Draft version of Chapter 6, "The Next Generation of Electric Power Unit Commitment Models", Intl. Series in Operations Research and Management Science, vol. 36, Kluwer Academic Publishers, Boston, April 2001 (editors: B. F. Hobbs et al)

CHAPTER

OPTIMAL SELF-COMMITMENT UNDER UNCERTAIN ENERGY AND RESERVE PRICES

R. Rajaraman (rajesh@lrca.com), L. Kirsch (laurence@lrca.com)
Laurits R. Christensen Associates

F. L. Alvarado (alvarado@engr.wisc.edu)
The University of Wisconsin

C. Clark (cclark@epri.com)
Electric Power Research Institute

Abstract This paper describes and solves the problem of finding the optimal self-commitment policy in the presence of exogenous price uncertainty, inter-product substitution options (energy versus reserves sales), different markets (real-time versus day-ahead), while taking into consideration inter-temporal effects. The generator models consider minimum and maximum output levels for energy and different kinds of reserves, ramping rate limits, minimum up and down times, incremental energy costs and start-up and shutdown costs. Finding the optimal market-responsive generator commitment and dispatch policy in response to exogenous uncertain prices for energy and reserves is analogous to exercising a sequence of financial *options*. The method can be used to develop bids for energy and reserve services in competitive power markets. The method can also be used for determining the optimal policy of physically allocating generating and reserve output among different markets (e.g., hour-ahead versus day-ahead).

1. INTRODUCTION

Unit commitment refers to the problem of deciding when to start and when to shutdown generators in anticipation of changing demand [1]. In traditional utility systems, the problem of unit commitment was formulated and solved as a multi-period optimization problem. In the traditional problem formulation, the anticipated demand was an input variable. The problem was solved for multiple generators, generally owned by the same entity (a utility). The start-up, shutdown and operating costs of the generators were assumed known. The standard way to analyze and solve this problem was by dynamic programming, and within this category of problems, the most popular solution method in recent years has been the use of Lagrangian relaxation [2,3]. Recently, a new method of decommitment has also been proposed [4,5].

In a deregulated market this changes. Generators generally have to optimally self-commit their units. Since in most power pools, no single merchant owns all the generating assets, the need to meet system load is replaced with the need to optimize profits of the merchant's generating plants based on the uncertain market prices at the location(s) where the generator(s) is(are) located. Of course, the forecasts of markets prices depend upon a number of factors, the most important of which include demand, system-wide generation availability and cost characteristics, and transmission constraints.

We pose the problem of finding the profit-maximizing commitment policy of a generating plant that has elected to self-commit in response to *exogenous but uncertain* energy and reserve price forecasts. Typically, one generator's output does not physically constrain the output of a different generator[1], so this policy can be applied to each generator in the merchant's portfolio separately and independently. Therefore, for ease of exposition, we assume the case of a single generator. Generator characteristics such as start-up and shutdown costs, minimum and maximum up and down times, ramping rates, etc., of this generator are assumed known. The variation of prices for energy and reserves in future time frames is known only statistically. In particular, the prices follow a *stochastic rather than deterministic* process. We model the process using a Markov chain. The method is applicable to multiple markets (e.g., day-ahead, hour-ahead) and multiple products (energy, reserves).

Other researchers have modeled the effect of energy price uncertainty on generator valuation. In [6], the author models the effect of the spark spread on short-term generator valuation. In [7], the authors propose mean reverting price processes and use financial options theory to value a generating plant. However, in both [6] and [7], the authors neglect the effect of realistic operating constraints such as minimum start-up and shutdown times. In [8], the authors improve upon this work to more realistically include the effect of operating constraints to find the short-term value of a generating asset; however, they neglect the effect of ramping rates. All these papers make assumptions about the risk-neutral price process in order to value the power plant.

---

[1] Exceptions include multiple hydroelectric units connected in series, and restrictions on aggregate emission levels from multiple generators within an area.

This paper focuses not on generator valuation, but on finding the general principles for generator self-commitment in the presence of exogenous price uncertainty and market multiplicity. The basic mathematical principles are derived from dynamic programming theory [10]. We consider energy and reserve markets, although the method can be extended to include additional market choices, such as day-ahead versus real-time markets. The paper is organized as follows. Section 2 defines the problem. Section 3 gives the dynamic programming solution to the problem. Section 4, 5 illustrate the features of the optimal commitment policy using simple illustrative examples. Section 6 gives a more detailed numerical result for a peaking generator. Section 7 concludes the paper. Appendix A is a technical section that solves the single-period optimal generator dispatch problem given exogenous prices of energy and reserves.

2. THE PROBLEM

We begin by describing the exogenous inputs to the problem.

Generator capability and cost characteristics. At any given time t, generator G is assumed to be in *state* x_t, where x_t is a member of a discrete set X={state 1, state 2, ..., state K} of possible states. *Inter-temporal constraints* are represented by *state transition rules* that specify the possible states that the generator can move to in time period t+1, given that the generator is in a state x_t at time t. Generally, there is a cost associated in moving between different states. In a simple representation, two states are sufficient: "in service" (or "up") and "out of service" (or "down"). However, in general many more states may be needed to represent the various conditions of the start-up and shutdown process.

The degree to which a generator can participate in providing reserves depends on its ability to respond to the reserve needs in a timely manner. For regulating and spinning reserves, the generator must be already be in service; the amount of MW of reserves that a generator can offer must be consistent with its ramping rate. Generators that are already at a maximum in terms of energy provision are unable to also participate in the reserves market. Thus, to participate in the reserves markets, the generator cannot simultaneously sell all of its capability in the energy market.

The parameters that describe a generator include:

- Minimum and maximum output levels
- Ramping rates
- Minimum up and down times for the generator
- Incremental energy costs and no-load costs
- Start-up, shutdown, and banking costs.

Generator states and state transitions. Generators can be in any of a number of several possible UP, DOWN or transitional states. For example, for a generator with total capacity of 200 MW, and ramp rate of 100 MW/hour, we could define 2 UP states, UP_1, and UP_2. The state UP_1 would cover the operating range [0 MW, 100 MW], while the state UP_2 would cover the operating range [101 MW, 200 MW]. Likewise, minimum down times can be enforced by defining multiple DOWN states. Only certain transitions among these states are permissible. Furthermore, transition between states generally involves a cost. For example, going from a cold shutdown to an online state will involve a start-up cost.

