Offer Stack Optimization for Price Takers in Electricity Markets

P. J. Neame, A. B. Philpott, G. Pritchard
University of Auckland
New Zealand.
email: p.neame@auckland.ac.nz

Abstract

We consider a small generator offering into an electricity spot market. For a given time period, they must submit an offer stack, consisting of a finite number of quantity-price pairs (or tranches). We assume that the generator is a price-taker — that is, it cannot offer in enough power to substantially affect the market price. We consider how to optimize the offer stack for various assumptions on the cost of production, and the probability distribution for market price.

1 Introduction

In this paper, we consider a class of problems facing generating companies who sell electricity into wholesale electricity markets. In recent years various forms of these markets have emerged throughout the world. The problems we study arise in markets that have a dispatch and pricing mechanism similar to that now implemented in the United Kingdom, Australia and New Zealand, where each generator produces an *offer stack*, consisting of a finite set of price-quantity pairs (*tranches*), indicating that they are willing to produce those quantities at the corresponding prices. The number of tranches that can be offered in any offer period depends on the market. (In New Zealand each generator is restricted to five tranches in each half-hour period). The prices in these markets are determined by a central authority (M-Co in New Zealand) that clears demand by systematically dispatching the cheapest supply offers until demand is met (allowing for transmission network losses and constraints). The marginal cost of supply at any node of the transmission network (i.e. the shadow price of the energy balance constraint) then defines the spot price of energy at this node.

Each generator in such a market seeks the best possible stack to offer at each time period. Their objective in the short term is to maximize their (possibly risk-adjusted) return taking into account any contract position. In this paper we choose to ignore contracts and deal with a model in which the generators are risk neutral and seek an offer stack in each dispatch period that will maximize their expected revenue minus expected costs. For simplicity we shall restrict attention to a generator offering a single offer stack at a single node of the transmission network.

We shall further assume that the generator is small enough that their actions have no effect on the spot price. This price-taking assumption is a considerable simplification of the price-making case where generators can influence the spot price by varying their offers. (This is the situation that is likely to pertain in practice for all but the smallest generators.) With perfect information on the behaviour of competitors and demand, a price-making generator can construct a single offer that maximizes their return, assuming that the behaviour of competitors does not alter. However such a policy might fail to be optimal after the potential responses of competitors are taken into account. In these circumstances the appropriate models will come from game theory.

A slightly simpler case occurs when the competitors' actions are unknown, but assumed to be independent of the policy of the generator we have in mind. This case has been studied by Anderson and Philpott [1] who represent the uncertainties in competitor behaviour and demand by a single *market distribution function* $\psi(q,p)$, defined to be the probability of a single-tranche offer (q,p) not being fully dispatched. This can be used to derive optimality conditions for offer stacks in this price-making context.

In the price-taking model that we consider, the probability $\psi(q,p)$ is a function of p only. (It is the probability that the clearing price is below p. For simplicity of exposition we restrict attention in this paper to cases where the price distribution in any offer period has a known density function with bounded support.) This means that much of our analysis is a special case of that in [1]. However our emphasis in this paper is on the development of procedures for computing optimal solutions. In this respect, studying the price-taking environment gives clues to how one might proceed in developing procedures for the more general case. Indeed, even with the simplifications described, the solution of the model is not elementary.

In the next section we formally define offer stacks and formulate the objective function that the generator seeks to maximize. Section 3 derives optimality conditions for offer stacks with finitely many tranches. In section 4 we give an example to demonstrate that there are often many local optima to these problems, and then describe a strategy for dealing with these. In section 5 we deal with the special case where the cost of generation is a piecewise linear function. This leads to a simple dynamic programming algorithm for computing a globally optimal offer stack.

