The electricity seller's Psi

G. Pritchard, A. Philpott, P. Neame, G. Zakeri
Dept. of Statistics
University of Auckland
New Zealand
g.pritchard@auckland.ac.nz

Abstract

A producer of electricity, or indeed any other commodity, may be faced with the problem of formulating an "offer curve" representing the quantity that it is willing to supply at any given price. We consider this problem as one of stochastic variational programming - find the offer curve which maximizes expected profit, given stochastic market conditions (demand and competitor behaviour). Sufficient information to solve the problem is contained in the market distribution (or "Ψ") function, which gives the probability that the market can absorb particular (quantity, price) pairs. We discuss ways to estimate Ψ functions, using both parametric and non-parametric models. More complex offer-curve optimization problems can arise when offering electricity into a transmission network. In such markets, there is usually a different price at every node. Even if a producer has only one generating plant, it will often have an interest in prices at other nodes, through financial contracts or retail customers located at those nodes. We present some examples of such problems, and a method for solving them.

1. Introduction

Consider the following problem faced by a producer in a commodity industry. The producer must formulate an *offer curve* (or *supply curve*) $\mathbf{s} = ((q(t), p(t))\ , 0 \leq t \leq T)$ representing the quantity $q(t)$ that it is willing to supply at a corresponding price $p(t)$. (The parametric description is chosen to allow the curve to have both horizontal and vertical segments.) Assume that this curve is required to be continuous, and $q(\cdot)$ and $p(\cdot)$ must both be non-decreasing in t.

A random *point of dispatch* (q, p) on the curve will then be chosen. This point is determined by events in the "rest of the market" and is outside the control of this producer, although its distribution is assumed to be known. A payoff $R(q, p)$ to the producer then results. The problem is to choose the offer curve $\mathbf{s}$ to maximize the producer's expected payoff.

This work was partly supported by New Zealand Public Good Science Fund grant UOA 803.

Key words and phrases: supply function, electric power, financial transmission rights, dynamic programming, stochastic programming

For some industries, the optimal-offer-curve problem belongs to a fairly abstract class of econometric models. However, the context of this paper is intended to be a pool-type electric power market. In the operation of such a market, each power generator must regularly submit an offer curve to a centralized coordinating authority, which then determines the point of dispatch on each curve in such a way as to optimize the operation of the overall system. The optimal-offer-curve problem thus has an immediate and concrete relevance.

The problem is one of stochastic optimization, made interesting by (among other things) the fact that the decision variable $\mathbf{s}$ is infinite-dimensional. This feature arises as a result of the stochastic element. In the special case where the point of dispatch is deterministic, the problem is essentially the standard monopolist's (or oligopolist's) profit-optimization problem from elementary economics (see e.g. [3]). In this case, it is not really necessary to consider the whole offer curve, but only the single point which will be the point of dispatch.

In another important special case, the producer is a price-taker – there is a random market price with a distribution unaffected by the offer curve submitted. If the producer has a convex cost function $C(q)$ giving the cost to produce each quantity q, then there is a straightforward solution: each $p(t)$ should be the marginal cost $C'(q(t))$ of production at quantity $q(t)$. Note that this is independent of the distribution of the market price.

Our main interest in this paper is with situations where the producer has some degree of market power, but still faces a random point of dispatch due to random behaviour by competitors and consumers.

The paper is laid out as follows. In section 2 we introduce the market distribution function as defined by Anderson and Philpott in [1], and show how it contains the necessary information to solve the optimal offer-curve problem. Then in section 3 we discuss the estimation of market distribution functions from historical price and volume data. The final section considers generalizations of the problem which occur in the context of an electricity market with locational pricing, giving different prices at different nodes in the network.

2. The market distribution function

The stochastic behaviour of the "rest of the market" can be succinctly described by the *market distribution function* Ψ, as defined in [1]. Let $\Psi(q,p)$ denote the probability that, if $\mathbf{s}$ passes through (q,p), the point of dispatch on $\mathbf{s}$ will lie below (q,p) (i.e. will be a point (q^*,p^*) with $q^* \leq q$ and $p^* \leq p$).

It is necessary to reassure ourselves that $\Psi(q,p)$ is well-defined. For this we must make the assumption:

(A) : If offer curves $\mathbf{s}$, $\mathbf{s}'$ both pass through (q,p), then the point of dispatch on $\mathbf{s}$ will lie below (q,p) if and only if the point of dispatch on $\mathbf{s}'$ does.

