Estimation and Optimisation in Electricity Pool Markets

Pei-Teh Hu
Department of Engineering Science
University of Auckland
New Zealand
p.hu@auckland.ac.nz

Abstract

In electricity pool markets, generators submit offers of quantity-price pairs to an independent system operator, who dispatches power to meet loads at the lowest cost using these offers. The market-clearing price at any location is the marginal cost of supplying electricity to this location. By withholding capacity or offering at a high price, generators can influence the market-clearing price, which is also affected by their competitors' behaviour. We represent the uncertainty in the competitor offers and the demand by a single probability distribution called the market distribution function. In this paper, we explore a methodology to estimate the market distribution function and construct an offer stack that gives the maximum expected profit with respect to this distribution. A case study based on historical data from Mighty River Power will be presented.

1. Introduction

In electricity pool markets, generators submit offers of quantity-price pairs to an independent system operator (ISO) for each trading period. These quantity-price pairs are known as offer tranches and together they form a non-decreasing piecewise function called an offer stack. Figure 1 is a plot of an example offer stack, which contains five tranches.

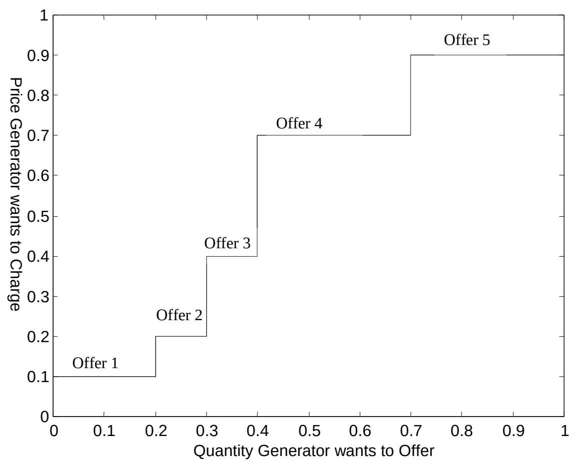

Figure 1. An example offer stack

The ISO dispatches power to meet loads at the lowest cost based on these offers. For any given trading period, an offer tranche of a generator can be either fully dispatched or partially dispatched. We say a tranche is *fully dispatched* if all of its quantity is cleared by the ISO in a trading period, and *partially dispatched* if some fraction of its quantity is cleared. The market-clearing price at any location is the marginal cost of supplying electricity to this location. That is, a generator which is partially dispatched is also the marginal generator in this trading period.

The market-clearing prices vary because of the uncertainty of the demand and generators' behaviour. Generators can make the market-clearing price go up by withholding capacity, or depress the clearing price by offering at low prices. It is possible to calculate optimal offer stacks that maximize the profit of a generator. In this paper, we address this problem for a single period.

2. Offer stack optimization conditions

Anderson and Philpott [1] observed that the uncertainty in the competitors' offers and the demand can together be represented by a single probability distribution, called the *market distribution function* $\psi(q,p)$. $\psi(q,p)$ is the probability of *not* being fully dispatched by the market if we offer generation q at price p. The market distribution function is monotonic in both its arguments q and p. If we keep the offering price p the same but increase (decrease) the offering quantity q, we will expect the probability of not being fully dispatched to increase (decrease). If we increase (decrease) the offering price p but keep the offering quantity q the same, we will expect the probability of not being fully dispatched to increase (decrease) as well. These can be expressed as follows:

$$\psi(q_1,p) \leq \psi(q_2,p), \text{ for } q_1 \leq q_2 , \tag{1}$$

$$\psi(q,p_1) \leq \psi(q,p_2), \text{ for } p_1 \leq p_2 . \tag{2}$$

By assuming that the probability of not being fully dispatched by the ISO is 0 if a generator offers quantity 0 at price 0, we obtain the first boundary condition:

$$\psi(0,0) = 0. \tag{3}$$

By assuming that the probability of not being fully dispatched by the ISO if we offer quantity q_M (the maximum quantity a generator can offer) at price p_M (the maximum price a generator will charge) is 1, we obtain the second boundary condition:

$$\psi(q_M,p_M) = 1. \tag{4}$$

Let $R(q,p)$ be the generator's profit if it offers quantity q at price p at any given trading period. Usually it can be expressed by

$$R(q, p) = qp - cq - Qp + fQ , \tag{5}$$

where c is the cost of one unit of electricity generation, Q is the contract position, and f is the price for one unit of contract generation. In practice, we do not include the term fQ in the optimisation problem since it is a constant.

