Some new approaches to forecasting the price of electricity: A Study of Californian Market

[1]**Eduardo F. Mendes** , [2]**Les Oxley** and [3]**Marco Reale**

[1]Department of Economics, University of Canterbury
E-Mail: eduardo.mendes@canterbury.ac.nz
[2]Department of Economics, University of Canterbury
E-Mail: les.oxley@canterbury.ac.nz
[3]Department of Mathematics and Statistics, University of Canterbury
E-Mail: marco.reale@canterbury.ac.nz

Keywords: *electricity time series; forecasting performance; semi- and non-parametric methods*

ABSTRACT

In this paper we consider the forecasting performance of a range of semi- and non-parametric methods applied to high frequency electricity price data. Electricity price time-series data tend to be highly seasonal, mean reverting with price jumps/spikes and time- and price-dependent volatility. The typical approach in this area has been to use a range of tools that have proven popular in the financial econometrics literature, where volatility clustering is common. However, electricity time series tend to exhibit higher volatility on a daily basis, but within a mean reverting framework, albeit with occasional large 'spikes'. In this paper we compare the existing forecasting performance of some popular parametric methods, notably GARCH ARMAX, with approaches that are new to this area of applied econometrics, in particular, Artificial Neural Networks (ANN); Linear Regression Trees, Local Regressions and Generalised Additive Models. Section 2 presents the characteristics of the data used which in this case are spot electricity prices from the Californian market 07/1999-12/2000. This period includes the 'crisis' months of May-August 2000 where extreme volatility was observed. Section 3 presents the results and ranking of methods on the basis of forecasting performance. Section 4 concludes. JEL CLASSIFICATIONS: C14, C45, C53

1 INTRODUCTION

Concerns over climate change, spiralling crude oil prices and security of electricity supply, have led to a resurgence of interest in energy-related issues. Electricity market modelling and forecasting has been given a particular boost following deregulation in many countries and the highly publicised Californian experiences of 2000 and more locally, the effects of two dry years in New Zealand, 2001 and 2003 and the Auckland cable failures of 1998 see, Weron (2006). Analysis of the electricity market has been facilitated by accessibility to high frequency load and price data which, in the case of New Zealand and Australia, is available at intervals as frequent as every 5 minutes.

Research into the operations and characteristics of electricity markets can be categorised into four main areas. Firstly, the modelling and forecasting of electricity load. This area has been predominately the domain of electrical or systems engineers concerned to minimise risks to supply. Modelling and forecasting here has involved both parametric, often simple multivariate regression, ARIMA time series approaches and smoothing methods see eg., and semi/non-parametric neural networks see eg., Hippert, Pedreira, and Souza (2001). Much of this literature has been published in electrical engineering outlets and has generally not entered the mainstream economics/econometrics literature. Secondly, there has been growing interest in forecasting spot and forward electricity prices. This interest has been fuelled by both the needs of a deregulated market to understand 'how the market works' and how best (most profitably) to respond to any systematic, forecastable events. Californian experiences, including widespread bankruptcy of some of the players see eg., Knittel and Roberts (2005), has added extra impetus. This area of research has typically been the realm of economists and econometricians who have used parametric time series tools from financial econometrics, and applied them to electricity data see eg., Worthington, Kay-Spratley, and Higgs (2005), Misiorek, Trueck, and Weron (2006), Conejo, Contreras, Espnola, and Plazas (2005), and Escribano, Peña, and Villaplana (2002). These methods typically comprise simple ARIMA or GARCH models whereas others have attempted to model some of the specific characteristics of electricity data. In particular, electricity price time series data tend to be highly seasonal, mean reverting with price jumps/spikes and time- and price-dependent volatility see Weron and Przybylowicz (2000), Huisman and Mathieu (2003), Goto and Karolyi (2003). Furthermore, as noted by Knittel and Roberts (2005) the data tend to exhibit large values of higher order moments relative to a Gaussian distribution which render models based on normality and log-normality of limited use. Mount, Ning, and Cai (2006) explain why price spikes are a typical feature of a deregulated market for electricity and argue in favour of a regime-switching model. Papers that have specifically considered the modelling of non-linearities and/or spikes include Huisman and Mathieu (2003) who argue that a regime jump process performs better in modelling jumps in combination with mean-reversion than a stochastic jump model. Moral-Carcedo and Vicens-Otero (2005) model the non-linearity of the response of demand to temperature using Smooth Transition (STR), Threshold Regression (TR) and Switching Regression (SR) models. They conclude that the Logistic Smooth Transition (LTSR) offers advantages over other models and is their model of choice when applied to Spanish electricity data.