Assumption (A) will hold whenever the market into which our producer sells has non-increasing marginal utility for the commodity. That is to say, the market implicitly chooses a random non-increasing function $U(q)$ giving its marginal utility (and hence willingness to pay) for quantity q from this producer. The actual construction of $U(\cdot)$ will depend on market structure, but $U(\cdot)$ is ultimately a representation of the consumer demand function and competing offers from other producers. The point of dispatch chosen for an offer curve $\mathbf{s}$ will then be the point where $\mathbf{s}$ intersects the graph of $p = U(q)$. (If the intersection is a horizontal or vertical line segment, an arbitrary choice is required; for definiteness take the left-hand or lower endpoint as the point of dispatch.)

Another case in which (A) can be verified occurs when the point of dispatch arises as the solution to a convex optimization problem which is (implicitly or explicitly) solved by the market. This occurs in a pool-type electricity market – see [1].

It is apparent that $\Psi(q,p)$ takes values in $[0,1]$, and is a non-decreasing function of both q and p.

Let $R(q,p)$ denote the payoff to the producer when the point of dispatch is (q,p). If this includes immediate revenue only, then $R(q,p) = qp$. If there is a cost $C(q)$ to produce q, then $R(q,p) = qp - C(q)$. It is common in the electric power industry to have some fixed quantity q_0 of output effectively pre-sold (through "contracts for differences"); in this case $R(q,p) = (q - q_0)p$. Another possibility is to use R to represent risk aversion.

Once a market distribution function Ψ and a payoff function R are specified, it is possible in principle to solve the optimal offer curve problem by an analytic procedure described in [1]. The expected payoff when the offer curve $\mathbf{s}$ is submitted is given by the line integral

$$V(\mathbf{s}) = \int_{\mathbf{s}} R(q,p)\; d\Psi(q,p).$$

The optimal offer curve problem can then be written

$$\begin{aligned} &\max V(\mathbf{s}) \\ &\text{ s.t. } \mathbf{s} = ((q(t), p(t))\,,\, 0 \le t \le T) \\ &\qquad q(\cdot), p(\cdot) \text{ non-decreasing} \\ &\qquad 0 \le q(t) \le q_M, \end{aligned}$$

where q_M is the maximum quantity that can be produced. This can be viewed as a problem in the calculus of variations. It can be shown that if $\mathbf{s}$ is stationary with respect to feasible variations, then it must at every point either be vertical or horizontal, or satisfy the first-order condition

$$\frac{\partial R}{\partial q}\frac{\partial \Psi}{\partial p} - \frac{\partial R}{\partial p}\frac{\partial \Psi}{\partial q} = 0.$$

2.1 Example.

Suppose we have a single-price market for the commodity, and that production q_1 by the "rest of the industry" is related to the market price p via the equation $q_1 = S(p)$, where $S(\cdot)$ is a deterministic, known function. Assume also that demand for the commodity is a random variable D with a probability density function f, and is inelastic with respect to price. Then

$$\Psi(q,p) = P\left(D < q + S(p)\right) = \int_0^{q+S(p)} f(\eta)\; d\eta$$

and so

$$\frac{\partial \Psi}{\partial p} = f(q + S(p)) \qquad \text{and} \qquad \frac{\partial \Psi}{\partial q} = f(q + S(p))S'(p).$$

If the payoff function is simply $R(q, p) = qp$, then

$$\frac{\partial R}{\partial q}\frac{\partial \Psi}{\partial p} - \frac{\partial R}{\partial p}\frac{\partial \Psi}{\partial q} \quad = \quad f(q + S(p))(pS'(p) - q).$$

Hence the optimal offer curve is described by the equation $q = pS'(p)$. (The solution is degenerate in regions where $f(q+S(p)) = 0$, since the point of dispatch can never lie in such a region.)

3. Estimation of Ψ.

In order for the market distribution function theory to be useful in practice, a reasonable estimate of the Ψ function must be available. There are several ways to estimate Ψ from historical data. Suppose observations (q_i, p_i) $(i = 1, \ldots, n)$ of past points of dispatch are available. It is not necessary for all the data to arise from the same offer curve; indeed, it is probably better that a variety of offer curves be represented, as this gives better coverage of the (q, p) plane. In real markets, the admissible offer curves are almost always step functions, i.e. they consist of a number of horizontal and vertical line segments. Let $g_i \in \{q, p\}$ be a variable which indicates whether the dispatch (q_i, p_i) occurred on a horizontal $(g_i = q)$ or vertical $(g_i = p)$ segment of its corresponding offer curve.

