ABSTRACT

Title of dissertation:	ESSAYS IN FINANCIAL ECONOMICS
	Carl J. Ullrich, Doctor of Philosophy, 2006
Dissertation directed by:	Professor Gurdip Bakshi Department of Finance

Essay 1, Constrained Capacity and Equilibrium Forward Premia in Electricity Markets, develops a refinement of the equilibrium electricity pricing model in Bessembinder and Lemmon (2002). The refined model explicitly accounts for constrained capacity, an important feature in electricity markets. Explicitly including a role for capacity allows the model to reproduce the price spikes observed in wholesale electricity markets. The refined model implies that the equilibrium forward premium, defined to be the forward price minus the expected spot price, is decreasing in spot price variance when the expected spot electricity price is low, but is increasing in the spot price variance when the expected spot electricity price is high. Further, the refined model implies that, ceteris paribus, the equilibrium forward premium is increasing in the ratio of the expected spot electricity price to the fixed retail price. The implications of this model are closer to reality.

How does currency return volatility evolve over time and what are the properties of volatility dynamics? Is the drift of currency return volatility non-linear? What forms of non-linearities are admitted in the drift and diffusion functions? The

purpose of essay 2, Estimation of Continuous-Time Models for Foreign Exchange Volatility, is to estimate a large class of volatility processes and explore these issues using weekly data on two currency pairs: U.S. dollar-British pound and Japanese Yen-U.S. dollar. The estimation approach is based on maximum-likelihood estimation that relies on closed-form density approximations (Aït-Sahalia 1999, 2002). Based on volatility implied by currency options, the constant elasticity of variance specification provides a reasonable characterization of the variance of variance function. Extending the diffusion function beyond the CEV specification does not improve the fit of the model, regardless of the assumed form of the drift function. Further, I find that certain types of non-linearities in the drift function improve the goodness of fit statistics, though no generalizations can be made.

ESSAYS IN FINANCIAL ECONOMICS

by

Carl J. Ullrich

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2006

Advisory Committee:
Professor Gurdip Bakshi, Chair/Advisor
Professor John Chao
Professor Steve Heston
Professor Mark Loewenstein
Professor Greg Willard

c Copyright by
Carl J. Ullrich
2006

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1

Constrained Capacity and Equilibrium Forward Premia in Electricity Markets

1.1 Introduction

Electricity prices are known to exhibit price spikes - episodes in which the price can grow to 30-40 times the mean price level, followed by a relatively quick (on the order of days, or even hours) return to normal price levels.[1] The main reason for such behavior in electricity prices is the fact that electricity cannot be stored in any meaningful way. Supply and demand must balance in real-time. There can be no inventory to smooth supply and/or demand shocks. Hence, the available supply of electricity (the available system capacity) must exceed the demand at all times in order to avoid a blackout.

In an important article, Bessembinder and Lemmon (2002, hereafter BL) develop an equilibrium model of electricity forward pricing and hedging. In their

[1] One seemingly counterintuitive aspect of electricity markets is that the price of electricity can be negative. Unlike traditional markets, the economic assumption of free disposal does not apply to the case of electricity. The reason for this has to do with the physical properties of power plants. Large nuclear and coal-fired plants cannot be easily shut down and restarted. A nuclear power plant that has been shut down cannot be restarted for a week or more, and then only at a significant cost. Thus it may be cheaper for the owner of electricity to pay others to take the output than to turn off their machines.

model, the spot price of electricity is determined by the level of demand and by the convexity of the supply curve. BL define the forward premium to be the difference between (i) the time $t = 0$ forward price for electricity to be delivered at time $t = 1$, and (ii) the time $t = 0$ expectation of the spot price that will prevail at time $t = 1$. BL then derive an expression for the equilibrium forward premium in terms of the moments of the expected spot electricity price.

This paper makes two contributions to the theory of electricity markets and our understanding of how risk premiums are determined in constrained capacity and zero inventory markets. First, I propose a refinement to the BL model to account explicitly for constrained capacity, a fact in electricity markets. In the refined model, the relevant quantity in determining the spot price of electricity is not the absolute level of demand, but the level of available supply in excess of demand. The refined model is designed so that it retains the tractability and ease of interpretation of the original BL model, but is better able to replicate the spikes seen in electricity spot and forward prices.

Second, in the presence of constrained capacity, I formally show that the forward premium depends on the level of the retail electricity price relative to the expected spot price. BL argue that the fixed retail price of electricity must exceed the expected wholesale (spot) price of electricity. While it is generally true that the long-term average wholesale price is less than the retail price, in the short term it can and does happen that expected spot price of electricity is greater than the retail price of electricity. This feature of electricity markets has important implications for the equilibrium forward premium, which the extant literature has not fully explored theoretically or empirically.

My main results are as follows. First, the equilibrium forward premium is decreasing in the variance of the spot price when the expected spot price is less than the fixed retail price. However, the premium is increasing in the variance of the spot price when the expected spot price exceeds the retail price. This relationship between the forward premium and the spot price variance reflects retailers' net hedging demands. High spot prices tend to occur in periods of high demand, so retail sales and the spot price are positively correlated. When the spot price is less than the retail price, retailers' profits are positively related to the spot price. In this case, retailers reduce forward purchases, thereby driving down the forward price. However, when the spot price exceeds the retail price, retailers' profits are negatively related to the spot price. In this case, retailers increase forward purchases to cap the price they must pay to meet their demand, and thereby smooth their profit stream. In either case, the effect increases with the variance of the spot price.

Second, while the original BL model predicts that the optimal forward premium is uniformly increasing in the skewness of the spot price, the refined model predicts that the premium increases in the skewness of the spot price for most levels of the expected spot price, but can decrease in spot price skewness for very high expected spot prices.

Third, the optimal forward purchase for a retailer of electricity increases in the ratio of the expected spot price of electricity to the retail price of electricity, even if the variance and skewness of the spot price are unchanged. Retailers are assumed to be mean-variance profit optimizers. The profits of electricity retailers become more variable as the expected spot price increases relative to the retail price. As the spot price increases relative to the retail price, retailers optimally would like to increase forward purchases to counteract the increase in the variance of their profit.

Fourth, ceteris paribus, the equilibrium forward premium increases in the ratio of the expected spot price of electricity to the retail price of electricity. While the optimal forward purchase for a retailer of electricity increases in the expected spot price-retail price ratio, producers are not directly exposed to the retail price of electricity. As the expected spot price increases relative to the retail price, retailers want to increase forward purchases to counteract the corresponding increase in the variability of their profit stream. However, producers' hedging demands are unchanged as long as the moments of the spot price are unchanged. The equilibrium forward premium must therefore increase.

This article proceeds as follows. Section II briefly presents the modeling paradigm adopted in Bessembinder and Lemmon (2002). Section III introduces the refined model and discusses the determinants of forward premia in the presence of constrained capacity. Section IV first briefly reviews the daily data for the Pennsylvania-New Jersey-Maryland (PJM) electricity market and aims to clarify the main empirical features driving the modeling approach using a simple three-state, discrete version of the model. The discrete model requires no approximations. Section V presents numerical simulation evidence designed to enable comparisons with the simulations in the BL paper. Section VI concludes and summarizes the contributions.

1.2 The Bessembinder and Lemmon (2002) Model

I begin with a brief discussion of the equilibrium electricity pricing model in Bessembinder and Lemmon (2002). BL first assume the wholesale market for electricity consists of $N_P + N_R$ total participants, who can be divided into two types - (1) N_P producers, i.e., owners of electric generating capacity, and (2) N_R retailers, i.e., firms that supply electricity to retail consumers. Their notation is followed as closely as possible.

Both producers and retailers are assumed to be mean-variance profit optimizers. Specifically, the utility function $U[.]$ for producers and retailers takes the form,

$$U[\pi] = E(\pi) - \frac{A}{2}\text{Var}(\pi) \tag{1.1}$$

where π is profit and $E(.)$ is the expectation operator. Both producers and retailers share the same dislike coefficient, $A \in \Re^+$, on the variance term in the utility function, and $\Re$ is the real line.

There are two times in the model, $t = 0$ and $t = 1$. At $t = 0$, producers and retailers of electricity may enter into forward contracts. At $t = 1$, (random) demand Q_{Rj} is realized for each retailer $j = 1, N_R$. Demand is the fundamental uncertainty driving the model and is exogenous. Most retail consumers face a fixed retail price, thus demand is unaffected by the wholesale price, at least to first-order.

Retailers can meet their $t = 1$ demand either by (i) purchasing electricity forward at $t = 0$ for delivery at $t = 1$, or, (ii) waiting until $t = 1$ and purchasing electricity spot. Adopt the convention that the spot and forward sales are positive and purchases are negative. For example, if retailer j elects to sell forward, which he is allowed to do, then $Q^F_{Rj} > 0$. If the amount purchased forward by retailer j, $Q^F_{Rj} < 0$ is insufficient to meet retailer j's demand, then retailer j must make up the difference by purchasing electricity in the spot market in the amount $(Q_{Rj} + Q^F_{Rj})$. Importantly, neither producers nor retailers may accumulate inventory. This means that at $t = 1$ supply and demand must balance and this restriction is binding.

BL assume that the total cost for producer i to produce electricity is given by

$$TC_i = F + \frac{a}{c}(Q_{Pi})^c , \tag{1.2}$$

where $F > 0$ are fixed costs, Q_{Pi} = output of producer i, and $c > 2$. The coefficient $a \in \Re^+$.

Begin at $t = 1$. The forward positions of each producer i, Q^F_{Pi}, and each retailer j, Q^F_{Rj}, are thus known. It is straightforward to show that the optimal wholesale spot sales for producer i, Q^W_{Pi}, is:

$$Q^W_{Pi} = \left(\frac{P_W}{a}\right)^{\frac{1}{c-1}} - Q^F_{Pi}, \tag{1.3}$$

and the equilibrium wholesale spot price, P_W, are given by,

$$P_W = a\left(\frac{Q^D}{N_P}\right)^{(c-1)} , \tag{1.4}$$

where

$$Q^D \equiv \sum_{j=1}^{N_R} Q_{Rj} \tag{1.5}$$

is the total system retail demand.

Given the optimal spot sales, one can move back to time $t = 0$ and determine the optimal forward positions for producer i,

$$Q^F_{Pi} \quad = \quad \frac{1}{A}\left(\frac{P_F - E(P_W)}{\mathrm{Var}(P_W)}\right) + \left(\frac{1}{a}\right)^{\frac{1}{c-1}}\left(1-\frac{1}{c}\right)\frac{\mathrm{cov}\left(P_W^{\frac{c}{c-1}}, P_W\right)}{\mathrm{Var}(P_S)}, \qquad (1.6)$$

and retailer j as,

$$Q^F_{Rj} \quad = \quad \frac{1}{A}\left(\frac{P_F - E(P_W)}{\mathrm{Var}(P_W)}\right) + \frac{P_R\,\mathrm{cov}(Q_{Rj}, P_W)}{\mathrm{Var}(P_W)} - \frac{\mathrm{cov}(Q_{Rj}P_W, P_W)}{\mathrm{Var}(P_W)}, \quad (1.7)$$

where P_R is the (fixed) retail price.

Using the fact that forward contracts are in net zero supply, defining $N \equiv (N_P + N_R)/A$, and defining $x \equiv \frac{1}{c-1}$, BL finally arrive at an expression for the forward premium, their equation (12),

$$P_F - E(P_W) = \frac{-N_P}{N\,a^x\,c}\left[c\,P_R\,\mathrm{cov}(P_W^x, P_W) - \mathrm{cov}(P_W^{x+1}, P_W)\right]. \qquad (1.8)$$

As BL explain, the forward price converges to the expected spot price if (i) the number of firms approaches infinity ($N_P + N_R \to \infty$), or (ii) the firms are risk neutral ($A = 0$). In either case, $N \to \infty$ and the right hand side of equation (1.8) goes to 0. Otherwise, the two covariance terms capture the hedging demands of producers and retailers. In their words (p.1358), "The forward price will be less than the expected spot price if the first term in brackets, which reflects retail revenue risk, is greater than the second term, which reflects production cost risk."

1.3 Accounting for Constrained Capacity in Electricity Markets

In this section I develop a refinement of the BL model that (i) explicitly accounts for constrained capacity, and (ii) is better able to produce price spikes than the BL model. The refined model allows me to derive the equilibrium spot price, the optimal spot and forward positions for producer i, and ultimately the equilibrium forward premium, $P_F - E(P_W)$. Derivations can be found in the Appendix. The section begins with a discussion of the characteristics of the supply curve in electricity markets, also known as the supply stack.

1.3.1 The Supply Stack

The production cost function assumed in the BL model, equation (1.2), captures one of the main features of electricity markets, supply curve convexity (recall that, in equation (1.2), $c > 2$), but it does not model explicitly the available supply of electricity. Because there can be no inventory, the spot price of electricity is determined not just by the level of demand, but also by the total amount of supply (electric generating capacity) available to meet that demand in real-time. On any given day, various power plants are unavailable due to scheduled maintenance and/or unscheduled, forced outages. Thus, not only does peak demand vary from day to day, but supply also has variations. For instance, suppose that on Tuesday, all the electric generating units in the system are available, while on Wednesday some fraction of the generating units are out of service, effectively shifting the supply curve to the left. Even if demand is the same on both days, it will be the case that the spot price is higher on Wednesday.

To focus on the defining characteristics of supply versus demand, Tables 1.1 and 1.22 present data for the Pennsylvania-New Jersey-Maryland (hereafter

PJM) market for the June 2000 through March 2003 time period, available at www.pjm.com. We refer the reader to Longstaff and Wang (2004) for a discussion of the PJM electricity market. All the data have units of Megawatts (MW). Table 1.1 displays summary statistics for available supply, i.e., capacity while Table 1.2 displays summary statistics for demand. From Table 1.1 it is clear that capacity varies from day-to-day, though the variation is much smaller than the variation in demand.

The supply curve in electricity markets, also known as the supply stack, results from the *stacking* of power plants in order from least expensive to most expensive. The least expensive 'baseload' units include hydroelectric, nuclear, and coal-fired plants. At the top of the supply stack are natural gas-fired 'peaking' units. These peaking units are much more flexible than baseload units - they can be started and stopped on short notice - but they are more expensive to operate. As a result, the supply stack looks nearly like a step function - it is relatively flat as long as demand does not exceed the capacity of baseload units, then it increases dramatically when demand is such that peaking units are needed. See, for example, Figure 1.2 in Eydeland and Wolyniec, 2003.

1.3.2 Production Cost Function

Motivated by the previous subsection where supply considerations play an important role, I propose a refinement to the BL model that formalizes the impact of available supply on risk premia and prices. In my model, the determinant of total cost is not just the demand, but the available supply minus demand. Consider the following production cost function for producer i that captures this feature,

$$TC_i = F_i + \frac{a}{c}\left(\frac{1}{H_{Pi} - Q_{Pi}}\right)^c, \qquad H_{Pi} - Q_{Pi} > 0, \tag{1.9}$$

where F_i = fixed costs, H_{Pi} = capacity (or available supply) of producer i, and, as before, Q_{Pi} = output of producer i. In the production function, assume $a \in \Re^+$ and $c \in \Re^+$. Thus, the presence of capacity makes the production cost function hyperbolic in Q_{Pi} as opposed to a power function in Bessimbinder and Lemmon (2002).

The benefit of this production cost function is that it ensures that, as demand approaches available supply, the spot price generated by the model increases dramatically. This fact is articulated further in the following subsection.

1.3.3 Optimal Producer Spot Sales and the Equilibrium Spot Price

Begin at $t = 1$ taking the optimal forward positions for producer i, Q^F_{Pi}, and retailer j, Q^F_{Rj}, as given. The ex-post profit for producer i is given by

$$\pi_{Pi} = P_F Q^F_{Pi} + P_W Q^W_{Pi} - F - \frac{a}{c}\left(\frac{1}{H_{Pi} - Q^W_{Pi} - Q^F_{Pi}}\right)^c, \tag{1.10}$$

where P_F is the forward price, P_W is the wholesale spot price, and (as before) Q^W_{Pi} is the spot sale for producer i. Notice that the output of producer i, Q_{Pi}, has been replaced by the sum of producer i's spot and forward sales, $Q^W_{Pi} + Q^F_{Pi}$. Via the steps of the proof in the Appendix, it can be seen that the optimal spot sale for producer

i is:

$$Q_{Pi}^{W} = H_{Pi} - Q_{Pi}^{F} - \left(\frac{a}{P_W}\right)^{\frac{1}{c+1}}. \tag{1.11}$$

Producer spot sales are increasing in the spot price, P_W, as expected.

Defining the total system supply as

$$H^S \equiv \sum_{i=1}^{N_P} H_{Pi}, \tag{1.12}$$

and using the fact that supply and demand must balance, the equilibrium spot price is given by

$$P_W = a\left(\frac{N_P}{H^S - Q^D}\right)^{c+1}. \tag{1.13}$$

The equilibrium spot price increases in total system demand Q^D and decreases in total system capacity H^S. At first glance it appears that an increase in producers N_P would increase the spot price. This is not the case, however, as a new supplier brings new capacity, thereby increasing H^S as well.

