Statistical Analysis of New Zealand Electricity Prices:

A Risk Manager's Perspective

by

Caroline Moy

A thesis

submitted to the Victoria University of Wellington

in partial fulfilment of the

requirements for the degree of

Masters of Science

in Stochastic Processes for Finance and Insurance

Victoria University of Wellington

2010

Abstract

This thesis considers the conventional SARIMA model and the EVT-GARCH model for forecasting electricity prices. However, we find that these models do not adequately capture the important characteristics of the electricity price data. A new model is developed, the EVT-SARIMA model, for forecasting electricity prices which is found to be the best at modelling the nature of the electricity prices. A time series of half-hourly electricity price data from the Hayward node in New Zealand is transformed into a daily average price series and using this resulting series, appropriate models are fitted for estimating and forecasting.

The new EVT-SARIMA model is used to simulate 1000 time series of daily electricity prices, over a 90 day period, to consider strategies for managing the risk associated with price volatility. The effects of different financial instruments on the cumulative distribution functions of predicted revenue obtained using our model are considered. Results suggest that different contracts have different effects on the predicted revenue. However, all contracts have the effect of reducing variability in the predicted revenue values and thus, should be used by a risk manager to reduce the range of probable revenue values. The quantity traded and which contracts to use is dependent on the objectives of the risk manager.

This is dedicated to my Father, who left us too soon ...

and my Mother, for her ongoing support and guidance.

Acknowledgements

Of the many people who have assisted in some way in the research work of this thesis, thanks are due in particular to:

Dr Leigh Roberts, my supervisor, for his advice and assistance throughout the course of this work;

Meridian Energy, in particular Chris Garr and Mike Roan, for their time and feedback and;

Friends for contributions by way of support and helpful discussions.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In recent years, we have seen extreme movements in electricity prices. For example, lately we have seen low lake levels lead to an increase in wholesale electricity prices and, as a result, an increase in retail electricity prices for the consumer. Similarly, increased investments by electricity generators have also led to increases in the wholesale electricity price. Such events demonstrate the importance of risk management in energy trading. The sensitivity of electricity prices to numerous factors requires suitable actions by risk managers to ensure the survival of any player in the electricity market.

Risk management in financial markets has been well studied and there has been much literature written on different strategies for managing risk in such markets. However, an energy trader faces risks that are unique to electricity trading. Electricity, as a commodity to be traded, is different from any other commodity. Electricity prices are highly sensitive to changes in weather, more so in New Zealand due to our dependence on hydro-power.

One also faces the difficulty of storing electricity, and the requirement to balance electricity supply and a relatively inelastic demand at all times. Thus, given the volatile nature of electricity prices, the methods for managing risk in other markets need to be adapted for risk management in the electricity market. Indeed, there exists methods for hedging risk in the NZ electricity industry. These include contracts available on *energyhedge.co.nz*, contracts specified between retailers and generators and more recently, trading on the ASX. However, compared to their international counterparts, players in the NZ electricity market have few tools available for hedging risk. For example, US firms can trade on NYMEX, the world's largest physical commodities futures exchange founded in 1882. Comparatively, it has only been recently that NZ firms have had options and futures available to them for the purposes of risk management. These are traded on the ASX, and have only been available since July 2009.

In this thesis, we develop a new model for forecasting electricity prices. We find that our model, the EVT-SARIMA model, appears to be the most appropriate model for forecasting electricity prices. This model incorporates the weekly seasonal pattern and provides relatively good predictions of the random jumps in price. We then consider possible risk management strategies from the point of view of an electricity generator. We demonstrate the effect of numerous contracts by observing the effect that these contracts have on the cumulative distribution function (CDF) of revenue. We obtain this CDF of the revenue curve using 1000 price series simulations using our EVT-SARIMA model.

The thesis has been set out as follows. In chapter 2, we provide a brief overview of the conventional SARIMA model and the theory behind this model. In chapter 3, we introduce relevant results from extreme value theory (EVT) and apply these results to the conventional SARIMA model to obtain our new model, the EVT-SARIMA model. We then provide a literature review and consider possible hedging strategies, and the implications of such strategies, in chapter 4. Finally, in chapter 5, we conclude with a summary of our findings and possible areas of further research.

Chapter 2

The Conventional Seasonal ARIMA Model

In this chapter we introduce the seasonal autoregressive integrated moving average (SARIMA) model. We discuss the various components of the SARIMA model and finally, fit a conventional SARIMA model to the daily average electricity price data. The reader should consult Taylor (2005) for further details of the models discussed in this chapter and the application of such models to financial time series.

2.1 ARMA(p,q) Models

We first consider Autoregressive Moving Average (ARMA) models. These models are a combination of an autoregressive (AR) model with p lags and a moving average model with q lags. For further discussion of ARMA models we suggest the reader see Brooks (2002) and, a more comprehensive look at these models is given in Bowerman, O'Connell & Koehler (2005).

2.1.1 AR(p)

Suppose we wish to model a process X_t, indexed by time t. An autoregressive model implies the current value of X depends only on the previous values of X and an error term. A general AR(p) model would be given by

$$\begin{aligned} X_t &= \phi_1 X_{t-1} + \phi_2 X_{t-2} + ... + \phi_p X_{t-p} + \epsilon_t \\ &= \sum_{i=1}^{p} \phi_i X_{t-i} + \epsilon_t \end{aligned}$$

where the $\phi_i' s$ are autoregressive parameters that must be estimated and ϵ_t follows a white noise process. ϵ_t has the following properties:

$$E[\epsilon_t] = 0, \qquad E[\epsilon_t^2] = \sigma_\epsilon^2, \qquad E[\epsilon_t \epsilon_{t+\tau}] = 0$$

for all t and $\tau \neq 0$. If we introduce the lag operator L and define

$$LX_t = X_{t-1}$$

then, we may represent the AR(p) model as

$$\phi_p(L) X_t = \epsilon_t$$

where $\phi_p(L) = 1 - \phi_1 L - \phi_2 L^2 - ... - \phi_p L^p$.

2.1.2 MA(q)

The moving average process is similar to the autoregressive process as they are both based on a white noise process. The MA(q) process may be written as

$$\begin{aligned} X_t &= \epsilon_t + \theta_1 \epsilon_{t-1} + \theta_2 \epsilon_{t-2} + ... + \theta_q \epsilon_{t-q} \\ &= \sum_{i=1}^{q} \theta_i \epsilon_{t-i} + \epsilon_t \\ &= \theta_q(L) \epsilon_t \end{aligned}$$

where ϵ_t is a white noise process and $\theta_q(L) = 1 + \theta_1 L^1 + \theta_2 L^2 + ... + \theta_q L^q$. A moving average model is a linear combination of white noise processes. Such a model implies that X_t depends only on the current and previous values of the white noise terms. Now, we can combine the AR(p) and MA(q) models to obtain an ARMA(p,q) model denoted

$$X_t = \sum_{i=1}^{p} \phi_i X_{t-i} + \epsilon_t + \sum_{i=1}^{q} \theta_i \epsilon_{t-i} \qquad or \qquad \phi_p(L) X_t = \theta_q(L) \epsilon_t$$

2.2 The ARIMA Model

ARMA(p,q) models are (weakly) stationary processes. However, if we were to model a non - stationary process, we may need to use an Autoregressive Integrated Moving Average (ARIMA(p,d,q)) model instead. We use such a model if the process we are modelling is non-stationary but becomes stationary after differencing. The number of times the process needs to be differenced before it becomes stationary is given by d in the model specification. After differencing the process d times, we then fit the appropriate ARMA(p,q) model to the resulting stationary process. We define the ARIMA(p,d,q) model as

$$\phi_p(L)(1-L)^d X_t = \theta_q(L)\epsilon_t$$

2.2.1 Seasonal ARIMA

The above model may be fitted to data with a seasonal trend. A seasonal ARIMA (SARIMA) model is an ARIMA model which includes moving average and autoregressive terms at a lag equal to the number of seasons (s). That is, a SARIMA model consists of two components. The first is an ARIMA model fitted to the original data series and the second component is an ARIMA model fitted to the series given by

$$\nabla_s X_t = X_t - X_{t-s}$$

which models the seasonal trend in the data. Ghysels & Osborn (2001) discuss the inclusion of the seasonal lags in the AR or MA polynomials. The multiplicative SARIMA model is defined as

$$\Theta_P(L^s)\phi_p(L)(1-L^s)^D(1-L)^d X_t = \Phi_Q(L^s)\theta_q(L)\epsilon_t$$

and we denote it by $ARIMA(p,d,q)(P,D,Q)_s$. Such a model contains two ARIMA components - one to model the general data series and another to model the seasonal trend in the data. For the rest of the paper, we refer to such a model as the SARIMA model.

In Sbai & Simpson (2008), both the SARIMA model and the Holt-Winters model are used to forecast electricity prices. These models are compared and results suggest that the SARIMA model is better than the Holt-Winters model for the purposes of forecasting nodal electricity prices. Thus, we also initially fit a SARIMA model to our electricity data. Then, in the following chapter we use this SARIMA model as the basis of our forecasting model for daily average electricity prices. For further details in model decomposition, the reader should see Bowerman et al. (2005). Ghysels & Osborn (2001) and Abraham & Ledolter (1983) provide a detailed description of the different methods one may use to appropriately fit a conventional SARIMA model to a data series.

2.3 Application to Daily Average Electricity Prices

Now, we consider modelling daily average electricity price series. We use half hourly electricity price data from the Haywards node. Our original price series consists of half-hourly electricity prices from 1 January 2006 to 31 August 2008. From this price series we calculate the daily average prices, which results in a series of 974 data points. We obtain the series given by Figure 2.1 and it is this series we use to fit an appropriate model.

From Figure 2.1 we can see that the electricity prices are quite volatile. The prices do not appear to be mean-reverting and exhibit random spikes as well as clusters of relatively high prices. The prices range from \$11.34 per MWh to \$450.49 per MWh, which is a large range of possible prices that one could face on a day to day basis. A suitable model for forecasting electricity prices would need to incorporate such properties.

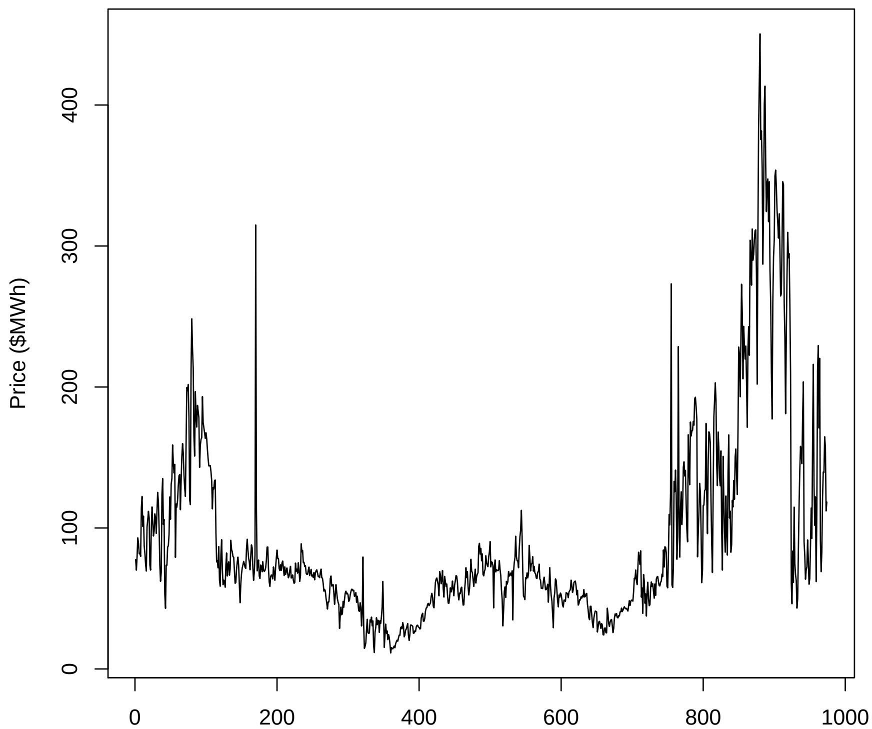

Figure 2.1: Daily Average Prices 1 Jan 2006 - 31 Aug 2008

Initially, to model any data series, we must consider which type of model is appropriate. We need to determine whether the process we are modelling is stationary or non-stationary. To do so, we use the autocorrelation function to confirm stationarity or non-stationarity.

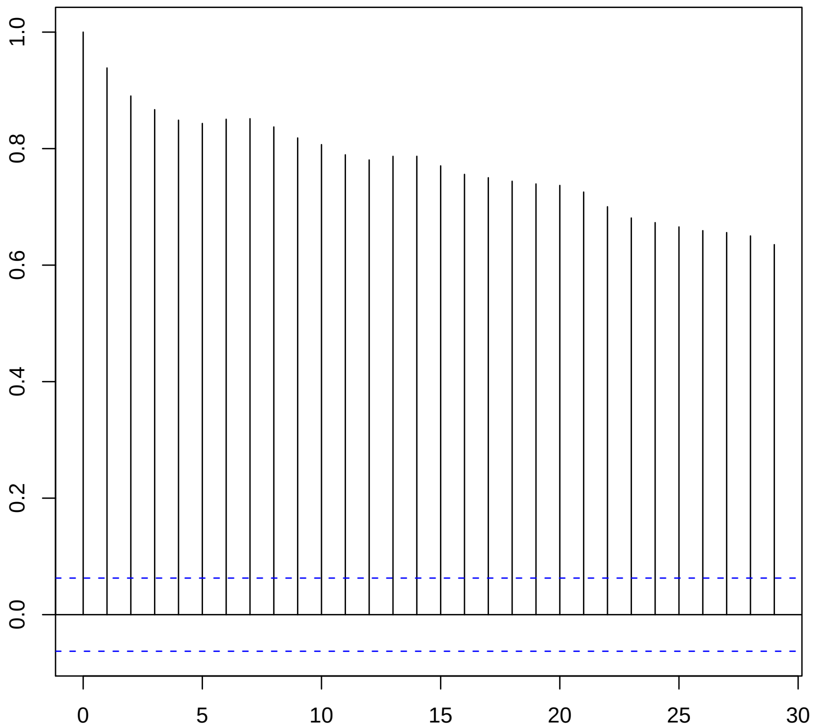

Figure 2.2: Autocorrelation Function of Daily Average Prices

From Figure 2.2, the autocorrelation function does not appear to cut off after any particular lag, which suggests that the daily average price series is a non stationary process. This non-stationarity may be due to seasonal effects or trends. As mentioned above, we may fit an ARIMA or SARIMA model if the series becomes stationary after differencing. We difference the daily average price series once and look at the autocorrelation function of the resulting series.

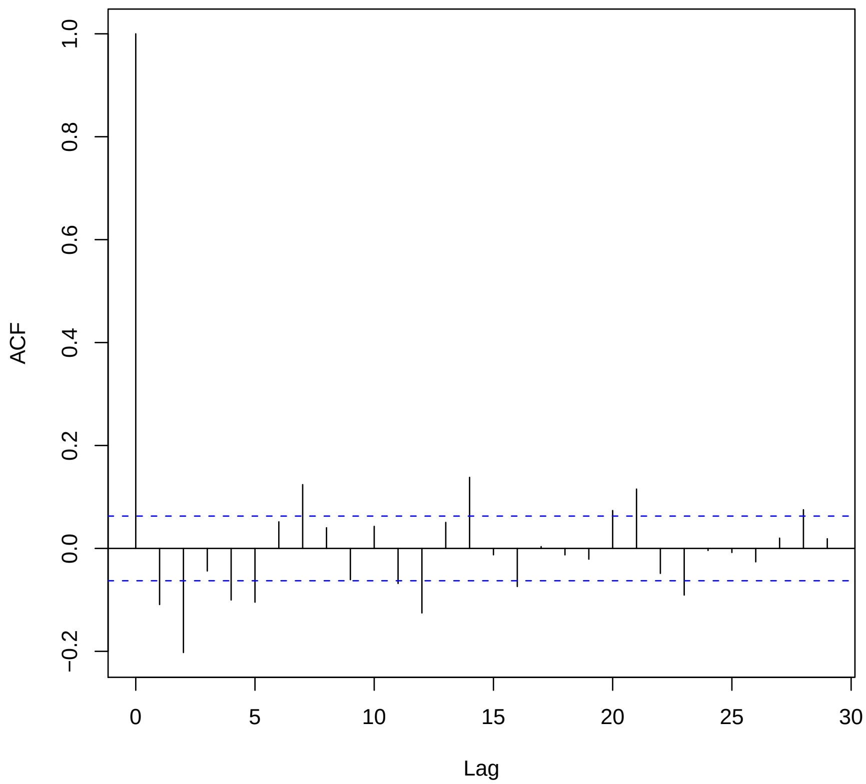

Figure 2.3: Autocorrelation Function of Daily Average Prices after One Difference

After differencing the series once, Figure 2.3 shows there are significant positive lags at multiples of 7, which suggests that a seasonal trend exists at a lag equal to 7. Thus, a SARIMA model may be an appropriate model for this data, with $s = 7$. This suggests that the daily average electricity price may follow a weekly cycle, which should be incorporated in the SARIMA model. To estimate the appropriate number of AR and MA terms for each component of the SARIMA model, we consider a number of different combinations of AR and MA polynomials and compare the AIC values for each resulting model. We find the "best" model, the model with the smallest AIC value, to be

$$\phi_2(L)(1 - L^7)X_t = \Phi(L^7)\theta(L)\epsilon_t$$

which may be denoted as $ARIMA(2, 0, 1)(0, 1, 1)_7$.