2 Offer stacks

It is convenient to model an offer stack as a continuous curve $\mathbf{s} = \{(q(t), p(t)), 0 \leq t \leq T\}$, in which the components $q(t)$ and $p(t)$ are monotonic increasing piecewise differentiable functions of t. Here $q(t)$ traces the quantity component of the offer curve and $p(t)$ traces the price component. This definition encompasses the case where the generator offers a continuous supply function $q(p)$, as well as the case where there are a fixed number of tranches — in this case the curve $\mathbf{s}$ will consist of horizontal and vertical sections. Let $f(p)$ denote the probability density function of spot price. We assume f has bounded support giving a bound

p_M on the spot price, so that $p(t) \leq p_M$, $0 \leq t \leq T$. Furthermore, we assume that $q(0) = 0$, $p(0) = p_0 = \inf \{p \mid f(p) > 0\}$, and that $q(t) \leq q_M$, $0 \leq t \leq T$, where q_M is the generation capacity of the generator. We also assume that $q(T) = q_M$, and $p(T) = p_M$, which includes the case where the final section of the curve is vertical at q_M.

We denote the cost of generating an amount q by $c(q)$. Suppose we are dispatched q at a clearing price of p. The return we make is then $qp - c(q)$. We are interested in maximising the expected return obtained by choosing a particular stack $\mathbf{s}$ to offer into the market. Then the expected return from offering $\mathbf{s}$ is

$$F = \int_{t=0}^{T} [q(t)p(t) - c(q(t))] f(p(t)) p'(t) dt.$$

If the generator can offer a quantity that is a continuous function $q(p)$ of the price p then the expected return is

$$F = \int_{p=0}^{p_M} [q(p)p - c(q(p))] f(p) dp.$$

Note that if $f(p) > 0$ for all $p \in [p_0, p_M]$, then F is maximized when

$$\frac{\partial (qp - c(q))}{\partial q} = 0,$$

giving $p - c'(q) = 0$. Hence the optimal strategy for the generator is to offer a curve $q(p)$ with the property that the price of each offered amount is its marginal cost.

In the remainder of this paper we focus on the case where the generator may choose no more than m tranches. In this case, $\mathbf{s}$ is a step function that changes direction at points $(0, p_1)$, (q_1, p_1), (q_1, p_2), ..., (q_m, p_m). If $p_m < p_M$ then there is a vertical section from (q_m, p_m) to (q_m, p_M). The expected return is now

$$F(p_1, \ldots, p_m, q_1, \ldots, q_m) = \sum_{i=1}^{m} \int_{p_i}^{p_{i+1}} [pq_i - c(q_i)] f(p) dp, \qquad (1)$$

where for convenience we identify p_{m+1} and p_M.

3 Optimality Conditions

The problem of computing an optimal offer stack with m or fewer tranches is now the following nonlinear programming problem.

$$\begin{aligned} P(m) \quad &: \quad \max F(p_1, \ldots, p_m, q_1, \ldots, q_m) \\ \text{such that } p_i \quad &\leq \quad p_{i+1},\ i = 1, \ldots, m-1, && (2) \\ q_i \quad &\leq \quad q_{i+1},\ i = 1, \ldots, m-1, && (3) \\ 0 \quad &\leq \quad p_i \leq p_M, i = 1, \ldots, m, && (4) \\ 0 \quad &\leq \quad q_i \leq q_M,\ i = 1, \ldots, m. && (5) \end{aligned}$$

Since c is continuous, F is continuous, and since the feasible region of $P(m)$ is compact, $P(m)$ has an optimal solution. We begin by considering first order conditions with respect to the variables $p_2, \ldots, p_m, q_1, \ldots, q_{m-1}$, then consider what happens at the ends of the stack, and then derive second order conditions. Observe that $P(m)$ is a constrained optimization problem. We assume that any candidate solution to $P(m)$ that we investigate satisfies the constraints (2) and (3) as strict inequalities. If this is not the case, then essentially the same solution (having fewer tranches) can be constructed, satisfying this condition.

