Estimation of electricity market distribution functions

E.J. Anderson * A.B. Philpott †‡

Abstract

In an electricity pool market the market distribution function gives the probability that a generator offering a certain quantity of power at a certain price will not be dispatched all of this quantity by the pool. It represents the uncertainty in a pool market associated with the offers of the other agents as well as demand. We present a general Bayesian update scheme for market distribution functions. To illustrate the approach a particular form of this procedure is applied to real data obtained from a New Zealand electricity generator.

Keywords: electricity markets, optimization, Bayesian statistics

*Australian Graduate School of Management, University of New South Wales, Australia.

†Department of Engineering Science, University of Auckland, New Zealand.

‡This work has been supported by the Australian Research Council and the New Zealand Public Good Science Fund.

1 Introduction

In recent years, wholesale markets for electricity generation and distribution have emerged in various regions around the world. Although implementations of these markets vary, they all endeavour to deliver electric power to consumers at a competitive price, and provide sensible signals for investment and new entry. For an overview of different market designs see e.g. [2], [6]. This paper focuses on a specific type of market structure called an *electricity pool*, characterized by a central dispatch and pricing mechanism. In a pool market the price of electricity in each trading period is determined by solving an optimization problem that matches supply and demand so as to minimize the total revealed cost of power delivery. The first electricity market of this type to be created was for England and Wales, but similar markets now operate in Australia, New Zealand, Scandinavia and some parts of Latin America and North America.

In this paper we consider the problem faced by a generator bidding into a wholesale electricity pool market in circumstances where the demand for electricity and the behaviour of other generators are to some extent uncertain. In a pool market, generators are required to submit offers to an independent system operator in the form of a supply function, specifying the amount that they are prepared to supply to the market at each price level. In many markets (Australia and New Zealand, for example) the supply functions are submitted in the form of a step function called an *offer stack*. The offer stack defines a finite set of quantities to be offered and the price per MWhr to be asked for each. It is convenient to model the offers as a continuous curve $\mathbf{s} = \{(x(t), y(t)),\ 0 \leq t \leq T\}$, in which the components $x(t)$ and $y(t)$ are monotonic increasing piecewise differentiable functions of t. Here $x(t)$ traces the quantity component of the offer curve and $y(t)$ traces the price component. In each trading period of a pool market, the independent system operator dispatches the offers of every generator in order of increasing price until demand is met. The electricity market clearing price is the offer price of the last generator dispatched.

A number of authors (see e.g. [3], [4], [5]) have developed models of electricity prices. The majority of these models are focused on derivative pricing, and so they are generally constructed from continuous-time stochastic processes that are modified to reflect the changing volatility and occasional price spikes observed in electricity spot prices. As observed in [5] most of the classical financial asset-pricing models are inadequate when dealing with electricity spot prices. A further weakness of these price models is that they assume that prices are exogenous, and so they cannot be used to optimize the offers of a generator that has sufficient market power to influence the price by its offering strategy.

One approach to modelling the effect of generator strategies on electricity prices is provided by the recent paper of Anderson and Philpott [1]. They define the *market distribution function* $\psi(q,p)$ for a generator G, to be the probability that G is not fully dispatched by the market if it makes a single offer to generate an amount q at price p. (G is not fully dispatched by the market if a plot of the residual demand function passes below and to the left of the point (p,q)). As shown in [1], if $\psi(q,p)$ is known then the generator can compute an offer stack $\mathbf{s}$ that maximizes expected profit by maximizing the line integral

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

where $R(q,p)$ is the profit when the generator is dispatched amount q at a price p. In practice, the function $\psi(q,p)$ is not known, and must be estimated from observations of market behaviour. In this paper we describe a general Bayesian estimation procedure for market distribution functions. The exact form that this procedure should take in practice will depend upon the market structure, and the amount of information that this reveals to participants as they trade. In practice we will expect both demand and bidding behaviour to vary with the time of day, and so our estimate of $\psi(q,p)$ will be with respect to a particular time (or set of times), and new information will only become available once a day.

The paper is laid out as follows. In the next section we derive a general Bayesian update formula for participants in pool markets. This model assumes that the generator knows its own dispatch quantity and clearing price following every offer, but does not have any other information. In section 3 we show how this model takes a special form when all generators are located at the same node, and all the uncertainty arises from the demand at this node. We then look at the situation in which participants are located at different nodes of a transmission network, so line losses and constraints alter the effect of their actions at the generator's node. In our Bayesian model the market distribution function will represent these effects in some probabilistic fashion. We show how a generator can construct a posterior estimate of the market distribution function from observations about its own dispatch. Finally we illustrate the updating procedure by applying it to some data provided by a New Zealand generating company.