Generator dispatch constraints. A generator may have additional dispatch constraints that restrict its operation. For example, the generator may need to be offline (in the "off" state) during certain periods for scheduled repairs. These restrictions are modelled as time dependent constraints on generator states.

Exogenous price forecast of energy. A discrete Markov process models the exogenous price p_t for energy. In each time period, a discrete price state represents a price range. The price at time $t+\Delta t$ is probabilistically related to the price at time t via a price transition matrix. That is, $prob(a_1 \leq p_{t+\Delta t} \leq b_1 \mid a \leq p_t \leq b)$ is a known quantity. One can think of a price forecast at any time t to be a baseline price point plus a random uncertainty around the baseline forecast

Since the exogenous price forecast is an important input of the problem, we digress a little to discuss how one may obtain estimates of this input. We consider two ways in which price forecasts can be made:

- One possible method is to use historical data. For example, to obtain a price forecast for next week, one could use past week data, and data from other weeks with similar load/weather patterns as that predicted for next week. This would be a statistical data-mining problem.
- Another way to forecast prices would be to use numerous Monte Carlo iterations of structural computer models (such as optimal power flow models and production cost models) to model the uncertainty of prices.

The physical spot markets for energy include real-time, hour-ahead, day-ahead, and possibly week-ahead markets[2]. Each of these spot markets is a different market, e.g., energy prices for a particular hour could be different in day-ahead and real-time markets and could have very different characteristics in terms of price volatility. A generator will often have a choice as to which market to use to sell its energy.

Exogenous price forecast of reserves. Operating reserves are distinguished by the speed with which they become available and the length

[2] It is doubtful whether forward prices quoted month-ahead (or more) will influence an individual generator's commitment and dispatch decisions.

of time that they remain available. In the nomenclature of the Federal Energy Regulatory Commission, the primary types of reserves (from fastest to slowest) are regulating, spinning, supplemental, and backup reserves. For example, regulating reserves need to be available for following moment to moment fluctuations in system demand, and can generally be offered by generators on Automatic Generation Control (AGC). As another example, to offer 10 MW of spinning reserves, a generator must be online and must be capable of producing 10 MW within 10 minutes. Supplemental reserves and backup reserves are slower reserve types. For generators who provide reserve services, there are two types of reserve costs (see [11] for more details). These are *reserve availability costs*, which are the costs of making reserves available even if they are not actually used, and *reserve use costs*, which are the costs incurred when the reserves are actually used. Generally, reserve use costs are compensated at the spot price of energy[3]. Reserve availability costs are the opportunity costs of generators, i.e., they include off-economic dispatch costs, and costs of starting up or shutting down generators. In California, New York, New England, and the Pennsylvania-Jersey-Maryland (PJM) system, there is currently a competitive market clearing process for setting the reserve availability costs of some or all of the above reserve types. Reserve availability prices are modeled similarly to energy prices, i.e., as a discrete Markov chain. However, we allow the reserve availability prices to be correlated with the energy prices. One particularly simple way to model reserve uncertainty is to assume perfect correlation between energy prices and reserve availability prices.

Exogenous fuel price forecasts. Fuel prices can be modeled similarly to reserve availability prices. However, since most practical commitment periods are less than a week and fuel prices typically show much less volatility than energy prices over this interval, it is a not a bad approximation to keep these inputs constant. Therefore, for ease of exposition, we assume fixed fuel prices throughout the paper. However, it is fairly straightforward to also include uncertainty in fuel prices; see, for example, [8].

Start and end time periods. We will assume that the start time period is at time 1, and that the end time period is at time T. For most practical problems, T will be between 1 day and 1 week.

Next we define some notation; explicit functional dependencies are often omitted for clarity. Given generator state x_t and a vector of energy and reserve availability prices, p_t, at time t,

---

[3] In general, reserve use costs could also include the wear-and-tear costs of ramping up and down to follow system demand. However the current custom is that, in competitive markets for reserve services, these costs are not directly compensated, and must somehow be internalized by the generators.

1. $u_t(x_t,p_t)$ (or simply u_t) defines the *commitment policy* for time t, i.e., it signifies a particular valid rule to move the generator from state x_t at time t to a new state $u_t(x_t,p_t)$ at time t+1.
2. $d_t(x_t,p_t)$ (or simply d_t) defines the *dispatch policy* for time t, i.e., it represents a particular dispatch of energy and reserves for the generator at time t, given that the generator is in state x_t.
3. $\pi_t(x_t, p_t)$ denotes *single-period* profits at time t[4], i.e., the profits realized by the dispatch $d_t(x_t,p_t)$.
4. $c_t(x_t,u_t)$ denotes the cost of transitioning (e.g., start-up or shutdown costs) from state x_t at time t to state $u_t(x_t,p_t)$ at time t+1 due to the commitment policy u_t used in time period t.

The problem can now be posed as:

[PROBLEM] Find the best commitment and dispatch policy (**u*, d***) that maximizes *expected total profits* $E(\Sigma(\pi_t(x_t,p_t)- c_t(x_t,u_t))$ over all possible commitment policies $\mathbf{u}=(u_1,u_2,u_3,....,u_T)$ and all possible dispatch policies $\mathbf{d}=(d_1,d_2,d_3,....,d_T)$, where E denotes the expected value over the uncertain price forecasts[5].

Before we proceed further, it is useful to summarize the essential features of the problem:

- **Inter-product substitutability**: market participants have a choice between sales in the energy versus sales in the reserve markets. These markets operate simultaneously (though the markets for energy and reserves may *clear* sequentially, as they currently do in California). Moreover, market participants have a choice of offering their products in different markets (e.g., day-ahead versus hour-ahead markets).
- **Price uncertainty**. The future prices of energy and of reserves at the location of the generator of interest are unknown but follow a known random process. The general characteristics of the random process are estimated by the generator wishing to self-commit.
- **Inter-temporal effects**: Inter-temporal constraints affect the generator's operations. This may lead to situations when a market participant can elect to remain on during certain periods when operation will be at a loss in return for likely (but not certain) profits in later periods.

[4] Current period profits *do not* include transition costs.

[5] This objective function assumes that the generator is risk-neutral. If the generator is risk-averse, the objective function should reduce the expected outcome according to some measure of risk. For example, the objective might be to maximize $E\{\pi\text{-}c\}\text{-}a*V\{\pi\text{-}c\}$where "V" is the variance of net profit and "a" is a risk aversion coefficient. As another example, the objective function could be an exponential utility function with constant relative risk aversion [12]. Now the objective function would be multiplicative in nature, but we can take *natural logarithms* to convert the objective function to the form shown in this paper.