Anderson and Philpott discuss generator offer curves in their paper [1] in detail. Using the expression in (5), the expected return $E(s)$ is the line integral

$$E(s) = \int_s R(q, p) d\psi(q, p), \quad (6)$$

if a generator offers in a curve s given that the market distribution function $\psi(q,p)$ is continuously differentiable everywhere in

$$\Psi = \{(q, p) \mid 0 < \psi(q, p) < 1\},$$

with partial derivatives ψ_q and ψ_p. Write s in parametric form so that

$$s = \{(x(t), y(t)) \mid 0 < t < T\}. \quad (7)$$

Substituting (7) in (6), we obtain $E(s)$ with the parameter t

$$E(s) = \int_0^T \{R(x(t), y(t))[\psi_q(x(t), y(t))x'(t) + \psi_p(x(t), y(t))y'(t)]\} dt. \quad (8)$$

By (1), (2), and (8), the problem maximizing $E(s)$ over curve s can be written as

(P1) Maximize $\int_0^T \{R(x(t), y(t))[\psi_q(x(t), y(t))x'(t) + \psi_p(x(t), y(t))y'(t)]\} dt$

Subject to
$$0 \le x(t) \le q_M$$
$$0 \le y(t) \le p_M$$
$$x'(t) \ge 0$$
$$y'(t) \ge 0$$

The problem can be seen as an autonomous nonlinear optimal control problem in which controls are $x'(t)$ and $y'(t)$, and any admissible trajectory must satisfy the state constraints. Anderson and Philpott [1] demonstrate that the curve s^* is a local optimal solution to (P1) if it satisfies the following condition

$$Z(q, p) = R_q\psi_p - R_p\psi_q = 0, \text{ for } (q, p) \in s^*. \quad (9)$$

3. Estimation of the market distribution function

There are several approaches that have been developed to estimate the market distribution function. Pritchard et al. [2] describes a maximum-likelihood estimation mechanism. In this paper, we use a Bayesian approach which is also described in [2].

In the Bayesian approach, we have experimented with a number of models. An approach that appears to give realistic results assumes that if a generator offers q at price 0, then the market-clearing price $P(q)$ has a log normal distribution so

$$\log(P(q)) \sim Normal(\beta - \alpha q, \sigma^2), \quad (10)$$

where (α, β) is distributed over some set A and σ is a constant.

Therefore, $P(q)$ has a density function f and a distribution function F given by

$$f(P(q)) = (2\pi\sigma^2)^{-\frac{1}{2}} e^{-\frac{(\log P(q)-(\beta+\alpha q))^2}{2\sigma^2}},$$
$$F(P(q)) = \int_{-\infty}^{\log p} e^{-\frac{(z-(\beta+\alpha q))^2}{2\sigma^2}} dz.$$

We define a family of market distribution function $\{\psi^{\alpha,\beta}, (\alpha,\beta) \in A\}$, where

$\psi^{\alpha,\beta}(q,p)$

= The probability that a generator is not being fully dispatched if it offers q at p

= The probability that the market-clearing price $P(q)$ is less than the offer price p

$= (2\pi\sigma^2)^{-\frac{1}{2}} \int_{-\infty}^{\log p} e^{-\frac{(z-(\beta+\alpha q))^2}{2\sigma^2}} dz.$

Then

$$\psi_p^{\alpha,\beta} = (2\pi\sigma^2)^{-\frac{1}{2}} (1/p) e^{-\frac{(\log(p)-(\beta+\alpha q))^2}{2\sigma^2}}, \quad (11)$$
$$\psi_q^{\alpha,\beta} = (2\pi\sigma^2)^{-\frac{1}{2}} \alpha e^{-\frac{(\log(p)-(\beta+\alpha q))^2}{2\sigma^2}}. \quad (12)$$