3.1 First Order Conditions

We first consider the case where F is a smooth function of its arguments, and obtain the stationarity conditions. From (1),

$$\frac{\partial F}{\partial q_i} = \int_{p_i}^{p_{i+1}} pf(p)dp - c'(q_i) \int_{p_i}^{p_{i+1}} f(p)dp,$$

and

$$\frac{\partial F}{\partial p_i} = \begin{cases} [q_{i-1}p_i - c(q_{i-1})]\, f(p_i) - [q_i p_i - c(q_i)]\, f(p_i) & , \quad i > 1 \\ [-q_1 p_1 + c(q_1)]\, f(p_1) & , \quad i = 1. \end{cases}$$

Now $\frac{\partial F}{\partial p_i} = 0$ implies that the horizontal sections (sections with fixed price) of an optimal stack satisfy

$$p_i = \frac{c(q_i) - c(q_{i-1})}{q_i - q_{i-1}}, \; i = 2, \ldots, m. \tag{6}$$

(We will discuss p_1 in Section 3.2.) That is, the generator offers in a given quantity at its average marginal cost.

If the quantities that can be offered in are determined (for example, the generator can only run at a few different output levels), and there are sufficiently many tranches allowed, then this is enough to determine the optimal tranches.

However, if the quantities offered can also be altered, then each of the vertical sections (sections with fixed quantity) can be shifted. Assuming $f(p_i) > 0$, the condition that $\frac{\partial F}{\partial q_i} = 0$ for $i < m$ yields

$$c'(q_i) = \frac{\int_{p_i}^{p_{i+1}} pf(p)dp}{\int_{p_i}^{p_{i+1}} f(p)dp} \text{ for all } i = 1, \ldots, m-1. \tag{7}$$

That is, the conditional expected price along the vertical section is equal to the marginal cost.

3.2 Boundary conditions

The above stationarity conditions are necessary for optimality. Note though, that we did not consider p_1, and q_m, which define the beginning and end of the stack. Obviously, it is possible that these cannot be moved in both directions. A negative quantity makes no sense, and we assume that negative prices are not allowed. A realistic model will also gives a maximum price and quantity that needs to be considered — however, if the model lacks either then the condition (7) applies to q_m.

We label $q_0 = 0$ and $p_0 = \inf\{p \mid f(p) > 0\}$. If $p_1 > p_0$, then the optimality condition is that $p_1 = \frac{c(q_1)}{q_1}$; otherwise, if $p_1 = p_0$, then the optimality condition is that $\frac{\partial F}{\partial p_1} \leq 0$, that is, $p_1 \geq \frac{c(q_1)}{q_1}$.

On the other boundary, if $q_m = q_M$, then

$$c'(q_i) = \frac{\int_{p_m}^{p_M} pf(p)dp}{\int_{p_m}^{p_M} f(p)dp};$$

otherwise, if $q_m = q_M$ then

$$c'(q_i) \leq \frac{\int_{p_m}^{p_M} pf(p)dp}{\int_{p_m}^{p_M} f(p)dp}.$$

3.3 Second Order Conditions

To determine whether a particular stack yields a maximum of the revenue function F, we need to check the second order conditions.

Firstly,

$$\frac{\partial^2 F}{\partial q_i^2} = -c''(q_i) \int_{p_i}^{p_{i+1}} f(p)dp.$$

Also, for $i > 1$,

$$\begin{aligned}\frac{\partial^2 F}{\partial p_i^2} &= q_{i-1}f(p_i) - q_{i-1}p_i f(p_i) - c(q_{i-1})f'(p_i) - (q_i f(p_i) + q_i p_i f'(p_i) - c(q_i)f'(p_i)) \\ &= (q_{i-1} - q_i)\, f(p_i),\end{aligned}$$

if the stack satisfies the first order conditions for optimality (6). Similarly,

$$\frac{\partial^2 F}{\partial p_1^2} = -q_1 f(p_1).$$

The non-diagonal elements of the Hessian matrix are

$$\frac{\partial^2 F}{\partial q_i \partial p_i} = \frac{\partial^2 F}{\partial p_i \partial q_i} = (c'(q_i) - p_i)\, f(p_i),$$

and for $i > 1$,

$$\frac{\partial^2 F}{\partial q_{i-1} \partial p_i} = \frac{\partial^2 F}{\partial p_i \partial q_{i-1}} = (p_{i-1} - c'(q_{i-1}))\, f(p_i).$$