3.1 Non-parametric estimation of Ψ

If the observations are independent, a likelihood function for the observed data is

$$\prod_{i=1}^{n} \frac{\partial \Psi}{\partial g_i}(q_i, p_i).$$

Attempting to find a maximum-likelihood estimator suggests the following construction. In the (q, p) plane, draw a vertical line through each point (q_i, p_i) with $g_i = q$, and a horizontal line through each point (q_i, p_i) with $g_i = p$. This divides the plane into a number of rectangular cells. Take Ψ to be constant on each cell. There are now only a finite number of Ψ values to be determined. Let γ_i^0 denote the value of Ψ on the cell to the left of (if $g_i = q$) or below (if $g_i = p$) the point (q_i, p_i). Let γ_i^1 denote the value of Ψ on the cell to the right of (if $g_i = q$) or above (if $g_i = p$) the point (q_i, p_i). Now choose the cell values to solve

$$\begin{aligned} \max \ & \prod_{i=1}^{n} (\gamma_i^1 - \gamma_i^0) \\ \text{s.t.} \ & \text{all cell values are in } [0, 1] \\ & \Psi \text{ is non-decreasing in both } q \text{ and } p \end{aligned}$$

Figure 1 shows the result of an application of this technique to a data set supplied by a large New Zealand electric power generation company, with shading showing the different values of the estimated Ψ function.

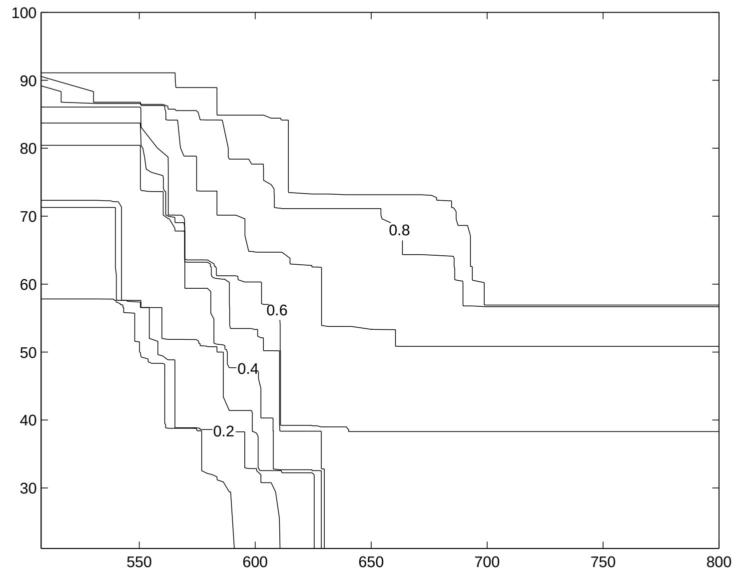

Figure 1: A Ψ function estimated by the non-parametric method.

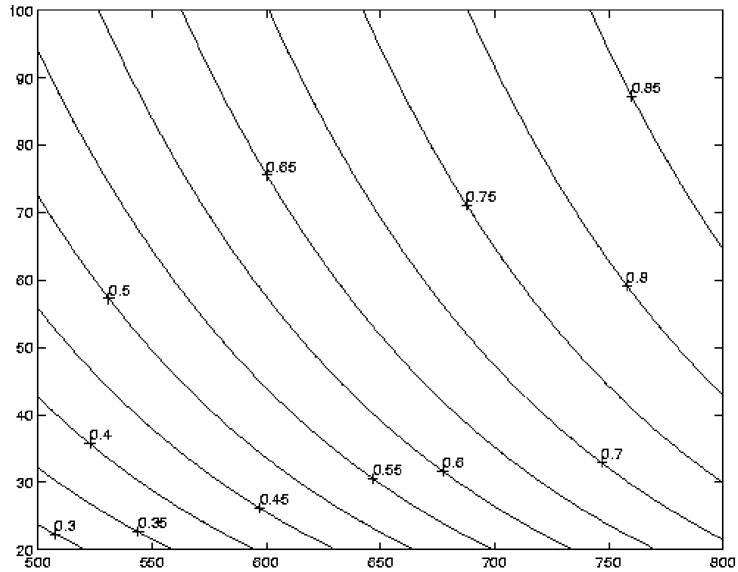

Figure 2: A Ψ function estimated by the Bayesian parametric method, from the same data as Figure 1.