2 Bayesian estimation of ψ

We shall describe a Bayesian approach to estimate the function $\psi(q,p)$. The first step is to decide on a parameterised family of market distribution functions, where the Bayesian updates will determine a distribution over the parameters. We write

$$\psi(q,p) = \int_A \psi^\alpha(q,p)\phi(\alpha)d\alpha \tag{1}$$

where $\phi(\alpha)$ is a (possibly multivariate) density over a parameter space A, and $\{\psi^\alpha(q,p)\}$ is a family of basis functions parameterised by $\alpha \in A$. We require that each $\psi^\alpha(q,p)$ has the properties of a market distribution function (i.e. it is monotonic nondecreasing in both arguments, and lies between 0 and 1.) To derive the Bayesian update formula suppose that at some period i the generator submits a step function offer stack $\mathbf{s}$, the price turns out to be p_i and the generator is dispatched q_i. Let E_δ be the event that (q_i, p_i) falls on the section of $\mathbf{s}$ between $(x(\tau^*), y(\tau^*))$ and $(x(\tau^*+\delta), y(\tau^*+\delta))$. Then

$$\begin{aligned}
\Pr[E_\delta|\alpha] &= \psi^\alpha(x(\tau^*+\delta), y(\tau^*+\delta)) - \psi^\alpha(x(\tau^*), y(\tau^*)) \\
&= \psi^\alpha(x(\tau^*+\delta), y(\tau^*+\delta)) - \psi^\alpha(x(\tau^*), y(\tau^*+\delta)) \\
&\quad + \psi^\alpha(x(\tau^*), y(\tau^*+\delta)) - \psi^\alpha(x(\tau^*), y(\tau^*)) \\
&= \delta\frac{\partial\psi^\alpha}{\partial q}(x(\tau^*), y(\tau^*))x'(\tau^*) + \delta\frac{\partial\psi^\alpha}{\partial p}(x(\tau^*), y(\tau^*))y'(\tau^*) + o(\delta^2)
\end{aligned}$$

Writing $\tilde{\phi}(\alpha)$ for the posterior density function for α, we have

$$\begin{aligned}
\tilde{\phi}(\alpha) &= \lim_{\delta\to 0} \Pr[E_\delta|\alpha]\frac{\phi(\alpha)}{\Pr[E_\delta]} \\
&= \frac{\phi(\alpha)g(\alpha)}{\int_{\alpha_0}^{\alpha_1}\phi(\alpha)g(\alpha)d\alpha}
\end{aligned}$$

where

$$g(\alpha) = \frac{\partial\psi^\alpha}{\partial q}(x(\tau^*), y(\tau^*))x'(\tau^*) + \frac{\partial\psi^\alpha}{\partial p}(x(\tau^*), y(\tau^*))y'(\tau^*).$$

Observe that the normalisation of $\tilde{\phi}$ (dividing by the integral) need only be applied at the end of a sequence of updates. For example, given observations (q_i, p_i), $i = 1, 2, \ldots, k$, we have

$$\tilde{\phi}(\alpha) \propto \phi(\alpha) \prod_{i=1}^{k} [\frac{\partial \psi^\alpha}{\partial q}(q_i, p_i) x'(\tau^*) + \frac{\partial \psi^\alpha}{\partial p}(q_i, p_i) y'(\tau^*)].$$

The posterior estimation of the function $\psi(q, p)$ as in (1) is just

$$\psi(q, p) = \int_A \psi^\alpha(q, p) \tilde{\phi}(\alpha) d\alpha.$$

3 Updates using a demand distribution

The procedure outlined in the previous section takes a simpler form if all participants are located at the same node, at which there is a demand h with known probability density function $f(h)$. As above we suppose that at some period i the generator submits an offer stack $\mathbf{s}$, the price turns out to be p and the generator is dispatched q. The first step is to decide on a parameterised family of market distribution functions, where the Bayesian updates will determine a distribution over the parameter. In the case where the rest of the market is modelled by a supply function S, and demand is inelastic, the market distribution function $\psi(q, p)$ can be expressed as $\Pr(h < q + S(p))$. (Although we confine attention to inelastic demand, our model can also represent the case where demand is defined by a known demand function $D(p)$ with random "shocks" h, giving $\psi(q, p) = \Pr(h < q + S(p) - D(p)))$. Since the distribution of h is assumed to be known it makes sense to use this distribution rather than estimate it (implicitly or explicitly) within the Bayesian framework.