Thus the Hessian is tri-diagonal. We use the abbreviations $c_i' = c'(q_i)$, $c_i'' = c''(q_i)$, $f_i = f(p_i)$ and $\int_{p_1}^{p_2} f = \int_{p_1}^{p_2} f dp$ to save space, and consider $-H$, as we are looking for conditions for a stack to be maximal.

$$-H = \begin{bmatrix} q_1 f_1 & (p_1 - c_1')f_1 & & & & \\ (p_1 - c_1')f_1 & c_1'' \int_{p_1}^{p_2} f & (c_1' - p_2)f_2 & & & \\ & (c_1' - p_2)f_2 & (q_2 - q_1)f_2 & (p_2 - c_2')f_2 & & \\ & \ddots & \ddots & \ddots & & \\ & & (c_{m-1}' - p_m)f_m & (q_m - q_{m-1})f_m & (p_m - c_m')f_m \\ & & & (p_m - c_m')f_m & c_m'' \int_{p_m}^{p_{m+1}} f \end{bmatrix}. \tag{8}$$

A sufficient condition for a candidate stack to be a local maximum is that $-H$ is positive definite and a necessary condition is that $-H$ is positive semi-definite.

4 Examples

We now consider two examples that demonstrate some features of this problem. Firstly, we consider a simple example, to gain some insight into what the stationarity conditions mean. Then we consider a more complex example which has multiple stacks satisfying the stationarity conditions.

4.1 A Simple Example

We consider the case where the price is uniformly distributed on $[0, 1]$, so

$$f(p) = \begin{cases} 1 & , \quad p \in [0, 1], \\ 0 & , \quad \text{otherwise}, \end{cases}$$

and $c(q) = q^2$ for all q.

From (7),

$$c'(q_i) = \frac{\frac{1}{2}p_{i+1}^2 - \frac{1}{2}p_i^2}{p_{i+1} - p_i},$$

yielding

$$p_{i+1} = 4q_i - p_i. \tag{9}$$

Similarly (6),

$$p_i = \frac{q_i^2 - q_{i-1}^2}{q_i - q_{i-1}}.$$

gives

$$q_i = p_i - q_{i-1}. \tag{10}$$

Note that if (p, q) lies on the line $p = 2q$ (the infinite tranche solution, with the price equal to the marginal cost), then neither (9) nor (10) suggests a new point to move to. If the point is below this line, then we can only move upwards; if it is above the line, we can only move across.

Thus, from any p_1, we can iteratively solve (9) and (10) to build an offer stack. In order to obtain the optimal m-tranche stack, it remains to satisfy an optimality condition on the boundary. As there is no maximum quantity, the only possible boundary condition is that $p_{m+1} = 1$.

For example, if $m = 3$, we obtain $p_1 = \frac{1}{7}$ giving the solution shown in Figure 1. The expected return from this solution is approximately 98% of the expected return from the optimal infinite tranche solution.

4.2 Example With Multiple Solutions

We now consider a more complicated example, in order to demonstrate that finding an optimal stack is not always straightforward.