To check whether our model is appropriate for the data we are modelling, we consider the autocorrelation function (acf) and partial autocorrelation function (pacf) for the residuals and the squared residuals of the model. The pacf indicates whether the model includes an appropriate number of autoregressive (AR) terms and the acf allows us to check whether a suitable number of moving average (MA) terms have been used.

In Figure 2.4 we can see that there is no significant autocorrelation or partial autocorrelation in the residuals or squared residuals, which suggests that our model includes a suitable number of parameters and is a relatively good fit to the data.

Figure **??** shows that even though the acf and pacf of both the residuals and squared residuals suggest that the SARIMA model is a good fit, the model does not capture the

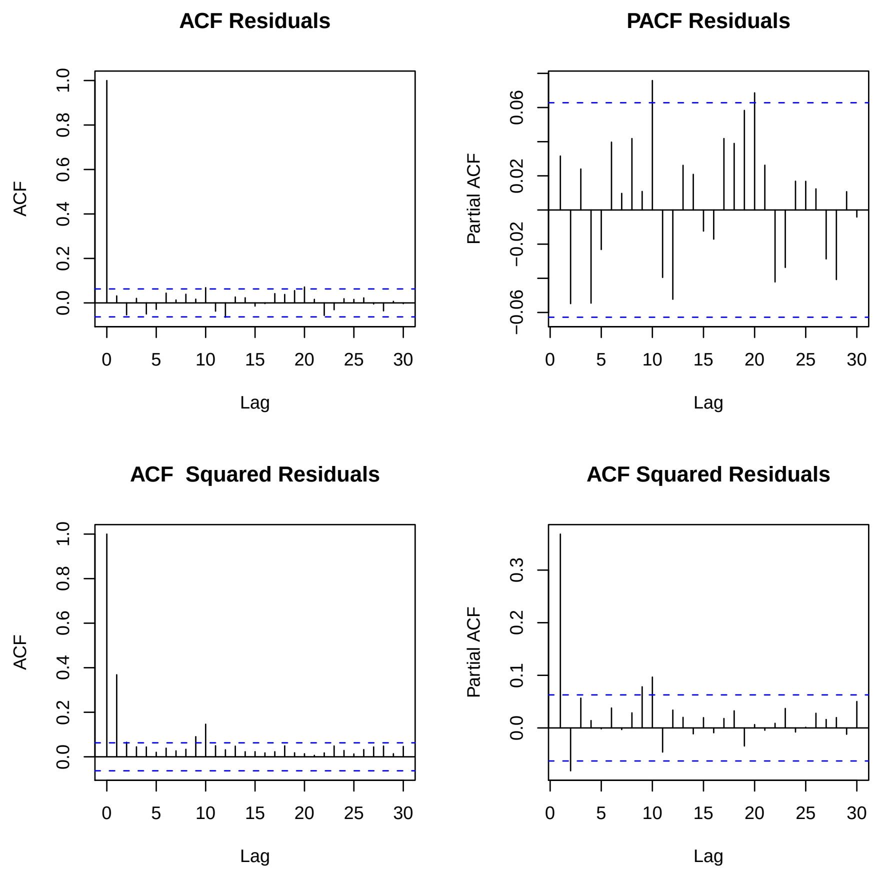

Figure 2.4: Autocorrelation Function (ACF) of Residuals and Squared Residuals. Partial Autocorrelation Function (PACF) of Residuals and Squared Residuals

volatility in the daily average electricity prices. Thus, the conventional SARIMA model would not be an adequate model for forecasting prices, as the model doesn't capture the important characteristics of the actual price series.

The model tends to largely underestimate the volatility in prices and suggests that prices constantly lie close to a level of just below $100 MWh, whereas the actual data series fluctuates a lot over time and does not stay at a constant level. Moreover, the model does not incorporate the random jumps in price seen in the actual data series. Hence, such a model would be particularly inappropriate for the purposes of risk management, where one is attempting to hedge against risk associated with price volatility.

In the following chapters, we attempt to find a model which captures the nature of the electricity price data series more accurately for the purposes of managing the risk related to volatile prices that an electricity generator may face.

Chapter 3

The EVT-SARIMA Model

In this chapter we introduce a new model, the EVT-SARIMA model, that we use to model the daily average prices. We first provide a brief overview of the results from extreme value theory that are relevant for our model. We apply these results to estimate the tail of the innovation distribution from the conventional SARIMA model. Then, the EVT-based estimations for the tail of the innovation distribution are applied to the conventional SARIMA model to obtain a new model, the EVT-SARIMA model, which we find makes better predictions of daily average electricity prices compared with other models.

3.1 Extreme Value Theory

In this section of the thesis we provide a brief overview of relevant results from extreme value theory and apply these results to our SARIMA model for daily average electricity prices from the previous section. For more theoretical detail we refer to the reader to Embrechts, Kluppelberg & Mikosch (2003) and for more detail with respect to the statistical methods when applying extreme value theory the reader should consult McNeil, Frey & Embrechts (2005).

In general, there are two main models for extreme values. There are the "block maxima" models which are very wasteful of data as we retain only the maximum losses in large blocks. The more commonly used models are the "points over threshold" models where observations which exceed a particular threshold are modelled. These models are generally considered more useful as they use data more efficiently and it is these models that we use in this thesis. For further details on either of these models we suggest the reader see McNeil et al. (2005).

3.1.1 The Generalised Extreme Value Distribution

In statistics, the normal distribution is an important limiting distribution for sample sums or averages which may be explained through the central limit theorem. Similarly, in extreme value theory, there is a family of important limiting distributions when studying the limiting behaviour of sample extreme values. This family of distributions can be represented under a single parameterisation known as the generalised extreme value distribution (GEV). We define the distribution function (d.f) of the standard GEV by:

$$H = \begin{cases} exp[-(1+\xi x)^{-1/\xi}] & \xi \neq 0 \\ exp(-e^{-x}) & \xi = 0 \end{cases}$$

where $1 + \xi x > 0$ and we call ξ the shape parameter. Within this family of distributions, there are three distributions that are special cases defined by ξ. They are the Frechet Distribution ($\xi > 0$, with shape parameter $\alpha = 1/\xi$)

$$\Phi_\alpha(x) = \begin{cases} 0 & x \leq 0 \\ exp(-x^{-\alpha}) & x > 0, \alpha > 0 \end{cases}$$

Gumbel Distribution ($\xi = 0$)

$$\Lambda(x) = exp(-exp^{-x})$$

Weibull Distribution ($\xi < 0$, with shape parameter α = -1/ξ)

$$\Psi_\alpha(x) = \begin{cases} exp[-(-x^{-\alpha})] & x \leq 0, \alpha < 0 \\ 1 & x > 0 \end{cases}$$

We can extend the family of distributions by introducing two more parameters, μ a location parameter and $\sigma > 0$, a scale parameter. Then, we can define

$$H_{\xi,\mu,\sigma}(x) = H_\xi((x - \mu)/\sigma)$$

such that $H_{\xi,\mu,\sigma}(x)$ is of the type H_ξ.

3.1.2 The Fisher-Tippett Theorem

The Fisher-Tippet Theorem is a major result in EVT and describes the limiting behaviour of normalised sample maxima. It suggests that the asymptotic distribution of the maxima belongs to one of the three distributions defined above.

Suppose we have a sequence of i.i.d random variables X_1, X_2, ... , X_n from an unknown common distribution F. Let the maximum of the first n observations be given by $M_n = \max(X_1,\ X_2,\ \ldots\ ,\ X_n)$. Now, suppose that we can find sequences of real numbers, $a_n > 0$ and b_n, such that the sequence of normalised maxima $(M_n - b_n)/a_n$ converges in distribution. Hence,

$$P\{(M_n - b_n)/a_n \leq x\} = F^n(a_n x + b_n) \to H_\xi, \qquad as \quad n \to \infty$$

for some non-degenerate distribution function H_ξ. We say that F is in the maximum domain of attraction of H if the above condition is satisfied and we denote this by $F \in MDA(H)$.

Fisher & Tippett (1928) showed that if $F \in MDA(H)$ then H is of the type H_ξ for some ξ. This implies that if the normalised maxima converge in distribution then the limit distribution will be some extreme value distribution with parameters ξ, μ and σ.

2.2 Peaks over Thresholds Method

Suppose we have a sequence of random variables X_1, X_2, ... , X_n which are independently and identically distributed (i.i.d) and come from a common distribution F. We choose a high threshold u and model the exceedances over this u. Given a particular value of u, the excess distribution over the threshold u is given by

$$F_u(x) = P(X - u \leq x | X > u) = \frac{F(x+u) - F(u)}{1 - F(u)}$$

for $0 \leq x < x_F - u$, where $x_F \leq \infty$ is the right endpoint of F.

To model the exceedances of u we use the Generalised Pareto Distribution (GPD). The d.f of the GPD is given by:

$$G_{\xi,\beta}(x) = \begin{cases} 1 - (1 + \xi x/\beta)^{-1/\xi} & \xi \neq 0 \\ 1 - exp[-x/\beta] & \xi = 0 \end{cases}$$

where the shape parameter $\beta > 0$ and $x \geq 0$ when the scale parameter $\xi \geq 0$ and $0 \leq x \leq -\beta/\xi$ when $\xi < 0$. We use the GPD to describe the distribution of the excesses over the threshold u.

In our attempt to model the exceedances using the GPD we use a result from Pickands, Balkema and De Haan which says that we can find a positive-measurable function β(u) such that

$$\lim_{u \to x_F} \sup_{0 \leq x < x_F - u} |F_u(x) - G_{\xi,\beta(u)}(x)| = 0$$

if and only if $F \in MDA(H_\xi)$, $\xi \in \mathbb{R}$. This suggests that a GPD is a natural limiting distribution for an excess distribution. Given the condition that $F \in MDA(H_\xi)$ is required in both the Fisher Tippet Theorem and the result above, the distributions where the normalised maxima converge to a GEV distribution are also distributions for which the excess distribution converges to the GPD as the threshold is raised. The shape parameter of the limiting GPD would be the same as the shape parameter for the limiting GEV distribution for the maxima. Thus, the GEV distributions are included within the parametrisation of the GPD. Hence, for a sufficiently high threshold u, we can approximate the distribution of the exceedances of u by $G_{\xi,\beta}(x)$ for some values ξ and β.

3.2 Application of EVT to the SARIMA Model

We now apply the results from the previous section to the tail of the innovation series from our conventional SARIMA model to obtain the EVT-SARIMA model for daily average electricity prices. In most cases when fitting the conventional SARIMA model to a data series, one generally uses the assumption that the innovations follow a normal distribution. We first check whether this is a reasonable assumption for our particular data series. To do so, we consider a Normal Q-Q plot of the residuals, which compares the quantiles of the Normal distribution and the distribution of the given data set. If the data follows a Normal distribution, the Q-Q plot would show a straight line with a positive slope.

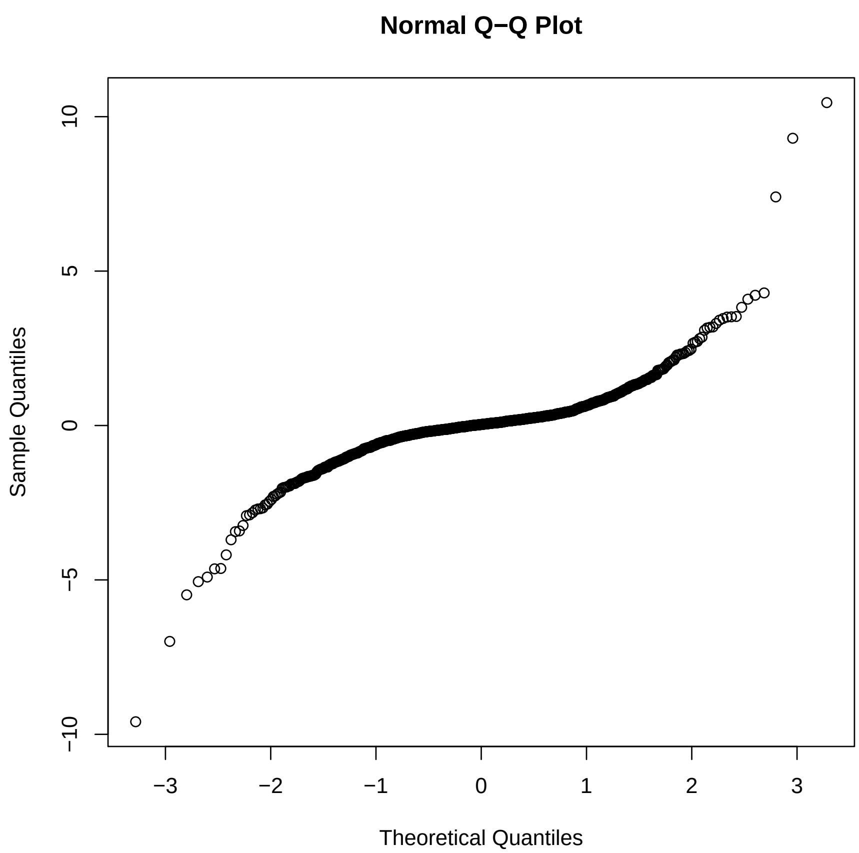

Figure 3.1: Normal Q-Q Plot for Residuals

The assumption of normality isn't appropriate as we can see from Figure 3.1, the points from the data distribution deviate from a positively sloped straight line. Moreover, the Q-Q plot suggests that the innovation distribution has fatter tails than the normal distribution. Hence, fitting a GPD to the tail distribution of the innovations may indeed be more appropriate than the normality assumption.

To fit a GPD to the innovation series, the approach we use is similar to the method given in McNeil & Frey (2000). Suppose we have a series of innovations e_i, $i = 1, ..., n$. We then order these innovations such that $e_{(1)} \geq e_{(2)} \geq ... \geq e_{(n)}$. We choose the first k data points from the ordered series, for our data we choose $k = 100$, and fit a GPD to the series given by

$$e_{(1)} - e_{(k+1)}, ..., e_{(k)} - e_{(k+1)}$$

In order to fit a GPD, we must first determine an appropriate threshold u. To estimate u, we use a mean excess plot, which is explained in Embrechts et al. (2003).

In Figure 3.2, the points of the mean excess plot appear to follow a reasonably straight line after a value of 2.75. Thus, we let the threshold u be 2.75. The upward direction of the points indicates that the data is heavy tailed and the GPD we fit will have a positive shape parameter. From our estimate of u we obtain parameter estimates $\xi = 0.5972$ and $\beta = 1.3413$ using the maximum likelihood estimation method discussed in Smith (1985).