Equation (1.13) shows analytically that the spot price can increase dramatically as demand approaches supply, i.e., as $H^S - Q^D$ approaches zero. In the BL model, supply curve convexity is captured by the fact that $c > 2$. Using daily data, BL estimate that c is approximately 4. The spot price expression that obtains in the BL model, equation (1.4), is unable to reproduce the tremendous increase in spot prices that occurs as demand approaches available supply, as demonstrated in Figure 1. Thus, having a model with constrained capacity overcomes some of the shortcomings of the BL model.

Figure 1 plots the spot prices generated by the refined model, equation (1.13), and the BL model, equation (1.4), as a function of the level of demand. The parameter values, discussed in detail in the simulation section below, in as much as possible are chosen to be the same as in BL. In particular, $c = 4$ for the BL model, $c = 1.5$ for the refined model, and available system supply $H^S = 240$ MW. Figure 1 demonstrates that, when demand approaches the available supply, equation (1.13) ensures that the spot price increases dramatically.

As discussed above, electricity spot prices reflect the nature of the supply stack - prices are relatively constant for low levels of demand, then increase dramatically at higher levels of demand. The refined model is able to reproduce this behavior using reasonable parameter values.

In the BL model, each producer sells (or buys) the same amount forward. The remainder of the realized demand is split equally amongst the producers. That is, substituting equation (1.4) into equation (1.3) results in $Q_{Pi}^W = \frac{Q^D}{N_P} - Q_{Pi}^F$. However, in the model with capacity each producer is characterized by the same total cost function, but producers are allowed to have differing amounts of capacity, H_{Pi}.

Returning to the refined model, substitute equation (1.13) into equation (1.11) which yields

$$Q_{Pi}^W = H_{Pi} - Q_{Pi}^F - \frac{H^S - Q^D}{N_P}. \qquad (1.14)$$

In this model, rather than each producer having the same spot sales, each producer is allowed to retain the same amount of unused capacity. Each producer's spot sales are equal to his available time $t = 1$ capacity net of his forward position $(H_{Pi} - Q_{Pi}^F)$

minus an equal fraction ($\frac{1}{N_P}$) of the system *excess* capacity ($H^S - Q^D$). Given the production cost function in equation (1.9), each producer's costs are determined by his excess capacity. In order to equalize marginal costs, each producer in equilibrium must retain the same excess.

1.3.4 Optimal Producer Forward Position

Stepping back to $t = 0$, the optimal forward position for producer i is shown in the Appendix to be:

$$Q^F_{Pi} = \frac{1}{A}\left(\frac{P_F - E(P_W)}{\mathrm{Var}(P_W)}\right) + H_{Pi} - \frac{1}{a^x(x+1)}\frac{\mathrm{cov}\left(P_W^{x+1}, P_W\right)}{\mathrm{Var}(P_W)}, \tag{1.15}$$

where $x \equiv \frac{-1}{c+1}$. The optimal forward position for retailer j is unchanged from the BL case, so that equation (1.7) above still applies.

Following BL, I use a second order Taylor series approximation to express the producer's optimal forward position in terms of the variance and skewness of the spot price.

$$Q^F_{Pi} \approx \frac{\mathrm{PREM}}{A\mathrm{Var}(P_W)} + H_{Pi} - \frac{H^S - Q^D}{N_P} - \frac{x}{2a^x}E\left(P_W\right)^{x-1}\frac{\mathrm{Skew}(P_W)}{\mathrm{Var}(P_W)}, \tag{1.16}$$

where $\mathrm{PREM} \equiv P_F - E(P_W)$. The first term on the right hand side is the well-known optimal response to a nonzero forward premium (e.g., Anderson and Danthine, 1980 and Hirshleifer and Subramanyam, 1993).

In the absence any forward premium, and given normally distributed spot prices, the first and last terms on the right hand side of equation (1.16) will be zero. In this case, each producer will sell forward all of his capacity minus an equal fraction of the overall system reserve, i.e., $H_{Pi} - \frac{H^S - Q^D}{N_P}$. Noting that $x < 0$, the optimal forward position for a producer increases in spot price skewness.

1.3.5 Equilibrium Forward Premium

Using the fact that forward contracts are in net zero supply, and defining $N \equiv (N_P + N_R)/A$, the forward premium is given by,

$$P_F - E(P_W) = \frac{N_P}{Na^x}\left[P_R \operatorname{cov}(P_W^x, P_W) - \left(\frac{x}{x+1}\right)\operatorname{cov}(P_W^{x+1}, P_W)\right]. \quad (1.17)$$

Using a second-order Taylor series approximation, the forward premium can be rewritten in terms of the anticipated variance and skewness of the spot price.

$$P_F - E(P_W) = \alpha \operatorname{Var}(P_W) + \gamma \operatorname{Skew}(P_W), \quad (1.18)$$

where

$$\alpha \equiv \frac{xN_P}{Na^x}E(P_W)^{x-1}\left[P_R - E(P_W)\right], \text{ and,} \quad (1.19)$$

$$\gamma \equiv \frac{xN_P}{2Na^x}E(P_W)^{x-2}\left[(x-1)P_R - xE(P_W)\right]. \quad (1.20)$$

To get some intuition for this result, consider the expression for the coefficient α, equation (1.19). BL argue that the retail price of electricity must exceed to expected spot price of electricity, $\frac{E(P_W)}{P_R} < 1$, otherwise no retailer would enter the market. (In their simulation exercises, BL set $P_R = 1.2E(P_W)$.) In this case, equation (1.19) implies that the forward premium decreases in anticipated spot price variance ($\alpha < 0$, recall that $-1 < x < 0$). BL point out that the profits of electricity retailers are positively exposed to wholesale spot prices, because more power is sold when the spot price is high. The lower is the expected spot price *relative* to the retail price, the less electricity retailers purchase forward, and the lower is the forward price. This effect is stronger the greater is the spot price variance.

When the forward contract in question is short-term, on the order of one day, the expected spot price can and often does exceed the retail price, $\frac{E(P_W)}{P_R} > 1$. To see this point, note that the retail price of electricity is on the order of \$80-\$100/MWh. In the PJM market for the June 2000 through March 2003 time period, there are 201 separate occurrences of hours for which the forward price and the spot price exceeded \$100/MWh, and 10 occurrences when both prices exceeded \$500/MWh. The forward risk premia derived in (1.20) has the following interpretation:

1. When the expected spot price exceeds the retail price then retailers' profits are *negatively* exposed to spot prices. The higher is the expected spot price relative to the retail price, the more electricity retailers purchase forward (to cap their losses), thereby driving up the forward price.

2. Conceptually, the forward premia effect is also stronger the greater is the spot price variance. Thus, the sign of α is negative when the expected spot price is less than the retail price, $\frac{E(P_W)}{P_R} < 1$, and positive when the expected spot price exceeds the retail price, $\frac{E(P_W)}{P_R} > 1$.

3. The sign of γ, Sgn(γ) can also change. Equation (1.20) implies that γ will be positive, and thus the equilibrium forward premium will be increasing in spot price skewness, if $\frac{E(P_W)}{P_R} < (c+2)$. However, the forward premium will be decreasing in spot price skewness if $\frac{E(P_W)}{P_R} > (c+2)$.

Table 1.3 summarizes the signs of α and γ and their possible impacts. Overall, allowing the expected spot price to exceed the retail price offers more flexibility to model forward premia in electricity markets.

1.4 State Space Model of Excess Demand

In this section I first briefly describe the data available from PJM. Then I develop a simple three-state model of excess demand to illustrate the main features of the model. The data help to guide my choices of parameter values in the state space model and in the simulations of the following section. These parameter values are chosen to be as similar as possible to those used in the BL simulations so as to highlight the more realistic features of the new model with constrained capacity.

This section demonstrates, (i) the implications of the model in an exact setting, with no need to resort to approximations or numerical simulations, and (ii) that the equilibrium forward premium increases in the ratio of the expected spot price to the retail price.

1.4.1 PJM Data

Longstaff and Wang (2004) document the distributional properties of hourly spot and forward electricity prices in the PJM market for the period June 2000 through November 2002. Panel A of Table 1.4 updates the distributional statistics for the hourly PJM spot market through May 2003 and Panel B of Table 1.4 provides the distributional statistics for the hourly PJM forward market through

May 2003. Table 1.4 confirms the result that electricity spot and forward prices are positively skewed and demonstrate extremely large excess kurtosis. These distributional characteristics reflect the occurrence of price spikes in the wholesale electricity market.

Under normal market conditions, prices remain at a relatively low level. Occasionally, prices spike to very high levels, then quickly return to normal. These spikes cause the distribution of spot prices to be right-skewed and to have fat tails. One of the contributions of the refined model is the ability to reproduce price spikes and, therefore, reproduce the distributional characteristics of wholesale electricity prices.

The demand for electricity is driven primarily by weather conditions. Thus electricity prices demonstrate pronounced seasonal patterns. The most extreme price behavior occurs in the summer months, conventionally defined to be June, July, and August. Succeeding seasons are defined to be succeeding three month periods. The sample mean for the summer (not shown in the tables) is \$38.80/MWh, but the maximum is over \$1,000/MWh. Also, the overall standard deviation for the summer months is \$65.52/MWh, nearly 170% of the mean. In the mild spring season, the standard deviation of spot prices is 'only' 67% of the mean, while the maximum price is more than four times the mean.

1.4.2 Estimated Exponent c in the Production Function

In order to gain insight into reasonable values of the production cost function exponent c, take logs on both sides of equation (1.13) to obtain

$$\log(P_W) = \log\left(aN_P^{c+1}\right) - (c+1)\log(H - Q^D). \tag{1.21}$$

Performing a simple OLS regression of $\log(P_W)$ on $\log(H - Q^D)$, by season, yields a range for the production function exponent c. The results are provided in Table 1.5. The value of c varies from approximately 0.75 to 1.5. Equation (1.13) does a good job of capturing variation in the spot price, particularly in the summer months ($R^2 = 73.8\%$).

1.4.3 Three-State Model

Consider the following three-state world indexed by $\omega = (\omega_1, \omega_2, \omega_3)$. The probability of each state is given by ϕ_k, i.e.,

$$\text{Prob}[\omega] = \begin{pmatrix} \phi_1 \\ \phi_2 \\ \phi_3 \end{pmatrix}, \tag{1.22}$$

where $0 \leq \phi_k \leq 1 \ \forall k = 1, 2, 3$, and $\sum_{k=1}^{3} \phi_k = 1$. By focusing on the three-state case, I solve for the forward premium, and the optimal producer forward position, in closed form, with no need to rely on approximations.

For the purposes of the discrete model, I take $H^S - Q^D$ to be the fundamental uncertainty. I assume $H^S - Q^D$ to be a symmetrical distribution[2] $\Big(\text{Skew}(H^S - Q^D) =$

[2]In order to ensure symmetry in the excess capacity distribution, in what follows I will set

$0\big)$, with $E\big(H^S - Q^D\big) = \widehat{HQ}$. $H^S - Q^D$ variability is controlled by the parameter σ,

$$\big(H^S - Q^D\big)[\omega] = \widehat{HQ}\begin{pmatrix} 1+\sigma \\ 1 \\ 1-\sigma \end{pmatrix}. \tag{1.23}$$

Thus, state 1 is the high excess capacity (and therefore low price) state; state 2 is the normal state; and state 3 is the low excess capacity (and therefore high price) state.

Convexity in the production cost function, equation (1.9), ensures that the distribution of spot prices is positively skewed even though excess demand is symmetrical,

$$P_W[\omega] = \begin{pmatrix} \psi_1 \\ \psi_2 \\ \psi_3 \end{pmatrix}, \tag{1.24}$$

where $\psi_1 \equiv a_1(N/\widehat{HQ})^{-1/x}(1+\sigma)^{1/x}$, $\psi_2 \equiv a_1(N/\widehat{HQ})^{-1/x}$, and $\psi_3 \equiv a_1(N/\widehat{HQ})^{-1/x}(1-\sigma)^{1/x}$. Given the spot price P_W, it follows that the mean $\mu_{P_W} \equiv E\big(P_W\big)$ and variance $\sigma^2_{P_W} \equiv E\big((P_W - \mu_{P_W})^2\big)$ of the spot price are given by

$$\mu_{P_W} = \sum_{k=1}^{3} \phi_k \psi_k, \tag{1.25}$$

$$\sigma^2_{P_W} = \sum_{k=1}^{3} \phi_k \psi_k^2 - \mu^2_{P_W}. \tag{1.26}$$

$\phi_1 = \phi_3 = \phi$ and $\phi_2 = 1 - 2\phi$. The assumption of symmetric excess demand is consistent with the actual PJM data. The overall sample has skewness = -0.58.

Further, the skewness, $\mathrm{Skew}(P_W) \equiv E\Big((P_W - \mu_{P_W})^3\Big)/(\sigma^2_{P_W})^{3/2}$, and kurtosis, $\mathrm{Kurt}(P_W) \equiv E\Big((P_W - \mu_{P_W})^4\Big)/(\sigma^2_{P_W})^2$, of the spot price are given by

$$\mathrm{Skew}(P_W) = \left(\frac{1}{\sigma^2_{P_W}}\right)^{3/2} \left[\Big(\sum_{k=1}^{3} \phi_k \psi_k^3\Big) - 3\mu_{P_W}\Big(\sum_{k=1}^{3} \phi_k \psi_k^2\Big) + 2\mu^3_{P_W}\right], \quad (1.27)$$

$$\mathrm{Kurt}(P_W) = \left(\frac{1}{\sigma^2_{P_W}}\right)^{2} \left[\Big(\sum_{k=1}^{3} \phi_k \psi_k^4\Big) - 4\mu_{P_W}\Big(\sum_{k=1}^{3} \phi_k \psi_k^3\Big) \right. \quad (1.28)$$

$$\left. +6\mu^2_{P_W}\Big(\sum_{k=1}^{3} \phi_k \psi_k^2\Big) - 3\mu^4_{P_W}\right] \quad (1.29)$$

The forward premium ($\mathrm{PREM} \equiv P_F - E(P_W)$) is given in equation (1.17). Define the ratio of the expected spot price to the retail price to be $\frac{E(P_W)}{P_R} \equiv \Theta$. Then Appendix B shows that the forward premium is given by

$$\mathrm{PREM} = \frac{N_P}{N\,a^x}\left\{\frac{\mu_{P_W}}{\Theta}\left[\Big(\sum_{k=1}^{3} \phi_k \psi_k^{x+1}\Big) - \mu_{P_W^x}\mu_{P_W}\right] - \right.$$

$$\left.\Big(\frac{x}{x+1}\Big)\left[\Big(\sum_{k=1}^{3} \phi_k \psi_k^{x+2}\Big) - \mu_{P_W^{x+1}}\mu_{P_W}\right]\right\}, \quad (1.30)$$

where $\mu_{P_W^x} \equiv E(P_W^x)$ and $\mu_{P_W^{x+1}} \equiv E(P_W^{x+1})$. This equation is exact; it requires no approximations.

Now take the partial derivative of the forward premium with respect to Θ, i.e., with respect to the ratio of the expected spot price to the retail price.

$$\frac{\partial\,\mathrm{PREM}}{\partial\theta} = -\frac{N_P\,\mu_{P_W}}{N\,a^x\,\Theta^2}\left[\Big(\sum_{k=1}^{3} \phi_k \psi_k^{x+1}\Big) - \mu_{P_W^x}\mu_{P_W}\right], \quad (1.31)$$

$$> 0. \quad (1.32)$$

Ceteris paribus, the equilibrium for premium is increasing in the ratio of the expected spot price to the retail price.

Substituting the expressions for the forward premium in equation (1.30) and the spot price variance in equation (1.26) into equation (1.15), the optimal forward position for producer i can be expressed as

$$Q_{Pi}^F = \frac{N_P}{A\,N\,a^x\Big(\sum_{k=1}^3 \phi_k \psi_k^2 - \mu_{P_W}^2\Big)} \Bigg\{ \frac{\mu_{P_W}}{\Theta} \Bigg[\Big(\sum_{k=1}^3 \phi_k \psi_k^{x+1} \Big) - \mu_{P_W^x} \mu_{P_W} \Bigg] - \Big(\frac{x}{x+1}\Big) \Bigg[\Big(\sum_{k=1}^3 \phi_k \psi_k^{x+2} \Big) - \mu_{P_W^{x+1}} \mu_{P_W} \Bigg] \Bigg\}. \tag{1.33}$$

It is easy to see that the optimal producer forward position, $Q_{P_i}^F$, also increases in the ratio of the expected spot price to the retail price, Θ. However, this response is due entirely to the increased forward premium. Producers *hedging* demands are unchanged. This point is important, and it is discussed further in the simulation section below.