We are interested in finding both the optimal *commitment* and the optimal *dispatch* policy. We stress that the problem is complicated by the fact that at the time the commitment decision is made, *future prices are uncertain.* The next section addresses this problem.

3. OPTIMAL COMMITMENT AND DISPATCH POLICY

We now present the optimal commitment and dispatch policy. The optimal dispatch policy is fairly straightforward: given an exogenous price forecast for time period t, the generator takes its current state x_t as given, and dispatches energy and reserves in an optimal manner for time period t, without regard to other time periods. Appendix A briefly describes the single-period optimal dispatch. However, the profit-maximizing *commitment* decision for transitioning to the next time period is more complicated because actions taken now affect future time periods.

The *backward dynamic programming (DP)* method for solving this problem starts at the final time period T and works backward using the following steps. The backward DP method (see [4]) is as follows[6]:

(Step 1) Let $J_T(x_T,p_T) = \max[\pi_T(x_T,p_T) - c_T(u_T,p_T)]$ over *all possible* commitment policies u_T, and dispatch policies d_T. Let the optimal dispatch policy be denoted by d^*_T, and the optimal commitment policy by u^*_T. $J_T(x_T,p_T)$ must be computed for each possible state x_T, and each possible price level p_T.

(Step 2) Let $J_{T-1}(x_{T-1},p_{T-1}) = \max[\pi_{T-1} - c_{T-1} + E[J_T(x_T,p_T)] \mid p_{T-1}]$ over all possible commitment policies u_{T-1}, and dispatch policies d_{T-1}. The expectation E is taken over all possible price levels p_T, *given* that the price in time T−1 is p_{T-1}; p_T is probabilistically related to p_{T-1} via the Markov chain. The state at time T, x_T is related to x_{T-1} by the commitment policy u_{T-1}. Let the optimal dispatch policy be $d^*_{T-1}(x_{T-1},p_{T-1})$, and the optimal commitment be $u^*_{T-1}(x_{T-1},p_{T-1})$. $J_{T-1}(x_{T-1},p_{T-1})$ must be computed for each possible state x_{T-1}, and each possible price level p_{T-1}.

(Step T) Let $J_1(x_1,p_1) = \max[\pi_1 - c_1 + E[J_2(x_2,p_2) \mid p_1]]$ over all possible commitment policies u_1, dispatch policies d_1, and price levels p_1. The expectation E is taken over all possible price levels p_2, *given* that the price in time 1 is p_1. The state at time 2, x_2 is related to the previous state x_1 by the commitment policy u_1, i.e., $x_2= u^*_1(x_1,p_1)$. Let the optimal dispatch policy be $d^*_1(x_1,p_1)$ and the optimal commitment policy be $u^*_1(x_1,p_1)$. $J_1(x_1,p_1)$ must be computed for each possible state x_1, and each price level p_1.

6 For ease of exposition, we have ignored the boundary condition at time T+1.

If the generator is in state x^*_1, and sees price level p^*_1 at time t=1, the profit-maximizing schedules are represented by the commitment actions **$[u^*_1(x^*_1,p^*_1), u^*_2(x^*_2,p^*_2), \ldots, u^*_T(x^*_T,p^*_T)]$** and the dispatch decisions **$[d^*_1(x^*_1,p^*_1), d^*_2(x^*_2,p^*_2), \ldots, d^*_T(x^*_T,p^*_2)]$**, where $x^*_{t+1}=u^*_t(x^*_t,p^*_t)$, for t=1,2,...,T-1, and p^*_{t+1} is related probabilistically to p^*_t via the Markov chain, and **the maximum expected profits are $J_1(x^*_1,p^*_1)$**. Actual profits and actual schedules depend upon actual price levels encountered in the different time periods.

The algorithm for finding the optimal commitment policy is similar to the problem of determining the value of an option using a tree approach [9]. Therefore the problem of finding an optimal commitment policy can be thought of as a generalized tree approach that values and exercises a sequence of complicated options in each time period. The options involve decisions such as whether to commit or not, whether to ramp up or ramp down, whether to participate in the energy or reserves markets, etc.

4. ILLUSTRATIVE EXAMPLE 1

This section illustrates the concept. For simplicity we assume only one product, energy, and one 3-period market, i.e., T=3.

Suppose that the generator parameters are as shown in Table 1. For this example, the generator at time t can be in one of two states, "UP" or "DOWN," i.e., X={UP,DOWN}. There are no inter-temporal constraints. The generator can move from any state at time t to any state at time t+1.

Table 1. Example parameters

Start-up costs, per start, in $	$60
Shutdown costs, per shutdown, in $	$12
Incremental costs, in $/MWh (when UP)	$10
Minimum MW (when UP)	5
Maximum MW (when UP)	50

A shutdown cost is incurred when the generator moves from an UP state to the DOWN state. A start-up cost is incurred when the generator moves from the DOWN state to the UP state. All other transitions result in zero costs.

In each time period, an exogenous price forecast p_t may be described by two possibilities: p_t is HIGH or p_t is LOW, each having a specified probability of occurrence. The HIGH and LOW prices in each time period are allowed to vary, as shown in Table 2. Further assume that the prices at time t+1 depend probabilistically upon the prices at time t. The probability of a HIGH-HIGH transition is α_1 and the probability of a LOW-LOW transition is $\alpha_{2.}$ These are exogenous, with assumed values α_1=0.8 and α_2=0.7.

The backward DP algorithm finds the profit-maximizing commitment and dispatch policy. The solution is shown in Tables 3(a)-(c). Columns in these tables correspond to time periods. The entry in Table 3(a) that corresponds to a "time period t" column and a "state/price" row represents the maximum total *expected profits* for time periods t to T, given that the state at time t is x_t and the price is p_t.

Table 2. Exogenous Energy Price Forecast ($/MWh)
(Number of Time Periods T=3)

Time Period	1	2	3
HIGH Price ($/MWh)	9	20	11
LOW Price ($/MWh)	2	8	1

Similarly, given state x_t and price p_t at time t, the corresponding entries in Tables 3(b)-(e) show respectively:

1. The optimal dispatch for time t.
2. The optimal commitment policy for time t, i.e., the next state to move to at time t+1.
3. The maximum profits obtained from the optimal dispatch at time t.
4. The cost of the optimal commitment policy, or the cost to move from the current state x_t to the new state at time t+1.

Table 3 illustrates the results of a standard Backward DP computation. These results are obtained, as is standard practice, by solving for the respective entries in the table from right to left.