Suppose that a generator submits an offer stack containing n_i tranches at period i, for $i = 1, 2, 3, \dots, m$. For each period i, we denote Q_{ij} as the last unit the generator is going to sell at price P_{ij} on the j^{th} tranche with $P_{i0} = Q_{i0} = 0$, for $j = 1, 2, \dots, n_i$. Let (q_i, p_i) be the observation at period i where q_i is the quantity the generator get dispatched at period i and p_i is the market-clearing price at period i, for $i = 1, 2, 3, \dots, m$. We can separate observations into two subsets: a partially dispatched set $\{(q_i, p_i) \mid i \in D\}$ and a fully dispatched set $\{(q_i, p_i) \mid i \in E\}$ where

$$D = \{i \mid p_i = P_{ij}, Q_{i(j-1)} < q_i < Q_{ij}, i = 1,\dots,m, j = 1,\dots,n_i\}, \quad (13)$$
$$E = \{i \mid P_{i(j-1)} < p_i < P_{ij}, q_i = Q_{ij}, i = 1,\dots,m, j = 1,\dots,n_i\}.$$

Given the prior density function $\phi_{i-1}(\alpha,\beta)$, we have the following relationships as shown in [3]:

$$\phi_i(\alpha,\beta) \propto \psi_q^{\alpha,\beta}(q_i, p_i)\phi_{i-1}(\alpha,\beta), \text{if } i \in D,$$
$$\phi_i(\alpha,\beta) \propto \psi_p^{\alpha,\beta}(q_i, p_i)\phi_{i-1}(\alpha,\beta), \text{if } i \in E. \quad (14)$$

This updating scheme assumes that all of the observations are equally weighted. Thus the order of observations with respect to time is not important. With the updated density function $\phi_i(\alpha,\beta)$ defined on $(\alpha,\beta) \in \mathrm{A}$, we obtain the market distribution function $\psi_i(\alpha,\beta)$ for period i to be

$$\psi_i(q,p) = (2\pi\sigma^2)^{-\frac{1}{2}} \int_A \int_{-\infty}^{\log p} e^{-\frac{(z-(\beta+\alpha q))^2}{2\sigma^2}} \phi_i(\alpha,\beta) dz d\alpha d\beta. \quad (15)$$

Once $\psi_i(q, p)$ is obtained, the optimal curve s^* for the expected return (6) can be found by condition (9).

In practice, the calculation of ψ_i is time consuming. Since $Z_i(q,p) = 0$ can be seen as the shape of optimal curve and it is what we are interested in, we can find Z_i directly from (11) and (12) while calculating ϕ_i.

4. Case study – Mighty River Power

Mighty River Power (MRP) is the second largest hydro generator in New Zealand (largest in North Island). It controls the water gate at Lake Taupo and owns 8 stations along Waikato River: Aratitia, Ohakuri, Atiamuri, Whakamaru, Maraetai, Waipapa, Arapuni, and Karapiro. Since these 8 stations are sitting at 8 different locations (or nodes), MRP submits a 5-tranche offer stack for each of them to the ISO at each half-hour trading period. At the end of the trading period, the ISO decides the market-clearing price for each node. In this case study, MRP provided dispatched quantities, market clearing prices, and 5-tranche offer stacks for each station at each trading period from December 2000 to June 2001.

A key feature of MRP's river system is that MRP has the right to arrange the amount of generation for each station to minimize the cost of delivering the power. This is as long as the total generation meets the sum of dispatched quantity of each station by the ISO for that period. For this reason, we can address the problem (P1) at a company level. Although 8 stations submit offer stacks separately, their tranches are at the same price level. Thus, we can put them together to form a company offer stack for each period. Also, we can have a company dispatch by adding up the dispatches from 8 stations and obtain a company market-clearing price by taking the average of market-clearing prices of 8 nodes for each trading period by assuming the differences of market-clearing prices among 8 nodes are small. Hence, the observation (q_i, p_i) at a company level sitting on the corresponding company offer stack is obtained.

4.1 Implementation of the original scheme

Here, we estimate the market distribution for weekdays peak-hour (from 8:30 a.m. to 8 p.m.) observations in June 2001 by assuming the market behaviour is similar during these periods. Assume we have a prior density function $\phi_0(\alpha, \beta)$ uniformly distributed on [0, 0.01] x [3,8] with a contract position of 400 MW, a generation cost of \$20/MW-hr, and standard deviation $\sigma = 4$. By (10) - (14), we set the normalized updating scheme as

$$\phi_i = f_i \phi_{i-1} = (\prod_{j=1}^{i} f_j)\phi_0, \text{ for } i = 1, 2, ..., m, \tag{16}$$