The discrete case expressions for the forward premium and the optimal producer forward position, equations (1.30) and (1.33), are exact and can be evaluated easily in a spreadsheet. In what follows, I take $\phi_1 = \phi_3 = 10\%$ and $\phi_2 = 80\%$. I then vary σ, the variability parameter from 5% to 50% and examine the forward premium and the optimal producer forward position.

I set the number of producers and the number of retailers, N_P and N_R, to 20. I take expected system capacity to be 200 MW. Each producer is assumed to be identical, with 10 MW of capacity. Expected demand to be 100 MW, thus the expected excess capacity is 100 MW. Because excess capacity is symmetrically distributed, expected excess demand is always equal to 100 MW. The exponent c is set equal to 1.5. Consistent with the BL simulations, I set $a = 30\left(\frac{100}{N_P}\right)^{(c+1)}$ thus ensuring that the spot price is \$30/MWh in the normal state, i.e., state 2.

Figure 3 plots the optimal producer forward position as a function of the demand standard deviation, σ, and the ratio of the expected spot price to the retail price, $\frac{E(P_W)}{P_R} \equiv \Theta$. Figure 4 plots the equilibrium forward premium as a function of the demand standard deviation, σ, and the ratio of the expected spot price to the retail price, $\frac{E(P_W)}{P_R} \equiv \Theta$. Figures 3 and 4 demonstrate, for fixed σ, both the optimal producer forward position and the equilibrium forward premium are increasing in Θ.

1.5 Simulated Spot Price Characteristics

BL illustrate their model through numerical simulations. In order to enable straightforward comparisons of the refined model presented in this paper with the original BL model, in this subsection I perform numerical simulations. I set parameters to be consistent with those chosen by BL as much as possible.

Referring again to Table 1.2, note that the standard deviation of demand as a percentage of the mean varies by month from approximately 10% - 24%. From Tables 1.1 and 1.2, mean demand as a fraction of mean capacity ranges from 36% to 62%. At a maximum, demand reaches almost 92% of available capacity.

As in BL, I first take the demand distribution to be truncated (at zero) normal, with a mean of 100 MW. The demand standard deviation is varied from 4% to 25%.

Second, the number of producers is set to $N_P = 20$. System capacity is set to 240 MW. Third, all producers are assumed to be identical, thus $H_{Pi} = 12$ MW, $\forall i = 1, \ldots, N_P$.

Similar to BL, I set $a = 30\left(\frac{140}{N_P}\right)^{(c+1)}$ thus ensuring that the spot price is \$30/MWh given demand of 100 MW. The intent of this assumption is to maintain comparability across values of c. c is fixed at 1.5 in the refined model and $c = 4$ in the BL model. The simulation exercises were repeated for different values of c, with no significant effect on the conclusions.

Finally, the coefficient of the variance term in the utility function, A, in equation (1.1) is set to $\frac{0.8}{2^c}$, as in BL. Each simulation consists of 25,000 trials.

Figures 2a through 2e plot various characteristics of the spot price as functions of the demand standard deviation, σ. The figures demonstrate the ability of the refined model to mimic actual spot price behavior.

Figure 2a plots the expected spot price as a function of σ, for the refined model and the BL model. The expected spot price generated by the two models is very similar. Owing to the choice of parameters values, for both models the expected spot price is of the same order as the actual PJM prices displayed in Panel A of Table 1.4. In both models, the expected spot price increases in σ, a direct result of supply curve convexity.

Figure 2b plots the spot price variance as a function of σ. Again the models are similar, with the BL model producing slightly higher variances. Notice that, as σ increases, the spot price variance in the refined model increases at a faster rate than in the BL model. This is another manifestation of the step function-like behavior of the supply curve in the refined model.

It can be argued that the refined model does a good job in reproducing price spikes. Consider Figure 2c, which plots the maximum spot price generated in the course of the 25,000 trials, as a function of σ. As σ increases, the maximum Q^D observed in the course of the trials increases, so that the minimum $H^S - Q^D$ decreases, thereby increasing the spot price from equation (1.13). The maximum spot price generated by the refined model was \$1,064/MWh. The BL model managed only \$264/MWh. The refined model is able to reproduce the price spikes seen in electricity markets.

Figures 2d and 2e plot for each model the spot price skewness and kurtosis, respectively, again as a function of σ. From Panel A of Table 1.4, the skewness in PJM spot prices varies from slightly below 2.0 to well over 10. Figure 2d shows that the refined model can easily generate this level of skewness (while still maintaining a reasonable level for the expected spot price). In no case does the BL produce skewness greater than 1.5. The PJM spot price kurtosis varies from around 8 to over 300. Figure 2e demonstrates that the refined model can produce kurtosis of this magnitude, while the BL model cannot.

The spot price behavior in the refined model is attributable directly to the form of the production cost function (equation (1.9)) and the spot price expression (equation (1.13)) that results. The step function-like behavior of the spot price from equation (1.13) mimics the actual supply curve in electricity markets. Occasionally, demand in wholesale electricity markets approaches the available supply. This can happen because (i) demand increases to a high level, (ii) supply decreases due to plant outages, or (iii) both. When demand approaches supply, there is no inventory to absorb the shock, so the shock is translated directly into spot prices. The result is that spot prices spike to extremely high levels. These high spot prices do not persist for long, rarely for more than a few days. This behavior shows up in the summary statistics as positive skewness and high excess kurtosis. Under reasonable assumptions, the refined model is able to reproduce the price spikes that occur in electricity markets. As in the actual market, these price spikes lead to positive skewness and large excess kurtosis in the distribution of spot prices.

1.5.1 Optimal Producer Forward Positions

Recalling that expected demand is 100 MW, equation (1.16) shows that, in the absence of any forward premium and with normally distributed spot prices, each producer will sell forward his entire capacity, $H_{Pi} = 12$ MW, minus an equal fraction of the system excess capacity, $\frac{H^S - Q^D}{N_P} = 7$ MW. Thus, each producer will sell forward 5 MW. I will refer to this as the *normal* case. Also from equation (1.16), producers will optimally decrease (increase) forward sales in response to a negative (positive) forward premium. Further, producers increase forward sales in response to positive spot price skewness.

Figure 5 plots the optimal forward position for a producer as a function of demand standard deviation, σ, and the ratio of expected spot price to the retail price, $\frac{E(P_W)}{P_R}$. When the retail price exceeds the expected spot price, i.e., $\frac{E(P_W)}{P_R} > 1$, then the forward premium is uniformly positive and producers optimally sell forward more than the 5 MW of the normal case. As σ increases, both the forward premium and spot price skewness increase. Both effects lead to higher forward sales by producers.

When σ is low and the retail price exceeds the expected spot price, i.e., $\frac{E(P_W)}{P_R} < 1$, the forward premium is negative. In this region, spot price skewness is positive but small in magnitude so that the response to the forward premium is the dominant effect and producers optimally reduce forward sales below the 5 MW of the normal case. As σ grows, both the forward premium and spot price skewness grow, hence producers' optimal forward sales grow. Notice that, even though the absolute magnitude of optimal producer forward sales is less for low expected spot prices than for high expected spot prices, the growth rate of the forward sales as σ increases is higher for low expected spot prices.

1.5.2 Equilibrium Forward Premium in Simulated Economies

It was emphasized above that the behavior of the equilibrium forward premium depends on the level of the expected spot price *relative* to the retail price. As summarized in Table 1.3 and discussed above, the forward premium decreases in the spot price variance when the expected spot price is less than the retail price, i.e., $\frac{E(P_W)}{P_R} < 1$, and increases in the spot price variance when the expected spot price exceeds the retail price, i.e., $\frac{E(P_W)}{P_R} > 1$.

Figure 6 plots the forward premium as a function of demand standard deviation, σ, and the $\frac{E(P_W)}{P_R}$ ratio. Consider first the cases with $\frac{E(P_W)}{P_R} < 1$. The premium first decreases and finally increases in σ. Recall that, in this region, the forward premium decreases with spot price variance and increases in spot price skewness. At low demand standard deviations, the negative relation of the forward premium to the spot price variance dominates. At higher demand standard deviations, spot price skewness increases dramatically (as shown in Figure 2d and discussed in the previous subsection). Here the positive relation between the forward premium and the spot price skewness is the dominant factor. The closer to 1 is $\frac{E(P_W)}{P_R}$, the sooner this reversal occurs.

At high expected spot prices, e.g., $\frac{E(P_W)}{P_R} \geq 1$, the forward premium is uniformly positive and increasing in σ. At higher expected spot prices, the forward premium is increasing in both the spot price variance and skewness.

For the purposes of examining the behavior of the forward premium, it is instructive to consider the hedging portions of the optimal forward positions for producers, equation (1.16), and retailers, equation (1.7). The first term on the right hand side of both equations (1.16) and (1.7), is the optimal response to a nonzero forward premium $\frac{PREM}{A\text{Var}(P_W)}$. The remainder of the right hand side of equation (1.16) is the forward position entered into by producers to hedge their profit stream. Similarly, the remainder of the right hand side of equation (1.7) is the forward position entered into by retailers to hedge their profit stream. Refer to these as the hedging demands of producers and retailers, respectively.

Figures 7a and 7b plot the hedging demands of producers and retailers as functions of demand standard deviation and the $\frac{E(P_W)}{P_R}$ ratio.[3] Consider first producers hedging demands, Figure 7a. The hedging demands of producers do not depend on the $\frac{E(P_W)}{P_R}$ ratio. Producers' hedging demands are invariant to the retail price. Changing the $\frac{E(P_W)}{P_R}$ ratio does not change the variance or the skewness of the spot price, thus it does not change producers' hedging demands. However, producers do increase forward sales in response to an increase in σ. Higher demand standard

[3] In creating these figures, I have assumed that all retailers are identical, and that the demand for each is simply an equal fraction of the overall system demand.

deviation results in an increase in the variance of producers' profits. In order to smooth their profit stream, producers optimally increase their forward sales. In the event that a high (low) spot price actually obtains, producers will make (lose) money on their spot sales and lose (make) on their forward sales.

Now consider retailers' hedging demands. Notice that, for any σ, as the $\frac{E(P_W)}{P_R}$ ratio increases, retailers optimally desire to increase their forward purchases. As the $\frac{E(P_W)}{P_R}$ ratio increases, the variance of retailers' profits increases. Retailers thus increase their forward purchases to smooth their profit stream. This observation explains the fact that, for any fixed σ, the forward premium increases in the $\frac{E(P_W)}{P_R}$ ratio, as shown in Figure 6. As $\frac{E(P_W)}{P_R}$ ratio increases, producers' hedging demands are unchanged, but retailers wish to increase forward purchases. This increase in demand for forward purchases drives up the equilibrium forward premium. It needs to be emphasized that this forward premium behavior occurs even though the anticipated variance and skewness of spot prices do not change.

The refined model produces a much larger range for the forward premium than does the BL model. Using similar parameter values the BL model produced a maximum forward premium of 30%, compared to over 170% in the refined model. As BL point out, the size of the premium can be made arbitrarily large or small by varying the risk preference coefficient A. However, regardless of the value of A, the refined model produces a much larger range for the forward premium.

In order to examine the behavior of the forward premium in Region III, the simulations were rerun with $\frac{E(P_W)}{P_R}$ varied from 3.2 to 3.8. Recalling that $c = 1.5$, the forward premium is decreasing in spot price skewness for $\frac{E(P_W)}{P_R} > 3.5$. However, at these extremely high spot prices, the magnitude of α, the coefficient on the spot price variance in equation (15), increases and this increase swamps the skewness effect.

1.6 Conclusions

This paper presents a refinement to the equilibrium electricity model in Bessembinder and Lemmon (2002). Unlike the BL model, the refined model explicitly accounts for constrained capacity. Because electricity cannot be stored, supply and demand must balance in real-time, thus the available supply (capacity) is relevant, as well as the level of demand.

The new model is able to reproduce the price spikes observed in wholesale electricity markets. Given the importance of price spikes to industry participants, it is crucial for hedging purposes to capture price spikes.

The equilibrium forward premium is shown to be decreasing in the variance of the spot price when the expected spot price is low, less than the retail price of electricity. However, when the expected spot price exceeds the retail price, the forward premium is increasing in spot price variance. These relationships arise because of the hedging demands of retailers. At low expected spot prices, retailers' profits are positively related to the spot price. At high expected spot prices, retailers' profits are negatively related to the spot price. In either case, retailers' hedging demands increase with spot price variance. More work is needed to study the impact of constrained capacity.

Table 1.1: Summary Statistics for PJM Daily Capacity (MW)

Month	Mean	Stdev	Min	Max
Jun-00	56,551	129	55,944	56,652
Jul-00	56,410	404	55,433	56,978
Aug-00	56,701	314	55,910	57,045
Sep-00	57,405	153	56,609	57,469
Oct-00	58,511	15	58,506	58,556
Nov-00	57,877	1	57,876	57,877
Dec-00	57,905	19	57,877	57,919
Jan-01	58,099	75	57,876	58,126
Feb-01	58,464	80	58,201	58,501
Mar-01	58,803	121	58,656	58,943
Apr-01	59,221	55	59,199	59,388
May-01	58,948	80	58,868	59,188
Jun-01	58,231	129	58,043	58,356
Jul-01	58,579	61	58,455	58,711
Aug-01	58,845	150	58,625	59,050
Sep-01	58,890	0	58,890	58,890
Oct-01	59,605	4	59,604	59,618
Nov-01	59,217	0	59,217	59,217
Dec-01	59,357	0	59,357	59,357
Jan-02	61,424	85	61,294	61,601
Feb-02	60,719	25	60,698	60,748
Mar-02	61,662	0	61,662	61,662
Apr-02	61,613	14	61,603	61,656
May-02	61,817	297	61,564	62,244
Jun-02	62,797	219	62,485	63,208
Jul-02	62,621	251	62,136	62,802
Aug-02	62,156	212	61,860	62,465
Sep-02	62,983	240	62,844	63,672
Oct-02	63,315	15	63,257	63,320
Nov-02	63,873	0	63,873	63,873
Dec-02	63,471	303	63,149	64,055
Jan-03	64,144	144	63,727	64,237
Feb-03	64,357	60	64,255	64,432
Mar-03	63,875	42	63,727	63,923
Apr-03	64,453	152	63,941	64,723
May-03	63,587	52	63,561	63,844

Table 1.2: Summary Statistics for PJM Daily Demand (MW)

Month	Mean	Stdev	Min	Max
Jun-00	32,605	6,608	19,525	49,305
Jul-00	31,810	5,896	20,667	47,958
Aug-00	32,929	6,842	19,996	49,462
Sep-00	29,662	5,584	19,021	45,021
Oct-00	27,119	4,134	18,955	35,917
Nov-00	28,779	4,006	19,972	38,083
Dec-00	33,041	3,908	22,023	41,489
Jan-01	32,709	3,800	24,769	41,476
Feb-01	31,081	3,619	21,537	41,150
Mar-01	29,964	3,586	22,759	38,238
Apr-01	26,989	3,941	18,549	35,345
May-01	27,466	4,751	18,790	40,647
Jun-01	32,946	7,469	19,532	50,157
Jul-01	32,676	7,029	20,601	52,132
Aug-01	36,468	7,334	21,909	54,030
Sep-01	29,112	5,398	19,023	43,184
Oct-01	27,338	4,056	19,121	34,789
Nov-01	27,201	3,924	19,031	34,848
Dec-01	29,523	4,169	19,985	38,743
Jan-02	31,076	3,984	21,320	40,002
Feb-02	30,239	3,896	22,010	39,201
Mar-02	28,875	4,113	19,247	38,883
Apr-02	28,530	4,842	19,533	44,537
May-02	28,083	4,929	19,360	43,566
Jun-02	33,588	7,626	19,954	52,938
Jul-02	38,051	8,540	21,984	55,690
Aug-02	37,995	8,396	21,327	55,934
Sep-02	31,282	6,041	19,950	47,186
Oct-02	28,820	4,751	19,702	41,859
Nov-02	29,418	3,811	20,466	37,169
Dec-02	32,563	4,374	23,985	42,386
Jan-03	34,985	4,592	22,332	46,420
Feb-03	34,042	3,783	24,960	41,961
Mar-03	30,193	4,421	20,581	41,538
Apr-03	28,314	4,085	20,103	37,098
May-03	27,189	4,017	19,414	34,104

Table 1.3: Model Forward Premium Predictions

Region	Expected Spot Price	α	γ
I	$\frac{E(P_W)}{P_R} < 1$	-	+
II	$1 < \frac{E(P_W)}{P_R} < (c+2)$	+	+
III	$(c+2) < \frac{E(P_W)}{P_R}$	+	-

The table presents the model's predictions regarding the signs of the coefficients α and γ from equation (1.18), $P_F - E(P_W) = \alpha \text{Var}(P_W) + \gamma \text{Skew}(P_W)$. The left hand side is the forward premium, defined to be the difference between the forward price, P_F, and the expected spot price, $E(P_W)$. P_R is the fixed retail price of electricity. c is the exponent in the production cost function, $TC_i = F + \frac{a}{c}\left(\frac{1}{H_{Pi} - Q_{Pi}}\right)^c$. The sign of alpha can change depending upon the level of the expected spot price relative to the retail price.