In Figure 3.3, the solid line in the Q-Q plots corresponds to the Normal and GPD quantiles respectively. The plot suggests that the fitted GPD with parameter value $\xi = 0.5972$ is a reasonable fit to the given series. We can see that in comparison to the normal distribution, a GPD is more appropriate for the data as the points from the sample distribution follow

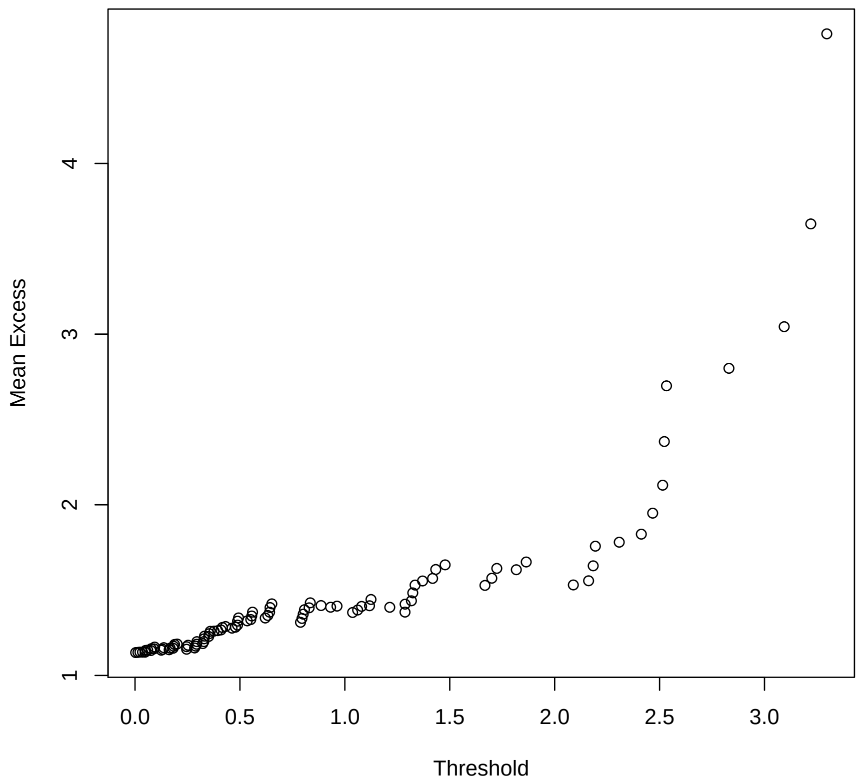

Figure 3.2: Mean Excess Plot

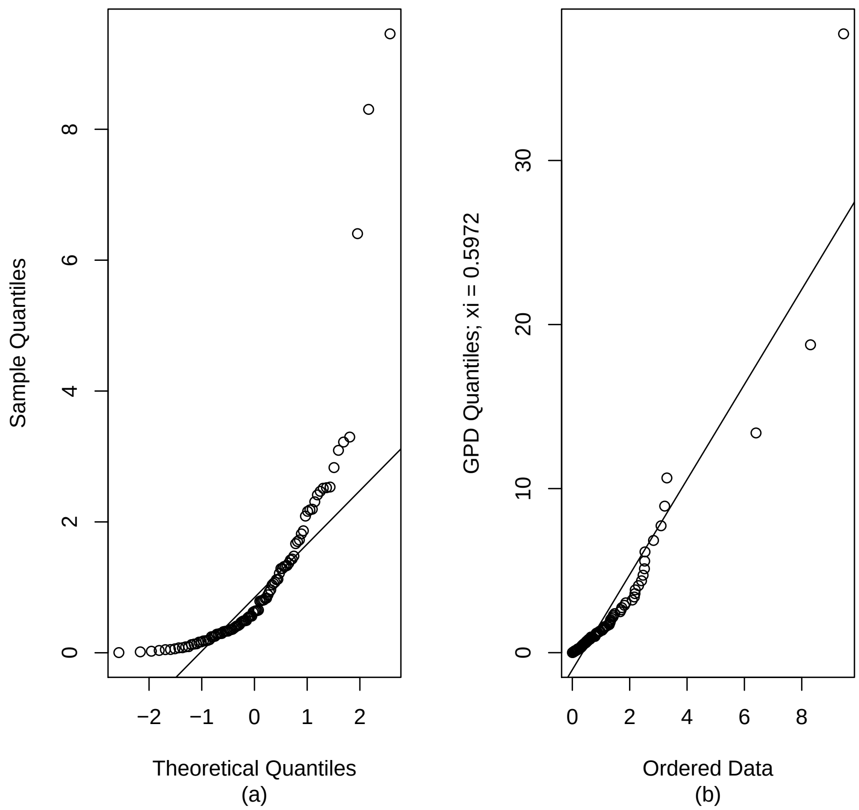

Figure 3.3: (a) Normal Q-Q Plot and (b) GPD Q-Q Plot

the GPD line more closely. We apply these results to the residuals of our SARIMA model to obtain the EVT-SARIMA model. We compare the EVT-SARIMA model to the conventional SARIMA model and the EVT-GARCH model in Table 3.1.

For the EVT-GARCH model, we fit a GARCH model to the data series and then apply EVT to the residuals, as above. We then combine the original GARCH model for the data series and an EVT based model for the residuals to obtain the EVT-GARCH model, similar to the methodology we used for the EVT-SARIMA model.

Table 3.1: Model Statistics

Price Series	Mean Price ($'s per MWh)	Variance	MSE
Actual	92.66	5135.64	-
SARIMA	92.69	2.46	5139.1
EVT-SARIMA	97.16	1030.41	3751.8
EVT-GARCH[1]	109.62	5737.93	10894.4

The conventional SARIMA model has the closest mean price to the actual data series whereas the EVT-SARIMA model and the EVT-GARCH model both appear to overestimate the mean price. The SARIMA and EVT-SARIMA model both underestimate the volatility of the actual price data, but the EVT-GARCH model and the actual price series have relatively close values for variance. This suggests that the EVT-GARCH model is the most appropriate for capturing the volatile nature of the data series. However, by visual comparison, we can see that the EVT-SARIMA model provides a closer estimate

[1] We also considered a conventional GARCH model for the data series. However, the EVT-GARCH model was found to be the more accurate model for the given data series compared with the conventional GARCH model and thus, we have only included the results from the EVT-GARCH model in our comparisons.

to the actual data, shown in Figure 3.4, than the EVT-GARCH model, given in Figure A.1.

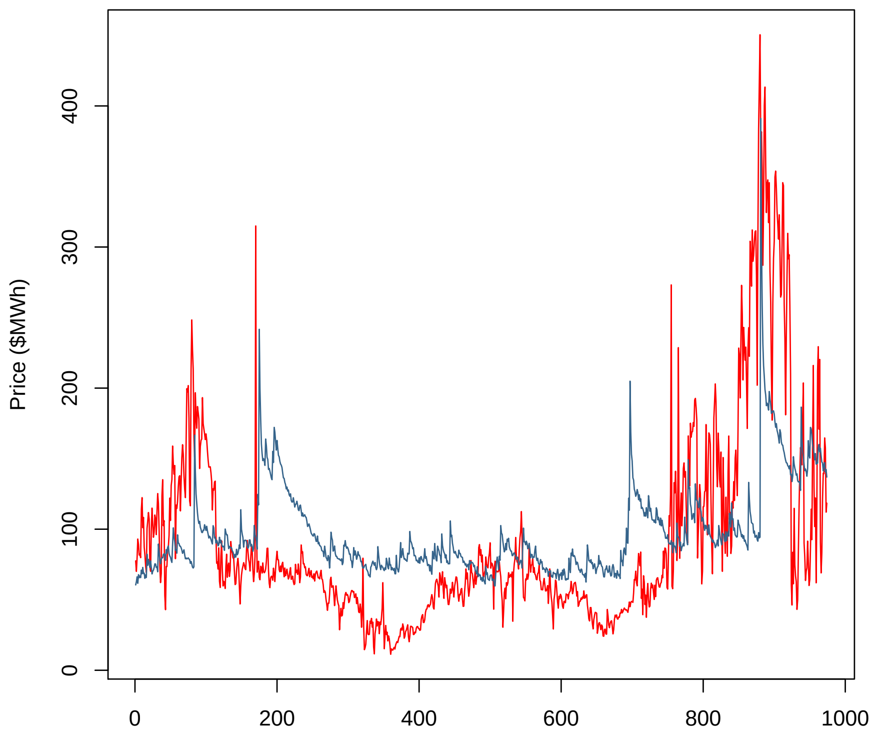

Figure 3.4: Forecasted Price Series using EVT-SARIMA Model (blue) and Actual price Series (red)

The mean squared error (MSE) supports this conclusion, as the EVT-SARIMA model has a much smaller MSE compared with that of the EVT-GARCH model. Indeed, the EVT-GARCH model may be a better predictor of variance, but it does not model the nature of the data series well. We can see in Figure A.1, that in periods of "low" volatility the EVT-GARCH model tends to overestimate the volatility in the electricity price. The random spikes in prices estimated by the EVT-GARCH model appear as single large jumps in price, which return to the lower price level immediately after. Moreover, the model tends to largely overestimate any random jump in the electricity price. This is not an accurate illustration of the characteristics of the actual electricity price data series.

From a risk manager's perspective, it is the random jumps that are of greater interest, as it is these random spikes in price that one generally wishes to hedge against. That is, the EVT-SARIMA model more appropriately estimates periods of low price volatility and random periods of high price volatility. Thus, we conclude that the EVT-SARIMA model is the most appropriate model for forecasting electricity prices for the purposes of risk management.

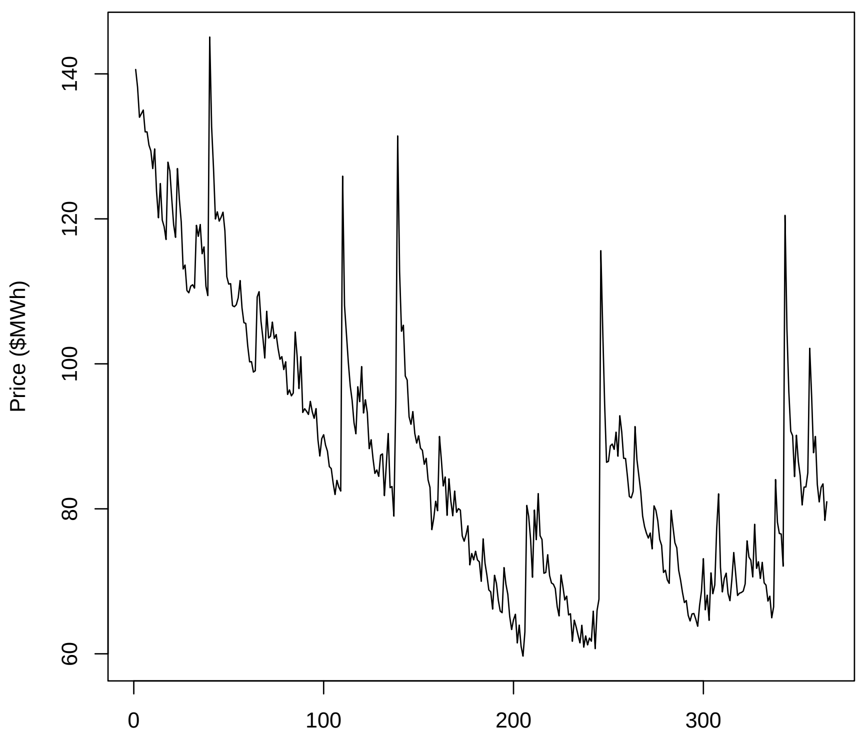

Figure 3.5: One Year Forecasted Electricity Price Series using EVT-SARIMA Model

Figure 3.5 is an example of a 1 year forecast of electricity prices that one may use to consider different hedging strategies. We obtain this series with our EVT-SARIMA model. We may also include the 95% upper and lower limits for our forecast series.

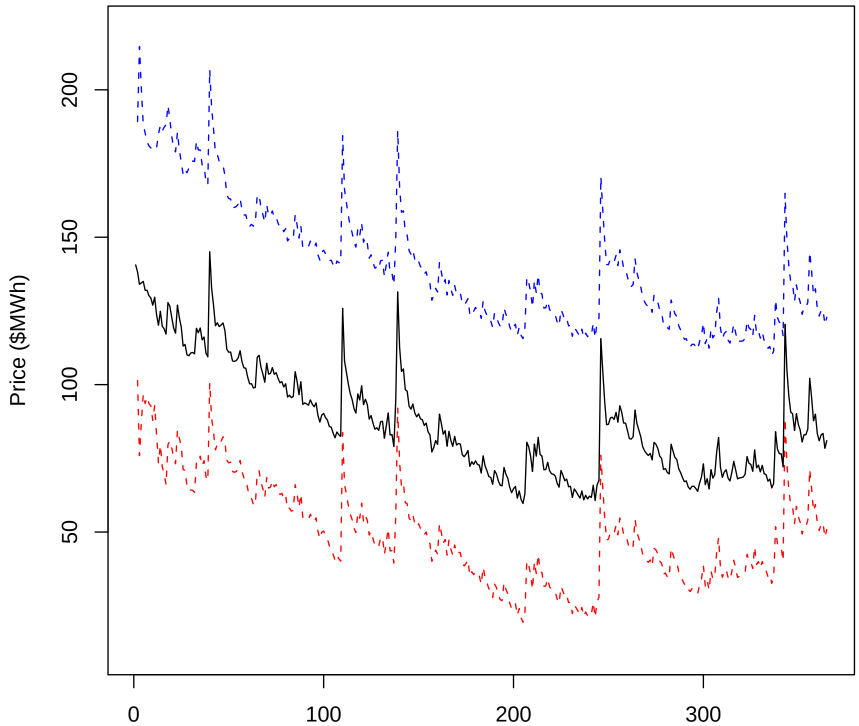

Figure 3.6: One Year Forecasted Electricity Price Series using EVT-SARIMA Model with 95% Upper and Lower Limits
Upper Limit (blue dashed line), predicted values (solid black line) and lower limit (red dashed line)

In the following chapter, we introduce a number of different strategies to hedge risk and observe the effect on revenue of these different types of contracts based on 1000 simulated price series, such as the one given in Figure 3.5 or Figure 3.6.

Chapter 4

Risk Management

In this chapter, we consider possible hedging strategies an electricity generator could use to manage the risk associated with price volatility. We base these strategies on the forecast prices from the model we have developed in the previous chapters.

We first provide a brief overview of the NZ electricity industry to gain an insight into the environment in which a generator is managing their risk. We then consider the relevant literature discussing the importance of risk management and the relationship between the hedging markets and the spot market. The chapter concludes with the application of different contracts and observing their effects on revenue obtained from our forecast prices.

Our methodology for analysing the effect of different contracts on revenue is to consider the cumulative distribution functions (CDFs) for revenue and to observe the changes in these CDFs as we introduce different contracts. To calculate the CDFs, we forecast 90 data points using our EVT-SARIMA model to obtain a simulated set of prices over a 90 day period. We then repeat this process 1000 times to obtain 1000 such simulations. From these simulated price series the 5th, 50th and 95th percentiles were calculated; and using a given set of generation quantities, these gave respectively the three CDFs for revenue plotted later in the chapter: sometimes we refer to these as the 5% CDF, 50% CDF and 95% CDF.

The lower 1st and 5th percentiles of revenue are also considered as inverse Value-at-Risk (VaR) measures. It would be preferable to subtract the 5th percentile of revenue from some benchmark revenue, thus calculating the conventional VaR. But, such a benchmark does not exist, whence we consider Revenue-at-Risk (RaR) instead of VaR.

4.1 The New Zealand Electricity Industry

The New Zealand electricity market consists of generators, transmission network owners, distribution network owners and retailers.

Generators offer different levels of electricity supply at different price levels. These offers are combined with offers from all the other generators and then stacked from lowest to highest price by the System Operator, Transpower. This is essentially the supply curve. Transpower then uses its forecast demand to determine how much electricity is required and sets the spot price from the bid stacks accordingly. Once the price is determined, all the electricity goes into a "pool" through the Grid Injection Points (GIPs). The generators are paid the price determined by Transpower at the GIPs. From the pool of electricity, the electricity flows through the national grid and exits at the Grid Exit Points (GXPs) and from here, distributed throughout New Zealand to end-users.

The Electricity Industry Reform Act (1998) *"prohibits cross involvement between electricity lines businesses with either electricity retail or generation activites"*[1]. This means that generators and retailers cannot own, or invest in, distribution companies. But, it is possible to have vertical integration between a generator and a retailer. There are, however, two exceptions to the legislation. Firstly, generation companies may own the lines that transport electricity from their power stations to the grid or local distribution network. Secondly, distribution companies can own a small amount of conventional generation capacity within their network.

The structure of the New Zealand electricity is similar to other electricity markets around the world in that it is based on bids made one day ahead. However, there are risks that generators face which are unique to New Zealand. For example, a large percentage of New Zealand electricity generation is hydro generation (approximately 60% of total generation). This means that if there is not enough water, supply would be highly affected and this would be seen through an increase in the wholesale electricity price. Moreover, New Zealand's hydro storage capacity is relatively small - New Zealand's hydro storage capacity would be about 6 weeks worth if we became completely dependent on hydro generation, compared with some Scandanavian countries with up to 2-3 years of storage capacity,[2] compounding the New Zealand electricity price susceptibility to weather changes.

In general, players in the electricity market face a highly variable electricity demand func-

[1] For a full copy of the Electricity Industry Reform Act go to www.legislation.govt.nz/act/public/1998/0088/latest/DLM428203.html

[2] Comparison from "Electricity Trading in the New Zealand Electricity Market", *Genesis Energy*

tion as demand is largely dependent on factors such as time of day, or weather. Generators need to ensure that they can supply enough electricity to satisfy demand at all times.

4.2 The Forward Market

The forward market enables players in the electricity market to manage the risk associated with the volatility of the electricity spot price. In general, there are two main types of forward contracts used in electricity markets. One is a conventional forward contract based on the physical delivery of electricity; and the other, a financial agreement on the strike price. In this paper, we only consider short positions in both types of forward contracts. However, any theory we discuss is relevant for a player taking a long position in these contracts, but will have the opposite outcomes.