where

$$f_i = \frac{\psi_q^{\alpha,\beta}(q_i, p_i)}{\int_3^8 \int_0^{0.01} \psi_q^{\alpha,\beta}(q_i, p_i)\phi_{i-1}(\alpha,\beta)\,d\alpha d\beta}, \text{if } i \in D, \tag{17}$$

$$f_i = \frac{\psi_p^{\alpha,\beta}(q_i, p_i)}{\int_3^8 \int_0^{0.01} \psi_p^{\alpha,\beta}(q_i, p_i)\phi_{i-1}(\alpha,\beta)\,d\alpha d\beta}, \text{if } i \in E.$$

After the first update we obtain a new estimated density function $\phi_1(\alpha, \beta)$ (shown in Figure 2). We can obtain the estimated market distribution function ψ_1 (Figure 3) by (15), and the optimal offering curve ($Z_1 = 0$) for next trading period (Figure 4) by (9).

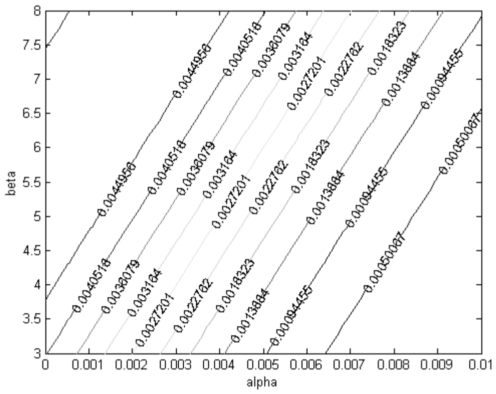

Figure 2. $\phi_1(\alpha, \beta)$

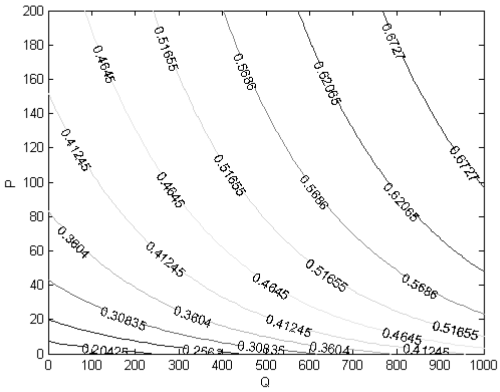

Figure 3. $\psi_1(p,q)$

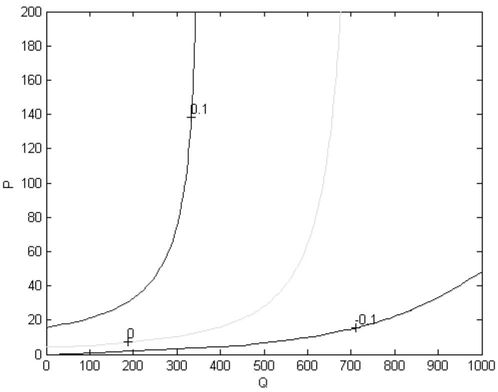

Figure 4. $Z_1 = 0$

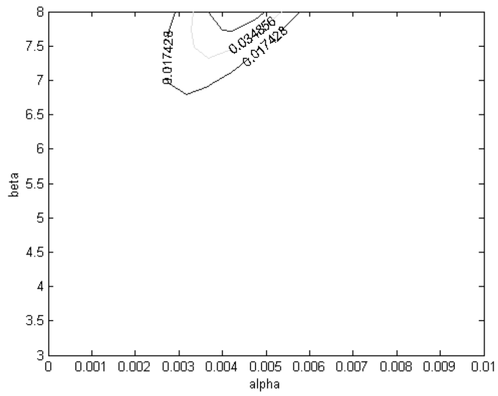

Figure 5. $\phi_{30}(\alpha, \beta)$

However, the density function ϕ converges to zero quickly if more observations are added in. This is because the range of small ϕ values increases along with the updating process. It is impossible to select the right range for α and β to prevent this happening. Figure 5 shows that the updated ϕ has only a small non-zero region left after 30 updates. In this case, ϕ converges to zero after 100 updates. To avoid this phenomenon, we introduce a smoothing factor k.