Our framework uses a family of continuous (rest-of-market) supply functions $S(\alpha, p)$, parameterised by $\alpha \in [\alpha_0, \alpha_1]$. Thus

$$\psi^\alpha(q, p) = \Pr(h < q + S(\alpha, p)),$$

where we are given a prior distribution over the parameter α with density function ϕ defined on $[\alpha_0, \alpha_1]$. Observe that assuming a continuous $S(\alpha, p)$ is simply a device for arriving at a continuous estimate of $\psi(q, p)$; the actual rest-of-market stack will conform to the particular market rules that apply, which may entail that it is a step function.

Now

$$\frac{\partial \psi^\alpha}{\partial q}(x(\tau^*), y(\tau^*)) = x'(\tau^*) f(x(\tau^*) + S(\alpha, y(\tau^*)))$$

$$\frac{\partial \psi^\alpha}{\partial p}(x(\tau^*), y(\tau^*)) = S'(\alpha, y(\tau^*)) y'(\tau^*) f(x(\tau^*) + S(\alpha, y(\tau^*)))$$

giving

$$\tilde{\phi}(\alpha) \propto \phi(\alpha) \prod_{i=1}^{k} (x'(\tau^*) + S'(\alpha, p_i) y'(\tau^*)) f(q_i + S(\alpha, p_i)).$$

The application of this formula becomes simpler when our offer stack is a step function. Here either $x'(\tau^*) = 0$ and $y'(\tau^*) = 1$ (on a vertical section), or $x'(\tau^*) = 1$ and $y'(\tau^*) = 0$ (on a horizontal section). Thus if we are dispatched at (q_i, p_i) on a vertical section then

$$\tilde{\phi}(\alpha) = \frac{\phi(\alpha) S'(\alpha, p_i) f(q_i + S(\alpha, p_i))}{\int_{\alpha_0}^{\alpha_1} \phi(\alpha) S'(\alpha, p_i) f(q_i + S(\alpha, p_i)) d\alpha}$$

and if we are dispatched at (q_i, p_i) on a horizontal section then

$$\tilde{\phi}(\alpha) = \frac{\phi(\alpha) f(q_i + S(\alpha, p_i))}{\int_{\alpha_0}^{\alpha_1} \phi(\alpha) f(q_i + S(\alpha, p_i)) d\alpha}.$$

These formulae provide a mechanism for updating ϕ. The application of the formulae as stated will only be possible when we have some analytic representation of the demand density f. The posterior estimate of the function $\psi(q, p)$ in terms of f and $\tilde{\phi}(\alpha)$ is now calculated to be

$$\psi(q, p) = \int_{\alpha_0}^{\alpha_1} \psi^\alpha(q, p) \tilde{\phi}(\alpha) d\alpha = \int_{\alpha_0}^{\alpha_1} \int_0^{q+S(\alpha,p)} f(h) \tilde{\phi}(\alpha) dh d\alpha.$$

4 Updates in a transmission system

In most markets the participants are located at nodes of a transmission system, and (along with other generators) supply electricity demand through this network. When there are no line losses or constraints in the network, the system can be modelled by aggregating the demand at the nodes of the network to form a single demand at a pool node at which all generation is assumed to be offered. The single-node analysis of the previous section can then be applied.

In the presence of losses and constraints the analysis of the previous section is difficult to apply. The effect of the actions of other participants on the price at each node depends on the location of those participants as well as the spatial variation of the demand in each trading period. To model this situation one might try to represent the network by an equivalent single node, or a small set of aggregated nodes. However this approach can become complicated for even the simplest of representations of the transmission network. Our approach is to create a model of $\psi(q,p)$ at a single node based only on observations at the node, in an effort to incorporate automatically the network effects along with the variation in demand and variation in competitor behaviour.

An important consideration in this model is the choice of the family $\{\psi^\alpha(q,p)\}$. This will depend on the particular characteristics of the market being studied. Our analysis in this section is representing the rest of the market as if they were acting as a single agent at the node where the generator is located. They can be thought of as offering an (unknown supply) function $S(\alpha,p)$ in conjunction with some effective nodal demand that is related to the multivariate nodal demands. One factor influencing the choice of $\{\psi^\alpha(q,p)\}$ will be the shape of $S(\alpha,p)$ (or its inverse T). A typical shape for the curve T in a single node model is shown in Figure 1 below.