The first type of contract is a simple fixed price forward contract for a commodity, it is an agreement to sell a particular amount of the commodity on the settlement date for a particular price, the strike price. A brief explanation of conventional forward contracts is provided below and for further detail we refer the reader to Hull (2004). Consider a forward contract between an electricity generator and a retailer. Suppose a generator agrees to supply electricity to a retailer, at a strike price K, and let T be the time to maturity of the contract. That is, at time T, the transaction will occur - the agreed amount of electricity will be delivered and paid for. In general, the initial forward price is given by F_0, where

$$F_0 = S_0 \times e^{rT}$$

with S_0 being the spot price and r is the risk-free interest rate. When entering such a contract, the strike price will be set as the current forward price. Thus, the value of a forward contract is 0 on the day it is first negotiated. We denote the value of a forward contract, for the seller, by

$$f = (K - F_0) \times e^{rT}$$

Hence, if the forward price on the settlement date is smaller than the strike price (the forward price on the date the contract was negotiated), then the forward contract will have a positive value from the seller's point of view. As a result of such a contract, the generator will receive $K \times Q$ at maturity, where Q is the amount of electricity that was agreed to be supplied.

A more common forward contract in the electricity markets is a swap, also known as a“contract for difference” (CFD). Similar to the above conventional forward contracts, such a contract specifies a settlement date, a strike price K and a quantity of electricity Q. However, once the settlement date is reached there is no delivery of electricity as a result of settling the contract. These contracts just specify a payment to be made between parties. That is, to settle the contract there are 2 possibilities: given the spot price on the settlement date S, if $K > S$ then, the buyer will pay the seller $(K - S) \times Q$. But, if the situation arises that $K < S$, then the seller must pay the buyer an amount given by $(S - K) \times Q$. Thus, such contracts are purely financial, whereas the simple fixed price forward contracts lead to physical delivery of electricity. CFDs ensure a certain price, the agreed upon strike price, for a specified quantity of electricity Q, without the generator actually having to deliver this quantity of electricity specified.

Conventional forward contracts and CFDs are both traded over the counter. Most forward trading in electricity markets is between generators and retailers for the purposes of hedging risk. Vertical integration of generators and retailers is also common for the purposes of hedging risk. That is, for a vertically integrated firm, a loss for the retailer as a result of an increase in the wholesale electricity price is offset by the gains received by the generator, and vice versa.

4.2.1 Other Financial Instruments

In financial markets, a common tool for hedging risk is futures contracts. Unlike forward contracts and CFD's, futures contracts are traded on an exchange. In this paper we also consider futures contracts as a means of managing the risk associated with electricity prices.

A futures contract is similar to a forward contract as it is an agreement to buy or sell an asset for a particular price at a particular date. However, most futures contracts do not result in the delivery of the underlying asset on the agreed upon settlement date. This is because most investors close out their positions before maturity of the contract. To do so, investors tend to enter into a trade opposite to the trade agreed upon in the futures contract. Thus, we will consider an investor taking a long position in a futures contract. Similarly to the contracts discussed above, a short position in such a contract will have opposite outcomes to those explained below.

The main difference between futures and forward contracts is that futures contracts are marked to market daily. Daily marking to market means that the holder of a futures contract will realise a profit or loss depending on the futures price at the close of trading every day up until the settlement date. Suppose an investor enters a futures contract at time 0 and we are now at time t ($t \neq 0$). On the close of trading today we have a futures price f_t. Then, an investor would make a profit/loss given by the difference in today's futures price and yesterday's futures price, that is, a profit/loss equal to $f_t - f_{t-1}$.

The daily gains or losses made as a result of marking to market are usually deposited into a margin account for the duration of the contract. In general, brokers allow investors to earn interest on the balance in their margin accounts. As a result, there exists a relationship between the futures price and the interest rate. Indeed, when the interest rate is deterministic, we would expect the futures and forward prices to be the same. However, as the interest rate changes unpredictably, the interest that one can earn on the balance of the margin account will also change unpredictably. Suppose the price S of the underlying asset increases when interest rates rise. If S increases, an investor holding a futures contract will make a profit which will be deposited in their margin account. The relationship between S and the interest rate implies that this increase in S will also mean an increase in the interest rate. Thus, the gain will be invested at a higher, more profitable, rate of interest. Conversely, a decrease in S will lead to a loss, but also a decrease in the interest rate implying that any interest earnings foregone will be small. Hence, when there is a strong positive correlation between S and the interest rate, we would expect futures prices to be higher than forward prices and vice versa. For further discussion, see Cox, Ingersoll & Ross (1981).

Throughout this chapter, we assume that the strike price of forward contracts and futures are the same given that our contracts are over a short period of 90 days, as discussed in Hull (2004). For the futures contracts, contract revenue includes the interest earned over the 90 day period on the balance in the margin account. After marking to market daily, the daily interest is calculated. The interest earned remains in the margin account and accumulates over the 90 day period. In considering the interest earned on any balance in the margin account from our futures contracts, we firstly assume a deterministic interest rate. We then introduce a stochastic interest rate instead and consider the affect on overall revenue of this stochastic interest rate, compared to using a deterministic interest rate.

4.3 The Forward and Spot Market: A Brief Literature Review

Now, we consider the effect the forward market has on the spot market. The availability of forward contracts is likely to lead to a reduction in the spot price, as mentioned in Anderson, Hu & Winchester (2007) and discussed further in Wolak (2000) and Green (1999). Wolak (2000) shows that if a firm exhibits high levels of contracting, it is rational for the firm to offer very aggressive bids to sell as much of its "left over" capacity as it can on the wholesale market - below, we refer to this "left-over" capacity as net residual demand. As a result of this aggressive bidding, the spot price will be driven downwards.

The existence of vertically integrated firms, where a generator is the owner of a retailer, must also be taken into consideration. McRae & Wolak (2009) suggest there are two opposing effects on a vertically integrated firms' profits when there is an increase in the spot price. The firm's profits from the spot market will increase and secondly, the firm's profits will decrease as it becomes more expensive to satisfy its retail demand. Thus, a firm may choose to withhold electricity to increase the spot price if the second effect is smaller than the first. McRae & Wolak (2009) finds that the larger the net residual demand a firm faces, the greater the generator's gain from a spot price increase. Using the notation given in McRae & Wolak (2009), $DR(p)$ is the quantity sold on the wholesale market, Q_R is the amount of retail electricity sold and Q_C is the quantity of electricity required for contractual obligations. Then, net residual demand $DR_F(p)$ is given by

$$DR_F(p) = DR(p) - (Q_R + Q_C)$$

This equation suggests that if the firm increases the amount sold through contracts (an increase in Q_C)), the net residual demand will decrease. Then, the firm's gain from a spot price increase will also decrease. Again, this would suggest that as the number of contracts a firm enters into increases, the firm will bid more aggressively on the wholesale market driving the market price down. This is due to a market price increase having little effect and thus, the firm having little incentive to drive the market price upwards.

In general, it is assumed that the price of a forward contract will be equal to the expected spot price. However, if a generator were to enter a contract their profits would be reduced as a result of this assumption. Thus, we are left questioning what incentive a generator has to enter into a forward contract. Wolak (2000) shows that contracts reduce the variability in profits, but this reduction in variability comes at a cost of a reduction in average profits. That is, there exists a trade-off between variability in profit and the average profit level when considering the appropriate amount of contract cover. The discussion in Wolak (2000) was based on the analysis of one firm's strategic decisions, and the effect of these decisions, given that all other firms in the market did not alter their bidding strategy. Green (1999) shows that a firm's decisions with regards to the amount of contract cover is dependent on what beliefs it has about their rivals behaviour. It is suggested that if a firm assumes that entering into a contract will reduce the contract cover of their rivals, then they are likely to enter into the contract.

In a report for the New Zealand Commerce Commission, Wolak concluded that *"the ability and incentive of large suppliers to exercise unilateral market power are important determinants of supply conditions that determine short-term wholesale prices, even after the impact of exogenous factors such as water availability and fossil fuel prices have been taken into account".*[3] Thus, the level of market power a generator has will affect the wholesale electricity price. For example, a generator with a large share of market power may have incentive to withhold electricity in storage or, offer higher prices in order to drive spot prices upwards.

When considering forward contracts, the assumption of equality between the forward contract price and the expected spot price is questionable. In Evans & Guthrie (2009), transaction costs are introduced into the storage process of a competitive equilibrium storage model. As a result of these transaction costs, they find that there will exist a difference in the spot price and the market value of the stored commodity. This price difference, or "friction-induced price divergence" leads to a convenience yield. A convenience yield is defined as a measure of the benefits from the ownership of the commodity that are not

[3] For full report, see www.comcom.govt.nz/BusinessCompetition/Publications/Electricityreport/DecisionsList.aspx

obtained by the owner of a contract for future delivery of the commodity. These benefits, for example, may be the reduction in costs of supplying varying output, or the ability to take immediate advantage of any price changes. Thus, in some cases, these benefits may outweigh any gains from immediate sale of the commodity and may induce the firm to hold the commodity. Hence, a premium may be required in order to entice generators into entering a contract, particularly to counter the negative impact of forward contracts on a generator's market power in the spot market, as suggested in Anderson & Hu (2005).

There are few tools available specifically for players in the electricity market to manage their risk. Generators and retailers can enter into private agreements, tailored for their individual requirements. But, few trading platforms or exchanges exist where players in the New Zealand electricity market have access to standardised contracts. The trading platform "energyhedge"[4] is for futures contracts which cover a period of approximately 5 years. Contracts are for a notional quantity of 250kW, using values from the Otahuhu, Haywards or Benmore grid reference points. However, this is not an authorised futures exchange and is not regulated under NZ law. More recently, futures and options have become available on the Australian Stock Exchange (ASX). These contracts use both Otahuhu and Benmore as the grid reference nodes.

If we compare the tools available on *energyhedge* and the ASX with exchanges such as NYMEX (New York Mercantile Exchange) or Energy Exchange Austria, we can see that the options available to New Zealand electricity market participants are relatively small. Both of the foreign exchanges are purely commodity based and have a comparatively large number of electricity based financial instruments available to traders.

[4] see www.energyhedge.co.nz"

4.4 Application

We are interested in observing the effect these contracts have on the revenue a generator may receive. To do so, we consider a cumulative distribution function (CDF) for the generator's revenue. We use a generation profile of daily average electricity generation[5] and throughout this paper we assume that the generation profile does not change. Using this generation profile and 1000 simulated price series of length 90 we obtain a CDF for revenue. Then, we see how different contracts affect the CDF.

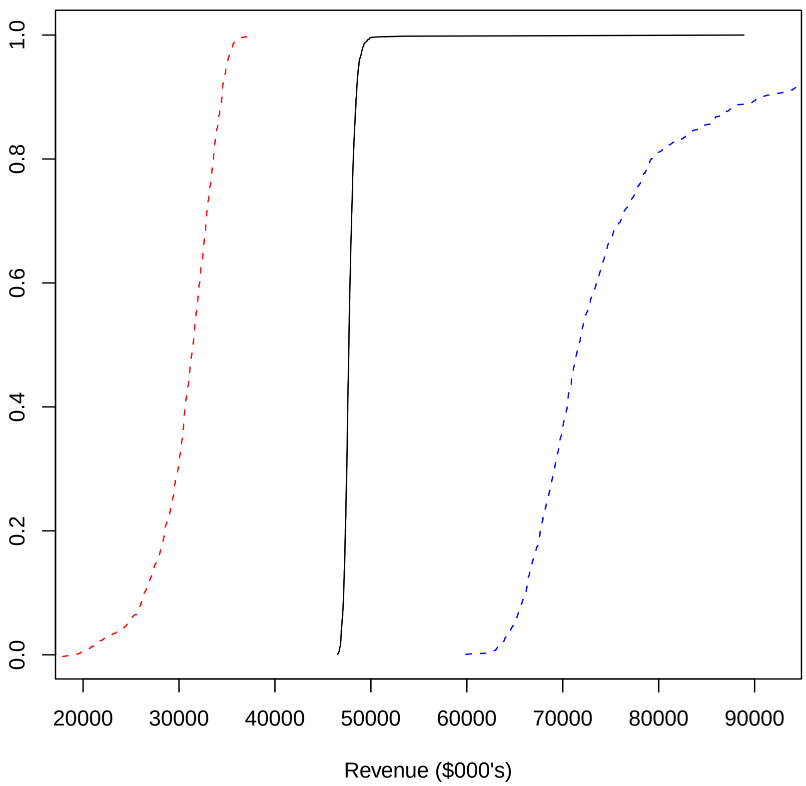

Figure 4.1: CDF for Revenue, no contracts
95% lower limit (dashed red line), predicted values (black) , 95% upper limit (dashed blue line)

[5]based on daily average generation from Benmore

There is a large range of possible revenues given no contractual obligations, as seen in Figure 4.1. If we consider the predicted values, the possible revenues range from \$46.5m to just under \$89m. We see that the lower tail of the CDF for revenue is very short, almost non existent. Thus, large increases in probability result from relatively small increases in revenue. This is further depicted by the very steep, almost vertical section of the CDF which suggests there is very little spread in the lower revenue values. A flat upper tail suggests that at larger revenue levels, there is a small increase in probability given a large increase in revenue. Hence, there is a relatively larger range of high levels of revenue available, compared with the possible range of low revenues. Given the overall range of possible revenues, it may be difficult to accurately budget for the future. Indeed, the large range of high revenues is favourable. However, a risk manager may be willing to forego the possibility of larger revenues for a smaller range of overall possible revenue values, thus creating more certainty in the level of revenue that will be obtained.

4.5 The Effect of Quantity Traded Through Contracts

We begin with a portfolio of conventional forward contracts negotiated for a 90 day period and assume the quantity of electricity generated is given over this period. We define total revenue as revenue from the spot market over the 90 day period (wholesale revenue) and contract revenue settled at the date of maturity. We take contract revenue as the amount of electricity traded multiplied by the agreed strike price.

We then consider a type of contract that is more commonly used in the electricity industry, contracts for differences (CFDs). As explained above, both parties agree on a particular strike price and on the settlement date, one party pays the other a difference between the strike price and the spot price for each MWh (or GWh) of electricity traded without any electricity being delivered as a result of the contract.

The method of calculating total revenue for CFDs is similar to that given for conventional forward contracts. However, now we have all electricity generated being sold on the spot market and contract revenue being the difference between the strike price and final spot price multiplied by the agreed upon quantity of electricity. Again, contract revenue is determined at the settlement date.

Finally, we consider futures contracts. To calculate the revenue from futures contracts, we again assume all generated electricity is sold on the wholesale market. We then mark to market at the close of each day and any profit or loss is immediately realised. We also calculate the interest earned on a daily basis on any balance in the margin account, accumulating this 'bank account' over time.

For all contracts, we begin by assuming an agreement is made to deliver a total of 3.5 GWh per day as a result of contractual obligations. We define the strike price for all contracts as $K = F_0$, as defined above for forward contracts. To determine F_0, we forecast 1000 price series of length 90 and use the very first price of each series as our value for S_0 and assume a constant rate of interest of $r = 5\%$ p.a. We then increase the amount of electricity supplied to contract holders to 4 GWh and then decrease it to 3 GWh per day and observe the effect of such changes on the CDFs for revenue. A summary of our results is given at the end of the chapter in Table 4.1.

Table 4.1: Summary of effect on revenue of different contracts given changes in amount of electricity traded through contracts

Contract	Amount Traded through Contracts (GWh)	Range of Predicted Overall Revenue from 50% CDF
Conventional	3	\$46.2m - \$82.8m (36.6)
Forward Contracts	3.5	\$46.1m - \$81.6m (35.5)
	4	\$46.0m - \$80.8m (34.8)
CFDs	3	\$198 000 - \$89.7m (89.5)
	3.5	-\$8m - \$89.9m (97.9)
	4	-\$16.1m - \$90m (106.1)
Futures	3	\$47m - \$83.7m (36.7)
	3.5	\$46.4m - \$81.8m (35.4)
	4	\$45.4m - \$80.2m (34.8)

For all conventional forward contracts and futures, the range of predicted revenue has decreased compared with having no contracts. This suggests that by using these contracts, one can decrease the variance in revenue obtained over the specified period. As a result, the ability to plan for future investments is greatly enhanced as the risk manager is more certain of the likely revenue they will receive. However, this comes at a cost. As mentioned above with respect to profits, we also see that there is a trade-off between variability in predicted revenue and the level of predicted revenue. That is, the introduction of contracts leads to a smaller range of predicted revenue values but also a reduction in overall predicted revenue.

This trade-off can further be seen when we increase the quantity of electricity traded through forward contracts and futures. As the amount traded through these contracts is increased, the range of predicted revenue values and the overall predicted revenue decreases further. We observe that futures contracts enable the largest reduction in the range of predicted revenue and generally maintains a similar level of predicted revenue to that of the forward contracts.