Table 1.4: Summary Statistics for PJM Hourly Spot Prices

Hour	Mean	Stdev	Skew	Kurt	Min	Max
1	$19.84	$11.08	2.93	15.89	$0.00	$107.84
2	$18.93	$12.22	3.47	24.23	-$16.40	$145.07
3	$17.04	$10.99	3.34	22.59	-$2.40	$127.43
4	$16.36	$10.40	3.25	23.27	-$2.42	$117.86
5	$17.48	$10.77	3.46	21.91	-$4.74	$114.60
6	$21.21	$12.79	2.76	15.15	$0.00	$125.42
7	$30.52	$22.35	2.05	8.92	$0.00	$182.13
8	$34.75	$25.08	1.92	8.03	$0.00	$207.46
9	$33.93	$20.54	2.01	8.66	-$1.92	$167.76
10	$38.05	$21.57	1.80	8.40	-$2.05	$186.04
11	$43.24	$24.84	1.87	10.55	$10.52	$249.68
12	$42.96	$40.97	10.54	174.88	$7.08	$846.50
13	$43.37	$53.36	11.18	167.87	$2.63	$1,005.53
14	$47.32	$64.82	10.51	141.00	$4.37	$1,020.28
15	$43.93	$67.43	10.56	139.34	$5.19	$1,019.97
16	$42.90	$70.43	10.60	134.65	$7.80	$1,019.72
17	$46.60	$62.95	10.83	150.61	$11.83	$1,019.74
18	$50.07	$54.44	11.22	184.76	$6.13	$1,019.75
19	$44.73	$44.01	9.82	153.00	$11.55	$801.55
20	$41.73	$30.56	7.89	142.43	$9.73	$645.32
21	$44.77	$39.19	13.95	321.50	$13.18	$994.98
22	$37.47	$22.89	3.59	37.22	$11.09	$352.38
23	$26.62	$13.95	2.43	11.42	$6.41	$116.32
24	$21.98	$11.81	3.91	29.37	$0.78	$157.24
Total	$34.41	$39.17	13.68	297.40	-$16.40	$1,020.28

The table contains the mean, standard deviation, skewness, and kurtosis for Pennsylvania, New Jersey, Maryland (PJM) market electricity spot prices for the three year period from June 2000 through May 2003. The units, for all values except skew and kurtosis, are $/MWh. Log prices are not reported since the price of electricity can be less than or equal to zero. 'Hour' refers to the hour *ending* at the given time. For example, Hour 1 refers to 12-1am, Hour 2 refers to 1-2am, Hour 17 refers to 4-5pm, etc.

Table 1.5: Summary Statistics for PJM Hourly Forward Prices

Hour	Mean	Stdev	Skew	Kurt	Min	Max
1	\$20.77	\$9.58	2.67	12.73	\$5.00	\$81.73
2	\$18.13	\$8.37	2.90	15.69	\$0.00	\$77.94
3	\$16.82	\$8.11	2.67	15.30	\$0.00	\$75.22
4	\$16.50	\$8.31	2.62	15.00	\$0.00	\$74.07
5	\$17.38	\$9.12	2.74	15.04	\$0.00	\$79.90
6	\$21.43	\$12.05	2.68	13.64	\$0.10	\$100.65
7	\$31.57	\$20.79	1.98	8.89	\$1.00	\$153.25
8	\$35.49	\$20.54	1.64	6.95	\$1.15	\$155.71
9	\$36.59	\$17.92	1.61	7.52	\$11.01	\$153.00
10	\$39.29	\$17.45	1.43	6.68	\$13.45	\$152.79
11	\$41.85	\$19.35	2.06	12.28	\$14.95	\$198.10
12	\$42.20	\$24.07	5.56	63.96	\$14.47	\$390.93
13	\$41.52	\$29.65	8.42	119.23	\$14.68	\$545.46
14	\$42.75	\$36.33	8.95	123.72	\$13.75	\$646.81
15	\$43.00	\$44.77	10.23	149.03	\$13.30	\$818.54
16	\$43.50	\$46.66	9.81	141.50	\$13.87	\$859.05
17	\$46.77	\$45.90	8.58	109.49	\$15.03	\$779.38
18	\$52.65	\$39.53	6.67	74.11	\$15.02	\$599.22
19	\$50.09	\$30.13	4.20	41.45	\$14.91	\$450.01
20	\$47.35	\$26.58	4.21	44.39	\$15.06	\$416.27
21	\$45.84	\$26.57	6.37	90.78	\$15.10	\$498.01
22	\$37.99	\$17.85	2.23	13.12	\$15.00	\$185.90
23	\$28.90	\$12.78	2.08	9.56	\$12.68	\$112.86
24	\$23.15	\$10.19	2.33	10.60	\$0.00	\$95.78
Total	\$35.06	\$28.18	8.68	168.43	\$0.00	\$859.05

The table contains the mean, standard deviation, skewness, and kurtosis for Pennsylvania, New Jersey, Maryland (PJM) market electricity forward prices for the three year period from June 2000 through May 2003. The units, for all values except skew and kurtosis, are \$/MWh. Log prices are not reported since the price of electricity can be less than or equal to zero. 'Hour' refers to the hour *ending* at the given time. For example, Hour 1 refers to 12-1am, Hour 2 refers to 1-2am, Hour 17 refers to 4-5pm, etc.

Table 1.6: Regression Results of log(Spot Price) on log(Capacity - Demand) for PJM Markets.

Season	NOBS	k	c	R^2
Summer	6,508	20.97	0.754	73.8%
		(0.131)	(0.013)	
Fall	6,521	24.63	1.073	39.6%
		(0.328)	(0.032)	
Winter	6,466	24.26	1.038	30.6%
		(0.391)	(0.038)	
Spring	6,575	30.02	1.564	39.5%
		(0.407)	(0.039)	

The following regression is performed,

$$\log(P_W) = k - (c+1)\log(H - Q_D),$$

Summer is defined to be June, July, and August. Succeeding seasons are defined to be succeeding three month periods. The data are taken from www.pjm.com.

Figure 1.1: Spot Price vs. Demand

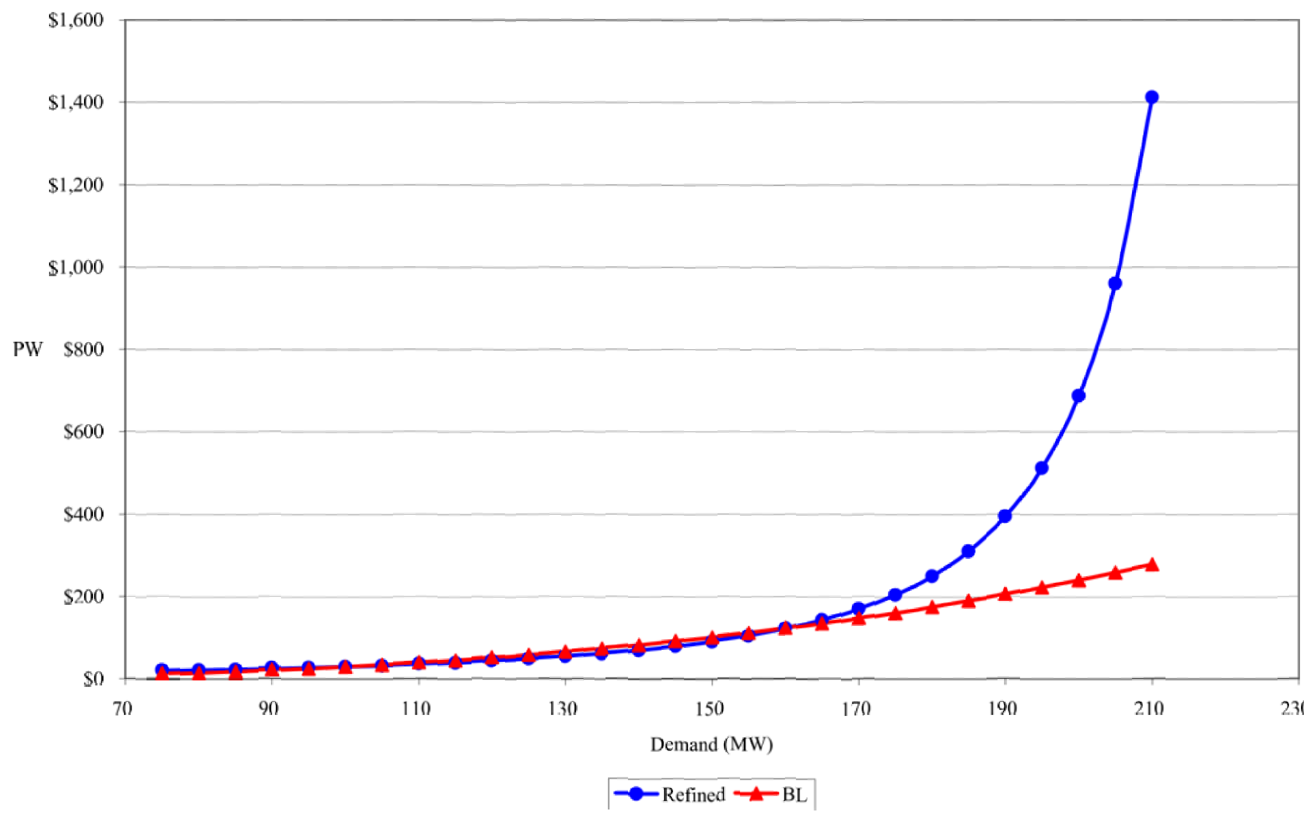

Figure 1.2a: Expected Spot Price vs. Demand Standard Deviation

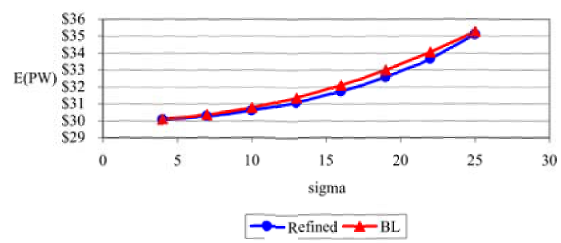

Figure 1.2b: Spot Price Variance vs. Demand Standard Deviation

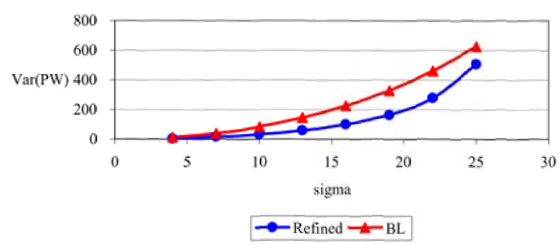

Figure 1.2d: Spot Price Skewness vs. Demand Standard Deviation

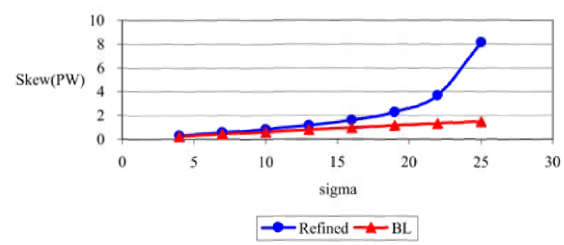

Figure 1.2c: Maximum Spot Price vs. Demand Standard Deviation

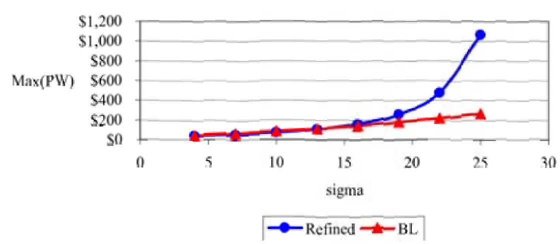

Figure 1.2e: Spot Price Kurtosis vs. Demand Standard Deviation

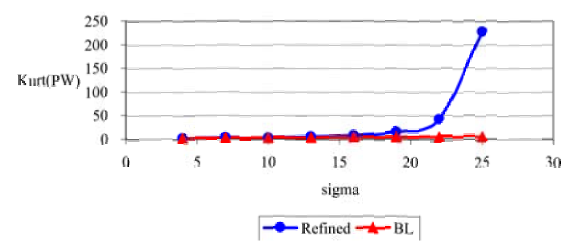

Figure 1.3: Optimal Producer Forward Sales from Discrete Model

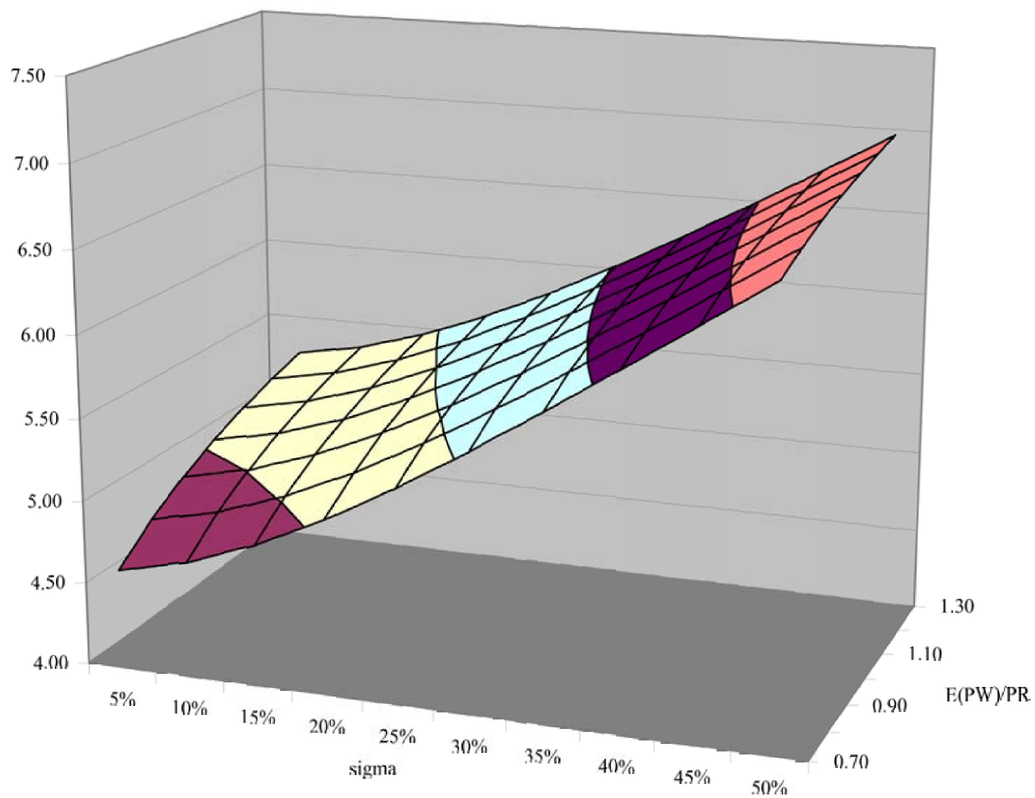

Figure 1.4: Forward Premium as a % of the Expected Spot Price
from Discrete Model