The third level of interest in electricity markets has come from those interested in modelling the behaviour of firms within a newly deregulated market. Here game theory and operations control methods have been applied with or without empirical validation see eg. Batstone (2000), Wolfram (1999), Newberry (1998) and Harvey and Hogan (2000).

Finally, there has been significant discussion of the legal and political implications of the consequences of deregulation particularly related to market power, volatile prices and security of supply see eg., Barton (2003).

In this paper we will present a comparison of the forecasting performance of a range of parametric and semi/non-parametric models as applied to spot electricity prices using data from the Californian market from 07/1999 to 12/2000. Results for the parametric models are taken from Misiorek, Trueck, and Weron (2006), where they found that including a GARCH component did not improve the forecast performance of the 'best' model – an autoregressive model formulated with exogenous variables. Here we take the ARX formulation as the best parametric specification for the California CalPX market clearing prices

The motivation for the paper comes from several sources. The first is to increase our understanding of the drivers and sources of predictability in electricity markets. The second is as a response to the current perceptions regarding the applicability of non-parametric methods to the forecasting of electricity prices. Misiorek, Trueck, and Weron (2006) state that:

> "AI-based models tend to be flexible and can handle complexity and non-linearity. This makes them *promising* (emphasis added) for short term predictions."

However, they then present a somewhat sceptical view on such methods when stating:

> "We have to note, however, that the advocated models have generally been compared *only to other AI-based techniques or simple statistical methods* (emphasis added)...The results of Conejo, Contreras, Espnola, and Plazas (2005), compared three time series specifications; a wavelet multivariate regression technique, and a multilayer perceptron with one hidden layer. ...the ANN technique was the worst of the five tested models....It would be interesting to evaluate representatives from both statistical and AI-based models. However, a comprehensive comparison of models, even from one class is a laborious task." (Misiorek, Trueck, and Weron 2006).

These same authors support regime switching models, which, "by construction should be well suited for modelling the non-linear nature of electricity prices" (Misiorek, Trueck, and Weron 2006).In this paper we seek to test whether this assumed inferiority of these techniques is supported by the data. We will compare the 'best; models for Misiorek, Trueck, and Weron (2006) with a range of semi- and non-parametric approaches discussed in section **??**. however, it is worth stressing at this point that the tournament, as it stands, should favor existing approaches as the 'best' parametric models were constructed to be just that. Here we are not attempting to create the (potentially) 'best' non-parametric alternative, but to take the covariates (and lags) found optimal for the parametric approach and find the best subset for each non-parametric formulation.

The plan of the paper is as follows. In section 2 we describe the data and the modeling strategy adopted for each of the models that will be used to compare to forecast (Californian) electricity price data. The specific models under scrutiny include a range of parametric models; ARIMA, and multiple regime (STAR)and non(semi)-parametric approaches including Artificial Neural Networks; Local regression; Linear Regression Trees and Generalised Additive Models. Section 3 presents the empirical results. Section 4 concludes.

2 CALIFORNIAN MARKET AND DATA

The California market was deregulated in 1998 and opened April 1st 1998. By May 1 2000 the market was in crisis which ended August 31, 2000. By that time Pacific Gas and Electric had gone bankrupt; the other two major power companies had amassed huge debts. Why did this happen? When the market was initially designed, two rules were put in place that left the utility companies unable to hedge against volatility. They were not permitted to sign long term contracts for wholesale electricity; retail rates were largely fixed and hence the companies were unable to pass-on any wholesale price increase onto customers.