The introduction of CFDs, compared with predicted revenue given no contracts, leads to an increase in the predicted level of revenue. However, such contracts also result in the possibility of negative revenues and thus, a much larger range of predicted revenue values. Hence, one would use CFDs for the purposes of increasing the level of predicted revenue rather than reducing variability in predicted revenue. As we increase the amount of electricity traded through CFDs, we increase the level of predicted revenue. But, we also increase the range of predicted revenue. So, in using CFDs as a means of managing risk one must balance the desire for higher revenue against the risk of possibly lower revenue values and hence, an increase in the overall range of predicted revenue values. In Figure 4.2, we magnify the tails of each CDF curve given different quantities of electricity being traded. From Figure 4.2(a), we can see that the contracts where a total of 4 GWh per day was traded resulted in a CDF with the lowest predicted revenue value. As we decrease the quantity of electricity traded, we see that the lowest predicted revenue value is no longer negative. Indeed, by decreasing the quantity traded, the minimum predicted revenue of \$-16m increases to \$198 000. However, if we consider the upper tails, we observe that the highest predicted revenue also results from trading the highest quantity of electricity. Thus, in using CFDs as a means of managing risk one must balance the desire for higher revenue against the risk of possibly lower revenue values and hence, an increase in the overall range of predicted revenue values.

Our results imply that there is scope with CFDs to obtain a much larger revenue than with no contracts, conventional forward contracts or futures. However, we see that this type of contract also means that large negative revenues may be faced by the seller over

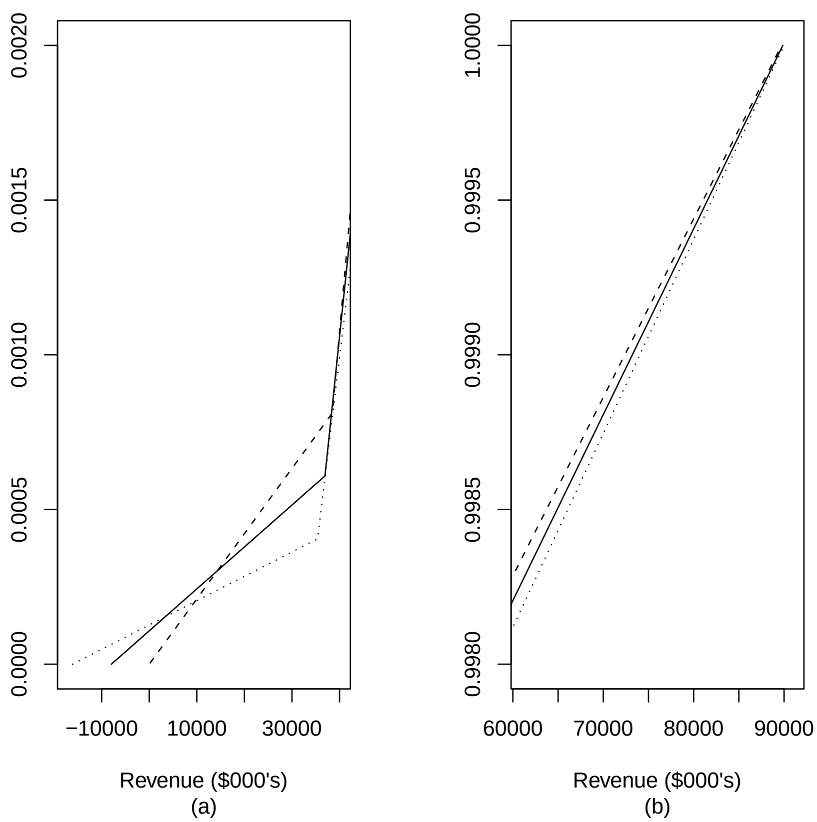

Figure 4.2: Tails of CDF for Revenue, with CFDs trading:
3 GWh (dashed line), 3.5 GWh (solid line), 4 GWh (dotted line) per day
(a) Lower Tail, (b) Upper Tail

the 90 day period. Thus, perhaps a risk manager would include CFDs in their portfolio for hedging risk to increase the possible levels of revenue, but would require some other hedging instrument to hedge against the possible negative revenue as a result of the CFDs. For example, one could use a mix of CFDs and futures so that the futures may be used to counter the increase in revenue spread caused by the CFDs and reduce the risk of negative revenue values. The CFDs would act as a means of increasing the larger predicted revenue values. We have not further investigated the profitability of using linear combinations of derivative contracts.

4.6 The Effect of the Interest Rate

In this section, we consider the effect the interest rate has on predicted revenue given a firm has entered into some type of contract. From Figure 4.3, we can see that as we increase the interest rate, the CDF for revenue moves to the right. Thus, the overall revenue we expect to obtain increases as a result of an increase in the interest rate. This can be seen further in Figure B.4, where we have magnified the CDFs for revenue given each particular interest rate. We observe that the CDF for revenue that is farthest to the right corresponds to the highest interest rate. This implies that a high interest rate leads to higher overall revenue given forward contracts. Hence, there appears to be a positive relationship between the interest rate and the electricity price. We discuss possible reasons for this relationship later in the section.

Now, we consider the effect of the interest rate on revenue given CFDs. In Figure 4.4, we can see that the interest rate has very little effect. Visually, it is difficult to see a

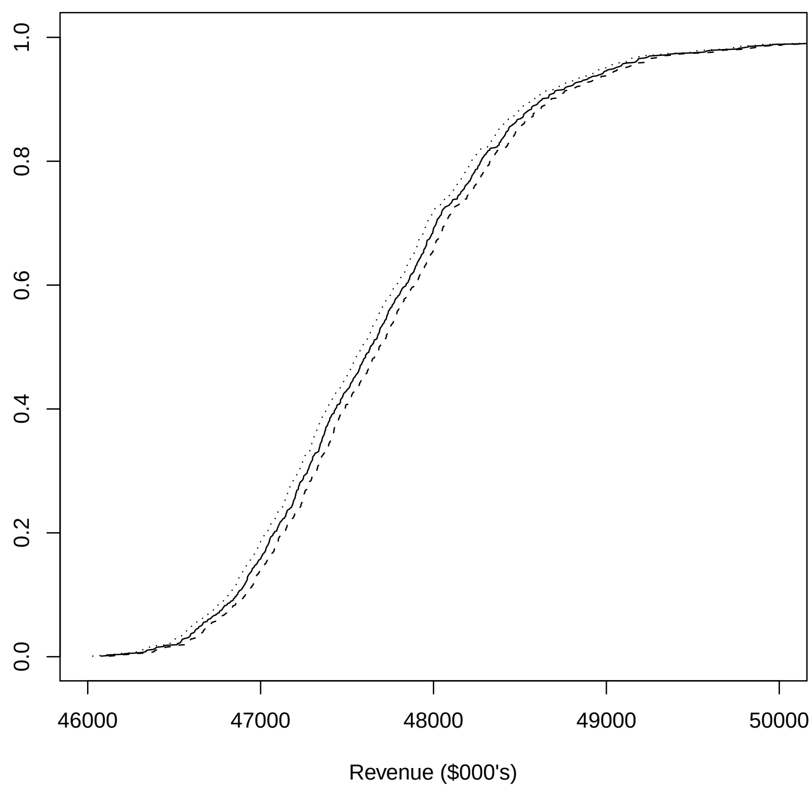

Figure 4.3: CDF for Revenue, with forward contracts trading 3.5 Gwh per day interest rate = 3% (dotted line), 5% (solid line), 7% (dashed line)

difference in the two CDFs for each interest rate, as seen in Figure 4.4. However, if we consider numerical values, the revenue values range from -\$8.0m to \$89.8 given an interest rate of 3%. If we consider an interest rate of 7%, the values range from -\$7.9m to \$89.9m. Thus, the interest rate has a relatively small effect on revenue. We note that as we increase the interest rate, there is a slight increase in the range of predicted revenue values. Hence, as the interest rate increases, the maximum predicted revenue value appears to increase slightly as well.

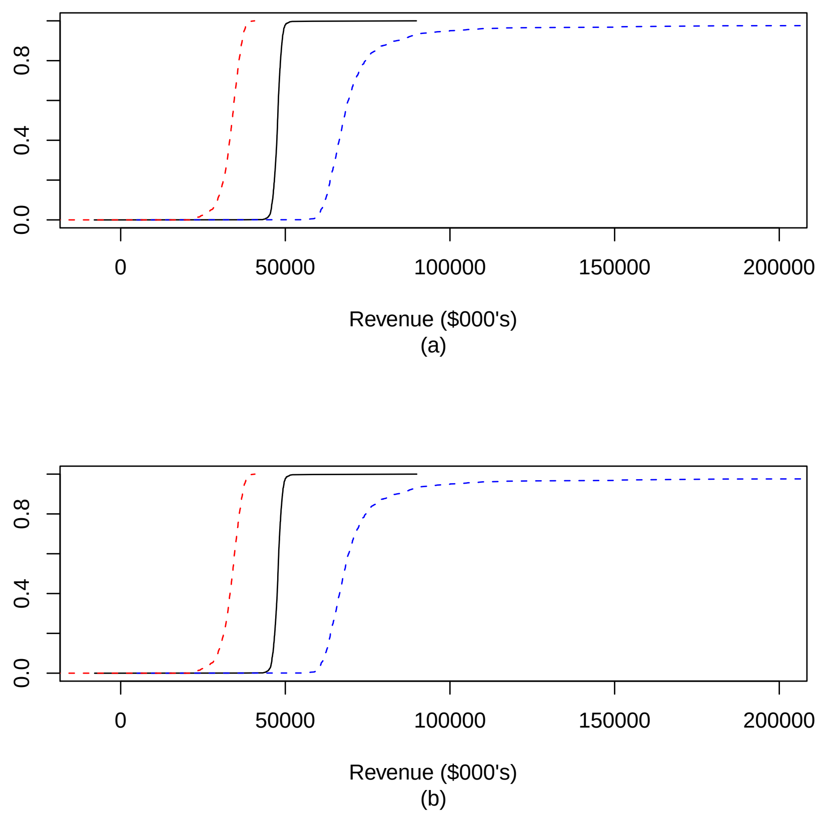

Figure 4.4: Tails of CDF for Revenue, with CFDs trading 3.5 GWh (solid line) (a) interest rate = 3%, (b) interest rate = 7%

We can see the CDFs for revenue curves at different interest rates in Figures D.1 and D.2. The effect of interest rates is very small, as shown in D.3, particularly given our

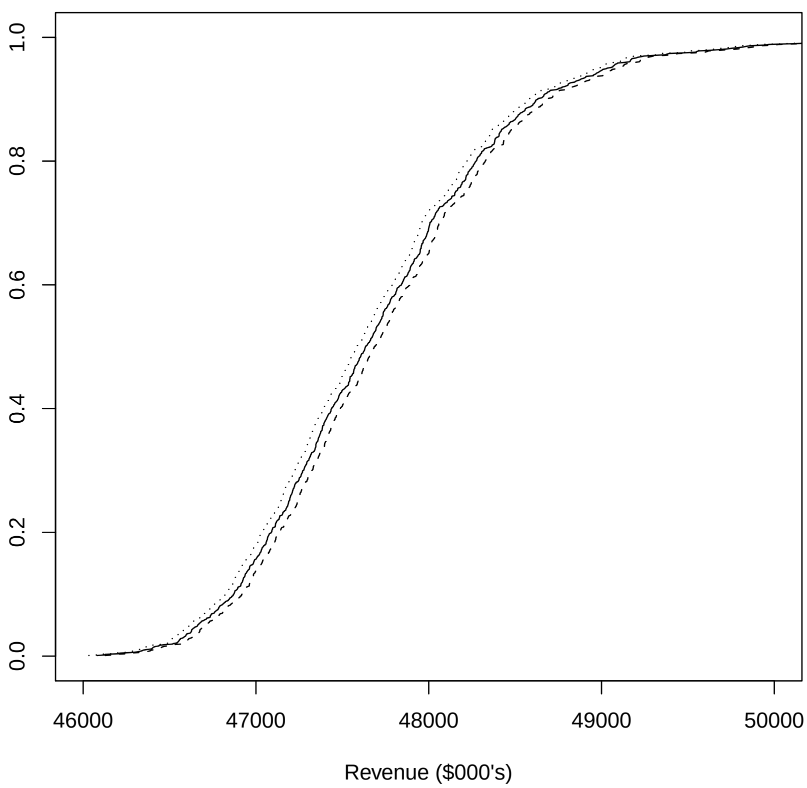

Figure 4.5: CDF for Revenue, with futures contracts trading 3.5 GWh per day with interest rate: 7%(dashed line), 5% (solid line), 3% (dotted line) per day

assumption that the initial futures price and the forward price are the same. That is, since we are considering contracts over a relatively short period, the possible interest that could be earned in a margin account is small and the effect of interest rates on the revenue from a futures contract appears minimal. However, if we were to consider contracts over a longer period, we would see that interest rates would have a more noticeable effect on revenue and so, we could no longer assume the forward and futures price to be equal.

In 4.5, we have magnified the different CDFs to see the effect of interest rates on the CDF for revenue. The small change in revenue that we do observe as a result of a change in interest rate is that as we increase the interest rate, the overall predicted revenue values increase. We also observe that the range of predicted revenues decreases. Our results suggest that electricity prices and interest rates may be positively correlated. Hence, as the interest rate increases, we would expect the futures price to be higher than the forward price.

4.6.1 Stochastic Interest Rate

So far, we have assumed that the interest rate has been deterministic. However, it's more realistic to consider a stochastic interest rate as we frequently observe unpredictable changes in the interest rate in the real world. Thus, we introduce a random component to our interest rate when forecasting revenue.

We include the random component of the interest rate by assuming the interest rate changes daily over the 90 day period. We define a sequence of random daily interest rates given by

$$r_{st} = r + 0.01 \times z$$

where z is a standard Normal random variable. We have taken the 'baseline' interest rate to be $r = 5\%$; and r_{st} is the stochastic interest rate per annum (p.a), compounded continuously.

We calculate the strike price and interest made on the balance of the margin account using this sequence of interest rates rather than using a deterministic rate of 5% throughout the duration of the contract.

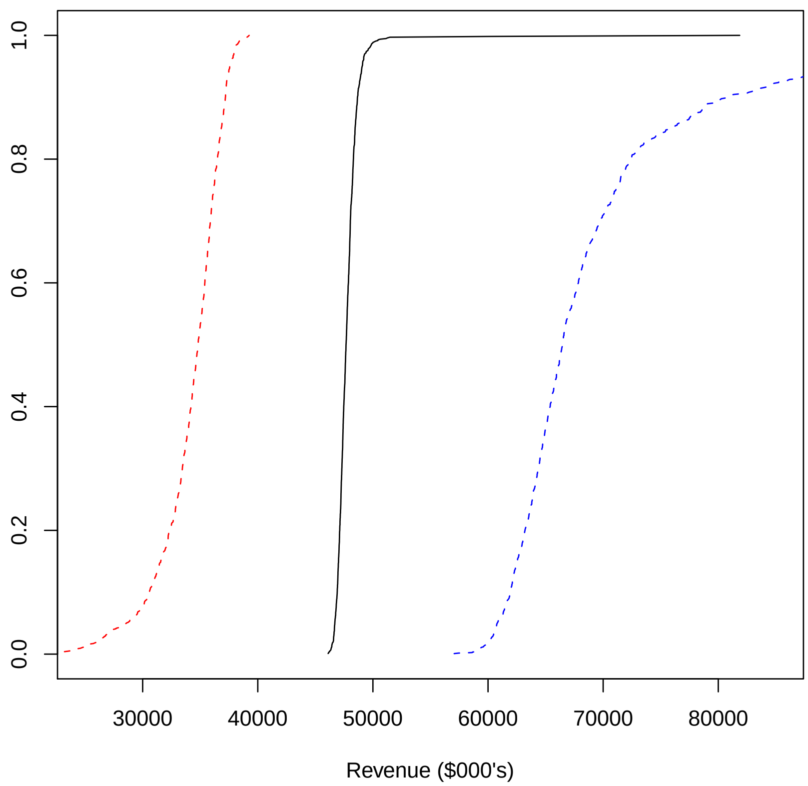

Figure 4.6: CDF for Revenue, with futures contracts trading 3.5 GWh per day with stochastic interest rate

We observe that over such a short period of time, the generated revenue given futures contracts with a stochastic interest rate is similar to the revenue using a deterministic interest rate. If we consider the range of predicted revenue values, we find that the introduction of the stochastic interest rate results in higher predicted values overall. The predicted revenue values range from \$46m to \$82m given a stochastic interest rate. However, if we assume a deterministic interest rate, the predicted revenue values range from \$36m to \$71m, which is lower overall compared with the values using a stochastic interest rate.

As mentioned previously, the CDFs for revenue given futures or conventional forward contracts are almost the same. However, when we consider a stochastic interest rate for the futures we can see a noticable difference in the two types of contract, compared with the predicted revenue using a deterministic interest rate. This can be seen in Figure 4.7.

From Figure 4.7 we see that the stochastic interest rate leads to larger revenue values overall. Indeed, we can see that in general the introduction of futures compared with forward contracts results in higher predicted revenue. Moreover, the positive deviation from the revenue CDF given forward contracts is larger for the revenue CDF given futures contracts and a stochastic interest rate compared with the CDF given a deterministic interest rate. This deviation can further be seen in the tails of each CDF, as shown in Figure 4.8.