Table 3(a). Values of optimal expected profits (in $)

(State, Price level)	Time Period 1	Time Period 2	Time Period 3
UP, HIGH	415.4	531.0	50.0
UP, LOW	103.9	–22.0	–45.0
DOWN, HIGH	360.4	0.0	0.0
DOWN, LOW	0.0	0.0	0.0

Table 3(b). Optimal dispatch policies d* (in MW)

(State, Price level)	Time Period 1	Time Period 2	Time Period 3
UP, HIGH	5	50	50
UP, LOW	5	5	5
DOWN, HIGH	0	0	0
DOWN, LOW	0	0	0

Table 3(c). Optimal commitment policies u*

(State, Price level)	Time Period 1	Time Period 2	Time Period 3
UP, HIGH	UP	UP	UP
UP, LOW	UP	DOWN	UP
DOWN, HIGH	UP	DOWN	DOWN
DOWN, LOW	DOWN	DOWN	DOWN

Table 3(d). Values of π^* (in $)

(State, Price level)	Time Period 1	Time Period 2	Time Period 3
UP, HIGH	–5.0	500.0	50.0
UP, LOW	–40.0	–10.0	–45.0
DOWN, HIGH	0.0	0.0	0.0
DOWN, LOW	0.0	0.0	0.0

Table 3(e). Values of c* (in $)

(State, Price level)	Time Period 1	Time Period 2	Time Period 3
UP, HIGH	0.0	0.0	0.0
UP, LOW	0.0	12.0	0.0
DOWN, HIGH	60.0	0.0	0.0
DOWN, LOW	0.0	0.0	0.0

Suppose that, at time t=1, the price level is HIGH and the generator is UP. The maximum expected profits are $415.4 over the three time periods. Since the current price level is HIGH and the generator is UP, the optimal commitment policy is to stay UP at t=2 (from Table 3(c)), *in spite of the possibility of net losses over the three periods*. The actual profits and the commitment policies at other times would however depend upon the actual price levels that occur in those time periods. For example, if the price level stays HIGH for both t=2 and t=3, the optimal commitment policy is to stay UP at t=3, realizing total profits of −5+500+50=$545. If, on the other hand, the price level becomes LOW for t=2, and LOW for t=3, the optimal commitment policy is to go DOWN at t=3, realizing profits of −5−10−22=−$37. In other words, the generator *loses* money under some price patterns, even with the optimal policy. However, the expected profits are maximized.

The optimal schedules given by the backward DP are not static. Instead, they depend upon the exogenous prices in each time period. Thus the DP method does not *merely* give an optimal schedule. Rather, it gives an optimal *scheduling policy* corresponding to different conditions.

5. ILLUSTRATIVE EXAMPLE 2

In this section, we describe the "*optionality*" features of the generator self-commitment problem, and show that it has features analogous to financial options. We also make additional three points:

1. Assuming a single average price forecast generally understates the value of the optionality, and could severely understate expected generator profits.
2. Running Monte Carlo methods without taking care to ensure that future prices are always uncertain (at the time the commitment is made) generally overstates the value of the optionality, and overstates expected generator profits.
3. Not considering reserve products (and multiple markets) tends to lower expected generator profits, because these additional products increase generator optionality.

First, we consider the optionality due to price uncertainty. Assume a single time-period horizon, and a single product --- the energy service. Assume that the generator for which we want to find the optimal commitment and dispatch policy has no start-up or shutdown costs, and no inter-temporal constraints. Assume that generator has a capacity of 100 MW, no minimum generation constraint, and constant incremental costs of $30/MWh over this range.

Table 4. Expected generator profits as a function of price volatility. Each price forecast is represented by a HIGH/LOW price, each with equal probability of occurrence.

	PRICE FORECASTS				
	1	**2**	**3**	**4**	**5**
HIGH price ($/MWH)	30	35	40	45	50
LOW price ($/MWh)	30	25	20	15	10
Expected price ($/MWh)	30	30	30	30	30
Std. Dev. ($/MWh)	0	7	14	21	28
Profits (HIGH price)	$0	$500	$1000	$1500	$2000
Profits (LOW price)	$0	$0	$0	$0	$0
Expected profits	$0	$250	$500	$750	$1000
Price volatility[7]	0%	17%	33%	50%	67%

Table 4 shows five different price forecast scenarios. Two possible price states, HIGH and LOW, each with a 50% probability of occurrence, represent each price forecast. The expected value of the prices for all price forecasts is $30/MWh. The optimal policy for the generator is to produce 100 MW whenever the energy price exceeds its incremental costs, and to produce 0 MW whenever the energy price is below its incremental costs. For example, for forecast #3, the generator will produce 0 MW when the price is LOW (or $20/MWh), and will make no profits. When the price is HIGH ($40/MWh), the generator will produce 100 MW and make a profit of 100*(40-30) = $1000. Since both price scenarios are equally likely for this forecast, expected profits are 0*0.5+1000*0.5 = $500.

Table 4 shows that expected generator profits *increase* with increasing price volatility[8]. This is analogous to the value of a financial option that increases in value when price volatility increases [9]. On the other hand, if one uses an average price of $30/MWh to find a commitment policy for the generators, then he/she will estimate that the generator will make no profit for any price forecast because the generator incremental costs will not exceed the expected energy price. *Therefore average price forecasts fail to calculate the value of optionality and understate generator profits.*

Next, we examine a potential pitfall associated with Monte Carlo methods. One possible way of finding the expected generator profits to capture the optionality value could be to generate a large number of random price scenarios for the time interval [0,T] by Monte Carlo methods. Using the ensemble of all generated price scenarios one could then use a deterministic model to solve for the optimal commitment and dispatch over this period for each member of the ensemble, and then average over

[7] Price volatility here is defined as the ratio of the standard deviation to the expected price, expressed in percent (and rounded).