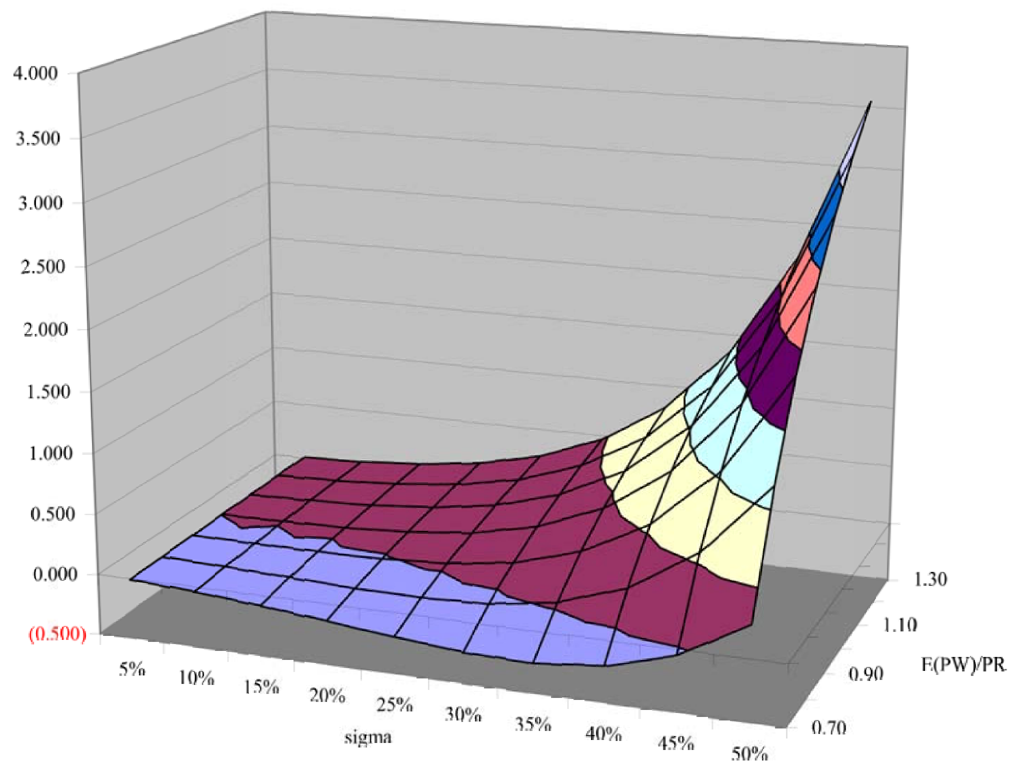

Figure 1.5: Optimal Producer Forward Sales from Simulations

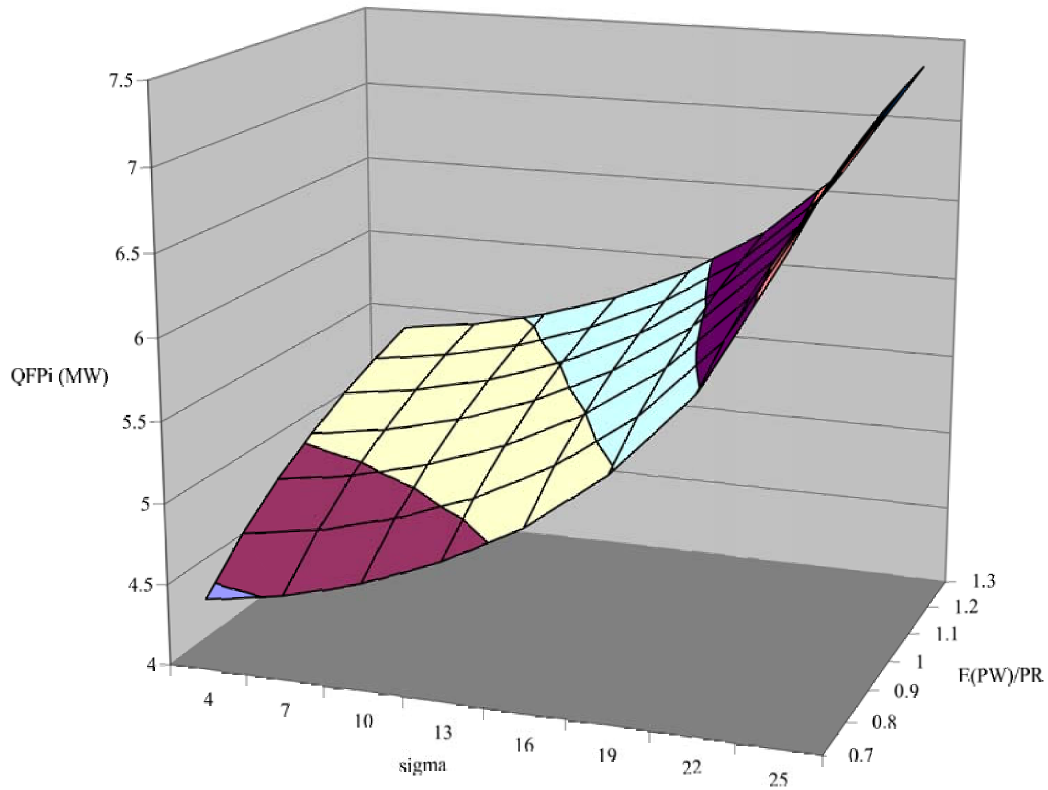

Figure 1.6: Forward Premium as a % of the Expected Spot Price
from Simulations

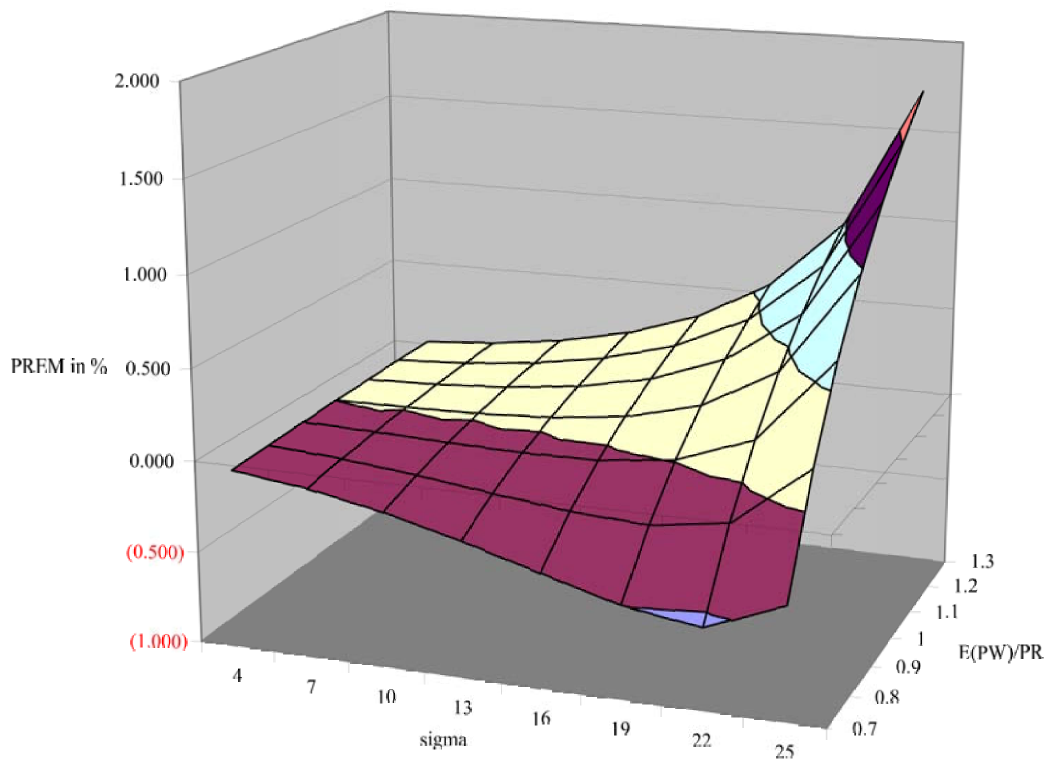

Figure 1.7a: Producer Hedging Demand

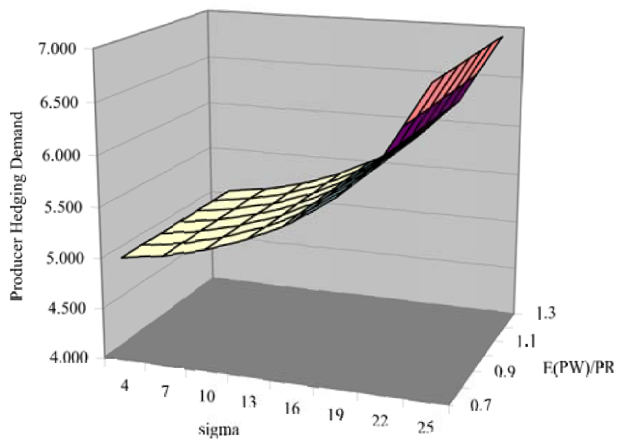

Figure 1.7b: Retailer Hedging Demand

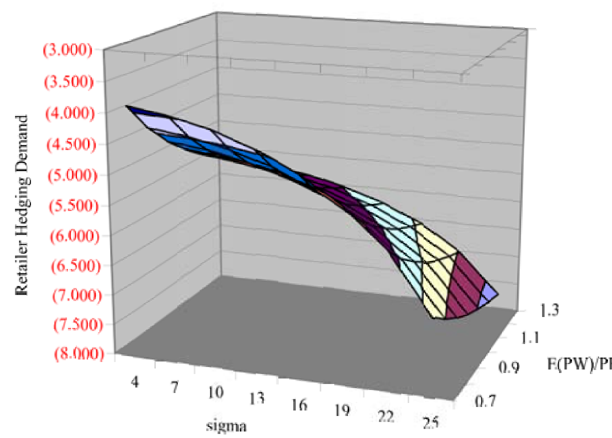

Chapter 2

Estimation of Continuous-Time Models of Foreign Exchange Volatility

2.1 Introduction

Substantial research effort has been directed towards understanding and modeling volatility (e.g., Engle (2004) and Andersen, Bollerslev, Diebold, Ebens (2001)) in financial markets. How does currency return volatility evolve over time and what are the properties of volatility dynamics? Is the drift of currency return volatility non-linear? What forms of non-linearities are admitted in the drift and diffusion functions? The purpose of this study is to estimate a large class of volatility processes and explore these issues using weekly data on two currency pairs: U.S. dollar-British pound and Japanese Yen-U.S. dollar. My estimation approach is based on maximum-likelihood estimation that relies on closed-form density approximations. I propose a general stochastic elasticity of variance specification for currency volatility that nests several plausible models of economic interest. Based on at-the-money volatility, I find that certain types of non-linearities in the drift function improve the goodness of fit statistics. Moreover, the constant elasticity of variance specification provides a reasonable characterization of the variance of variance function. The results pose challenges to existing paradigms and approaches to modeling currency returns and currency distributions (e.g., Bates (1996), Brandt and Santa-Clara (2002), and Bakshi, Carr, and Wu (2005)).

The general class of continuous-time models of currency return volatility can be written as a stochastic differential equation of the type:

$$dX_t = \mu[X_t; \Theta]\, dt \,+\, \sigma[X_t; \Theta] dW_t, \tag{2.1}$$

where W_t is a Brownian motion, $\mu[X_t; \Theta]$ is the risk neutral drift function, and $\sigma[X_t; \Theta]$ is the diffusion function. $\mu[X_t; \Theta]$ and $\sigma[X_t; \Theta]$ are assumed to be functions of the state variable X_t and some unknown parameter vector $\Theta \in \Re^{\mathrm{K}}$. If the preferred method for estimating the parameters Θ is maximum likelihood, then a dilemma arises as the transition density function is not known for the entire class, but for only a few simple cases. Many schemes have been proposed to deal with this issue, including numerical simulation, Monte Carlo Markov Chains, and various nonparametric approaches. In particular, Aït-Sahalia (1999, 2002) develops a method to approximate the transition density based on Hermite expansion. Bakshi, Ju, and Ou-Yang (2005) extend his approximation methodology to a broader class of drift and diffusion functions.

Using the density approximation approach of Aït-Sahalia (1999, 2002) and Bakshi, Ju, and Ou-Yang (2005) I estimate the stochastic elasticity of variance class of models where $\sigma[X_t; \Theta] = \sqrt{\beta_0 + \beta_1 X_t \,+\, \beta_2 X_t^{\beta_3}}$. I also consider a class of $\mu[X_t; \Theta] = \alpha_0 + \alpha_1 X_t + \alpha_2 X_t^2 + \alpha_3 / X_t$ as in Aït-Sahalia (1996). The advantage of this class of drift and diffusion functions is that it allows flexible specifications of mean-reversion in volatility and exponent parameter β_3.

Guided by Engle (2004), I proxy volatility using at-the-money implied volatility from currency call options. Even though currency return volatility is far less than volatility in equity markets, the exponent parameter β_3 is estimated to be approximately 1.26 for YEN-USD volatility. The exponent parameter β_3 is less than unity for the full sample (1996-2004) of USD-GBP return volatility, but subperiod estimation reveals that this result does not hold in the most recent (2000-2004) data.

As opposed to a large literature that advocates non-linearities in the currency returns and currency return volatility, I find that adding a nonlinearity to the drift function through $\alpha_2 X_t^2$ improves the fit of the model for only the USD-GBP currency pair. There is little change in log-likelihood criterion function with the YEN-USD currency pair.

One important implication of the results of this paper is that generalizing the constant elasticity of variance (CEV) diffusion function does not substantially improve the fit of the model for any currency pair. For instance, the contribution of the $\beta_1 X_t$ term in $\sigma[X]$ is not that large. The values of β_1 are consistently small but nonetheless statistically significant.

This article proceeds as follows. Section II details the specific forms of the drift function, $\mu[X_t; \Theta]$, and the diffusion function, $\sigma[X_t; \Theta]$ to be estimated. Section III reviews the estimation procedure. Section IV describes the data, provides the estimation results, and briefly compares results for FX and equity volatility. Section V concludes.

2.2 Currency Volatility Models in the Stochastic Elasticity of Variance Class

In what follows, suppose S_t^{fh} is the currency h price of currency f at time-t, with h being the home currency. Let $R_{t,t+1}^{fh} \equiv \log\left(S_{t+1}^{fh}/S_t^{fh}\right)$ denote the logarithmic return on the currency.

Let X_t denote the instantaneous volatility of currency returns. Consider the following time-homogeneous continuous-time process for X_t,

$$dX_t = \mu[X_t;\Theta]\,dt + \sigma[X_t;\Theta]dW_t \tag{2.2}$$

where the class of drift and diffusion functions, respectively denoted as $\mu[X_t;\Theta]$ and $\sigma[X_t;\Theta]$, for foreign exchange volatility, are due to Aït-Sahalia (1996)

$$\mu[X_t;\Theta] = \alpha_0 + \alpha_1 X_t + \alpha_2 X_t^2 + \alpha_3/X_t, \tag{2.3}$$

$$\sigma[X_t;\Theta] = \sqrt{\beta_0 + \beta_1 X_t + \beta_2 X_t^{\beta_3}}. \tag{2.4}$$

This model, the stochastic elasticity of variance - nonlinear drift model, is referred to henceforth as SEV-ND. Notice that the SEV-ND allows for nonlinear drift and diffusion and nests various well-known continuous-time models:

- When $\alpha_2 = \alpha_3 = 0$ and $\beta_2 = 0$ in (2.3)-(2.4), the SEV-ND model reduces to the affine model (hereby AFF).

$$dX_t = (\alpha_0 + \alpha_1 X_t)dt + \sqrt{\beta_0 + \beta_1 X_t}dW_t. \tag{2.5}$$

The volatility model in Heston (1993) is a restricted case with $\beta_0 = 0$ in (2.5);

- Setting $\beta_0 = \beta_1 = 0$, I get the constant elasticity of variance (CEV) diffusion model with,

$$\sigma[X_t; \Theta] \;=\; \beta_2 X_t^{\beta_3}. \tag{2.6}$$

Three versions of the CEV model, each with differing drift functions, are estimated. The models, referred to in the following as CEV-CD (constant drift), CEV-LD (linear drift), and CEV-ND, respectively, are given by

$$dX_t \;=\; \begin{cases} \alpha_0\, dt + \beta_2\, X_t^{\beta_3}\, dW_t & \textbf{CEV-CD} \\ (\alpha_0 + \alpha_1 X_t)\, dt + \beta_2\, X_t^{\beta_3}\, dW_t & \textbf{CEV-LD} \\ (\alpha_0 + \alpha_1 X_t \;+\; \alpha_2\, X_t^2 + \alpha_3/X_t)\, dt \;+\; \beta_2\, X_t^{\beta_3} dW_t. & \textbf{CEV-ND.} \end{cases} \tag{2.7}$$

These models also are studied empirically in Aït-Sahalia (1996) and Durham (2004) for interest rates, and Bakshi, Ju, and Ou-Yang (2005) for equity index volatility.

- Restricting $\mu[X]$ and $\sigma[X]$ provides two additional versions of the SEV model, corresponding to the specifications of the drift function $\mu[X_t; \Theta]$, respectively denoted as SEV-CD and SEV-LD:

$$dX_t \;=\; \alpha_0\, dt + \sqrt{\beta_0 + \beta_1 X_t + \beta_2 X_t^{\beta_3}} dW_t, \tag{2.8}$$

$$dX_t \;=\; (\alpha_0 + \alpha_1\, X_t)dt + \sqrt{\beta_0 + \beta_1\, X_t + \beta_2\, X_t^{\beta_3}}\, dW_t, \tag{2.9}$$

The model in (2.9) permits a linear drift for foreign exchange volatility while (2.8) accommodates a constant drift.[1]

[1]Chan et. el. (1992) extend the diffusion function by allowing the exponent to be a part of the parameter set Θ, rather than being fixed at $\frac{1}{2}$. This results in the constant elasticity of variance (CEV) model specified in equation (2.6).

The general model is driven by an unknown parameter vector,

$$\Theta \equiv (\alpha_0, \alpha_1, \alpha_2, \alpha_3, \beta_0, \beta_1, \beta_2, \beta_3), \tag{2.10}$$

which needs to be estimated. Unfortunately, the density function for the stochastic elasticity of variance model is not known in exact form. However, it can be approximated using the methods in Aït-Sahalia (1999, 2002) and Bakshi, Ju, and Ou-Yang (2005). In the remainder of the paper I present the density approximation based estimation methodology and the insights it enables regarding the dynamic evolution of foreign exchange volatility. These results have important implications for building models of currency returns, option prices, and risk management.