Again, we consider the price distribution to be uniform, although this time distributed over $[0, 5]$, so

$$f(p) = \begin{cases} \frac{1}{5} & , \quad p \in [0, 5], \\ 0 & , \quad \text{otherwise}. \end{cases}$$

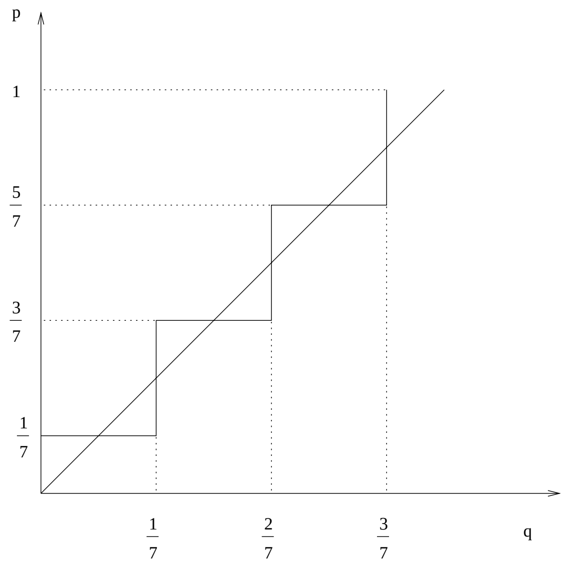

Figure 1: Optimal tranches for simple example

We consider the cost function

$$c(q) = \frac{5}{16}\left(q^2 - \sin^2 q\right).$$

In Figure 2, we plot the optimal continuous stack $p = c'(q)$ and three stacks with four tranches, each of which satisfies the stationarity conditions.

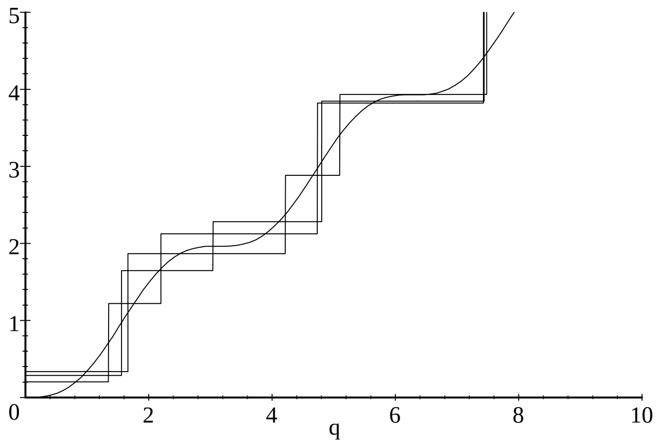

Figure 2: Stack satisfying stationarity conditions

It seems to be rather challenging to find an algorithm that will automatically find the absolute best offer stack for more complicated cost functions. One option is to search along possible p_1 values, in search of an offer stack that ends on the boundary ($p_{m+1} = p_M$ or $q_m = q_M$).

However, the stacks plotted in Figure 2 appear to be quite sensitive to the initial value p_1, and are only clearly distinct on the middle tranches. This suggests an alternative approach, which is to consider a stack of say $\lfloor \frac{m}{2} \rfloor$ tranches starting from an unknown p_1 and a stack of $\lceil \frac{m}{2} \rceil$ tranches starting from q_m. Thus we would replace a global optimization problem over one variable, with a function sensitive to that variable, with a two variable problem, with less sensitivity to the variables.

5 Piecewise Linear Cost Functions

We now consider the special case where the cost function is piecewise linear. This would be the case in a model that has a number of turbines that can be switched on or off, and the marginal cost of electricity is constant for a constant number of turbines switched on.

Note that as the cost function is nonsmooth, the condition (7) is generalized and for the piecewise linear case yields

$$2c'_-(q_i) - p_i \le p_{i+1} \le 2c'_+(q_i) - p_i,$$

where $c'_-(q_i)$ is the left derivative of c at q_i and $c'_+(q_i)$ is the right derivative of c at q_i. As before, $p_1 = \frac{c(q_i)-c(q_{i-1})}{q_i-q_{i-1}}$.