[8] All other factors (e.g., expected energy price) remaining constant.

different Monte Carlo runs. *This is, however, not always correct*. To see this, consider the following example. Assume a two-period system, with each period having a 50% chance of HIGH price ($35/MWh) and a 50% chance of having a LOW price ($10/MWh), regardless of the previous period. The generator has the same characteristics as the above example, except that there is a minimum generation limit of 90 MW, and an additional inter-temporal constraint that, once online, the generator has to stay online for two consecutive periods. The boundary condition is that the generator is offline initially, and must be offline at the end of two periods. It can be verified that the optimal policy is not to commit the generators regardless of what the period 1 price is[9]. *Therefore the expected profit under the optimal commitment policy is zero.* On the other hand, suppose that we *first* generate all the price scenarios (using a Monte Carlo method), and *then* run a deterministic optimal unit commitment on each possible price sequence. The four equally possible price sequences in the two periods are {HIGH, HIGH}, {HIGH, LOW}, {LOW, LOW}, and {LOW, HIGH}. If we make four deterministic unit commitment runs on these four price sequences, the deterministic unit commitment will only run the generator at maximum output (100 MW) for both time periods when the price sequence is {HIGH, HIGH}. The profit for this price sequence is $1000. For all other price sequences, the generator will not run, and the profit will be zero. *Hence expected profits using this method will be 0.25*1000+0.75*0 = $250, which clearly overstates the expected profits of zero under the true optimal commitment policy.* The reason for this is that in each Monte Carlo run, the generator "peeped ahead" and "knew" the future prices and therefore chose the commitment accordingly. Models for commitment based on the traditional approach are likely to follow a variant of this deterministic optimization method. This approach would result in *overestimation* of generator profits.

Monte Carlo methods are very efficient when one needs to simulate a large number of different random outcomes and find the expected value (or some other statistic) of some function based on these random outcomes. They are much more complicated to implement[10], and prohibitively

[9] If period 1 price is HIGH, and the generator decided to commit, then the generator would produce 100 MW in period 1 to make a profit of $500. However, there is a 50/50 chance that the period 2 price is HIGH or LOW. If the period 2 price is HIGH, the generator's two-period profit will be $1000. If period 2 had a LOW price, the generator would produce the minimum 90 MW and lose 90*20=$1800 in period 2 for a net two-period loss of $1300. Therefore, if the generator commits to be online when the period 1 price is HIGH, the expected two-period payoff is 1000*0.5 − $1300*0.5 = ($150), for an expected loss of $150.

[10] See reference [8] for the correct way to implement Monte Carlo methods for such problems. However, in [8], the Monte Carlo method becomes prohibitively expensive as

expensive, when the value of a function at any given time t itself depends on what may happen *in the future*, as in optimal commitment policy problems that have inter-temporal constraints[11]. For example, using the finance analogy, Monte Carlo methods are used for European style options and for those other types of options when one does not have to worry about *when* it is optimal to exercise the option. However, American style options are much more difficult to solve for using Monte Carlo methods [12, pp. 685].

Consider a final example to show the optionality value of multiple products (or multiple markets). Assume a single time-period problem and a single reserve product. Consider a generator whose incremental cost is $30/MWh, maximum capacity is 100 MW, minimum capacity is 0 MW, and maximum reserve capacity of 30 MW. Assume that the price of energy is $45/MWh and reserve availability costs of $20/MW/h. If the generator offers 100 MW of energy only, it will make profits of 100*(45–30) = $1500, on revenues of 100*45=$4500. On the other hand, if it maximizes its profits[12] and offers 30 MW of reserve and 70 MW of energy, its profits are 30*20+70*(45–30)=$1650 on revenues of 30*20+70*45=$3750. (Shifting more of the generator output to reserves increases profits, even though total revenues decrease relative to the energy-only sale. This is typical.)

In summary, the more optionality that a generator has, the higher its expected profits will be. Conversely, the more the operational constraints reduce this optionality (e.g., inter-temporal constraints), the lesser will be its expected profits, all other factors being equal.

6. ILLUSTRATIVE EXAMPLE 3

We now consider a slightly more realistic example[13]. Assume a peaking generator with the characteristics illustrated in Table 5. The start-up and shutdown times for the generator are assumed to be zero. Figure 1 shows the baseline energy price forecast. Figures 2 and 3 illustrate the anticipated

---

the number of generator states and price uncertainty states increase. See also Section 6 for how to use Monte Carlo methods *once the optimal self-commitment policy is known*.

[11] The Monte Carlo method discussed in this section will work *correctly* on the problem described in Table 4, because the optimal self-commitment policy does not have inter-temporal features.

[12] Appendix A shows how one may approach the problem when there are more than one reserve type. In this example, while there is profit to be made on the sale of both energy and reserves, the generator sees a higher profit margin in reserves, and so maximizes the sale of reserves (30 MW). The remainder is offered as energy (100-30=70 MW).

[13] The results of this section were derived using EPRI's PROFITMAX model.

baseline prices for 4 types of reserves: regulating, spinning, supplemental, and backup[14]. The total number of time periods is one week (168 hours).

Table 5. Generator characteristics

Start-up costs, per start, in $	100
Shutdown costs, per shutdown, in $	0
Production cost, constant coefficient ($ per hr)	100
Production cost, linear coefficient ($ per MWhr)	37
Production cost, quadratic coeff. ($ per MWhr^2)	37
Minimum MW (when UP)	10
Maximum MW (when UP)	110
Maximum regulating reserves (MW)	10
Maximum spinning reserves (MW)	20
Maximum supplemental reserves (MW)	40
Maximum backup reserves (MW)	110

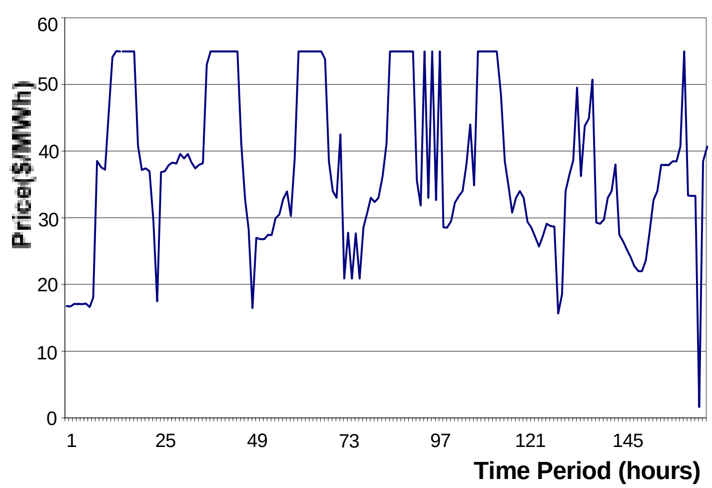

Figure 1: Price transition. Predicted baseline energy prices. This is the expected pattern of future prices (“baseline” prices) at time t=0.

[14] Spinning reserves are defined to be the capability that can be offered in 10 minutes (if online). Supplemental reserves are the amount of MW available in 20 minutes, and backup reserves are the amount available in 60 minutes.

