompages," *Revue E*, vol. 5, pp. 235–245, 1967.
- [22] R. Ferrero, J. Rivera, and S. Shahidepour, "Application of games with incomplete information for pricing electricity in deregulated power pools," *IEEE Transactions on Power Systems*, vol. 13, pp. 184–189, Feb. 1998.
- [23] P. Morris, *Introduction to Game Theory*. New York: Springer-Verlag, 1994.
- [24] S. Ross, *Stochastic Processes*, 2nd ed., ser. Wiley Series in Probability and Statistics. John Wiley and Sons, inc., 1996.
- [25] K. Murphy, "The Bayes Net Toolbox for Matlab," *Computing Science and Statistics: Proceedings of the Interface*, vol. 33, 2001.