We will now prove that in this case, we need only consider stacks where the tranche quantities lie at the breakpoints between the linear sections of the cost function. We then give an illustrative example in Section 5.1 and then present a dynamic programming approach to solving this special class of problems in Section 5.2.

Lemma 5.1 *If $c(q)$ is a continuous, piecewise linear function, then $P(m)$ has a solution $\{(q_1, p_1), \ldots, (q_m, p_m)\}$, in which for all $i = 1, \ldots, m$, q_i lies on a breakpoint of $c(\cdot)$ (that is $c'(q_i)$ does not exist).*

Proof. Since c is continuous, there exists an optimal stack $S = \{(q_1, p_1), \ldots, (q_m, p_m)\}$. Suppose for some i that q_i is not a breakpoint of $c(\cdot)$. Choose the largest such q_i.

Since S is optimal,

$$c'(q_i) = \frac{\int_{p_i}^{p_{i+1}} pf(p)dp}{\int_{p_i}^{p_{i+1}} f(p)dp}. \tag{11}$$

We now consider an alternative stack S', where (q_i, p_i) is replaced by (q^+, p_i), with q^+ a breakpoint and

$$c(q^+) = c(q_i) + c'(q_i)(q^+ - q_i). \tag{12}$$

Thus, q^+ is the next breakpoint, and q_i was the largest quantity not at a breakpoint, so (if $i < m$) $q^+ \le q_{i+1}$ and $q^+ \le q_M$, so S' is feasible for $P(m)$. Let the total return from stack S be $R(S)$ and likewise the return from stack S' be $R(S')$. Then

$$R(S) - R(S') = \int_{p_i}^{p_{i+1}} q_i p - c(q_i) f(p)dp - \int_{p_i}^{p_{i+1}} q^+ p - c(q^+) f(p)dp.$$

From (11), $\int_{p_i}^{p_{i+1}} pf(p)dp = c'(q_i) \int_{p_i}^{p_{i+1}} f(p)dp$, so using this, and (12),

$$\begin{aligned} R(S) - R(S') &= -(q^+ - q_i) \int_{p_i}^{p_{i+1}} pf(p)dp + \int_{p_i}^{p_{i+1}} c(q^+) - c(q_i) f(p)dp. \\ &= -(q^+ - q_i)c(q_i) \int_{p_i}^{p_{i+1}} f(p)dp + (q^+ - q_i)c(q_i) \int_{p_i}^{p_{i+1}} f(p)dp \\ &= 0. \end{aligned}$$

Thus S' has the same total return as S. By applying the same argument to S' a finite number of times, it can be seen that there exists a stack of at most m tranches, with all q_i lying on breakpoints of $c(\cdot)$, yielding the same return as S, and hence solving $P(m)$. ∎

A look at the Hessian (8) reveals that most stacks satisfying the stationarity conditions with any tranche quantity q_i not at a breakpoint must be a saddle point. We note that $c''(q_i) = 0$, and hence the negative of the Hessian matrix contains a principal minor

$$\begin{bmatrix} (q_i - q_{i-1}) f(p_i) & (p_i - c'(q_i)) f(p_i) \\ (p_i - c'(q_i)) f(p_i) & 0 \end{bmatrix},$$

which has determinant $-[(p_i - c'(q_i)) f(p_i)]^2$. We have assumed that $f(p) > 0$ for all p, so if $p_i < c'(q_i)$, $-H$ has a principal minor with negative determinant and thus cannot be positive semi-definite. So any stack satisfying the stationarity conditions, but containing q_i cannot be a local maximum. Note that if $p_i = c'(q_i)$, then higher order conditions would be required, but the proof above still holds.

5.1 An Example

We now consider an example where the cost function is piecewise linear. We consider an example where the price is uniformly distributed over [0,5], $q_M = 4$ and the cost function is defined by:

$$c(q) = \begin{cases} q & , \quad 0 \le q \le 1 \\ 2q - 1 & , \quad 1 \le q \le 2 \\ 3q - 3 & , \quad 2 \le q \le 3 \\ 4q - 6 & , \quad 3 \le q \le 4 \end{cases}$$

An exhaustive search yields the following solutions satisfying the optimality conditions.