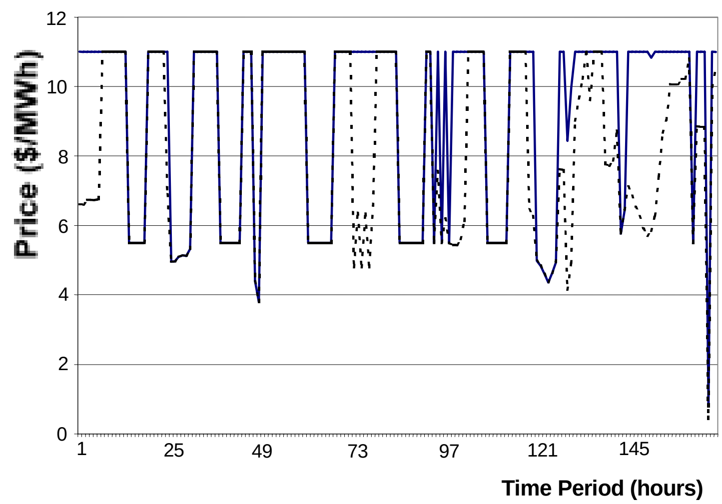

Figure 2. Baseline reserve availability prices, spinning and regulation. These are the anticipated future patterns of prices ("baseline" prices) at time t=0. Solid lines show regulating reserves while dashed lines show spinning reserves.

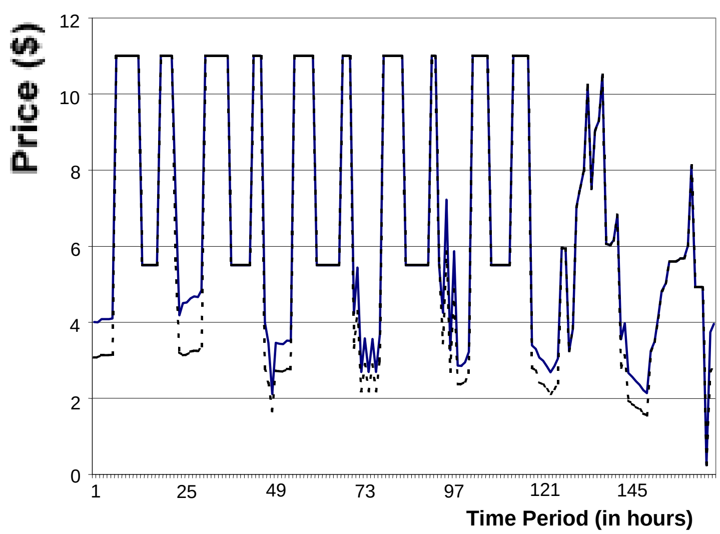

Figure 3. Baseline reserve availability prices, supplemental (solid lines) and backup (dashed lines). These are the anticipated patterns of future prices ("baseline" prices) at t=0.

We model price uncertainty using a three-node Markov process. We assume that prices can be at one of three levels: HIGH, BASELINE, or LOW. The BASELINE prices are as illustrated. The HIGH price for energy and reserves is 115% of the corresponding baseline price, and the LOW price is 85% of the baseline price. We assume that there is perfect correlation between energy and reserve availability prices; e.g., when the energy price is HIGH, so are the reserve availability prices. For simplicity, we assume that the probability of transition between any one price state to any other price state is 1/3. That is, it is equally likely for the price to change states regardless of the present state of prices.

Using the backward DP methods described in Section 3, we derived the optimal commitment and dispatch policy. The optimal commitment and dispatch policy at any time t is a function of the state that the generator is in, and a function of the uncertain price forecast for future time periods as observed at time t. *Future prices are always considered uncertain*. We then applied the optimal policy in numerous Monte Carlo runs to simulate different profit outcomes. From the Monte Carlo prices, we then calculated the actual eventual profits and generator outputs. Using these outcomes, we then illustrate the probability distribution of different generator outputs: generator profits, optimal generator dispatch of energy and reserves, etc.

When the optimal policy found by the DP is used to simulate a number of possible states, one obtains a distribution of possible outcomes. Figure 4 illustrates this distribution. There is no assurance that the distribution will be neatly “bell shaped” as in this example. In other examples it is possible to have distributions that are skewed or bimodal, particularly when start-up and shutdown costs are considered.

The effects of the optimal commitment policy on the state transitions are shown in Figure 5. For a large number of scenarios, all transitions between states that result from the optimal policy are recorded. For some times, the state is unique (either UP or DOWN), but for other times the system could end up in either state, depending on the price sequence. This is because it is not known which state the generator will be at a future time t. Figure 5 shows that, depending on the prices that are actually realized, the generator could be in any of the different states at a future time. The ball “size” represents the probability of ending up in a particular state. The “thickness” of a transition line indicates the likelihood that the particular transition will take place.

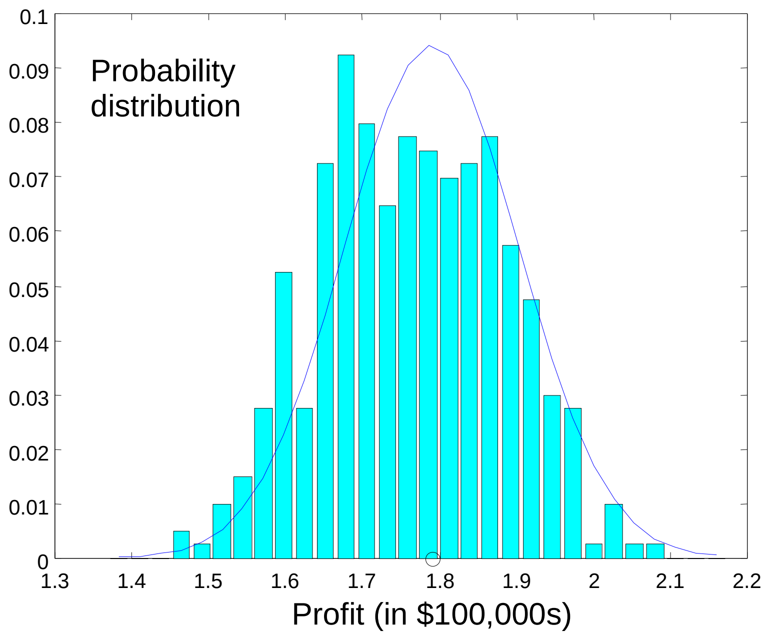

Figure 4. Distribution of anticipated profits. The optimal policy maximizes expected profits, but there is variance around this mean value even with an optimal policy.