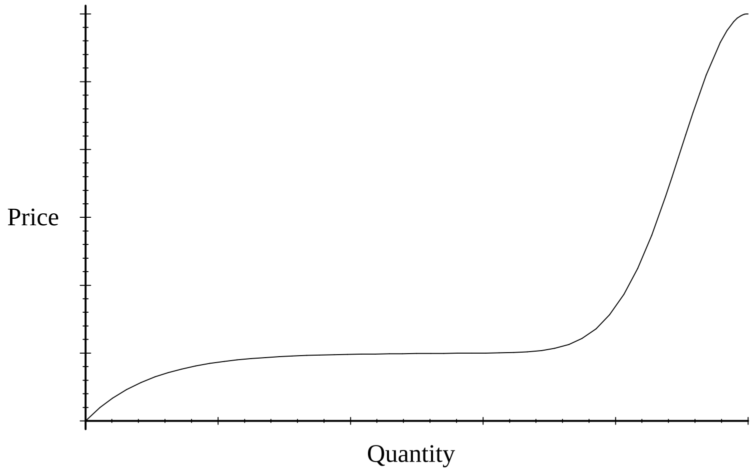

Figure 1: Typical shape of rest-of-market offer curve

For an effective nodal demand of h, $T(h-q)$ will be the anticipated clearing price at the node if the generator were to offer an amount q into the node at no cost. We would expect this function to take a large value when $q=0$ (the clearing price at the node if the generator withdraws all its supply) and decrease monotonically to zero for large q.

The approach we follow is to construct $\psi^{\alpha}(q,p)$ by specifying for each q a density function for p having a distribution that varies with q in such a way that the contours of $\psi^{\alpha}(q,p)$ have a similar shape to the expected shape to $T(h-q)$. Formally we set

$$\psi^{\alpha}(q,p)=\int_{-\infty}^{p} g^{q,\alpha}(z)dz.$$

where $g^{q,\alpha}(p)$ has the interpretation as being the probability density function of the clearing price at the generator's node if they were to offer an amount of electricity of q at price 0.

As stated above, the choice of $\{g^{q,\alpha}(p)\}$ in this model should reflect the expected shape of $T(h-q)$. A model that appears to work well in practice is to let $\alpha=(\alpha,\beta)$ be bivariate, and choose g to be a lognormal density, so that $\log p$ has a normal distribution with mean $\mu = \beta - \alpha q$. The variance of this normal distribution must be chosen with some care, since a fixed variance will give a lower spread of outcomes (in (q,p) space) for large α values than we get for small α values. To correct for this, we choose the variance to be $\sigma^2(1+\alpha^2)$ for some fixed constant σ. This gives

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

Given a prior density $\phi(\alpha,\beta)$ on the parameters $\alpha = (\alpha,\beta)$ we obtain a market distribution function

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

The Bayesian update is defined by

$$\begin{aligned}
\frac{\partial \psi^\alpha}{\partial q} &= \frac{1}{\sqrt{2\pi(1+\alpha^2)}\sigma} \int_{-\infty}^{\log p} \frac{\partial}{\partial q} e^{-\frac{(z-\beta+\alpha q)^2}{2(1+\alpha^2)\sigma^2}} dz \\
&= \frac{1}{\sqrt{2\pi(1+\alpha^2)}\sigma} \int_{-\infty}^{\log p} \alpha \frac{\partial}{\partial z} e^{-\frac{(z-\beta+\alpha q)^2}{2(1+\alpha^2)\sigma^2}} dz \\
&= \frac{1}{\sqrt{2\pi(1+\alpha^2)}\sigma} \alpha e^{-\frac{(\alpha q-\beta+\log p)^2}{2(1+\alpha^2)\sigma^2}},
\end{aligned}$$

and

$$\frac{\partial \psi^\alpha}{\partial p} = \frac{1}{p\sqrt{2\pi(1+\alpha^2)}\sigma} e^{-\frac{(\alpha q-\beta+\log p)^2}{2(1+\alpha^2)\sigma^2}}.$$

Thus given a dispatch of quantity q_i at price p_i we update $\phi(\alpha,\beta)$ by multiplying it by $\frac{\alpha}{\sqrt{(1+\alpha^2)}} e^{-\frac{(\alpha q_i-\beta+\log p_i)^2}{2(1+\alpha^2)\sigma^2}}$ if the generator is dispatched on a horizontal section of its stack, and update $\phi(\alpha,\beta)$ by multiplying it by $\frac{1}{p\sqrt{(1+\alpha^2)}} e^{-\frac{(\alpha q_i-\beta+\log p_i)^2}{2(1+\alpha^2)\sigma^2}}$ if the generator is dispatched on a vertical section of its stack. The final posterior distribution $\tilde{\phi}(\alpha,\beta)$ is then normalised (by dividing through by $\int \tilde{\phi}(\alpha,\beta) d\alpha d\beta$ at the completion of this sequence of updates.)