Soln 1 : $q_1 = 1, p_1 = 1, p_2 = 3, q_2 = 4 : F = 2.8$

Soln 2 : $q_1 = 2, p_1 = 1.5, p_2 = 3.5, q_2 = 4 : F = 2.9$

Soln 3 : $q_1 = 3, p_1 = 2, p_2 = 4, q_2 = 4 : F = 2.8$

Soln 4 : $q_1 = 1, p_1 = 1, p_2 = \frac{5}{2}, q_2 = 3 : F = 2.85$

Soln 5 : $q_1 = 2, p_1 = \frac{3}{2}, p_2 = 3, q_2 = 3 : F = 2.85$

Soln 6 : $q_1 = \frac{3}{2}, p_1 = \frac{4}{3}, p_2 = \frac{8}{3}, q_2 = 3 : F = 2.8\dot{3}$

Note that Soln 6 has tranche quantities that do not lie on the breakpoints of the cost function, but does not yield the maximum return. The optimal stack is given by Soln 2.

5.2 A Solution Technique

We have determined that piecewise linear cost functions can be solved by considering combinations of finitely many quantities. We now present a method for finding the optimal stack for this class of problems.

Suppose that the breakpoints of the piecewise linear cost function are at $0 = \theta_0$, θ_1, θ_2, ..., $\theta_N = q_M$. If $m \geq N$ then the optimal offer stack is the marginal cost function, as long as this is monotonic.

If the marginal cost function is not monotonic then we seek a set of m points from the breakpoints to form the breakpoints of the optimal offer stack. For $j = 1, \ldots, N$, let $\Pi(j)$ be the set of possible prices at which we could offer θ_j in a stationary stack. Formally

$$\Pi(j) = p_M \cup \{\frac{c(\theta_j) - c(\theta_i)}{\theta_j - \theta_i} \mid 0 \leq i < j\}.$$

To construct an optimal stack we list the elements of the set $\cup_j \Pi(j)$ in order of increasing size to give

$$\cup_j \Pi(j) = \{\pi_k \mid k = 1, 2, \ldots, K\}.$$

Define $V_n(j)$ to be the maximum extra contribution one could make to the expected return by constructing a stationary stack that started at the point (θ_n, π_j). The optimal stack will be that which maximises $V_0(j)$. This can be computed using the following dynamic programming recursion:

$$V_N(k) = \begin{cases} \int_{\pi_k}^{p_M} [p\theta_N - c(\theta_N)]\, f(p) dp & , \quad \pi_k \in \Pi(N), \\ -\infty & , \quad \text{otherwise,} \end{cases}$$

$$V_n(j) = \begin{cases} \max_{k>j, r>n} \int_{\pi_j}^{\pi_k} [p\theta_j - c(\theta_j)]\, f(p) dp + V_r(k) & , \quad \pi_j \in \Pi(n), \\ -\infty & , \quad \text{otherwise.} \end{cases}$$

This recursion can be applied when there are no restrictions on the number of tranches. If $m < N$, the recursion must be modified to include an extra state variable that records the number of tranches used and sets $V_n(j)$ to be $-\infty$ for any partial solution that has more than m tranches. An implementation of this recursion and its application to a number of simple examples is described in more detail in [2].

Acknowledgement

The authors would like to acknowledge the support of FRST under contract UOA803.

References

[1] A. Philpott, E. Anderson and J. Kaye, "Optimal participant behaviour in electricity markets," Proceeding of the ORSNZ Conference, 1998.

[2] N. Roberts, "Multi-period optimisation of hydro-electric generator offers," Year 4 Project Report, Department of Engineering Science, University of Auckland, 1999.