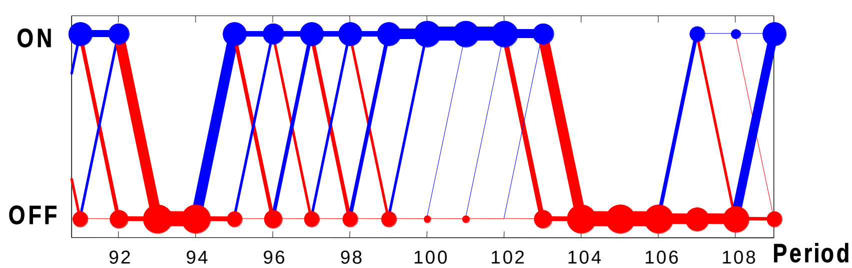

Figure 5. Optimal commitment policy with uncertainty. Only hours 91 to 109 are illustrated. During this interval the optimal policy has the option to take several transitions, depending on the actual price outcome realized. Only hours 93, 94 and 104 to 106 have a certain state (OFF in this case). For other periods the relative probability of being in either state is represented by the size of the "ball" and the relative transition probability is represented by the thickness of the transition "line." During certain periods there is both an up transition probability and a down transition probability in the optimal policy.

For comparison, Figure 6 illustrates the optimal commitment policy when there is no uncertainty, i.e., the baseline price forecast is the perfect forecast. It can be seen that the effect of price uncertainty recognizes the possibility that the generator will begin producing during certain periods if the prices become high (see, in particular, time periods 120-140). That is, a quick-responding generator has the luxury of producing when the prices are high, and going offline when the prices are low. Thus high price volatility tends to be beneficial for the *expected* profits of the peaking generator. For prices higher than its incremental costs, the peaking generator will maximize its output (and increase profits), while for prices below its incremental costs, the generator will shut down (and have zero profits). Since profits are bounded from below at zero, and monotonically increase as a function of price (above the generator's incremental costs) the generator's expected profits will increase.

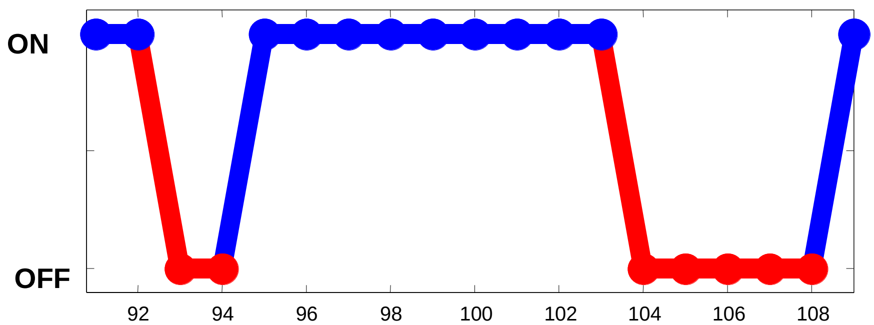

Figure 6. Optimal Commitment Policy (no uncertainty). All "balls" are the same size, corresponding to 100% probability of being in the corresponding state.

Figure 7 illustrates this notion. It shows that when expected prices are held constant among scenarios, expected profits increase as price uncertainty increases. This again shows that the problem of committing and dispatching a generator is analogous to that of exercising a financial option. The option value generally increases when uncertainty increases[15]. This is because a peaking generator is able to follow changes more readily than other types of generators. Thus, price volatility is beneficial to peaking units, a result that may be familiar to many readers. In Figure 11, we are able to precisely quantify just how large is this benefit.

[15] An intermediate or cycling generator has less optionality features because inter-temporal constraints affect its profits, and has features similar to *Asian* options [9]. A baseload generator has even less optionality features (excepting for the important question of which market to sell into), usually because its incremental costs are generally well below the market prices and is analogous to *forward* contracts [9].

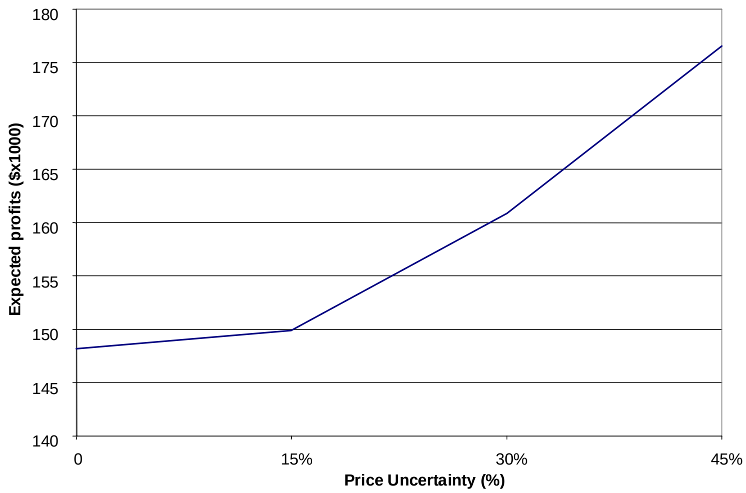

Figure 7. Variation in expected profit as a function of energy price uncertainty. As energy price uncertainty increases, the expected profits increase for this generator.

Figure 8 illustrates the expected sales of energy and reserves as a function of time. Note the striking fact that the generator in question offers energy only during certain periods, but derives income from making reserves available during many more periods.