4.2 The smoothing factor *k*

From (16), we can see every updating factor f_i is equally weighted during the process. That is, each observation is equally contributing to the estimated ϕ. As $i \to \infty$, the estimated density function ϕ should converge to a static function, that is, a real density function that represents true market behaviour. However, one may claim that is not an appropriate representation in practice. Since the market behaviour is changing with time, we shall select our offer strategy using the latest information. To incorporate this idea, we put more weight on recent observations and eliminate the effects from old observations as time passes. Here, we introduce the smoothing factor $k \in (0,1]$ to modify the updating scheme (16)

$$\phi_i = (f_i \phi_{i-1})^k, \text{ for } i = 1,2,... \tag{18}$$

This has the same effect in an updating scheme as repeating the most recent observation $(\frac{1}{k})$ times in comparison to the penultimate observation. For example, if k = 0.5, then we count the current observation twice as many times as the previous one. Also, k de-emphasizes the effect on the posterior density function caused by the observation and therefore keeps the element of updated density function away from zero. When k has a value close to zero, a large weight is put on recent observations, which is appropriate for circumstance in while the offers of competitors are known to be changing frequently over a day.

By using the same assumption as in section 4.1 but including the updating formula (18) with the smoothing factor k = 0.8, we obtain the 1st updated density function ϕ_1 (Figure 6), and the 30th updated density function ϕ_{30} (Figure 7). As we can see, ϕ_1 in Figure 6 displays a similar shape to ϕ_1 in Figure 2.

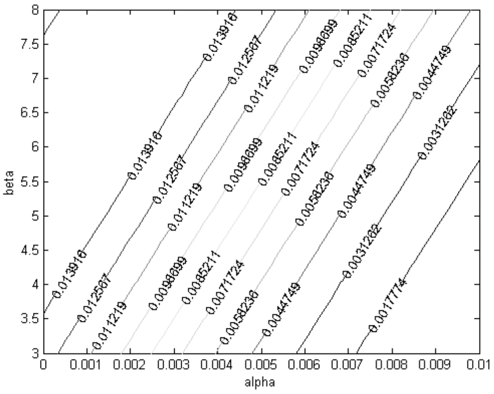

Figure 6. ϕ_1 with $k = 0.8$

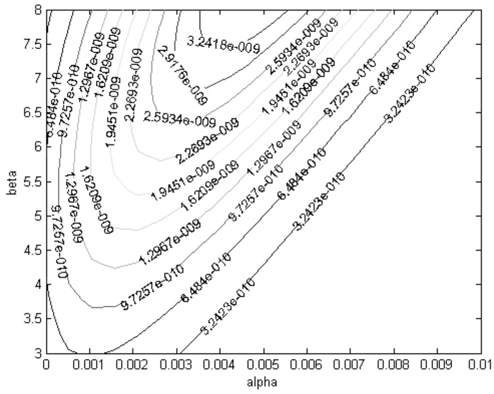

Figure 7. ϕ_{30} with $k = 0.8$

4.3 Fully Dispatched and partially dispatched observations

As the updating procedure continues, we observe the market distribution changes significantly while the nature of the updating observation is altered from a full dispatch to a partial dispatch. For example, the fact that the 19th observation is partially dispatched indicates that the previous offer can by made more aggressive for the current market. Thus, one should either withhold the capacity or increase the offering price in order to increase returns, that is, achieving a full dispatch for the next trading period. As we can see in Figure 8 and 9, our estimated market distribution has detected the change. Notice that ψ_{19} shifts towards the left compared with ψ_{18}. Moreover, the optimal offer curve Z_{19} = 0 (Figure 11) shifts towards the left compared with Z_{18} = 0 (Figure 10) as well, and hence the optimal offer becomes more aggressive.