2.3 Estimation Approach

Several techniques for estimating the parameters of the general diffusion given in equation (2.1) have been proposed in the literature. Gouriéroux, Monfort, and Renault (1993) and Gallant and Tauchen (1996) use simulation methods. Hansen and Scheinkman (1995) and Kessler and Sorensen (1999) develop methods to estimate the parameters using the GMM. Bollerslev and Zhou (2002) apply the GMM specifically to the case of FX volatility, using high frequency returns to generate realized volatility. Aït-Sahalia (1996) matches the underlying transition density nonparametrically, while Eraker (1997) and Jones (1997) resort to Bayesian approaches. Ruiz (1994) and Harvey and Shephard (1994) estimate a discrete-time stochastic volatility model using a quasi-maximum likelihood approach. Jacquier, Polson, and Rossi (1994) and Eraker (2001) resort to Markov Chain Monte Carlo

(MCMC) methods. For reasons discussed in Aït-Sahalia (2002) and Bakshi, Ju, and Ou-Yang (2005) I derive the closed-form density approximation and estimate the continuous-time model via maximum likelihood.

2.3.1 Likelihood Function

Without any loss of generality suppose one observes the discrete-time observations on foreign exchange volatility or variance X_t, $\{t = i\Delta | i = 0, n\}$ at equally spaced intervals Δ. Let $p_X[\Delta, X_{i\Delta}|X_{(i-1)\Delta}; \Theta]$ denote the conditional transition density of observing $X_{i\Delta}$ at time $i\Delta$ given that $X_{(i-1)\Delta}$ was observed at time $(i-1)\Delta$. An estimate of the parameter vector Θ can be found by maximizing the log likelihood function, defined to be,

$$L[\Theta] \equiv \sum_{i=1}^{n} \log\Bigg(p_X[\Delta, X_{i\Delta}|X_{(i-1)\Delta}; \Theta]\Bigg). \tag{2.11}$$

The estimated parameter vector $\hat{\Theta}$ is the set of values that maximize the log likelihood,

$$\hat{\Theta} = \text{argmax}_{\Theta}\ L[\Theta]. \tag{2.12}$$

As pointed out by Aït-Sahalia (1996) the exact transition density $p_X[\Delta, X_{i\Delta}|X_{(i-1)\Delta}; \Theta]$ is unknown in closed-form for the SEV-ND. However, based on Aït-Sahalia (1999, 2002) one could develop a closed-form, accurate approximation to the transition density and thus the likelihood function. I estimate the parameters Θ using an extension of the Aït-Sahalia (1999, 2002) approximation method proposed by Bakshi, Ju, and Ou-Yang (2005). The key features of this approximation approach are outlined below.

2.3.2 Density Approximation Approach

Aït-Sahalia (1999, 2002) relies on a transformation of the process $X(t)$ given in equation (2.1) to a process Y with unit variance. Given the transformation

$$Y \equiv \gamma(X;\Theta) = \int^{X} \frac{du}{\sigma(u;\Theta)}, \tag{2.13}$$

Ito's Lemma implies that

$$dY_t = \mu_Y[Y_t;\Theta]dt + dW_t, \tag{2.14}$$

where W_t is a Brownian motion, and the new drift function is given by

$$\mu_Y[Y_t;\Theta] = \frac{\mu(\gamma^{-1}(y;\Theta);\Theta)}{\sigma(\gamma^{-1}(y;\Theta);\Theta)} - \frac{1}{2}\frac{\partial\sigma}{\partial x}\Big(\gamma^{-1}(y;\Theta);\Theta\Big). \tag{2.15}$$

Aït-Sahalia (2002) then utilizes another transformation in order to convert the process Y into yet a third process Z,

$$Z \equiv \Delta^{-1/2}\Big(Y - y_0\Big) \tag{2.16}$$

which more closely resembles a normal random variable and thus is easier to approximate. He shows that it is possible to construct a convergent series of approximations to the transition density for Z, $p_Z(\Delta, z|y_0;\Theta)$. He then works backwards, using the Jacobian formula to obtain an approximation, $p_X^{(J)}$, to the original transition density of interest, $p_X[\Delta, X_{i\Delta}|X_{(i-1)\Delta};\Theta]$.

Importantly, in order for the approximation scheme to work in practice, Y_t and $\gamma^{-1}(y;\Theta)$ must be analytical in form. Further, the approximation results in recursion relations for expansion coefficients c_j which contain high dimensional integrals for

$j \geq 3$. These requirements limit the class of drift and diffusion functions to which the approximation method can be applied.

Bakshi, Ju, and Ou-Yang (2005) extend the methodology to a broader class of functions $\mu[X_t;\Theta]$ and $\sigma[X_t;\Theta]$ making it unnecessary to know the closed form expressions for Y_t and $\gamma^{-1}(y;\Theta)$. Specifically, Bakshi, Ju, and Ou-Yang (2005) begin by (see their Proposition 1, 2 and 3) defining the function

$$f[X] \equiv \frac{\mu[X]}{\sigma[X]} - \frac{\sigma'[X]}{2}, \tag{2.17}$$

where $\sigma'[X] \equiv \frac{\partial\sigma[X]}{\partial X}$. The density approximation can then be expressed in terms of $f[X]$, $\mu[X]$, and $\sigma[X]$. Up to K^{th} order, the density approximation is given by:

$$p_X^{(K)} \approx \frac{\Delta^{-1/2}}{\sigma[X;\Theta]}\phi\left[\frac{1}{\Delta^{1/2}}\int_{x_0}^{x}\frac{du}{\sigma[u]}\right]\exp\left(\int_{x_0}^{x} f[u]\frac{du}{\sigma[u]}\right)\sum_{k=0}^{K} c_k\left[\gamma[x]|\gamma[x_0];\Theta\right]\frac{\Delta^k}{k!}, \tag{2.18}$$

where $\phi[z] \equiv \frac{e^{-z^2/2}}{\sqrt{2\pi}}$. This approximation is shown to be highly accurate for small values of K. The function $f[X]$ and the coefficients $c_k, k = 1,4$ are given explicitly in Bakshi et al (2005) and are repeated here for convienience.

$$c_1[y] = \frac{1}{y-y_0}\int_{y_0}^{y}\lambda[w]dw, \tag{2.19}$$

$$c_2[y] = c_1^2[y] + \frac{1}{(y-y_0)^2}\left(\lambda[y]+\lambda[y_0]-2c_1[y]\right), \tag{2.20}$$

$$\begin{aligned} c_3[y] = {} & c_1^3[y] + \frac{3}{(y-y_0)^2}\left(c_1[y]\left(\lambda[y]+\lambda[y_0]\right)-3c_2[y]\right) + \\ & \frac{3}{(y-y_0)^3}\left(\frac{\lambda'[y]-\lambda'[y_0]}{2}+\int_{y_0}^{y}\lambda^2[w]dw\right), \end{aligned} \tag{2.21}$$

and,

$$\begin{aligned} c_4[y] = {} & c_1^4[y] + \frac{3}{(y-y_0)^2}\left(2c_2[y]\lambda[y]-8c_3[y]+2\lambda[y_0]c_1^2[y]\right) + \\ & \frac{12c_1[y]}{(y-y_0)^3}\left(\frac{\lambda'[y]-\lambda'[y_0]}{2}+\int_{y_0}^{y}\lambda^2[w]dw\right) + \end{aligned}$$

$$\frac{3}{(y-y_0)^4}(3\lambda^2[y]+5\lambda^2[y_0]+4\lambda[y]c_1[y]-$$
$$12c_2[y]+\lambda^{''}[y]+\lambda^{''}[y_0]), \tag{2.22}$$

where,

$$\lambda[y] = -\frac{1}{2}\Big(f^2[x]+f^{'}[x]\sigma[x]\Big), \tag{2.23}$$
$$\lambda^{'}[y] \equiv \frac{\partial\lambda[y]}{\partial y} = -\frac{1}{2}\sigma[x]\Big(f^2[x]+f^{'}[x]\sigma[x]\Big)^{'}, \tag{2.24}$$
$$\lambda^{''}[y] \equiv \frac{\partial^2\lambda[y]}{\partial y^2} = -\frac{1}{2}\sigma[x]\Big(\Big(f^2[x]+f^{'}[x]\sigma[x]\Big)^{'}\sigma[x]\Big)^{'}, \tag{2.25}$$
$$y-y_0 = \int_{x_0}^{x}\frac{du}{\sigma[u]}, \tag{2.26}$$

2.3.3 Density Approximation for the SEV-ND Model

An estimate of the values for the parameter vector Θ are found by replacing the actual transition density $p_X[\Delta, X_{i\Delta}|X_{(i-1)\Delta};\Theta]$ in equation (2.12) with its K^{th} order approximation $p_X^{(K)}[\Delta, X_{i\Delta}|X_{(i-1)\Delta};\Theta]$ from equation (2.18). In all estimation exercises, K is set to $K=4$ to achieve the desired degree of accuracy.

Equations (2.19) through (2.26) are straightforward (if tedious) to calculate and code for the case of equations (2.2), (2.3), and (2.4). The exact expressions required for the SEV-ND case are given below. Let $f^{'} \equiv \frac{\partial f[x]}{\partial x}$, $f^{''} \equiv \frac{\partial^2 f[x]}{\partial x^2}$, etc. Then

$$f[x] \equiv \frac{\mu[x]}{\sigma[x]} - \frac{\sigma^{'}[x]}{2}, \tag{2.27}$$
$$f^{'}[x][x] = \frac{\mu^{'}[x]}{\sigma[x]} - \frac{\mu[x]\sigma^{'}[x]}{\sigma^2[x]} - \frac{\sigma^{''}[x]}{2}, \tag{2.28}$$
$$f^{''}[x][x] = \frac{\mu^{''}[x]}{\sigma[x]} - \frac{1}{\sigma^2[x]}\Big(2\mu^{'}[x]\sigma^{'}[x]+\mu[x]\sigma^{''}[x]\Big) -$$

$$\frac{\mu[x]\,(\sigma'[x])^2}{\sigma^3[x]} - \frac{\sigma'''[x]}{2}, \tag{2.29}$$

$$\begin{aligned} f'''[x][x] &= \frac{\mu'''[x]}{\sigma[x]} - \frac{1}{\sigma^2[x]}\Big(3\mu''[x]\sigma'[x] + 3\mu'[x]\sigma''[x] + \mu[x]\sigma'''[x]\Big) + \\ & \frac{6}{\sigma^3[x]}\Big(\mu'[x]\,(\sigma'[x])^2 + \mu[x]\sigma'[x]\sigma''[x]\Big) \\ & -\frac{6}{\sigma^4[x]}\,\mu[x]\sigma'[x] - \frac{\sigma''''[x]}{2}, \end{aligned} \tag{2.30}$$

where

$$\mu[x] = \alpha_0 + \alpha_1 x + \alpha_2 x^2 + \alpha_3 x^{-1}, \tag{2.31}$$

$$\mu'[x] = \alpha_1 + 2\alpha_2 x - \alpha_3 x^{-2}, \tag{2.32}$$

$$\mu''[x] = 2\alpha_2 + 2\alpha_3 x^{-3}, \tag{2.33}$$

$$\mu'''[x] = -6\alpha_3 x^{-4}, \tag{2.34}$$

and

$$\sigma[x] = \Big(\beta_0 + \beta_1\, x \;+\; \beta_2\, x^{\beta_3}\Big)^{1/2}, \tag{2.35}$$

$$\sigma'[x] = \frac{1}{2\sigma[x]}\Big(\beta_1 \;+\; \beta_2\beta_3\, x^{\beta_3 - 1}\Big), \tag{2.36}$$

$$\sigma''[x] = \frac{1}{\sigma[x]}\Big(- \Big(\sigma'[x]\Big)^2 + \frac{1}{2}\beta_2\beta_3(\beta_3 - 1)\, x^{\beta_3 - 2}\Big), \tag{2.37}$$

$$\begin{aligned} \sigma'''[x] &= \frac{1}{2\sigma[x]}\beta_2\beta_3(\beta_3 - 1)\, x^{\beta_3 - 3}\Big(\beta_3 - 2 - \\ & \frac{\sigma'[x]}{\sigma[x]}x\Big) + \frac{\sigma'[x]}{\sigma[x]}\Big(\frac{(\sigma'[x])^2}{\sigma[x]} - 2\sigma''[x]\Big), \end{aligned} \tag{2.38}$$

$$\begin{aligned} \sigma''''[x] &= \frac{1}{2\sigma[x]}\beta_2\beta_3(\beta_3 - 1)\, x^{\beta_3 - 4}\Big((\beta_3 - 2)(\beta_3 - 3) - 2(\beta_3 - 2)x\frac{\sigma'[x]}{\sigma[x]} \\ & -x^2\frac{\sigma''[x]}{\sigma[x]} + 2x^2\Big(\frac{\sigma'[x]}{\sigma[x]}\Big)^2\Big) + \frac{1}{\sigma[x]}\Big(- 2(\sigma''[x])^2 - \\ & 2\sigma'[x]\sigma'''[x] - \frac{5}{\sigma[x]}(\sigma'[x])^2\sigma''[x] - \frac{2}{\sigma^2[x]}(\sigma'[x])^4\Big). \end{aligned} \tag{2.39}$$

Given these expressions, log likelihood values can be computed from the approximate conditional densities and the various models can be compared to one another.

2.4 Comparison of Volatility Models

This section presents estimation results of variance dynamics for the USD-GBP and YEN-USD currency pairs. For conciseness I divide the results into four parts containing a discussion of (i) estimation results for the drift function, $\mu[X_t; \Theta]$, (ii) estimation results for the diffusion function, $\sigma[X_t, \Theta]$, (iii) robustness tests, and (iv) a brief comparison of FX volatility results with those for equity volatility. The section begins with a brief discussion of the proxy for unobservable volatility and the data used in the estimation process.

2.4.1 Data

The volatility of logarithmic changes in spot currency is unobservable. Unobservable volatility is proxied using at-the-money, Black-Scholes implied volatility. Volatility is the square root of the variance. In the subsequent estimations, I follow the literature and estimate parameters for variance, i.e., $X(t) \equiv V^2$, where V is volatility.

Data used in this study include weekly (taken on Wednesday) FX call option implied volatilities for two currency pairs - U.S. dollar-British pound (USD-GBP) and Japanese yen-U.S. dollar (YEN-USD). The options used in the estimation all have one month until maturity.

The option prices are quoted for a fixed option delta, not for a fixed strike price. The option delta is (roughly) the probability that the option will expire in-the-money. For each date the data contain prices, quoted in the form of implied volatilities, for a 50-delta call, for each currency pair. Given the option delta, it is straightforward to calculate the strike price of the option. These 50-delta calls have spot-to-strike price ($\frac{S}{K}$) ratios that range from 1.005 to 0.998 with a mean of approximately 1.002. The data cover January 24, 1996 through January 28, 2004, a total of 415 weekly observations for each currency pair.

Summary statistics for the USD-GBP and YEN-USD volatility are found in Table 2.1, while time series plots are presented in Figure 1 and 2. Notice that YEN-USD volatility is nearly 40% higher than USD-GBP volatility. Further, YEN-USD volatility is more than twice as volatile as USD-GBP volatility (3.12% vs. 1.47%). As is detailed below, the key to modeling YEN-USD volatility is the specification of $\sigma[X_t; \Theta]$.

2.4.2 Estimation Results

For all the models and for both currency pairs, the coefficients β_0 and α_3 (see equations (2.3) and (2.4)) were found to be very near zero and to have little or no effect on the log likelihood value. Coefficients β_0 and α_3 were set to zero for all models.

Consider first the estimation results for the USD-GBP variance in Table 2.2. The table presents parameter estimates, standard errors (in parentheses), log likelihood values, and $-\frac{n}{2}$ times the Akaike Information Criteria (AIC). The models can be ranked using the value of $-\frac{n}{2}$AIC, which penalizes a more general model for the loss of degrees of freedom that results from the increased number of parameters. Notice that the estimation algorithm could not be made to converge for the SEV-LD model, hence results are not reported for that model. According to the AIC criteria, the CEV-ND model is slightly better than the SEV-ND model (2270.54 vs. 2269.75), though the difference is small. Evidently, adding a square root term to the standard CEV diffusion function does not improve the fit of the model.

In order to compare nested models, I employ the log likelihood ratio test statistic, $L^* \equiv -2.0\Big(L[\Theta_R] - L[\Theta_U]\Big)$, where $L[\Theta_R]$ is the log likelihood of the restricted model and $L[\Theta_U]$ is the log likelihood of the unrestricted model. The test statistic is distributed $\chi^2[df]$, where df is the number of exclusion restrictions. Table 2.4 presents the results of the log likelihood ratio test for the USD-GBP and YEN-USD currency pairs. The table shows that, for the USD-GBP currency pair, the CEV-ND cannot be rejected against the SEV-ND alternative. The test statistic is 1.78, while the 95% criterion value is 3.84. The CEV-CD and CEV-LD models can be rejected at the 1% significance level.