Knittel and Roberts (2005) fit a range of traditional models to an hourly time series of real-time Californian electricity prices and find that the forecasting performance of traditional models is 'poor' and can be improved when they address "the unique features of electricity data in particular, volatility clustering and higher order autocorrelation". Contreras, Espnola, Nogales, and Conejo (2003) utilise an ARIMA model to forecast Californian next-day electricity prices for the week of April 3, 2000, being the week is prior to the beginning of the dramatic price volatility period that took place May-August 2000. Their preferred ARIMA model predicts price better before the May-August crisis and only requires the previous 2 hours of data and three differentiations. Average errors in the pre-crisis period were around 5%, whereas they jump to 11% when this volatile period is included. For more on the California market, see also Moulton (2005) and Weron (2006).

In this study we forecast the day-ahead and week-ahead hourly Californian market clearing prices from the period preceding and including the market crisis cited above. We split the dataset into estimation and evaluation sets. The estimation set comprises the period from July 5, 1999 to April 2, 2000; the day before starting the crisis. Consequently, the period from April 3, 2000 to December 3, 2000 is used for evaluation purposes. The test scheme is the same used by Misiorek, Trueck, and Weron (2006) for the linear model; however, we specify the models weekly to capture changes in the model specification, i.e. a new regime.

The variable set used to forecast the prices is: last two days log-price (p_{t-24} and p_{t-48}), last week log-price (p_{t-168}), dummy variable for Saturday (d_{Sat}), Sunday (d_{Sun}) and Monday (d_{Mon}), the logarithm transformation of the next day forecasted load (l_t) and the minimum of previous day's 24 hourly log-prices (mp_t). The logarithm transformation of price and load is used to attain more stable variances.

We forecast the clearing price in a day-based framework (24 hours of the day in a turn) and we re-estimate the models every day, re-specifying[1] the models every week. Note that we use the model estimated on Sunday to forecast the whole week to evaluate the week-ahead performance.

As noticed by Misiorek, Trueck, and Weron (2006) and Cuaresma, Hlouskova, Kossmeier, and

[1] In re-specify the model we mean grow the model when it is needed, e.g. number of regimes in a multiple regime models.

Obersteiner (2004), modeling each hour of the day separately performs better than one specification for whole day. Then, we decide to model each hour of the day separately for all classes of models.

The variable selection procedure[2] used for the parametric models consist in select the subset of variables which minimizes the Bayesian Information Criteria (BIC). For the linear and tree-based models, the selected set contains all the variables. For the non-linear models with smooth transition (i.e. Multiple STAR and Artificial Neural Networks) we select the variables using a technique proposed by Rech, Teräsvirta, and Tschernig (2001). The idea is to approximate the non-linear model by a polynomial of sufficient high order and then apply some well-know variable selection technique to this approximation. We select all variables as they are significant for most models.

For the local regression and GAM we choose a subset of real-valued variables, which seems to present a non-linear relationship with p_t or a local behavior, to model non-parametrically. The selection of these variables is carried out using the an information criteria (Hastie, Tibshirani, and Friedman 2001, Hastie and Tibshirani 1990, Eubank 1988), where the effective number of parameters is given by the trace of the hat matrix[3]. We choose the Corrected AIC (AICc) (Hurvich, Simonoff, and Tsai 1998) which are not affected by significant problems of over-fitting (Manzan 2004). The AICc is shown below.

$$AICc = \log SSE + \frac{N + df}{N - df - 2}, \qquad (1)$$

where SSE is the sum of squared errors, N is the sample size and $df = Tr(H)$ is the effective number of parameters.

For both GAM and local regression, we calculated the AICc of a number of models and select the one which minimizes the information criteria. The estimated models were the following: the dummy variables modeled linearly and all the models with 1, 2, …, 5 non-parametric responses. In the GAM selected all five real-valued variables are modeled non-parametrically. For the local regression model the selected variables were only p_{t-24}, p_{t-168} and l_t.

Following Misiorek, Trueck, and Weron (2006) and Conejo, Contreras, Espnola, and Plazas (2005), we use the naive method as a benchmark for all models. The naive method can be described as follows: the price on hour t on Sundays, Mondays and Saturdays are equal to the same hour of the previous week; the price on hour t on Tuesdays to Fridays are equal to the same hour of the previous day. For the week-ahead forecast, the price is the same as last week. The naive test is passed if the errors for the model are smaller than the errors obtained for the naive method.