We again observe in Figure 4.8 that the CDF for revenue given futures contracts and a stochastic interest rate lies to the right of the CDF given conventional forward contracts. Thus, futures contracts with a stochastic interest rate appear to result in higher overall revenue values compared with revenue given conventional forward contracts or futures

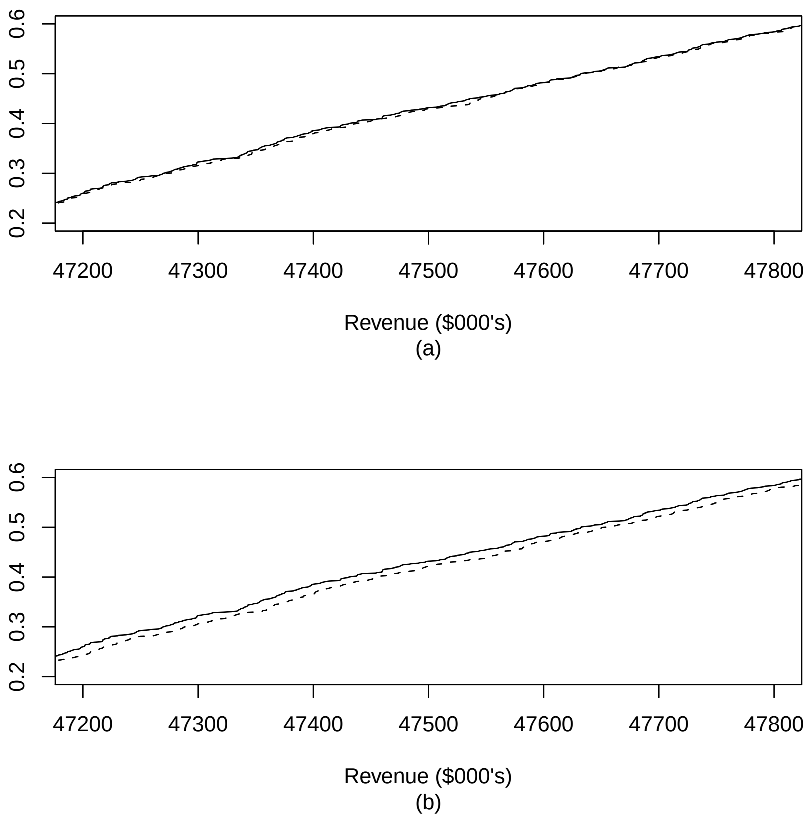

Figure 4.7: CDF for Revenue, with futures contracts (dashed line) trading 3.5 GWh per day compared with conventional forward contracts (solid line):
(a) with deterministic interest rate, (b) with stochastic interest rate

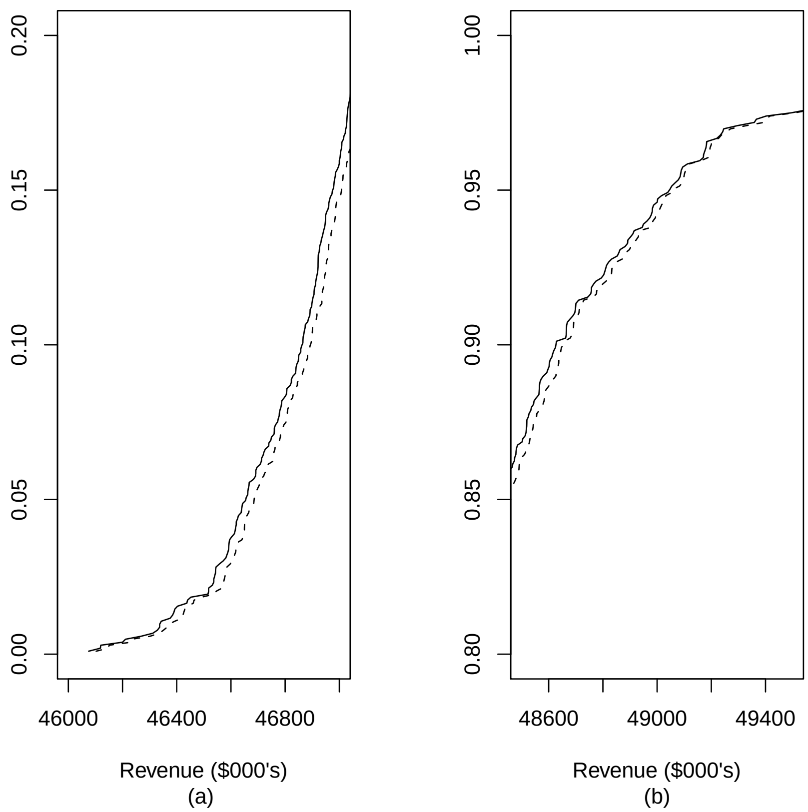

Figure 4.8: CDF for Revenue, with futures contracts and stochastic interest rate (dashed line) trading 3.5 GWh per day compared with forward contracts (solid line): (a) lower tail, (b) upper tail

contracts and a deterministic interest rate. Further discussion in the use of hedging with futures in the electricity industry is provided in Tanlapco, Lawarree & Liu (2002).

There are other advantages to entering futures rather than conventional forward contracts besides the higher predicted revenue. One main advantage is that there is virtually no credit risk with futures contracts due to the margin account and marking to market daily. However, with forward contracts there is some credit risk. Hull (2008) discusses further the differences of futures and forward contracts.

For the rest of the paper, when we draw comparisons between the different contracts we will use our results using a stochastic interest rate when considering the effect of futures contracts.

4.6.2 The Relationship Between Profitability and the Interest Rate

As we have observed empirically, from Figure 4.7 on page 52, the futures contract is more profitable than the conventional forward contract. This indicates that the interest rate and the futures price are positively correlated. Our results suggest that if we were to observe an increase in the interest rate, we would also expect an increase in the electricity price. Moreover, we have seen that our simplistic model of a stochastic interest rate led to more dramatic results in the level of revenue predicted. Indeed, we observed a relatively large increase in predicted revenue as a result of introducing the stochastic interest rate.

There does not appear to be any obvious reason for the existence of this relationship. A possible reason for the apparent positive relationship between the profitability of a contract and the interest rate may be due to our assumption that the generator's behaviour is given. If we were to model the generator's behaviour we may find that this 'relationship' disappears as the amount generated and sold also changes over time.

One could also argue that both the interest rate and the electricity price act as an indicator of the current economic conditions and thus may appear to be related to one another. That is, generally in times of high inflation we observe an increase in the rate of interest as the interest rate is used as a tool for reducing inflation. Generally, increases in the price of commodities such as electricity contribute towards higher inflation levels, as electricity is considered a necessity by most consumers. Thus, in such economic conditions, one would observe a high interest rate and high electricity prices. Conversely, in periods of low inflation, we would expect relatively low electricity prices. The interest rate would also be lower given less need to curb inflation through the interest rate and possibly the need to stimulate the economy. This finding is in accord with our empirical observations, which apparently indicate a positive correlation between the interest rate and the futures price.

There is scope for further research to determine whether there does indeed exist a relationship in the electricity price and the interest rate and if so, what causes this relationship.

4.6.3 Comparison of Contracts

As mentioned above, a mix of different contracts may be the most appropriate method of hedging the risk associated with variances in revenue, or spot price volatility. Given no contracts, we can see that overall revenue values are generally higher than with contracts.

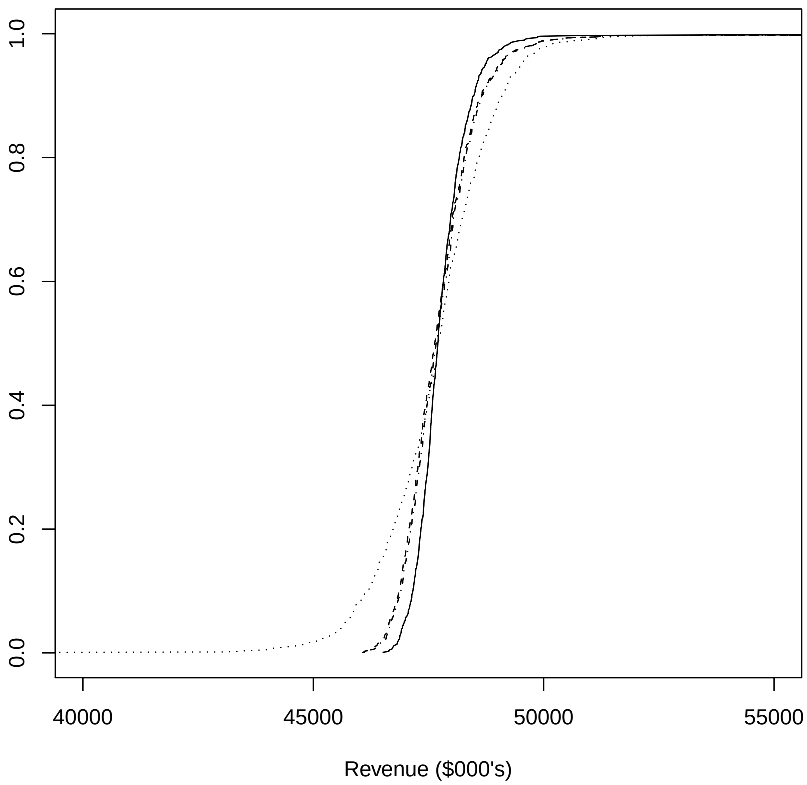

Figure 4.9: Comparison of CDFs for Revenue, no contracts (solid line), forward contracts (dashed line), CFDs (dotted line) both with r=5% and futures contracts with stochastic interest rate (dash-dot line) all contracts trading 3.5 GWh per day

The introduction of either conventional forward contracts or futures contracts result in a decrease in overall revenue, but we also obtain a more compact CDF. That is, we observe a decrease in the variability of revenue values. We observe that the difference in the CDFs for revenue given conventional forward contracts or futures contracts is relatively small. In Figure 4.9, we can see that the predicted revenue given futures contracts or conventional forward contracts appears identical. If we consider the mean revenue though, we have \$47.757m given futures and \$47.746m given conventional forward contracts. Thus, futures contracts result in slightly higher predicted revenue than conventional forward contracts.

Both conventional forward and futures contracts enable a risk manager to be more certain about the level of revenue that will be obtained, as they reduce the range of predicted values. However, this comes at the cost of higher revenue values. The possibility of high levels of revenue is also enticing for a risk manager. Indeed, it may be appropriate to minimise the variance in revenue obtained but one is also looking to ensure that the firm maximises profit. The CDF for revenue given CFDs has the largest range of values and has both the smallest and largest predicted revenue value. We see that it is only with CFDs that there is a possibility of obtaining a negative revenue.

A combination of contracts appears to be the most appropriate strategy. One could use CFDs to try and maximise revenue levels but balance their portfolio with forwards or futures contracts such that a minimum level of revenue is maintained. These contracts will also offset the increase in the range of predicted revenue values resulting from the CFDs. Hence, one could obtain a portfolio which reduced variability in revenues, yet ensured there was an opportunity to obtain high levels of revenue while still maintaining a minimum level of revenue through a combination of the contracts we have discussed.

As we have mentioned, certainty in revenue obtained is important to a risk manager. Thus, it is favourable to manipulate the CDF for revenue so that the range of predicted outcomes is small. One may also consider the 95% upper and lower limits and see what effect contracts have on these boundaries. That is, if we introduce trading of electricity through contracts and the upper and lower limit values aren't as different to the predicted values compared with no contracts, then this would suggest the level of certainty of the predicted values has increased as a result of the addition of contracts. Hence, if the difference between the upper and lower limits is small, we can predict values with greater certainty as we will have reduced the range of likely revenue values.

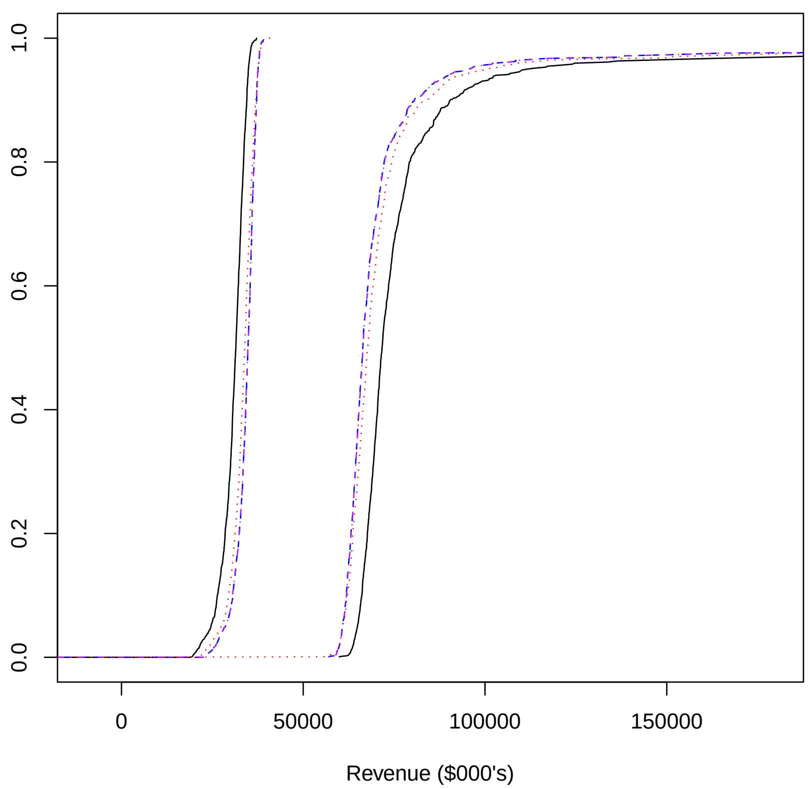

Figure 4.10: 95% Lower and Upper Limits of CDFs for Revenue, all trading 3.5 GWh per day no contracts (solid line, black), forward contracts (dashed line, blue), CFDs (dotted line, red) all with r=5% and futures contracts (dash-dot line, purple)

We find that in all cases, with the introduction of contracts, the 95% limits are closer together. This can be seen in Figure 4.10, where each lower and upper limit lies in between the limits for the CDF for revenue with no contracts. We also observe from Figure 4.10 that the difference in the upper and lower boundaries given conventional forward or futures contracts appears to be the smallest. This suggests that conventional forwards and futures are the most efficient at minimising the value of any possible deviation from the predicted revenue value, but also result in lower overall predicted revenue values compared with predicted revenue given no contracts. Thus, conventional forward and futures contracts allow for greater confidence in the predicted revenue values, as they reduce the variability in the overall likely revenue values and predicted revenue values.

A summary of the effect of these different types of contracts on the predicted revenue is given in Table 4.2.

Table 4.2: Summary of Different Contracts

Contract	Median Difference between 95% Limits[6]	Median of 50% CDF	RaR (99%)	RaR (95%)
No Contracts	\$40.08m	\$47.69m	\$46.73m	\$46.97m
Forward Contracts	\$31.47m	\$47.63m	\$46.34m	\$46.64m
CFDs	\$33.63m	\$47.69m	\$44.16m	\$45.66m
Futures	\$31.54m	\$47.66m	\$46.38m	\$46.69m

We can see that in all cases, the median difference between the 95% upper and lower limits decreases as a result of using contracts. The smallest decrease is from the introduction of CFDs. However, we notice that CFDs have almost no effect on the median level of revenue, whereas the other two types of contract result in a decrease in median revenue.

[6] We calculate this median difference by subtracting the median of the 5% CDF from the median of the 95% CDF

It appears that conventional forward contracts lead to the largest reduction in both the median difference in the 95% limits and median revenue compared with the other contracts.

We have shown that there does indeed appear to be a trade off between reducing variability in revenue and higher levels of revenue. We now consider the percentage changes for each, given the different types of contract in Table 4.3.

Table 4.3: Effect of Contracts on Revenue

Contract	% Change in Median Revenue from 50% CDF	% Change in Median Difference of 95% Limits
Forward Contracts	-0.155	-21.5
CFDs	-0.017	-16.1
Futures	-0.090	-21.3

Again, we see that conventional forward contracts lead to the greatest change in both median revenue and median difference of the 95% limits, and CFDs the smallest percentage change. Both conventional forward and futures contracts lead to a similar percentage change in median difference of 95% limits. However, futures contracts result in a smaller change in median revenue. Thus, one reduces predicted revenue by 0.065% for a 0.2% decrease in the change in difference of the limits when using forward contracts rather than futures contracts. Moreover, using conventional forward contracts instead of CFDs amounts to a 0.138% decrease in median revenue and a 5.4% change in the median difference of the 95% limits. However, using futures contracts instead of CFDs will result in a 0.073% change in median revenue and a 5.2% decrease in the difference of the limits.

Our results suggest that futures contracts result in the smallest proportional trade-off between a reduction in revenue and a reduction in the variability of revenue. That is, for a relatively large decrease in variability, we forgo the least amount of revenue given futures contracts compared with the other contracts. Indeed, the median change in revenue is smallest with CFDs but, the change in variability is also relatively small.