Table 2.3 shows that, unlike the USD-GBP case, for the case of YEN-USD variance the CEV-LD model is ranked first by the AIC criteria. However, the log likelihood values are very close for all but the AFF model. Log likelihood ratio test statistics reported in Table 2.4, which shows that the CEV-LD model cannot be rejected at any reasonable significance level when compared to the CEV-ND, SEV-LD, or SEV-ND models. Notice that, similar to the USD-GBP currency pair, generalizing the CEV diffusion function does not improve the fit of the model. However, whereas the USD-GBP currency pair benefited from adding a nonlinear term to the drift function, such a generalization does not improve the fit of the model for the YEN-USD currency pair.

Notice further that all the models have a harder time fitting the YEN-USD data than the USD-GBP data. (The log likelihood values are consistently lower for the YEN-USD case than for the USD-GBP case.) Given the summary statistics in Table 2.1, this result is not surprising. YEN-USD volatility is higher and more volatile than its USD-GBP counterpart. Thus, it is unsurprising that they require different model specifications to capture volatility dynamics.

Drift Function $\mu[X_t; \Theta]$ Estimates

Table 2.2 demonstrates that the drift function, $\mu[x_t; \Theta]$, for USD-GBP FX rate variance series is strongly mean reverting, with $\alpha_2 = -935.04$ (standard error $= 366.75$) for the preferred CEV-ND model. The large negative coefficient on the square of the variance, α_2, has the effect of quickly reducing the value of the variance when it becomes large and thereby forcing the variance back down to lower levels. Notice that, in moving from the CEV-ND model to the SEV-ND model, the estimates of the drift coefficients change very little.

The fact that the constant term in the drift function, α_0, is estimated to be negative means that, for volatility (i.e., $\sqrt{\text{variance}}$) values less than approximately 2.7%, the drift term for the USD-GBP variance can be *negative*. While this result would seem to indicate the counterfactual that it is possible for the variance to become less than zero, the sample data has no observations less than 3.61% for the USD-GBP volatility. Thus, the estimated drift function is likely not reliable in this low volatility region.

Nonlinearities in the drift function improve model fit for USD-GBP volatility. From Table 2.4, the CEV-CD and CEV-LD models are rejected at the 1% significance level. However, the specification of the drift function is not as important for YEN-USD volatility. Table 2.4 shows that the CEV-CD and CEV-LD models cannot be rejected versus the CEV-ND alternative.

Table 2.3 also shows that, for the preferred CEV-LD model of YEN-USD variance, the coefficient on the linear drift term, $\alpha_1 = -2.23$, is negative, thereby ensuring that the series is stationary. Further, unlike the USD-GBP case, the drift function for the YEN-USD variance does not become negative for small values. However, precise estimation of the drift function is difficult. The standard error for α_1 reported in Table 2.3 results in $t-$statistic of approximately 1.70 for the YEN-USD variance series. Table 2.3 shows that estimation of the drift function parameters is difficult for all the models. As discussed below, part of the reason for this difficulty lies in the time series pattern of YEN-USD volatility in the sample data. The volatility trended upward in the first half of the sample data, then dropped to a lower overall level and stayed relatively flat in the second half of the sample data. (See Figure 2.)

Diffusion Function $\sigma[X_t; \Theta]$ Estimates

For the USD-GBP ATM implied variance the estimate of β_3 is less than one for all the model with the CEV diffusion function. However, sub-period estimation (see the following subsection) reveals that this result no longer obtains for the most recent data.

Comparison of individual model specifications demonstrates some important points. Moving from the AFF model to the CEV-LD model increases the $-\frac{n}{2}$AIC value substantially, from 2263 to 2268. The square root specification for the diffusion function is inadequate to capture volatility dynamics. However, comparison of the CEV-CD and CEV-ND models with the corresponding SEV models shows clearly that moving beyond the constant elasticity of variance specification for the diffusion function does not improve model performance, regardless of the drift specification.

The diffusion functions for the YEN-USD ATM implied variance series is similar to the values found by Bakshi et al. (2005) for equity variance (see their Table 2).[2] Tables 2.2 and 2.3 make clear that the additional flexibility of the diffusion function for the SEV class of models as compared to the CEV class of models does not improve the fit. CEV-LD (CEV-ND) cannot be rejected versus the SEV-LD (SEV-ND) alternative. Indeed, when compared to the SEV-ND model, neither the CEV-CD nor the CEV-LD models can be rejected. As with the USD-GBP currency pair, generalizing the CEV diffusion function does not improve the fit of the model.

[2]While the drift function estimated in this paper is also similar to the drift function for equity volatility found by Bakshi et al., notice that,owing to a much larger data set, those authors are able to estimate α_0 with much greater precision.

For the YEN-USD currency pair, moving from the AFF model to the CEV-LD model results in a *dramatic* increase the $-\frac{n}{2}$AIC value, from 1850 to 1924. The square root diffusion function is not appropriate for modeling FX volatility. Notice that, moving from the CEV-LD (CEV-ND) model to the SEV-LD (SEV-ND) model causes the common parameter estimates to change very little, while the new parameter introduced in the SEV diffusion function statistically indistinguishable from zero.

The fact that all the models except AFF have similar log likelihood values serves to emphasize that the key to modeling YEN-USD volatility is the specification of the diffusion function $\sigma[X; |theta]$. Option pricing models and risk management strategies that rely on a square root diffusion function are misspecified.

Robustness Tests

In order to check the robustness of these findings, the data were split into two subperiods, corresponding to the time (1) before and (2) after 2000. While this choice is admittedly ad hoc, it serves to divide the data roughly in half. Subperiod (1) covers January 24, 1996 through the end of 1999 (NOBS=206), while subperiod (2) covers the beginning of 2000 through the end of the data, January 28, 2004 (NOBS=209).

AIC rankings and log likelihood tests for both subperiods (not reported) indicate that, for USD-GBP volatility, CEV-ND is still the preferred model, though there are interesting differences in the estimation results across the two subperiods.

In particular, the elasticity coefficient, β_3, increases dramatically in subperiod two. While $\beta_3 = 0.915$ for the entire sample and $\beta_3 = 0.831$ in subperiod (1), for subperiod (2) $\beta_3 = 1.220$. This result is consistent with the estimate for the YEN-USD currency pair and for the Bakshi et al. (2005) equity volatility result. The phenomena reported above whereby the elasticity coefficient β_3 is less than one for the ATM implied USD-GBP variance is no longer observed in the most recent data.

As with the USD-GBP currency pair, for the purposes of robustness tests the YEN-USD data were split into pre-2000, subperiod (1), and post-2000, subperiod (2). YEN-USD FX rate volatility is much higher and more volatile in subperiod (1) than in subperiod (2). YEN-USD volatility reaches values of 28.5% on October 7, 1998 and 22.3% on September 22, 1999. After January 1, 2000, the maximum observed value is only 14.50%, on January 17, 2001. Further, the standard deviation of ATM implied YEN-USD volatility is 3.81% in subperiod (1), but only 1.50% in subperiod (2), more in-line with USD-GBP volatility. Evidently, the level and the volatility of YEN-USD FX rate volatility have decreased dramatically since the turn of the century.

Recall that, for the case of YEN-USD variance, the CEV-LD model is ranked first by the AIC criteria. In subperiod (2), the CEV-LD model continues to fit the data better than the alternatives, but in subperiod (1) the CEV-CD model cannot be rejected versus the CEV-LD alternative. The coefficient on the linear drift term, α_1, is much higher in subperiod (2), indicating that the speed of mean reversion is faster in the post-2000 time frame. However, the drift function cannot be estimated with precision in the post-2000 time frame. This result is not surprising given the time series pattern in the data. In the post-2000 time period, YEN-USD volatility remained essentially flat. Owing to the less volatile behavior of YEN-USD volatility in subperiod (2), all of the models do a better job of fitting the data, despite the fact that the drift function cannot be estimated with precision.

Comparison of FX Volatility and Equity Volatility

Bakshi, Ju, and Ou-Yang (2005) report results for the CEV and SEV model estimations for equity volatility. They use market index volatility, as proxied by VIX. While this article reports that the CEV diffusion function is sufficient for FX volatility, they report that, for market index volatility, the SEV diffusion function fits best.

First, note that equity volatility is consistently higher than its FX counterpart for both currency pairs. The sample mean volatility value for the USD-GBP currency pair is 8.21%. The sample mean VIX volatility value in the Bakshi, Ju, and Ou-Yang (2005) data is 18.91%. Further, VIX volatility is itself more volatile than FX volatility. The sample standard deviation for the spanned USD-GBP volatility is 1.47%, while that for the VIX data is 5.85%. Equity volatility is higher and more volatile than FX rate volatility. Not surprisingly, equity volatility requires a more general diffusion function in order to fit the data.

The similarity of the estimates for β_3 for (i) VIX volatility, reported in Table 2 of Bakshi Ju, and Ou-Yang (2005), (ii) the volatility of the YEN-USD currency pair reported in Table 2.3 herein, and (iii) the volatility of the USD-GBP currency pair beginning in 2000 is striking. For the CEV-LD model (the highest ranked model for (ii)) Bakshi et al report $\beta_2 = 4.70$ and $\beta_3 = 1.27$, while the corresponding estimates in my context are $\beta_2 = 4.41$ and $\beta_3 = 1.26$. While I estimate $\beta_3 = 0.915$ for the volatility of the USD-GBP currency pair for the full sample, subperiod estimation reveals that $\beta_3 = 1.22$ since the turn of the century.

2.5 Conclusions

This article estimates univariate continuous-time diffusions for the case of FX volatility for three currency pairs, U.S. dollar-British pound, Japanese yen-U.S. dollar, and Japanese yen-British pound. Volatility is proxied by simple ATM implied volatility from FX options. All the options used in this article have one month until maturity.

A nonlinear term in the risk neutral drift function improved the fit of the model only for the USD-GBP currency pair. This is in contrast to the interest rate result of Aït-Sahalia (1996) and the equity result of Bakshi, Ju, and Ou-Yang (2005). For all cases, the estimated parameters of the drift function ensure that the volatility mean reverts and thus that the diffusion is stationary.

In none of the estimations reported in this paper does allowing for a more general diffusion function improve upon the constant elasticity of variance model.

One extension of this article involves studying the higher-order implications of triangular-arbitrage restriction on cross FX rates. The no-triangular-arbitrage relationship requires that $S_t^{YEN,GBP} = S_t^{YEN,USD} \times S_t^{USD,GBP}$, where $S^{f,h}$ is the spot

FX rate in terms of currency f per unit of currency h. Investigating the restrictions that this relationship implies for covariances may be worth pursuing since volatility can be recovered in a model-free manner (Barndorff-Nielson and Shephard (2004)).

Table 2.1: Basic Features of the Volatility Series

	USD-GBP Volatility	**YEN-USD** Volatility
NOBS	415	415
Mean	8.21%	11.46%
Stdev	1.47%	3.12%
Skewness	-0.11	1.37
Kurtosis	3.03	5.87
Autocorrelation: lag 1	0.9045	0.9167
Autocorrelation: lag 2	0.8200	0.8694
Autocorrelation: lag 3	0.7445	0.8343
Autocorrelation: lag 4	0.6796	0.8071
Autocorrelation: lag 5	0.6323	0.7726
Autocorrelation: lag 6	0.5910	0.7405

The table contains the sample mean, standard deviation, skewness, kurtosis, and the first six autocorrelations for USD-GBP and YEN-USD currency return volatility. Volatility is proxied by at-the-money implied Black-Scholes volatility. The data are sampled at a weekly frequency for the period January 24, 1996 through January 28, 2004.

Table 2.2: Estimation Results for US Dollar-British Pound FX Rate

model	L	$-\frac{n}{2}$ AIC	α_0	α_1	α_2	β_1	β_2	β_3
AFF	2266.00	2263.00	0.0330	-4.6953		0.0090		
			(0.0075)	(1.1532)		(0.0000)		
CEV-CD	2267.64	2264.64	0.0065				0.7982	0.9300
			(0.0023)				(0.0433)	(0.0902)
CEV-LD	2272.18	2268.18	0.0241	-3.3189			0.4910	0.8311
			(0.0068)	(1.1382)			(0.0255)	(0.0936)
CEV-ND	2275.54	2270.54	-0.0054	8.1108	-935.037		0.7547	0.9149
			(0.0125)	(4.4996)	(366.754)		(0.0422)	(0.0946)
SEV-CD	2267.64	2263.64	0.0065			0.0000	0.6372	1.8599
			(0.0023)			(0.0000)	(0.0691)	(0.1804)
SEV-LD								
SEV-ND	2275.75	2269.75	-0.0048	7.8463	-911.425	0.0040	22.8115	2.7033
			(0.0136)	(4.8213)	(390.871)	(0.0035)	(9.0927)	(1.4611)

The SEV-ND model, which nests the SEV-CD and SEV-LD models, is $dX_t = (\alpha_0 + \alpha_1 X_t + \alpha_2 X_t^2 + \alpha_3 X_t^{-1})dt + \sqrt{\beta_0 + \beta_1 X_t + \beta_2 X_t^{\beta_3}} dW_t$. The CEV-ND model, which nests the CEV-CD and CEV-LD models, is $dX_t = (\alpha_0 + \alpha_1 X_t + \alpha_2 X_t^2 + \alpha_3 X_t^{-1})dt + \beta_2 X_t^{\beta_3} dW_t$. The AFF model is $dX_t = (\alpha_0 + \alpha_1 X_t)dt + \sqrt{\beta_0 + \beta_1 X_t} dW_t$. In all cases, β_0 and α_3 were found to be zero and were thus eliminated from the estimation. The series is Black-Scholes implied variance for 1 month at-the-money options. The data are sampled at a weekly frequency for the period January 24, 1996 through January 28, 2004.

Table 2.3: Estimation Results for Japanese Yen-US Dollar FX Rate

model	L	$-\frac{n}{2}$ AIC	α_0	α_1	α_2	β_1	β_2	β_3
AFF	1853.11	1850.11	0.0791	-5.6115		0.0359		
			(0.0158)	(1.2540)		(0.0000)		
CEV-CD	1927.30	1924.30	0.0173				5.3298	1.3055
			(0.0045)				(0.2023)	(0.0691)
CEV-LD	1928.78	1924.78	0.0360	-2.2310			4.4161	1.2620
			(0.0120)	(1.2993)			(0.1743)	(0.0712)
CEV-ND	1929.16	1924.16	0.0227	0.3661	-101.771		4.6721	1.2747
			(0.0192)	(3.2665)	(118.345)		(0.1833)	(0.0720)
SEV-CD	1927.38	1923.38	0.0172			0.0012	40.2130	2.7060
			(0.0045)			(0.0031)	(9.9416)	(0.2854)
SEV-LD	1928.89	1923.89	0.0367	-2.2935		0.0016	29.5270	2.6392
			(0.0124)	(1.3241)		(0.0034)	(7.7708)	(0.2888)
SEV-ND	1929.19	1923.19	0.0239	0.1632	-95.121	0.0008	28.8387	2.6071
			(0.0202)	(3.4203)	(123.256)	(0.0036)	(7.3183)	(0.2933)

The SEV-ND model, which nests the SEV-CD and SEV-LD models, is $dX_t = (\alpha_0 + \alpha_1 X_t + \alpha_2 X_t^2 + \alpha_3 X_t^{-1})dt + \sqrt{\beta_0 + \beta_1 X_t + \beta_2 X_t^{\beta_3}} dW_t$. The CEV-ND model, which nests the CEV-CD and CEV-LD models, is $dX_t = (\alpha_0 + \alpha_1 X_t + \alpha_2 X_t^2 + \alpha_3 X_t^{-1})dt + \beta_2 X_t^{\beta_3} dW_t$. The AFF model is $dX_t = (\alpha_0 + \alpha_1 X_t)dt + \sqrt{\beta_0 + \beta_1 X_t} dW_t$. In all cases, β_0 and α_3 were found to be zero and were thus eliminated from the estimation. The series is Black-Scholes implied variance for 1 month at-the-money options. The data are sampled at a weekly frequency for the period January 24, 1996 through January 28, 2004.

Table 2.4: Pairwise Log Likelihood Ratio Tests

R model	U model	*df*	USD-GBP	YEN-USD
AFF	CEV-LD	1	12.361 (0.000)	151.342 (0.000)
CEV-CD	CEV-ND	2	16.719 (0.000)	3.711 (0.156)
CEV-LD	CEV-ND	1	11.803 (0.001)	0.906 (0.686)
CEV-LD	SEV-ND	1	7.143 (0.028)	0.816 (0.665)
CEV-CD	SEV-CD	1	0.000 (0.999)	0.151 (0.697)
CEV-LD	SEV-LD	1	NA	0.320 (0.572)
CEV-ND	SEV-ND	1	1.782 (0.182)	0.061 (0.802)

The table presents the results of the log likelihood ratio test statistic for comparing nested models and the significance level at which the restricted model can be rejected. The test statistic is $-2\Big(L[\hat{\Theta}_R] - L[\hat{\Theta}_U]\Big)$, where $\hat{\Theta}_R$ is the estimated parameter vector for the restricted model and $\hat{\Theta}_U$ is the estimated parameter vector for the unrestricted model. The test statistic is distributed $\chi^2[df]$, where *df* is the number of exclusion restrictions.