3.2 Bayesian parametric estimation of Ψ

The market distribution function could also be modelled parametrically as belonging to a family $(\Psi^\alpha : \alpha \in A)$ of possible market distribution functions. Apart from the natural maximum-likelihood estimators, this framework also lends itself well to a Bayesian approach, in which the estimated Ψ function can be successively updated as new observations become available. Given a prior density $(\phi(\alpha) : \alpha \in A)$, this density can be updated with the non-normalized Bayes factor

$$g(\alpha) = \prod_{i=1}^{n} \frac{\partial \Psi^\alpha}{\partial g_i}(q_i, p_i).$$

Anderson and Philpott propose in [2] the two-parameter model

$$\Psi^{\alpha,\beta}(q,p) = P\left(\pi^{\alpha,\beta}(q) < p\right),$$

where $\pi^{\alpha,\beta}(q)$ (the maximum price at which quantity q can be sold) is lognormally distributed, with $\log \pi^{\alpha,\beta}(q)$ having a normal distribution with mean $\beta - \alpha q$ and constant variance σ^2.

Figure 2 shows the result of applying this model to the same data set as in the subsection above.

4. Market distribution functions for network problems.

More complex offer-curve optimization problems can arise when offering electricity into a transmission network. In such markets, there is usually a different price at each node of the network. Even if a power-generator has only one plant, it will often have an interest in prices at other nodes, through financial contracts or retail customers located at those nodes. Such a situation requires a generalization of the concept of the market distribution function.

Suppose that our producer has a payoff function $R(q,p,p')$, depending not only on the point of dispatch (q,p), but also on the price p' prevailing at some other node in the network. Define functions $\Psi_{qp'}$ and $\Psi_{pp'}$ as follows. Let $\Psi_{qp'}(q,p,p')\ dq\ dp'$ be the probability that the point of dispatch would fall along an infinitesimal horizontal segment of offer curve from (q,p) to $(q+dq,p)$, while the price at the other node would fall between p' and $p'+dp'$, were such a horizontal segment present. Similarly, let $\Psi_{pp'}(q,p,p')\ dp\ dp'$ be the probability that the point of dispatch would fall along an infinitesimal vertical segment from (q,p) to $(q,p+dp)$, while the price at the other node would fall between p' and $p'+dp'$, were such a vertical segment present. These functions are analogous to mixed partial derivatives of the Ψ function as originally defined.

The expected payoff from an offer curve $\mathbf{s}$ can then be written

$$V(\mathbf{s}) = \int_{-\infty}^{\infty} \int_{\mathbf{s}} R(q,p,p')\ (\Psi_{qp'}\ dq\ +\ \Psi_{pp'}\ dp)\ dp',$$

while the first-order condition for optimality becomes

$$\int_{-\infty}^{\infty} \left(\frac{\partial R}{\partial q} \Psi_{pp'} - \frac{\partial R}{\partial p} \Psi_{qp'} \right)\ dp'\ =\ 0.$$

4.1 Example.

Consider a three-node grid in which "we" are the only generator large enough to influence prices.

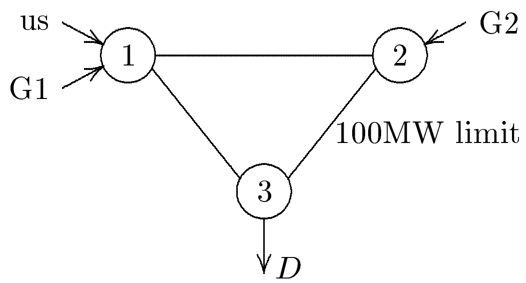

A competitive fringe consisting of many smaller generators provides additional supply "G1" and "G2"; which is offered to the market via fixed aggregated offer curves. For simplicity, assume that these take the forms $q = \beta_1 p$ and $q = \beta_2 p$ at nodes 1 and 2 respectively. Demand D is located at node 3 and is random, with a probability density function f. Transmission lines are lossless, of unlimited capacity (except for the 100MW limit on the line between nodes 2 and 3), and have equal admittances. (The significance of this last point is that 1/3 of any power injected at node 1, and 2/3 of any power injected at node 2, must flow via the limited-capacity line to reach the load.) Suppose that we can generate power at no marginal cost, and that we own a financial transmission right – a type of financial contract – which will pay us q_f megawatts times the price difference between nodes 1 and 3. Our payoff function is thus $R(q,p,p') = pq + q_f(p'-p)$, where p denotes the price at the local node 1, and p' the price at the consuming node 3.