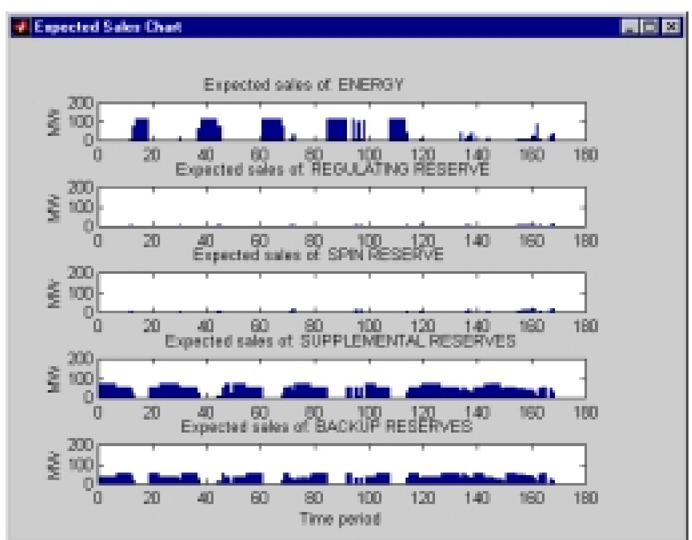

Figure 8. Optimal sales of energy and reserves

7. CONCLUSIONS

The main contribution of this paper is to describe the multi-period multi-market uncertain framework within which decisions for unit commitment and dispatch will have to take place for many units that operate in a deregulated market. The paper applies directly to the problem of optimal generator self-commitment. This paper describes a method for finding the most profitable market-responsive commitment and dispatch policy that takes into full account the optionality available to a generator: reserve market opportunities, multiple markets, price uncertainty, and inter-temporal constraints. The model uses backward dynamic programming. The algorithm used by the model can be thought of as a generalized tree that values and exercises a sequence of complicated options. This algorithm can be used to obtain optimal power market bids for energy and reserve services in markets that integrate both these needs. The method can also be used to profitably allocate output in different physical forward markets, e.g., hour-ahead versus day-ahead.

ACKNOWLEDGEMENTS

We thank EPRI PM&RM for support of this work. We thank Victor Niemeyer of EPRI for his support and encouragement. In particular, the work on the PROFITMAX model developed by Christensen Associates was sponsored by EPRI and initiated in 1997. We thank Blagoy Borissov of Christensen Associates for his help in simulating some of the results in this paper. We also thank Fritz Schulz for his help with the original implementation of PROFITMAX.

REFERENCES

1. A.J. Wood and B.F. Wollenberg. *Power Generation Operation and Control*, Wiley 1984.
2. A. Merlin and P. Sandrin. A New Method For Unit Commitment at Electricite De France. IEEE Transactions on Power Apparatus and Systems, PAS-102(5): 1983.
3. D.P. Bertsekas et al. Optimal Short-Term Scheduling of Large-Scale Power Systems. IEEE Transactions on Automatic Control, AC-28(1): 1983.
4. C.A. Li, R.B. Johnson and A.J. Svoboda. A New Unit Commitment Method. IEEE Transactions on Power Systems, 12(1): 1997.
5. C L. Tseng, S.S. Oren, A J. Svoboda and R.B. Johnson. A Unit Decommitment Method in Power System Scheduling. Electrical Power and Energy Systems, 19(6): 1997.
6. M. Hsu. Spark Spread Options are Hot!. Electricity Journal, 11(2): 1998.

7. S. Deng, B. Johnson, A. Sogomonian. Exotic Electricity Options and the Valuation of Electricity Generation and Transmission Assets. Research Report #PSERC 98-13, 1998 (also available in the limited access http://www.pserc.wisc.edu/index_publications.html).
8. C.L. Tseng and G. Bartz. Short Term Generation Asset Valuation. Research Report #PSERC 98-20, 1998 (also available in the limited access website http://www.pserc.wisc.edu/index_publications.html).
9. J.C. Hull. *Options, Futures, and Other Derivatives*, Prentice Hall, 1997.
10. D.P. Bersekas. *Dynamic Programming and Optimal Control*, volumes 1 and 2, Athena Scientific, June 1995.
11. L. Kirsch and R. Rajaraman. Profiting from Operating Reserves. Electricity Journal, March 1998.
12. H.R. Varian. *Microeconomic Analysis*, W.W. Norton & Company, third edition, 1992.
13. P. Wilmott. *Derivatives: The Theory and Practice of Financial Engineering*, John Wiley & Sons, 1998.

APPENDIX A

This appendix describes how a profit-maximizing generator would dispatch energy and reserve availability services *given* exogenous market prices for a given hour, and given that it is committed to be online.

Assumptions

Consider the output choice faced by a generator that can offer, in any given hour, energy and four reserve services. Assume that the energy price is PE, and that the availability prices for the four reserves are PR1, PR2, PR3, PR4 respectively. Suppose that the generator has maximum output level ME, and that, because of ramping limitations, the generator can provide the maximum quantities of the reserve XR1 for reserve 1, XR2 for reserve 2, XR3 for reserve 3, and XR4 for reserve 4. Further suppose that the generator's production cost function is:

$$\text{COST} = a\ \text{XE}^2 + b\ \text{XE} + c \qquad \text{(A-1)}$$

where a, b, and c are constants and XE is the generator's quantity of offered energy. What quantities of energy and reserves should the generator offer if it is maximizing profits?

The Optimization Problem

The generators problem is to maximize profits[16]:

$$\pi = \text{PE*XE} + \text{PR1*XR1} + \text{PR2*XR2} + \text{PR3*XR3} + \text{PR4*XR4} - (a\ \text{XE}^2 + b\ \text{XE} + c) \qquad \text{(A-2)}$$

[16] When a generator offers reserves, there is a certain probability of these reserves being *called*. One can easily include this effect in the objective function (A-2).

subject to the constraints that all energy and reserve quantities must respect maximum limits:

$$XE + XR1 + XR2 + XR3 + XR4 \leq ME \qquad \text{(A-3a)}$$

$$XR1 \leq MR1 \qquad \text{(A-3b)}$$

$$XR2 \leq MR2 \qquad \text{(A-3c)}$$

$$XR3 \leq MR3 \qquad \text{(A-3d)}$$

$$XR4 \leq MR4 \qquad \text{(A-3e)}$$

For simplicity and without loss of generality, we ignore the constraints that all reserve quantities must be non-negative, and that the energy quantity have a minimum limit. Although equations (A-2) and (A-3) make the reserves all appear to be mathematically identical, they are not because we assume (reasonably) that reserve prices have a particular order[17]:

$$PR1 \geq PR2 \geq PR3 \geq PR4 \geq 0 \qquad \text{(A-4)}$$

Because production costs (A-1) depend only upon energy output, the generator will prefer to sell Reserve 1 first and Reserve 4 last.

The solution to the above problem is:

$$XE = \frac{PE - b - \mu}{2a} \qquad \text{(A-5)}$$

If Reserve 1 is (optimally) offered in a positive quantity that is less than its limit, then μ=PR1. But if Reserve 1 *is* (optimally) offered to its limit, then μ<PR1. In general:

- if Reserve j is offered at all, then all reserves <j are at their limits;
- if Reserve j is offered in a positive quantity that is less than its limit, then μ=PRj and all reserves >j are not offered at all.

[17] Because of market imperfections, this relationship is not always obeyed; e.g. see historical ancillary service prices from the California ISO website (http://www.caiso.com).