[2] All the model/variable selection procedures were carried out using only the in-sample observations.

[3] The hat matrix H is defined as $\hat{y} = Hy$, where $\hat{y}$ is the forecasted outcome and y the actual outcome.

3 RESULTS

To assess the forecasting performance of each model, we use different statistical measures. This performance can be evaluated once the true market prices are available. For every day and all the weeks three types of average prediction errors (typically used in the electricity price forecasting literature, see e.g. Weron (2006)) were computed: one corresponding to the 24 hours of each day and two to the 168 hours of each week.

The Mean Daily Error (MDE) is computed as

$$MDE = \frac{1}{24} \sum_{h=1}^{24} \frac{|p_h - \hat{p}_h|}{\bar{p}_{24}}, \qquad (2)$$

where p_h and $\hat{p}_h$ are respectively the actual price and the forecasted price for hour h and $\bar{p}_{24}$ is the mean hourly price for a given day. The use of $\bar{p}_{24}$ avoid the adverse effect of prices close to zero.

Analogous to the MDE, the Mean Weekly Error (MWE) is computed as:

$$MWE = \frac{1}{168} \sum_{h=1}^{168} \frac{|p_h - \hat{p}_h|}{\bar{p}_{168}}, \qquad (3)$$

where p_h and $\hat{p}_h$ are respectively the actual price and the forecasted price for hour h in the week and $\bar{p}_{168}$ is the mean hourly price for a given week. Additionally, we compute the Weekly Root Mean Square error (WRMSE). The WRMSE is calculated as the square root of the 168 square differences between the actual and forecasted price:

$$MDE = \sqrt{\frac{1}{168} \sum_{h=1}^{168} (p_h - \hat{p}_h)^2}. \qquad (4)$$

The WRMSE puts more weight to differences in the high-price range than MDE and MWE. Such measures are important because price spikes may lead to financial losses in electricity trading. However, both measures are not robust against outliers.

3.1 Forecast Results

Table 1 below, refer to the daily forecasts presented as a weekly measure. The entry "*Linear*" refers to the preferred model from Misiorek, Trueck, and Weron (2006) and entries for this model in Tables **??**, **??**

and **??** replicate his results for this approach and likewise for "*Nave*". The other entries (GAM; Local Regression; ANN and Tree) presented below and in the Appendix are new. Table 1, below, summarises the results and demonstrate the following; for the MWE both GAM and Local Regression dominate Linear with Nave fourth. For WRMSE GAM followed by Local Regression with Linear and Nave joint third. Looking at the 'calm' (weeks 1-10) versus 'volatile periods' (weeks 11-35); Linear seems to forecast better in the early periods, less so in the more volatile episodes.

Table 1. WEEKLY BEST MODEL SUMMARY - DAY AHEAD FORECAST

Model	MWE	WRMSE
Local Regression	8	10
GAM	9	7
ANN	1	2
Naive	6	6
Tree	4	4
Linear	7	6

This table contains the weekly 'best model' (model with smallest error) summary in a Day-Ahead forecasting framework. For each model we show how many times it was be best model in each error measure (MWE and WRMSE), where 'GAM' aims for the Generalised Additive Model, 'ANN' for the Artificial Neural Networks, 'Naive' the Naive method for week ahead forecast, 'Local Regression' for Local Regression Model, 'Tree' the Linear Regression Tree model and 'Linear' the ARX model.

Table 2 below, refer to week ahead forecasts. These are new including the columns headed "*Linear*" and "*Naive*". Table 2 summarise these results and shows that Local Regression and GAM dominate all other approaches for both MWE and WRMSE. Also new is Table 3 which relate to day-ahead forecasts. Table 3 shows the models with smaller MDE sorted by day of week. The naive method is the best "forecasting method" in a day ahead forecast, followed by GAM and Local Regression. The linear, ANN and linear regression tree models have the worst performance.

4 CONCLUSION

Interest in modelling and in particular, forecasting, electricity prices is growing globally. Much interest has been focussed on modelling a small number of key markets including CalPX and NordPool, with British, Spanish and Australasian markets being included as high frequency data becomes available.