Observe that the methodogogy we describe uses no specific information about the (multivariate) distribution of load. This is incorporated implicitly to some degree in the estimation of $\tilde{\phi}(\alpha,\beta)$. The main reason for adopting this approach is its simplicity – the only data required are observed dispatch quantities and prices at previous periods.

Of course it is possible to incorporate more information into the estimation procedure, For example, since there is a strong correlation between total system demand and price, we might improve the estimation by incorporating system demand observations in our model. How this is done depends on the effect that a change in demand has on other participants. The simplest model assumes that their offers do not vary with demand. In this case if h is the system demand for a particular trading period (having density function $f(h)$ and expectation $\bar{h}$) then we let $\{\psi^\alpha(q,p)\}$ represent the family of market distribution functions that would pertain with demand $\bar{h}$. Now offering (q,p) into such a market with demand $\bar{h}$ gives the same set of dispatch outcomes as offering $(q+\bar{h}-h,p)$ into this market with demand h, so given a sequence of dispatches (q_i,p_i) and observed demand levels h_i,we obtain the following update formula.

$$\tilde{\phi}(\alpha) \propto \phi(\alpha)f(h_i)\prod_{i=1}^{k}[\frac{\partial\psi^\alpha}{\partial q}(q_i+\bar{h}-h_i,p_i)x'(\tau^*)+\frac{\partial\psi^\alpha}{\partial p}(q_i+\bar{h}-h_i,p_i)y'(\tau^*)].$$

The posterior market distribution function is then

$$\psi(q,p)=\int\int_{\alpha_0}^{\alpha_1}\psi^\alpha(q+\bar{h}-h,p)\tilde{\phi}(\alpha)f(h)d\alpha dh.$$

5 An Industrial Application

We have tested the model above on a data set provided by a hydro-electricity generator in New Zealand. We were provided with their offer stack at each half-hour trading period over the month of February 2001, as well as their actual dispatch and the clearing price. From this we could determine whether each dispatch was on a vertical section of the stack or on a horizontal section of the stack. Since offer behaviour changes with the time of day, we focus on a single trading period, namely 8:00-8:30 pm. The observations of dispatches in this trading period in February are shown in the scatterplot in Figure 2.

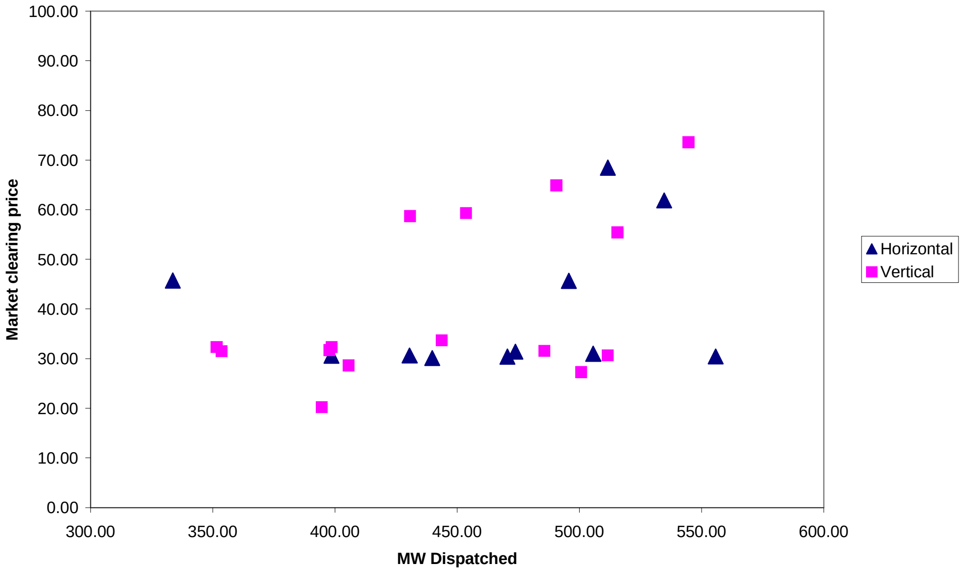

Figure 2: Electricity dispatches for period 8:00-8:30pm in February 2001

As mentioned in the previous section for each q we model $\log p$ with a normal distribution with mean $\mu = \beta - \alpha q$, and $\sigma = 0.4$. As prior distribution for (α, β) we choose ϕ to be a uniform density on $[0, 0.01] \times [3, 8]$. This prior generates the market distribution function with contours plotted in Figure 3.