In Table 4.2 we also consider Revenue at Risk (RaR). We use this measure instead of conventional Value at Risk (VaR) as noted in the preamble to this chapter. VaR is a widely used risk measure and is generally used to measure the riskiness of a portfolio of financial assets. Usually, VaR measures the possible loss in value of a portfolio over a set period for a specified confidence interval. The most commonly used VaR is the 95% VaR and for further detail on VaR, the reader should consult Hull (2004). Given that we are considering revenue, however, we use RaR as a measure of the possible low revenue values which one may face given a specified confidence level. That is, revenue will be equal to or below the RaR value, given a specified probability and period. For example, using the 99% RaR for revenue, we predict that there is a 1% chance that revenue will be equal to or below this value over the specified period.

We consider both the 99% and 95% RaR values. From Table 4.2 on page 59, we can see that the 99% and 95% RaR are largest given no contracts. We predict there is a 1% chance that revenue will be \$46.73m or less over a 90 day period and, a 5% chance that revenue will be less than or equal to \$46.97m. The introduction of forward contracts leads to the smallest RaR and futures allow for the largest RaR. This is a good indicator for a risk manager as to the level of capital required to cover these periods of low revenue if they were to occur. Thus, one could prepare for such situations by ensuring a certain level of capital is available to them. Hence, any costs that aren't covered as a result of a period of low revenue may be settled using the capital available. This suggests that a higher VaR is favourable, as less extra capital would be required to cover any costs. So, of all the contracts, if a risk manager's main objective was to increase RaR one would use futures contracts.

The choice of which contract to use, or what mix of contracts is required, is dependent on what the risk manager's particular objective is when managing risk. That is, do they want to increase maximum possible revenue, reduce the range of predicted revenue values, reduce the range of likely revenue values or increase the RaR. As we have seen, the introduction of any type of contract appears to reduce the size of possible deviations from the predicted values. Our observations imply that any type of contract allows a risk manager to have greater confidence in the predicted values they base their choice of hedging strategy on. Thus, it would appear to be appropriate for a risk manager to have some form of contract included in their portfolio, physical or purely financial, at least for the purposes of creating more certainty surrounding any revenue estimates.

Chapter 5

Conclusion

We have shown that conventional statistics models are not necessarily appropriate for the purposes of modelling electricity prices. Such models do not adequately capture the volatile nature of the prices. Rather than incorporate the random large spikes in price that electricity prices exhibit, these models tend to suggest that electricity prices have smaller, less variable jumps situated around a constant mean price. That is, in both the cases of the conventional SARIMA and the EVT-GARCH model, the prices have been modelled such that they appear to be mean reverting. However, the EVT-GARCH model does predict the variance of electricity prices most accurately, but does not model the randomness, nor the magnitude, of price changes well. Moreover, the "less" volatile periods were not captured by the EVT-GARCH model. Thus, it was necessary for us to develop a new model that would incorporate these properties into a simulated price series. By applying EVT to the residuals of the conventional SARIMA model we obtained a new model, the EVT-SARIMA model.

The EVT-SARIMA model appeared to predict the random price changes well, and had similar characteristics to the actual data series. That is, periods of low volatility combined with random periods of high volatility were exhibited in the forecast price series using the EVT-SARIMA model, which were similar to the corresponding periods in the actual price data. Our model did not exhibit a particular mean level of prices and the random spikes in price were of similar magnitude to those in the actual price series. Hence, we decided that our EVT-SARIMA model was the most appropriate model for forecasting prices.

Using our model, given generation quantities based on actual quantities traded at Benmore over a 90 day period, we then considered possible levels of revenue an electricity generator may receive over a 90 day period. From 1000 simulations we obtained a cumulative distribution function (CDF) for revenue and proceeded to alter this CDF through the use of different financial instruments. We observed the effects of conventional forward contracts, contracts for differences and futures contracts on the CDF. We found that conventional forward contracts and futures both reduced the variability in revenue that a firm may face. However, this reduction in variability came at a cost of receiving a lower level of revenue overall.

Conversely, CFDs increased the range of possible revenue values but allowed for higher revenue values compared with those given no contracts. But, CFDs may also result in negative revenue values. Thus, we concluded that the most appropriate strategy would be for a risk manager to have a portfolio of contracts consisting of a mix of the different types of contracts. That is, a mix of contracts would enable the risk manager to maintain a certain minimum level of revenue, reduce variability and possibly enable the opportunity to obtain higher levels of revenue.

In deciding on which particular mix of contracts is required, a number of things must be taken into consideration. Which contracts and the amount of electricity traded through contracts is dependent on how risk averse the firm, or risk manager, is. One would expect more electricity to be traded through contracts for a firm who is more risk averse as contracts result in greater certainty in the predicted level of revenue. In other words, as the quantity traded through contracts increases, the variability in revenue decreases and there will be greater certainty in the level of revenue that will be attained. This would also lead to less exposure to volatile spot prices in the wholesale market. Moreover, if one were to enter into forward contracts, increasing the quantity traded through contracting would ensure a certain amount of generated electricity is sold. Thus, only the "excess" electricity generation would be sold on the wholesale market and exposed to the volatile prices.

The interest rate also appears to play a role in the value of a contract. Our results suggest that there is a positive correlation between the electricity spot price and the interest rate. In particular, the introduction of a simplistic model for a stochastic interest rate further illustrated this correlation. Further research could be done to determine whether a relationship between the interest rate and the electricity spot price does exist. Moreover, if there is indeed a relationship, further work could be done to understand the underlying reasons for the existence of such a relationship.

We observed that all contracts reduced the 95% confidence intervals for our predicted values. This suggests that the introduction of contracts allows for greater confidence in the revenue predictions. One could also consider other financial instruments that are available. For example, given the recent introduction of options on the ASX, one may wish to observe the effects of different options on predicted revenue.

Other areas of research may be to estimate a model for electricity generation. Given that we assumed the generation profile was given over the 90 day period, our results may be affected by the current economic or seasonal conditions for that particular generation profile. That is, during different times of the year, more or less electricity may be necessary to satisfy demand thus affecting the amount of electricity generated. As a result, different contracts may be necessary for different times of the year. Perhaps during periods of low demand, it may be appropriate to trade more electricity through forward contracts to ensure a certain amount of electricity generated will be sold and to also provide greater certainty in how much electricity should be generated.

One could also consider contracts of varying length. For example, in the short term one may be interested in increasing revenue and therefore may use a mix of CFDs and forward contracts or futures over a short period. However, in the long term, one may be more concerned in reducing variability in revenue and so would consider contracting using forward contracts or futures instead of CFDs. Thus, a risk manager may create a portfolio consisting of CFDs, futures and forward contracts where a large proportion of the CFDs may be over short periods. Futures and forward contracts would also be included, all over varying lengths of time. The futures and forward contracts would act to reduce variability in revenue overall.

Our results suggest that a firm should use a variety of contracts to trade a portion of electricity. The actual combination of contracts required would be dependent on the risk manager's specific objectives and how averse they are to risk. However, the introduction of any contracts will enable the risk manager to reduce the variability in predicted revenue and to enable greater confidence in predicted revenue values.

Appendix A

EVT-GARCH model

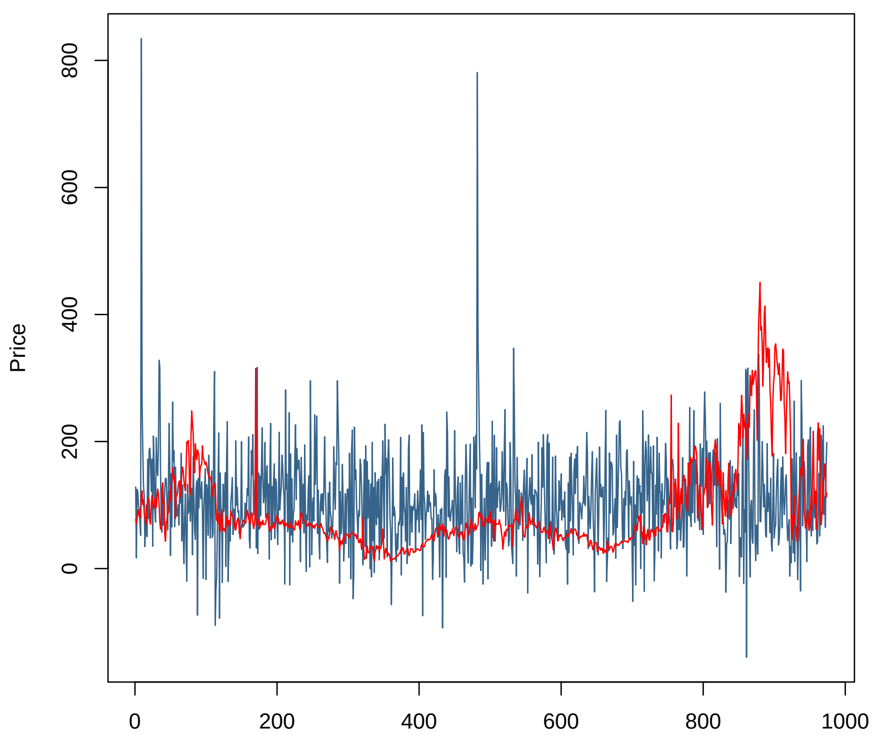

Figure A.1: EVT-GARCH model (blue) and Average Daily Prices (red)

Appendix B

Forward Contracts

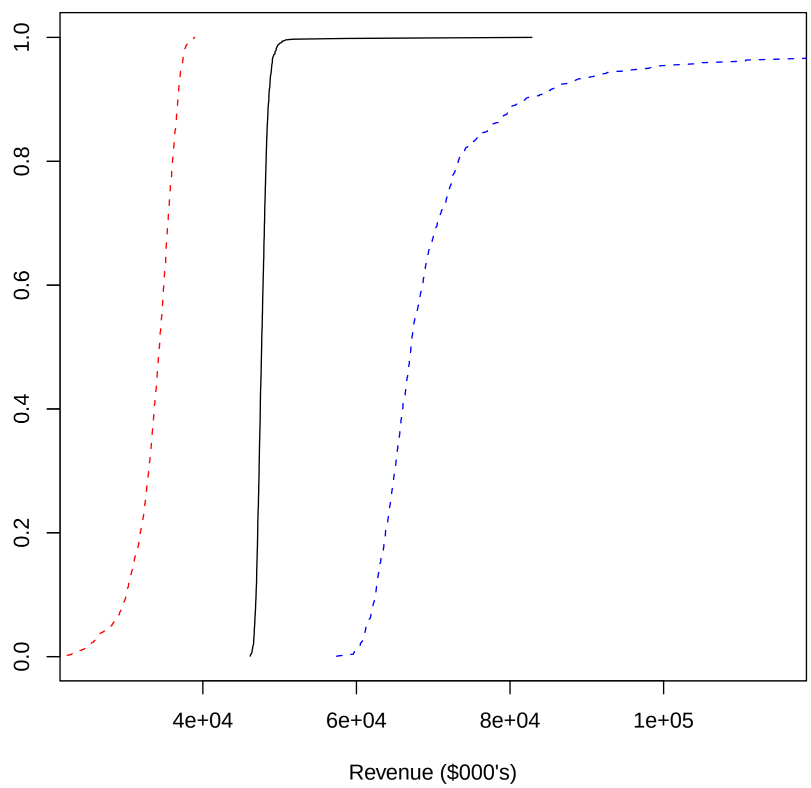

Figure B.1: CDF for Revenue given forward contracts trading 3 GWh per day, interest rate = 5%

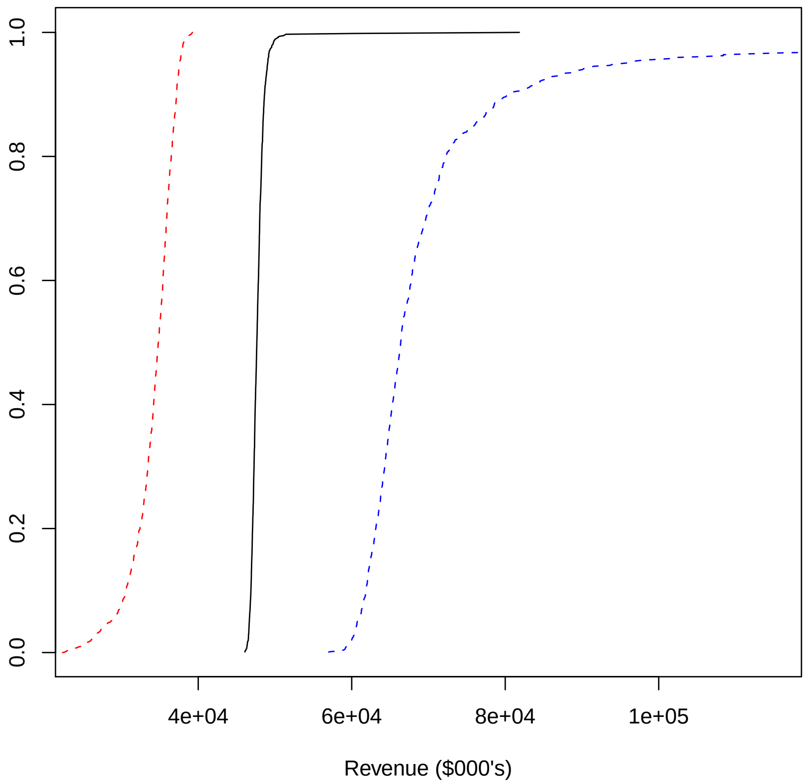

Figure B.2: CDF for Revenue, with forward contracts trading 3.5 GWh per day interest rate = 5%

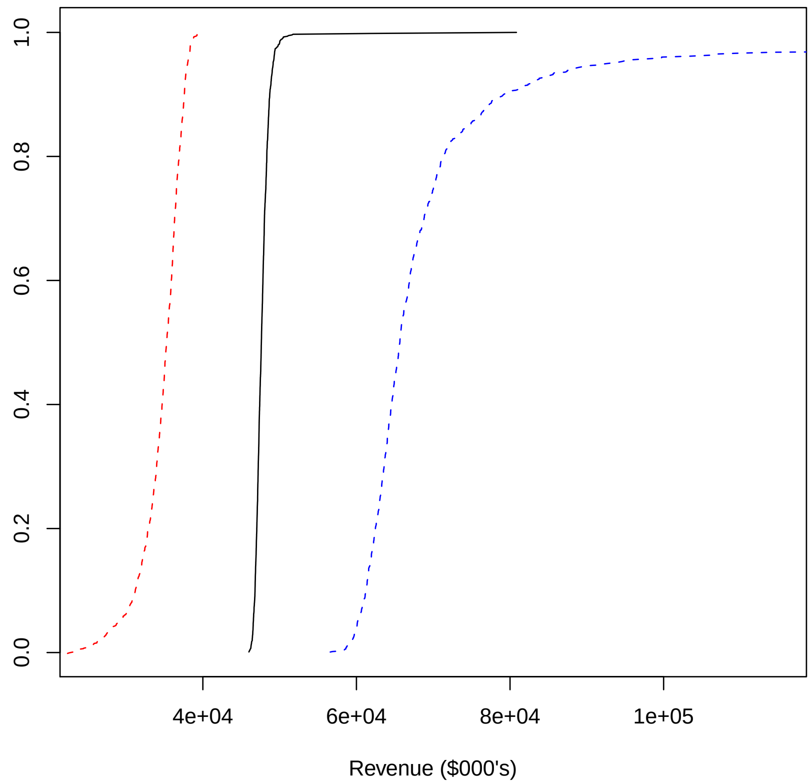

Figure B.3: CDF for Revenue given forward contracts trading 4 GWh per day, interest rate = 5%

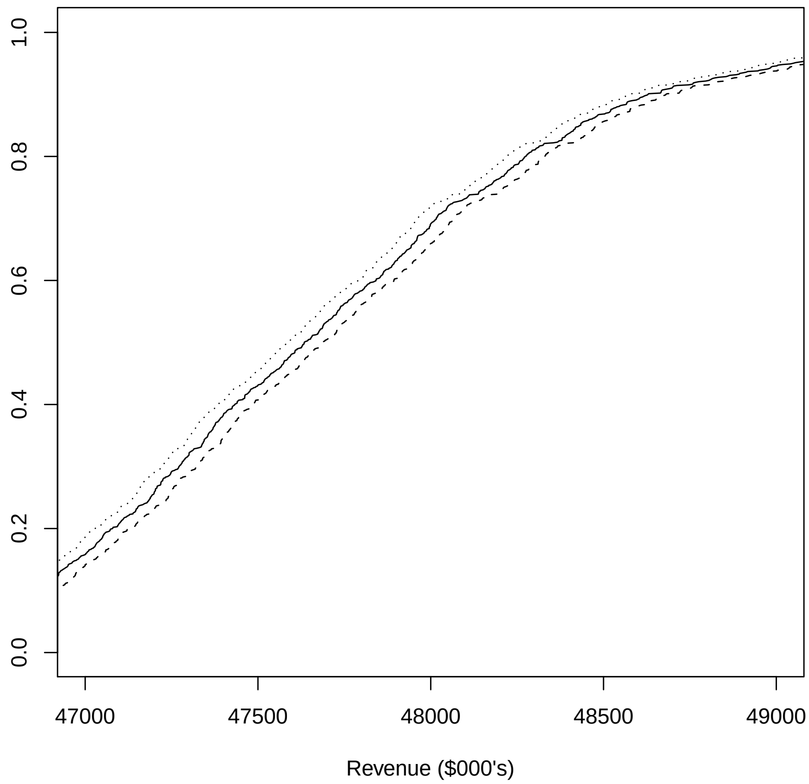

Figure B.4: CDF for Revenue given forward contracts trading 3 GWh per day interest rate = 3% (dotted line), 5% (solid line) and 7% (dashed line)