In this study we have analysed the CalPX data as a precursor to a more wide-ranging testing programme. In particular, we have sought to formally investigate the potential for using a range of non- (semi-) parametric methods that have proven useful in other areas of applied statistics. The particular nature of the

Table 2. WEEKLY BEST MODEL SUMMARY - WEEK AHEAD FORECAST

Model	MWE	WRMSE
Local Regression	14	16
GAM	19	16
ANN	2	2
Naive	0	1
Tree	0	0
Linear	0	0

This table contains the weekly 'best model' (model with smallest error) model summary in a Week-Ahead forecasting framework. For each model we show how many times it was be best model in each error measure (MWE and WRMSE), where 'GAM' aims for the Generalised Additive Model, 'ANN' for the Artificial Neural Networks, 'Naive' the Naive method for week ahead forecast, 'Local Regression' for Local Regression Model, 'Tree' the Linear Regression Tree model and 'Linear' the ARX model.

Table 3. DAILY BEST MODEL - DAY-AHEAD FORECAST

Model	**Total**
Local Regression	20.8%
GAM	21.6%
ANN	11.4%
NAIVE	29.0%
TREE	8.6%
LINEAR	8.6%

This table contains summary of the models with smallest MDE in each day of week. 'GAM' aims for the Generalised Additive Model, 'ANN' for the Artificial Neural Networks, 'Naive' the Naive method for week ahead forecast, 'Local Regression' for Local Regression Model, 'Tree' the Linear Regression Tree model and 'Linear' the ARX model.

electricity data eg., highly seasonal, mean reverting with occasional jumps/spikes and time- and price-dependant volatility, appears on the face of it to be a prima facie case for using a range of parametric methods developed for financial data. Weron et. al. (various) have demonstrated with the CalPX data the apparent dominance of linear ARX and 'Nave' forecasting methods. Incorporating GARCH-type effects apparently does not enhance the performance of these simple methods (see Weron (2006)) although these results relate to a small range of cases and would appear to contrast with those of Garcia, Contreras, van Akkeren, and Garcia (2005) who used both CalPX and Spanish data.

In this study we have contributed to the literature by formally testing the proposition on page 4 that casts doubt on the assumed poor performance of AI-based techniques. Our results fall into two groups. For the experiments undertaken by Weron; daily forecasts - weekly measure - the Linear (ARX) model performs well, but is dominated by Local Regression and GAM. ANN does not perform well, as postulated by Misiorek, Trueck, and Weron (2006), nor do Trees. However, 'Nave' works very well - the simplest and often 'best' way to forecast electricity prices is to assume your forecast tomorrow is simply informed by the same hour of the previous day (or week for Saturday, Sunday and Monday)! New results presented here, however, are somewhat more encouraging for the benefits of using non-parametric methods. Week ahead forecasts are dominated by Local Regression and GAM formulations and day ahead forecasts show a strong role for these two approaches and also the ANN. Linear models have somewhat less success. In addition it must be stressed that these comparisons were made within the constraints of the best linear model formulation where the covariates were chosen to maximise the performance of that formulation.

Future work in this area will involve; application of the parametric and non- (semi-) parametric methods to a range of alternate data sets and to include a number of other co-variates, i.e., NordPool and New Zealand data sets and the inclusion of weather and hydrological data.

5 REFERENCES

BARTON, B. (2003): "Does electricity market liberalization contribute to energy sustainability?," in *Energy Law and Sustainable Development*, ed. by A. J. Bradbrook, and R. L. Ottinger, pp. 217–231. Cambridge:IUCN Publications.

BATSTONE, S. (2000): "Risk managment for dergulated electricity market: simulation results from a hydro management model," in *Proceedings of the* 33^{rd} *Annual Conference*, ed. by A. Henderson, and A. Philpott. ORSNZ.

CONEJO, A. J., J. CONTRERAS, R. ESPNOLA, AND M. PLAZAS (2005): "Forecasting electricity prices for a day-ahead pool-based electric energy market," *International Journal of Forecasting*, 21, 435–462.