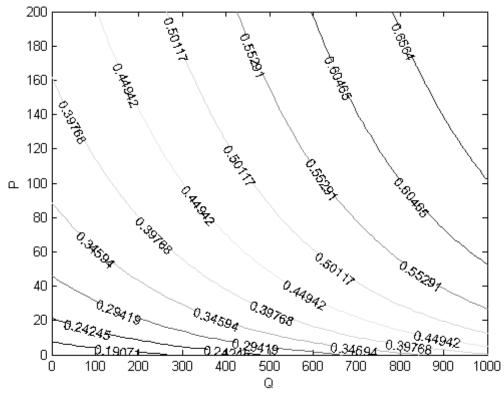

Figure 8. ψ_{18} with $k = 0.8$

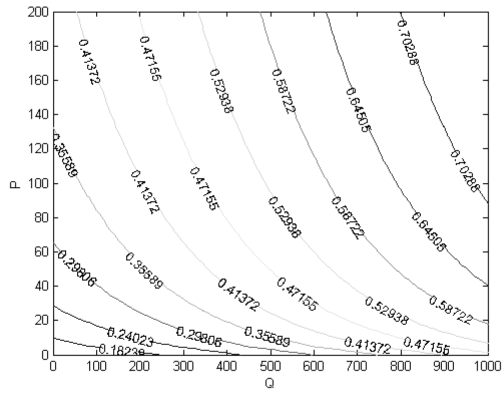

Figure 9. ψ_{19} with $k = 0.8$

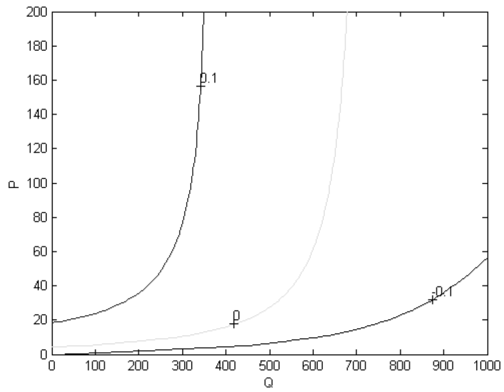

Figure 10. $Z_{18} = 0$ with $k = 0.8$

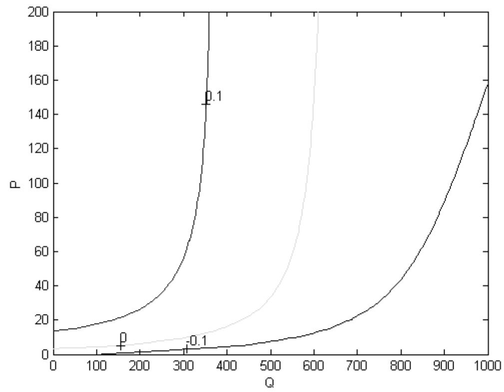

Figure 11. $Z_{19} = 0$ with $k = 0.8$

5. Conclusions

In this paper we have demonstrated a methodology to estimate the market distribution function by a Bayesian approach. In the case study, we see the drawback of putting equal weight on each observation while estimating the density function ϕ. A smoothing factor k is introduced to solve this numerical difficulty and hence assign appropriate weights according to the order of the observations with respect to time. However, this is a very crude approximation to model the dynamics of the rest of market. A more adequate method may be needed but this will require further study.

Another unfavorable approximation shown in this paper is that the family of ψ is simply constructed by assumption (10). To obtain a more decent ψ family, we can model other participants explicitly in the dispatch model.

A further issue that requires attention is that an unbiased estimation of the market distribution function requires a spread of observations. However, observations which have small dispatched quantities are very expensive to obtain in practice. In order to obtain such an observation, a generator needs to decrease their offer quantity, and hence this results in a decrease in revenue. This will not usually be supported by a company.

In further research, we will adapt this methodology following the approach of Neame et al.[4] to a hydro-generator with a single reservoir. For a hydro-generator, the amount of water kept in its reservoir can be considered as the amount of power it can generate, or an asset of the generator. The basic strategy for such a generator is to generate as much electricity as possible when the market-clearing price is high but hold capacity when the market-clearing price is low. This results in a more complicated multi-period stochastic optimization problem.

Acknowledgement

I am grateful to Assoc. Prof. A.B. Philpott on whose work this paper is largely built. This research is funded by Mighty River Power.

Reference

[1] E.J. Anderson and A.B. Philpott, “Optimal offer construction in electricity markets”, forthcoming in *Math of OR*, 2001.

[2] G. Pritchard, A.B. Philpott, P. Neame, G. Zakeri, “The electricity seller’s psi”, proceeding of the *ORSNZ*, 2001.

[3] E.J. Anderson, and A.B. Philpott, “Estimation of electricity market distribution functions”, submitted to *Annals of OR*, 2001.

[4] P.J. Neame, A.B. Philpott, and G. Pritchard, “Long-term hydro scheduling in a pool market”, proceeding of the *ORSNZ*, 2001