Appendix C

Contracts For Differences (CFDs)

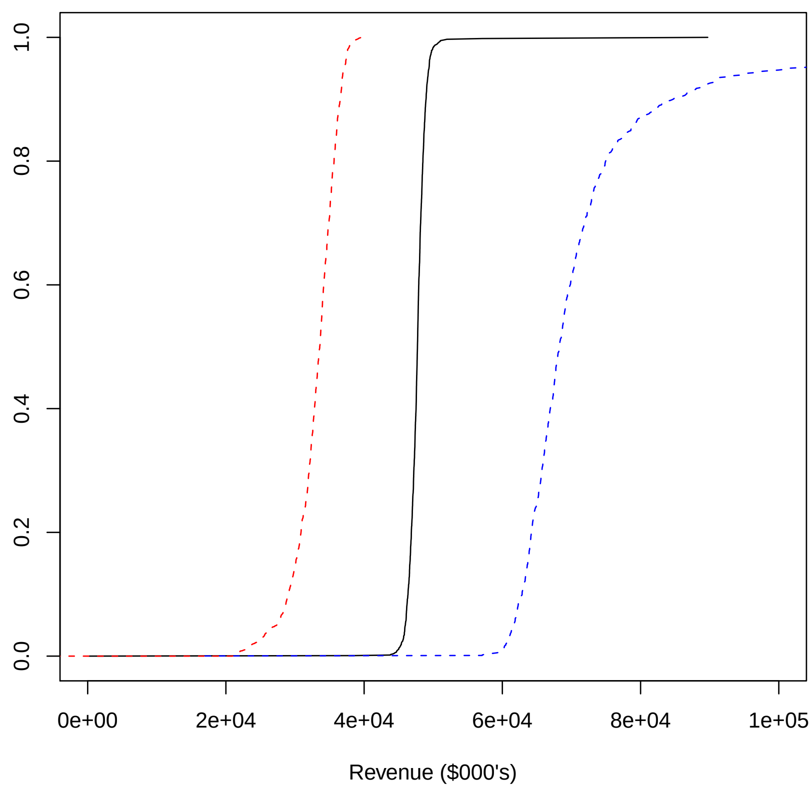

Figure C.1: CDF for Revenue given CFD trading 3 GWh per day, interest rate = 5%

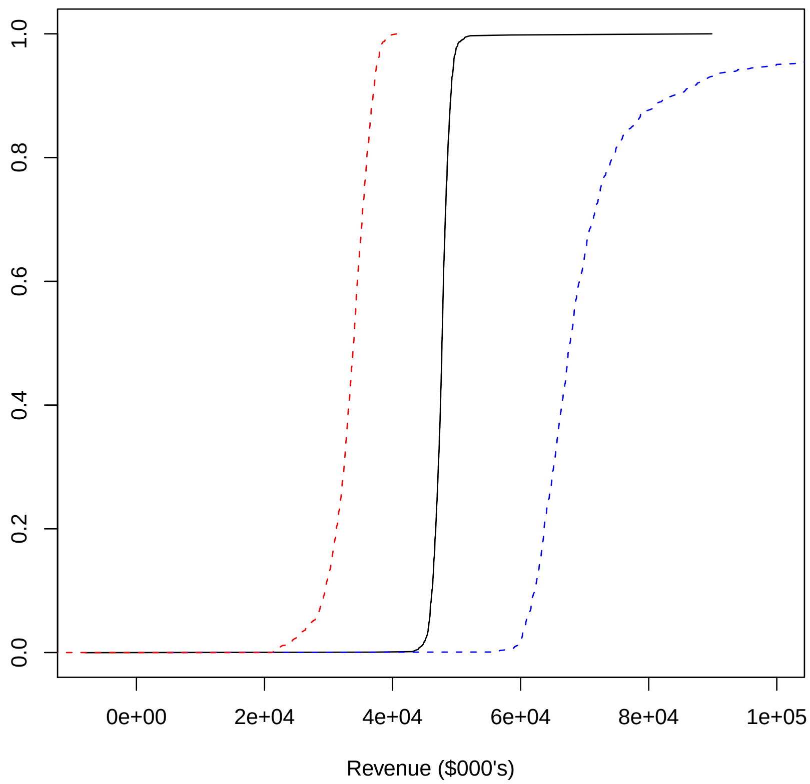

Figure C.2: CDF for Revenue, with CFDs trading 3.5 GWh per day

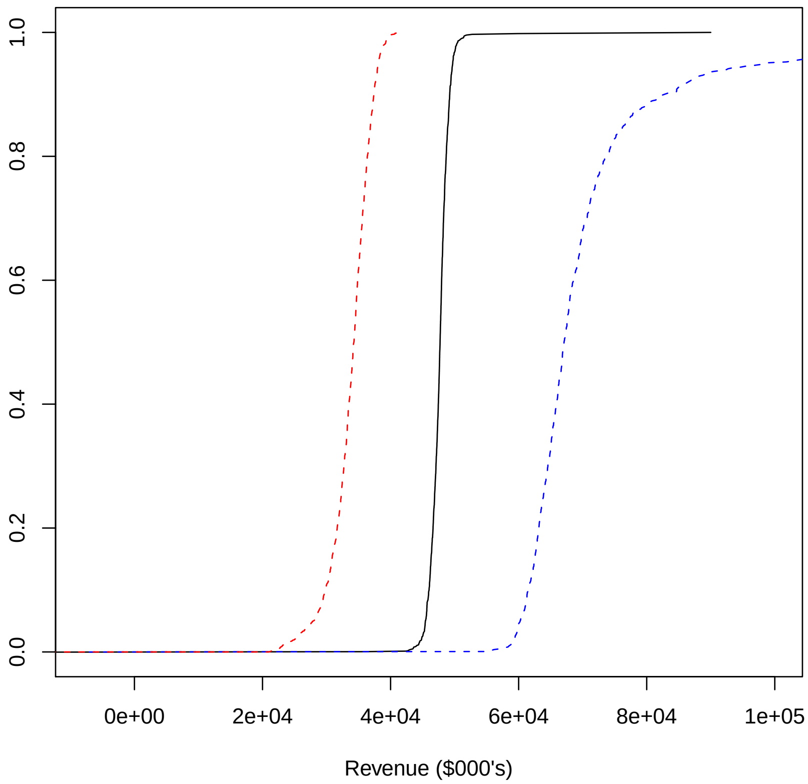

Figure C.3: CDF for Revenue given CFD trading 4 GWh per day, interest rate = 5%

Appendix D

Futures

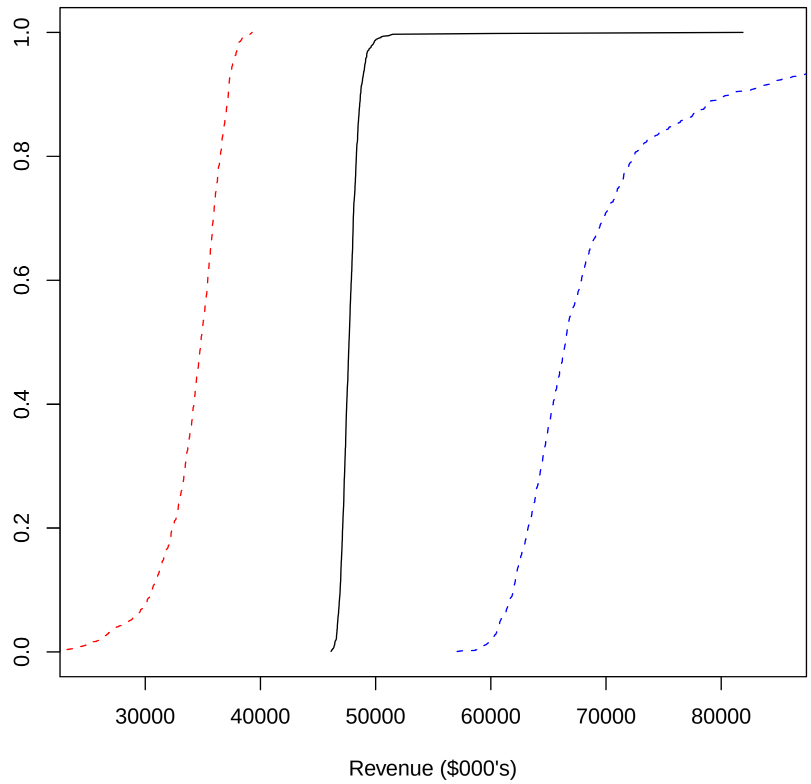

Figure D.1: CDF for Revenue given futures contracts trading 3.5 GWh per day, interest rate = 7%

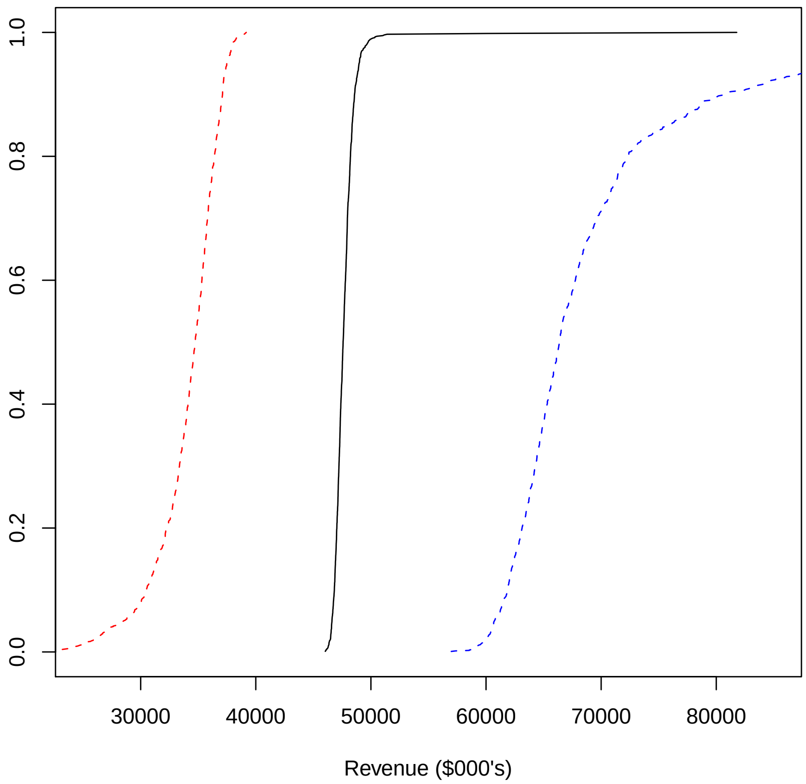

Figure D.2: CDF for Revenue given futures contracts trading 3.5 GWh per day, interest rate = 3%

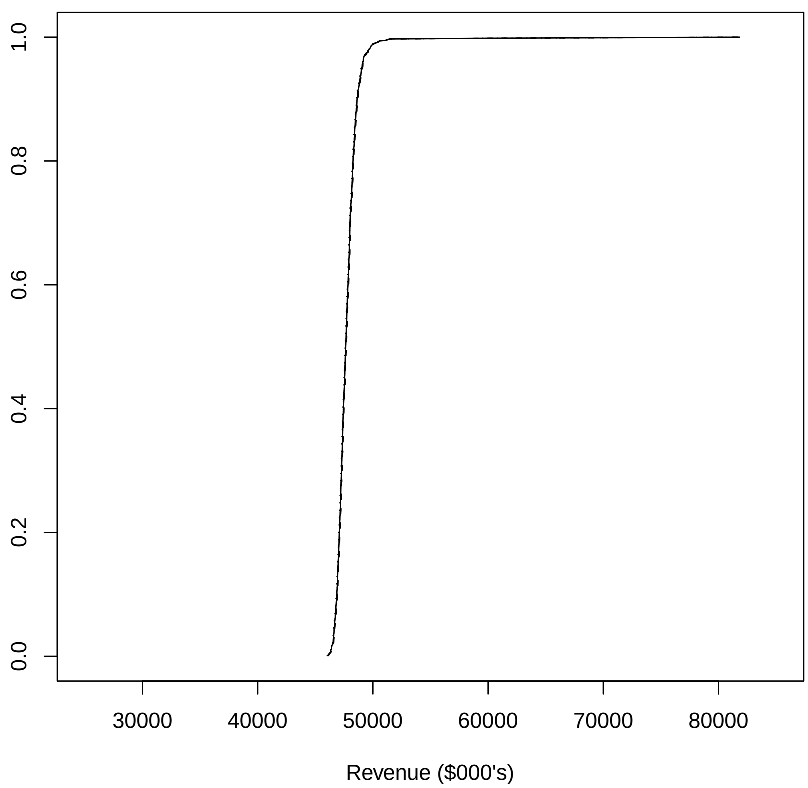

Figure D.3: CDF for Revenue given futures contracts trading 3.5 GWh per day, interest rate = 7% (dashed line), 5% (solid line) and 3% (dotted line)

Bibliography

Abraham, B. and Ledolter, J. (1983). *Statistical Methods for Forecasting*, Wiley Series in Probability and Mathematical Statistics, John Wiley and Sons, Inc.

Anderson, E. J. and Hu, X. (2005). Strategic retailer behaviour in an electricity contract market, Australian Graduate School of Management, Working paper. The University of New South Wales.

Anderson, E. J., Hu, X. and Winchester, D. (2007). Forward contracts in electricity markets: the Australian experience, *Energy Policy* **35**(5): 3089–3103.

Bowerman, B. L., O'Connell, R. T. and Koehler, A. B. (2005). *Forecasting, Time Series, and Regression*, Duxbury Applied Series, 4th edn, Thomson Brooks/Cole.

Brooks, C. (2002). *Introductory Econometrics for Finance*, Cambridge University Press.

Cox, J. C., Ingersoll, J. E. and Ross, S. A. (1981). The relation between forward prices and futures prices, **9**: 321–346.

Embrechts, P., Kluppelberg, C. and Mikosch, T. (2003). *Modelling Extremal Events for Insurance and Finance*, number 33 in *Applications of Mathematics*, stochastic modelling and applied probability edn, Springer.

Evans, L. and Guthrie, G. (2009). Commodity price behaviour resulting from transaction cost.

Fisher, R. A. and Tippett, L. H. C. (1928). Limiting forms of the frequency distribution of the largest or smallest member of a sample, *Proceedings of the Cambridge Philosophical Society* **24**: 180 – 190.

Ghysels, E. and Osborn, D. (2001). *The Econometric Analysis of Seasonal Time Series*, themes in modern econometrics, Cambridge University Press.

Green, R. (1999). The electricity contract market in England and Wales, *The Journal of Industrial Economics* **41**(1): 107–124.

Hull, J. C. (2004). *Fundamentals of Futures and Options Markets*, 5th edn, Prentice Hall Finance.

Hull, J. C. (2008). *Options, Futures and Other Derivatives*, 7 edn, Prentice Hall Finance.

McNeil, A. and Frey, R. (2000). Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach, *Journal of Empirical Finance* pp. 271–300.

McNeil, A., Frey, R. and Embrechts, P. (2005). *Quantitative Risk Management*, Princeton Series in Finance, Princeton University Press.

McRae, S. and Wolak, F. A. (2009). How do firms exercise unilateral market power? Evidence from a bid-based electricity market, EUI Publications, RSCAS Working Papers.

Sbai, E. and Simpson, M. (2008). Short-term forecasting of nodal electricity demand in NZ.

Smith, R. L. (1985). Maximum likelihood estimation in a class of nonregular cases, *Biometrika* **72**(1): 67–90.

Tanlapco, E., Lawarree, J. and Liu, C. (2002). Hedging with futures contracts in a deregulated electricity industry, *IEEE Transactions on Power Systems* **17**(3): 577–582.

Taylor, S. J. (2005). *Asset Price Dynamics, Volatility and Prediction*, Princeton University Press.

Wolak, F. A. (2000). An empirical analysis of the impact of hedge contracts on bidding behaviour in a competitive electricity market, *International Economic Journal* **14**(2).