CONTRERAS, J., R. ESPNOLA, F. J. NOGALES, AND A. CONEJO (2003): "ARIMA models to predict next-day electricity prices," *IEEE Transactions on Power Systems*, 18(3), 1014–1020.

CUARESMA, J., J. HLOUSKOVA, S. KOSSMEIER, AND M. OBERSTEINER (2004): "Forecasting electricity spot prices using linear univariate time series," *Applied Energy*, 77, 87–106.

ESCRIBANO, A., J. PEÑA, AND P. VILLAPLANA (2002): "Modelling electricity prices: international comparisons," Working paper, Universidad Carlos III.

EUBANK, R. L. (1988): *Spline Smoothing and Nonparametric Regression*. Marcel Dekker, Inc.

GARCIA, R., J. CONTRERAS, M. VAN AKKEREN, AND J. GARCIA (2005): "A GARCH forecasting model to predict day-ahead electricity prices," *IEEE Transactions on Power Systems*, 20(2), 867–874.

GOTO, M., AND G. A. KAROLYI (2003): "Understanding electricity price volatility within and across markets," Working paper, Ohio State University.

HARVEY, S., AND H. HOGAN (2000): "Californian electricity prices and forward market hedging," Technical report, Harvard University.

HASTIE, T., AND R. TIBSHIRANI (1990): *Generalized Additive Models*. Chapman and Hall.

HASTIE, T., R. TIBSHIRANI, AND J. FRIEDMAN (2001): *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. Springer.

HIPPERT, H. S., C. E. PEDREIRA, AND R. C. SOUZA (2001): "Neural networks for short-term load forecasting: a review and evaluation," *IEEE Transactions on Power Systems*, 16(1), 44–55.

HUISMAN, R., AND R. MATHIEU (2003): "Regime jumps in electricity prices," *Energy Economics*, 25, 425–434.

HURVICH, C. M., J. S. SIMONOFF, AND C. L. TSAI (1998): "Smoothing parameter selection in nonparametric regression using an improved Akaike Information Criterion," *Journal of the Royal Statistical Society B*, 60, 271–293.

KNITTEL, C. R., AND M. R. ROBERTS (2005): "An empirical examination of restructured electricity prices," *Energy Economics*, 27, 791–817.

MANZAN, S. (2004): "Model selection for nonlinear time series," *Empirical Economics*, 29, 901–920.

MISIOREK, A., S. TRUECK, AND R. WERON (2006): "Point and interval forecasting of spot electricity prices: lineas vs non-linear time series models," *Studies in Nonlinear Dynamics & Econometrics*, 10(3), Article 2.

MORAL-CARCEDO, J., AND J. VICENS-OTERO (2005): "Modelling the non-linear response of Spanish electricity demand to temperature variations," *Energy Economics*, 27, 477–494.

MOULTON, J. S. (2005): "California electricity futures: the NYMEX experience," *Energy Economics*, 27, 181–194.

MOUNT, T. D., Y. NING, AND X. CAI (2006): "Prediction price spikes in electricity markets using a regime-switching model with time-varying parameters," *Energy Economics*, 28, 62–80.

NEWBERRY, D. (1998): "Competition, contracts and entry in the electricity spot market," *RAND Journal of Economics*, 29(4), 762–749.

RECH, G., T. TERÄSVIRTA, AND R. TSCHERNIG (2001): "A simple variable selection technique for nonlinear models," *Communications in Statistics - Theory and Methods*, 30, 1227–1241.

WERON, R. (2006): *Modelling and forecasting electricity loads and prices: a statistical approach*. Wiley, Chichester.

WERON, R., AND B. PRZYBYLOWICZ (2000): "Hurst analysis of electricity price dynamics," *Physica A*, 283, 462–468.

WOLFRAM, C. (1999): "Measuting duopoly power in the Britsh electricity spot market," *American Economic Review*, 89(4), 805–826.

WORTHINGTON, A., A. KAY-SPRATLEY, AND H. HIGGS (2005): "Transmission of prices and proce volatility in Australian electricity markets: a multivariate GARCH analysis," *Energy Economics*, 27, 337–350.