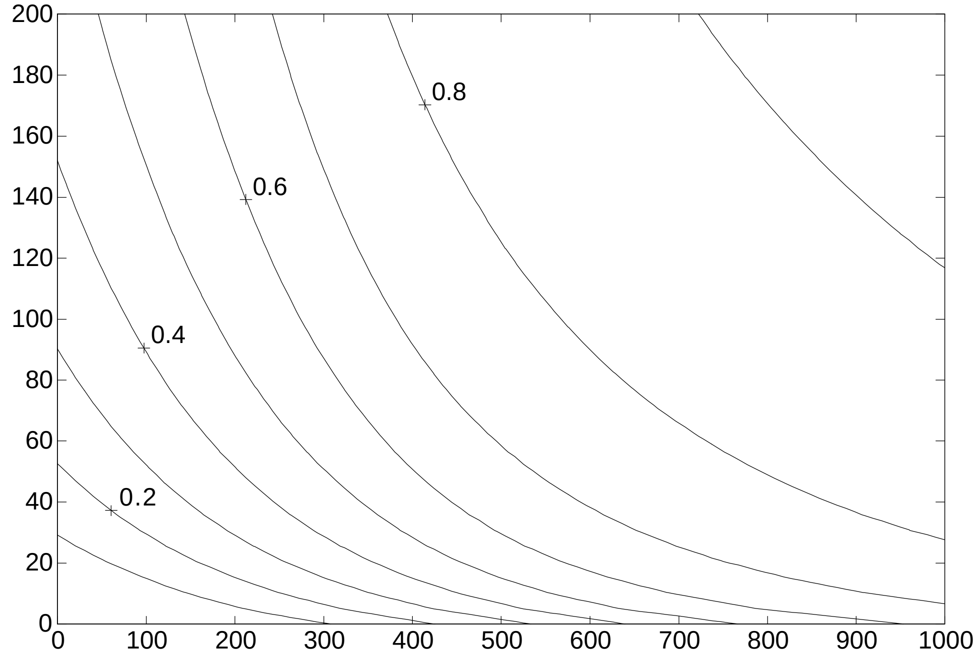

Figure 3: Contours of $\psi(q, p)$ using (prior) ϕ

Updating ϕ using the dispatch data to give a posterior density $\tilde{\phi}$ gives a market distribution function with contours as shown in Figure 4.