It can be shown by analysis of this network that

$$\Psi_{qp'} = \begin{cases} f(q+(\beta_1+\beta_2)p)\ \delta_p(p'), & \text{if } q+(\beta_1+2\beta_2)p < 300 \\ f((q+\beta_1 p+300)/2)\ (1/2)\ \delta_{2p-(300-q-\beta_1 p)/(2\beta_2)}(p'), & \text{if } q+(\beta_1+2\beta_2)p \geq 300 \text{ and } q+\beta_1 p \leq 300 \end{cases}$$

and

$$\Psi_{pp'} = \begin{cases} f(q+(\beta_1+\beta_2)p)\ (\beta_1+\beta_2)\ \delta_p(p'), & \text{if } q+(\beta_1+2\beta_2)p < 300 \\ f((q+\beta_1 p+300)/2)\ (\beta_1/2)\ \delta_{2p-(300-q-\beta_1 p)/(2\beta_2)}(p'), & \text{if } q+(\beta_1+2\beta_2)p \geq 300 \text{ and } q+\beta_1 p \leq 300 \end{cases}$$

Our offer curve must therefore pass through two distinct regions, labelled I and II on the diagram below. In region II, the transmission line between nodes 2 and 3 will be at capacity; in region I, it will not. The first-order optimality condition reduces to $q = (\beta_1+\beta_2)p$ in region I and $q = \beta_1 p - q_f$ in region II, independently of f. Allowing for the monotonicity constraint, this gives the optimal offer curve as the upper curve in the diagram below. (The value of q_T will in general depend on f.) Also shown, for comparison, is the optimal offer curve when $q_f = 0$, i.e. in the absence of the transmission right. It is interesting to note that ownership of the right requires a more aggressive offering strategy.

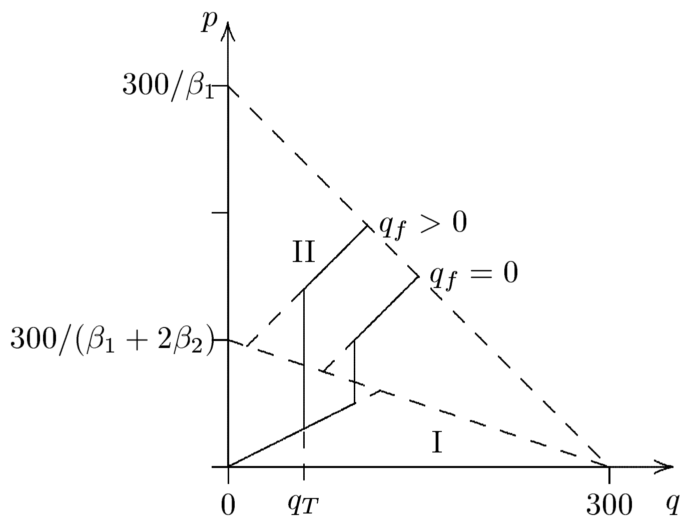

5. Optimal offer curves by dynamic programming.

In larger electric-power offering problems, it may be impractical to formulate an explicit Ψ function as in the previous section. It is still possible, however, to take a simulation approach to the optimal offer curve problem. In this section we outline such an algorithm.

Let us begin by subdividing the p-q plane with a finite rectangular grid, and restricting the class of admissible offer curves to those which follow the edges of this grid (i.e. an offer curve must intersect each cell of the grid only in its boundary). Suppose that for each edge $\mathbf{e}$ of a cell, we know the expected payoff $V(\mathbf{e})$ due to points of dispatch occurring on $\mathbf{e}$ that would be realized if $\mathbf{e}$ were part of the offer curve. That is, we know the quantity $\int_{\mathbf{e}} R(q,p)\, d\Psi(q,p)$, or a higher-dimensional equivalent of this if the payoff function involves prices at other nodes.