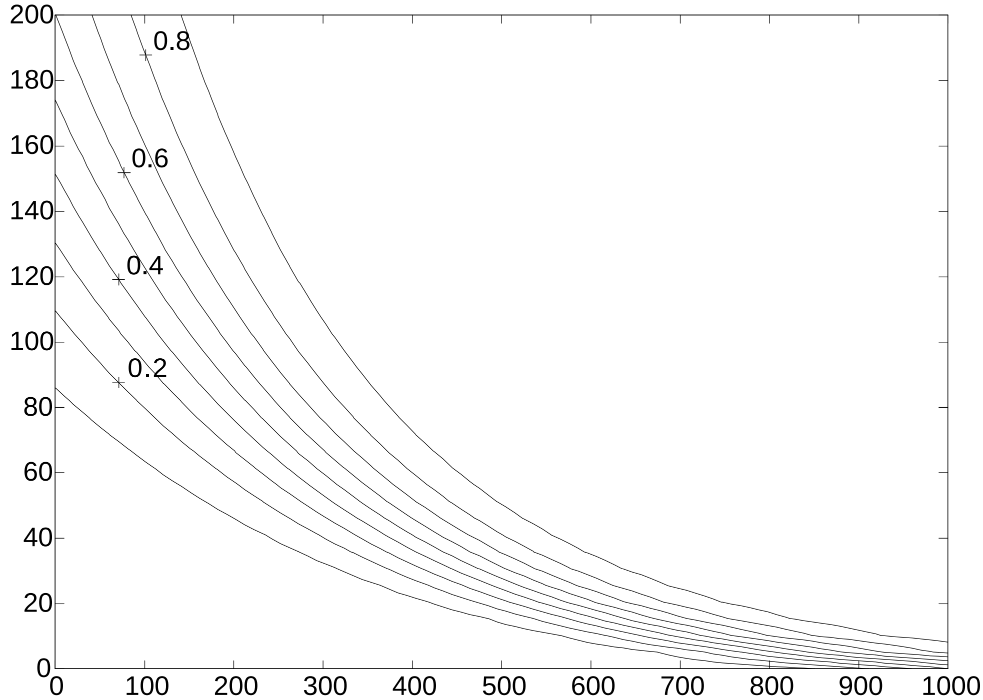

Figure 4: Contours of $\psi(q, p)$ using (posterior) $\tilde{\phi}$

The changes in the contours of ψ give some indication of the effect that the observed dispatches have had on the estimated market distribution function. The contours have become more closely spaced in the neighbourhood of the dispatches as one might expect.

The effect of the difference between the prior and posterior estimates of the market distribution function can be seen from computing an optimal supply curve for each case. This is done by computing

$$Z(q,p) = \frac{\partial}{\partial q}R(q,p)\frac{\partial}{\partial p}\psi(q,p) - \frac{\partial}{\partial p}R(q,p)\frac{\partial}{\partial q}\psi(q,p).$$

It can be shown (see [1]) that the zero contours of $Z(q,p)$ are extremals for the problem of maximizing the expected profit $V(\mathbf{s})$ over supply curves $\mathbf{s}$. (When the market rules require $\mathbf{s}$ to be a stack with a finite number of tranches it is possible to use $Z(q,p)$ to compute an optimal $\mathbf{s}$, but we shall not make use of this here.)

For our example we shall use a profit function corresponding to a marginal cost of \$20 per MWhr (corresponding to a fixed marginal value for water) and a contract level of 500MW. This gives a profit (ignoring the fixed term involving the contract quantity and strike price) of

$$R(q,p) = qp - 20q - 500p.$$

Figure 5 shows the respective optimal supply curves. The top curve is optimal for the prior estimate of ψ and the bottom curve is optimal when the posterior estimate is used.

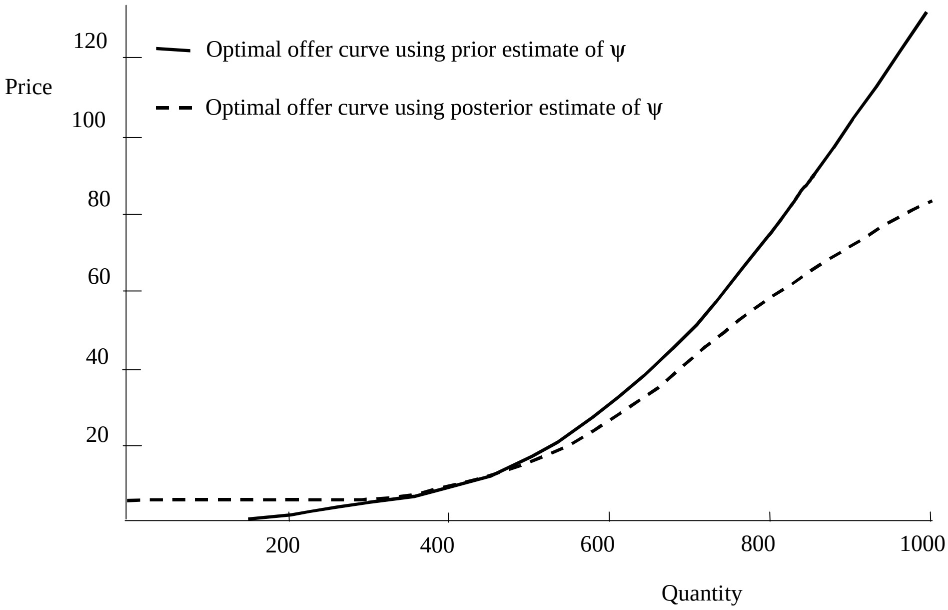

Figure 5: Optimal supply functions for prior and posterior ψ estimates

6 Conclusions

This paper has briefly described a Bayesian updating scheme for estimating a market distribution function for an electricity pool market. A central feature of our approach is the representation of this function using a parameterised family of functions. Clearly an inappropriate specification of this family will lead to poor results. Essentially the choice of a family of functions amounts to guessing the form of the supply functions that are offered by other participants in the market. Markets will vary in regard to the information available on historical bids of other generators.

The estimation of a market distribution function is problematic: not only is this a two dimensional continuous function, but it is constrained to be monotonic in both arguments. A maximum likelihood approach will produce a discontinuous estimate. The Bayesian approach has some significant advantages in this context. It can deal easily with the complexity of the model, and by working with a parameterised family of functions it will always produce an estimate with reasonable properties. Moreover we have demonstrated through the application of this technique to the New Zealand market data that it can be effective in practice. However it has some significant disadvantages: it requires a somewhat arbitrary choice of a family of functions, as discussed above, and the estimate produced does not have easily specified statistical properties

One practical difficulty in estimating market distribution functions using a Bayesian method is that the observations obtained from generators are clustered in a region of (q, p) space in which the generator has been placing its offer stacks. When these offers are earning good profits there is little incentive for a generator to experiment by offering in a different region by withholding capacity. However in the absence of such experimental observations there is considerable potential for errors in the estimated $\psi(q, p)$, especially in regions that are far away from the current offers, where the estimates can be sensitive to the particular parameterised family $\{\psi^{\alpha}(q, p)\}$ chosen.

The Bayesian updating scheme we have used assumes a stationary underlying market distribution function for the trading period considered. The data that are used to update the prior distribution are collected over a period of time (perhaps a month or more), so the more recent data should give a more accurate representation of current market conditions. In order to model this we might give a greater weight to more recent observations in the update procedure. This is done at each stage by putting the prior distribution (the posterior from the previous update) to some fractional power k before applying the update formula, giving

$$
\tilde{\phi}(\alpha,\beta) = \begin{cases} \frac{\alpha}{\sqrt{(1+\alpha^2)}} e^{-\frac{(\alpha q_i - \beta + \log p_i)^2}{2(1+\alpha^2)\sigma^2}} \phi(\alpha,\beta)^k, & \text{if the generator is dispatched at } (q,p) \text{ on} \\ & \text{a horizontal section of its stack,} \\ \frac{1}{p\sqrt{(1+\alpha^2)}} e^{-\frac{(\alpha q_i - \beta + \log p_i)^2}{2(1+\alpha^2)\sigma^2}} \phi(\alpha,\beta)^k, & \text{if the generator is dispatched at } (q,p) \text{ on} \\ & \text{a vertical section of its stack.} \end{cases}
$$

For example a choice of $k = 0.5$ has the effect of performing the most recent update twice as many times as its predecessor, which in turn is applied twice as many times as its predecessor, and so forth.

The approach we have outlined above uses only dispatch (and in the single-node case) demand information. In some situations (such as in Australia), the market clearing authority provides the complete market stack for each trading period after it has been dispatched. Since we can remove our own offer from this to obtain a stack representing the rest of the market, this gives us a direct way of estimating the function $\psi(q,p)$, namely given a total of N observations $S_1, \ldots, S_N$, we can take the estimate

$$
\psi(q,p) = \frac{1}{N} \sum_{i=1}^{N} F(q + S_i(p)).
$$

This direct approach has some disadvantages, especially when N is small. A step in S_i at some price p_i will translate into a discontinuity in ψ at p_i, a property that may not be desirable when carrying out stack optimization. Moreover in the Bayesian setting where we update a distribution over a family of continuous supply functions S_i, $i = 1, 2, \ldots, N$, then any observation of a step function gives a zero (posterior) probability for any S_i. In this case it is necessary to devise a framework in which any offer stack we observe can be generated (perhaps with a low probability) from the family of continuous functions we adopt. This is a subject of our ongoing research.

References

[1] Anderson, E.J. and Philpott, A.B. (2002). Optimal offer construction in electricity markets. *Math of OR, 27,1, 82-100..*

[2] Chao, H-P. and Huntington, H. (1998). *Designing Competitive Electricity Markets.* Kluwer Academic.

[3] Clewlow, L. and Strickland, C. (2000). *Energy Derivatives: Pricing and Risk Management.* Lacima Publications.

[4] Deng, S. (2000). Stochastic models of energy commodity prices and their applications: mean reversion with jumps and spikes. *POWER report PWP-073*, UCEI, Berkeley.

[5] Knittel, C.R. and Roberts, M.R. (2001). An empirical examination of deregulated electricity prices. *POWER report PWP-087*, UCEI, Berkeley.

[6] Wolak, F.A. (1997). Market design and price behaviour in restructured electricity markets: an international comparison. *POWER report PWP-051*, UCEI, Berkeley.