The problem of finding the best admissible offer curve is then a finite one, which can be solved by a dynamic programming algorithm. Suppose our grid covers the region $0 \le q \le q_{\max}$, $p_{\min} \le p \le p_{\max}$, and that all possible points of dispatch lie within this region. For each vertex x of the grid, let $W(x)$ denote the maximal expected payoff, due to points of dispatch above x, of any offer curve which passes through x. Then

$$W(x) = \begin{cases} 0, & \text{if } x \text{ is on } q = q_{\max} \text{ or } p = p_{\max} \\ \max(V(e_r(x)) + W(\nu_r(x)), V(e_u(x)) + W(\nu_u(x))), & \text{otherwise,} \end{cases}$$

where $\nu_r(x)$ and $\nu_u(x)$ are, respectively, the vertices adjacent to x to the right and above, and $e_r(x)$ and $e_u(x)$ the edges linking to those neighbours. It is thus straightforward to successively compute $W(x)$ for each vertex x of the grid, and hence to determine the optimal admissible offer curve. This can be taken as an approximation of the optimal offer curve for the original problem.

Rather than compute the edge values $V(\mathbf{e})$ as exact integrals, we can estimate them by a sampling procedure. Take n randomly chosen scenarios $\omega_1, \dots, \omega_n$, where each "scenario" is a realization of the random elements of the problem (competitors' offer curves, loads, outages, etc.). For each scenario i and each edge $\mathbf{e}$, let $V_i(\mathbf{e})$ be the payoff that results if ω_i results in a point of dispatch along $\mathbf{e}$, or 0 if no such dispatch can occur. Finally, approximate each $V(\mathbf{e})$ by $\hat{V}(\mathbf{e}) = n^{-1}\sum_{i=1}^{n} V_i(\mathbf{e})$. Note that $\hat{V}(\mathbf{e})$ is a consistent unbiased estimator of $V(\mathbf{e})$. Solving the optimal offer curve problem for the sampled edge values $\hat{V}(\mathbf{e})$ is a form of sample-path optimization.

To illustrate this, the diagram below shows the optimal offer curves for a version of the three-node problem of the previous section. The parameters here are $\beta_1 = 5$ MW^2/\$, $\beta_2 = 15$ MW^2/\$, a 50MW financial transmission right, and demand distributed normally with mean 200MW, standard deviation 30MW. The curves shown are the optimal offer curve for the dynamic programming approximation (stepped curve) and for the original problem.

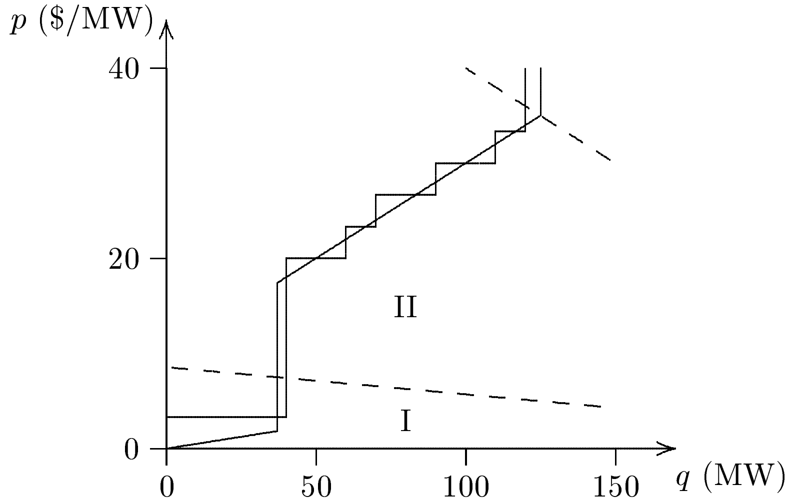

6. Conclusion.

We have considered offer-curve optimization problems from both a theoretical and a more applied perspective. The market distribution function provides a sound framework for solving such problems analytically; however, it is possible to carry this through only in the simplest cases.

In practice, most electricity generation firms will also have their own retail customers; they are therefore both producers and consumers of electricity, and must account for different prices at different nodes in a transmission grid. In addition, they will usually have financial hedge contracts of various kinds contributing large terms to the payoff function. We have indicated how the Ψ function can be generalized to cover such situations, by allowing for more than one price.

If a sufficiently detailed model of the grid and the competition is available, the sampling and dynamic-programming technique of section 5 allows offer-curve optimization problems to be solved. The algorithm would probably be successful on quite large problems, although the authors have yet to try this.

References

[1] E. Anderson and A. B. Philpott, *Optimal offer construction in electricity markets,* To appear in Mathematics of Operations Research.

[2] E. Anderson and A. B. Philpott, *Estimation of electricity market distribution functions,* In preparation.
[3] D. R. Fusfeld, *Economics,* Heath, 1976.