Figure 2.1: Dollar-Pound (USD-GBP) Volatility

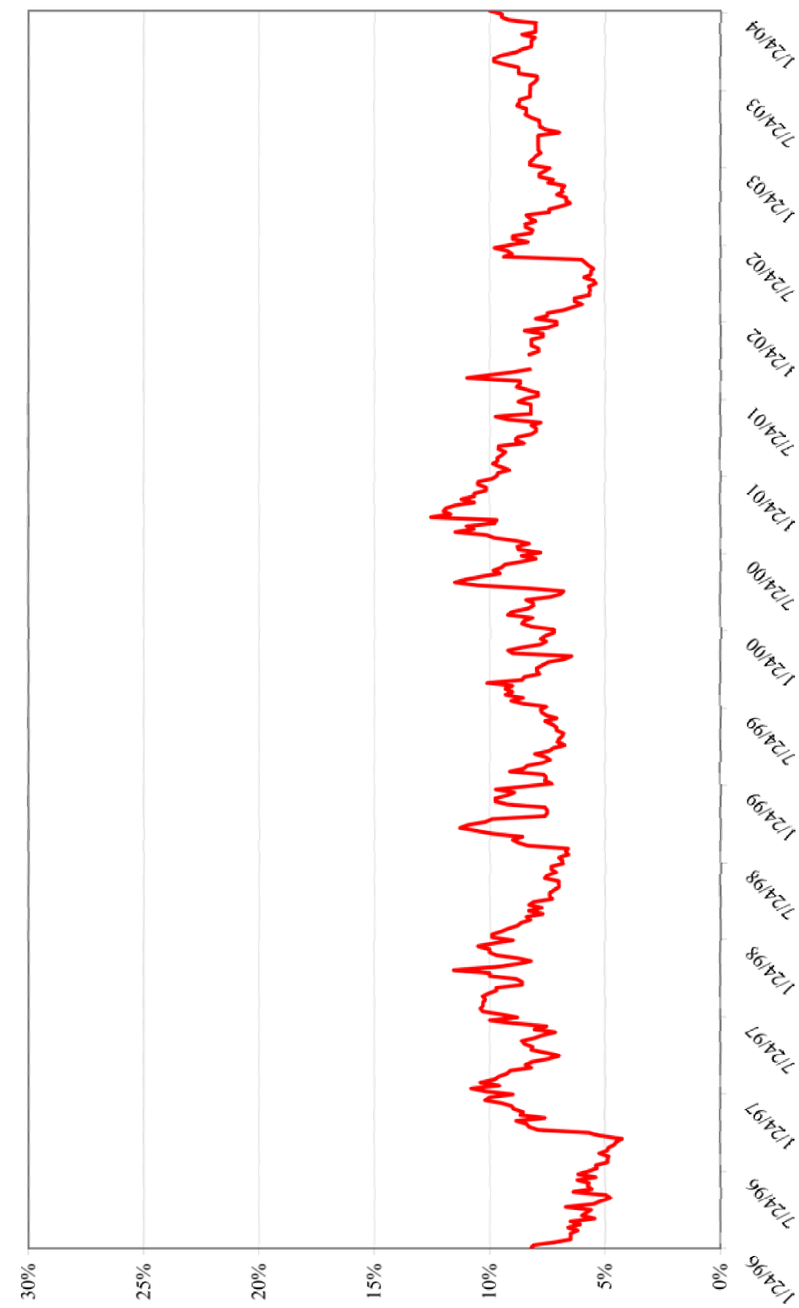

Figure 2.2: Yen-Dollar (Yen-USD) Volatility

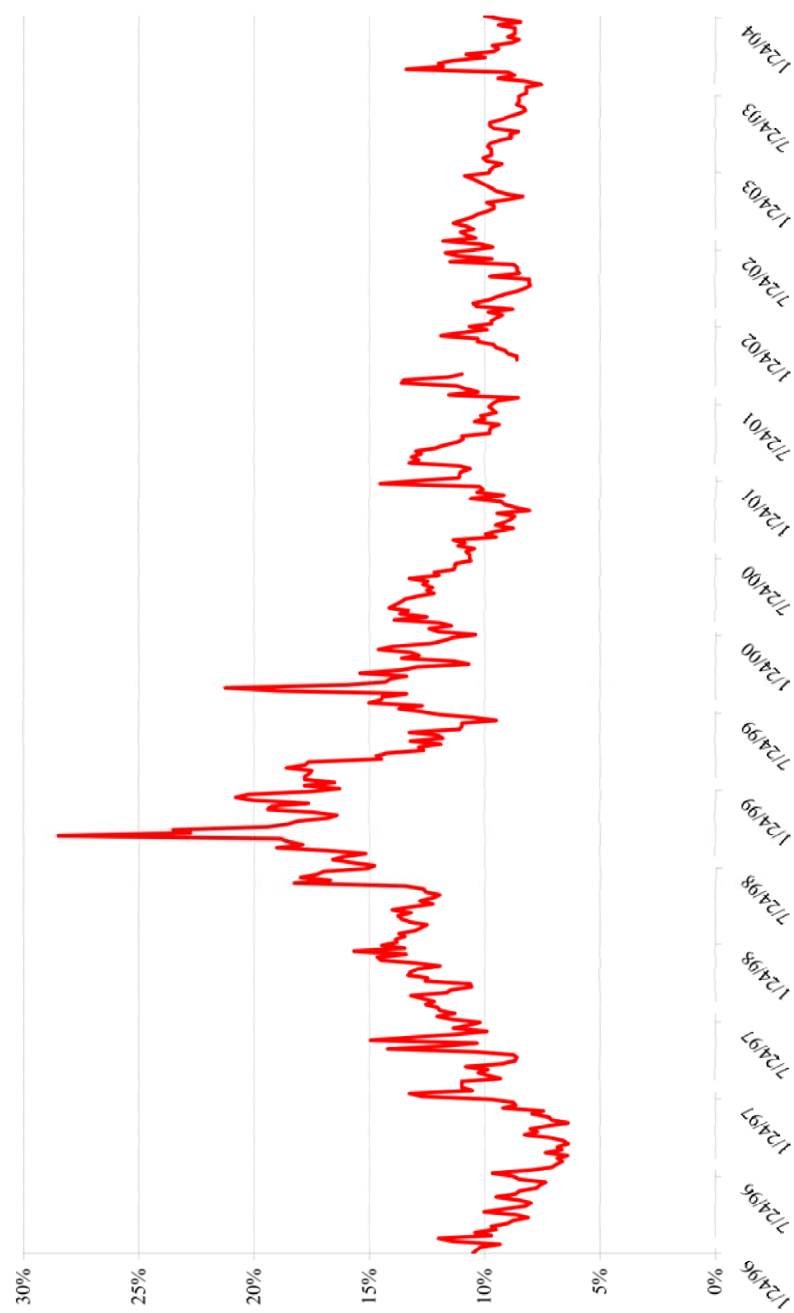

Appendix A

Derivations

The total cost function for producer i (equation (1.9) in the text) is given by

$$TC_i = F + \frac{a}{c}\left(\frac{1}{H_{Pi} - Q_{Pi}}\right)^c, \tag{A.1}$$

where H_{Pi} = capacity (or, available supply) of producer i, and Q_{Pi} = output of producer i.

The ex-post profit of producer i (equation (1.10)) is

$$\pi_{Pi} = P_F\, Q^F_{Pi} + P_W\, Q^W_{Pi} - F - \frac{a}{c}\left(\frac{1}{H_{Pi} - Q^W_{Pi} - Q^F_{Pi}}\right)^c, \tag{A.2}$$

where P_F is the forward price, P_W is the wholesale spot price, Q^F_{Pi} is the forward sale for producer i, and Q^W_{Pi} is the spot sale for producer i. Notice that the output for producer i, Q_{Pi}, has been replaced by the sum of producer i's spot and forward sales, $Q^W_{Pi} + Q^F_{Pi}$.

At time $t = 1$, Q^F_{Pi} is known. Differentiating equation (A.2) with respect to spot sales yields

$$\frac{\partial \pi_{Pi}}{\partial Q^W_{Pi}} = P_W - a(H_{pi} - Q^W_{Pi} - Q^F_{Pi})^{(-c-1)}. \tag{A.3}$$

Setting this expression equal to zero and solving for the optimal spot sales results in equation (1.11),

$$Q^W_{Pi} = H_{Pi} - Q^F_{Pi} - \left(\frac{a}{P_W}\right)^{\left(\frac{1}{c+1}\right)}. \tag{A.4}$$

At time $t = 1$, supply and demand must balance,

$$\sum_{j=1}^{N_R} Q_{Rj} = \sum_{i=1}^{N_P} Q_{Pi}. \tag{A.5}$$

Note that the left hand side is the total system demand, Q^D, and on the right hand side replace $Q_{Pi} = Q^W_{Pi} + Q^F_{Pi}$ to get

$$Q^D = \sum_{I=1}^{N_P} \left(Q^W_{Pi} + Q^F_{Pi} \right). \tag{A.6}$$

Now substitute equation (A.4) into the right hand side (A.6) to get

$$Q^D = \sum_{I=1}^{N_P} \left[H_{Pi} - \left(\frac{a}{P_W} \right)^{\frac{1}{c+1}} \right] \tag{A.7}$$

$$= H^S - N_P \left(\frac{a}{P_W} \right)^{\frac{1}{c+1}}. \tag{A.8}$$

Solving for the spot price then leads to equation (1.13):

$$P_W = a \left(\frac{N_P}{H^S - Q^D} \right)^{c+1}. \tag{A.9}$$

Finally substituting equation (A.9) into equation (A.4) results in equation (1.14)

$$Q^W_{Pi} = H_{Pi} - Q^F_{Pi} - \frac{H^S - Q^D}{N_P}. \tag{A.10}$$

Producer i is assumed to be a mean-variance profit optimizer with utility function given by

$$U_{Pi}[\pi_{Pi}] = E(\pi_{Pi}) - \frac{A}{2} \mathrm{Var}(\pi_{Pi}), \tag{A.11}$$

where $E(\pi_{Pi})$ is the time $t = 0$ expectation of time $t = 1$ profits.

Substituting equation (A.4) into equation (A.2) gives

$$\pi_{Pi} = P_F Q^F_{Pi} + P_W \left(H_{Pi} - Q^F_{Pi} \right) - F - a^{\frac{1}{c+1}} \left(1 + \frac{1}{c} \right) P_W^{\frac{c}{c+1}}. \tag{A.12}$$

Equation (A.11) can then be rewritten as

$$\begin{aligned} U_{Pi}[\pi_{Pi}] &= P_F Q^F_{Pi} + \left(H_{Pi} - Q^F_{Pi}\right) E(P_W) - F - a^{\left(\frac{1}{c+1}\right)} \left(1 + \frac{1}{c}\right) E\left(P_W^{\left(\frac{c}{c+1}\right)}\right) \\ & -\frac{A}{2}\left(H_{Pi} - Q^F_{Pi}\right)^2 \mathrm{Var}(P_W) - \frac{A}{2}\left[a^{\left(\frac{1}{c+1}\right)}\left(1+\frac{1}{c}\right)\right]^2 \mathrm{Var}\left(P_W^{\left(\frac{c}{c+1}\right)}\right) \\ & +A\left(H_{Pi} - Q^F_{Pi}\right) a^{\left(\frac{1}{c+1}\right)}\left(1+\frac{1}{c}\right) \mathrm{cov}\left(P_W, P_W^{\left(\frac{c}{c+1}\right)}\right). \end{aligned} \tag{A.13}$$

Producer i's time $t = 0$ problem is then to maximize equation (A.13) w.r.t. Q^F_{Pi},

$$\begin{aligned} \frac{\partial U_{Pi}[\pi_{Pi}]}{\partial Q^F_{Pi}} &= P_F - E(P_W) + A\left(H_{Pi} - Q^F_{Pi}\right)\mathrm{Var}(P_W) - \\ & \quad A\, a^{\frac{1}{c+1}}\left(1+\frac{1}{c}\right)\mathrm{cov}\left(P_W, P_W^{\left(\frac{c}{c+1}\right)}\right). \end{aligned} \tag{A.14}$$

Setting this equal to zero and solving for Q^F_{Pi} finally gives,

$$Q^F_{Pi} = \frac{1}{A}\left(\frac{P_F - E(P_W)}{\mathrm{Var}(P_W)}\right) + H_{Pi} - \frac{1}{a^x(x+1)}\frac{\mathrm{cov}\left(P_W^{x+1}, P_W\right)}{\mathrm{Var}(P_S)}, \tag{A.15}$$

where $x \equiv \frac{-1}{c+1}$.

The expression for the retailer j's optimal forward position (text equation (1.7)) is unchanged from the BL case and will not be derived here, merely repeated for convenience.

$$Q^F_{Rj} = \frac{1}{A}\left(\frac{P_F - E(P_W)}{\mathrm{Var}(P_W)}\right) + \frac{P_R\,\mathrm{cov}(Q_{Rj}, P_W)}{\mathrm{Var}(P_W)} + \frac{\mathrm{cov}(Q_{Rj}P_W, P_W)}{\mathrm{Var}(P_W)}, \tag{A.16}$$

where P_R is the (fixed) retail price. Notice that producers and retailers are assumed to have the same coefficient A for the variance term in their utility functions.

Forward contracts are in zero net supply. Imposing market clearing,

$$\sum_{j=1}^{N_R} Q^F_{Rj} + \sum_{i=1}^{N_P} Q^F_{Pi} = 0. \tag{A.17}$$

Define $N \equiv \frac{N_R + N_P}{A}$ and substitute equations (A.15) and (A.16) into equation (A.17) to get

$$
\begin{aligned}
0 &= N\left[P_F - E(P_W)\right] + H^S \mathrm{Var}(P_W) - \frac{N_P}{a_1^x (x+1)} \mathrm{cov}(P_W, P_W^{x+1}) \\
&\quad P_R \, \mathrm{cov}(Q^D, P_W) - \mathrm{cov}(P_W Q^D, P_W).
\end{aligned} \tag{A.18}
$$

Consider the last two terms in equation (A.18), i.e., the covariance terms containing the total system load, Q^D. Using the expression for Q^D from equation (A.8), we have,

$$
\mathrm{cov}(Q^D, P_W) = -\frac{N_P}{a_1^x} \mathrm{cov}(P_W^x, P_W), \tag{A.19}
$$

$$
\mathrm{cov}(P_W Q^D, P_W) = H^S \mathrm{Var}(P_W) - \frac{N_P}{a_1^x} \mathrm{cov}(P_W^{x+1}, P_W). \tag{A.20}
$$

Substituting these expressions in equation (A.18) finally yields the expression for the equilibrium forward premium:

$$
P_F - E(P_W) = \frac{N_P}{N a^x} \left[P_R \, \mathrm{cov}(P_W^x, P_W) - \left(\frac{x}{x+1} \right) \mathrm{cov}(P_W^{x+1}, P_W) \right]. \tag{A.21}
$$

As in BL, expanding P_W^x around the expectation $E(P_W)$, the covariance term $\mathrm{cov}(P_W^x, P_W)$ can be approximated to second order by

$$
\mathrm{cov}(P_W^x, P_W) \approx x \left[E(P_W)\right]^{(x-1)} \mathrm{Var}(P_W) + \frac{1}{2} x(x-1) \left[E(P_W)\right]^{(x-2)} \mathrm{Skew}(P_W) \tag{A.22}
$$

Applying this approximation for $\mathrm{cov}(P_W^x, P_W)$ and $\mathrm{cov}(P_W^{x+1}, P_W)$ implies equations (1.18), (1.19), and (1.20). □

Appendix B

Discrete Model

Consider again the expression for the equilibrium forward premium, equation (1.17), reproduced here for convienience.

$$P_F - E(P_W) \;=\; \frac{N_P}{Na^x}\left[P_R \operatorname{cov}(P_W^x, P_W) - \left(\frac{x}{x+1}\right)\operatorname{cov}(P_W^{x+1}, P_W)\right]. \quad \text{(B.1)}$$

This expression involves two covariances, $\operatorname{cov}(P_W^x, P_W)$, and $\operatorname{cov}(P_W^{x+1}, P_W)$. Focus on the former.

$$\begin{aligned} \operatorname{cov}(P_W^x, P_W) &\equiv E\Big((P_W^x - \mu_{P_W^x})(P_W - \mu_{P_W})\Big), && \text{(B.2)} \\ &= E\Big(P_W^x P_W\Big) - \mu_{P_W^x}\mu_{P_W}, && \text{(B.3)} \\ &= \sum_{k=1}^{3} \phi_k \psi_k^{x+1} - \mu_{P_W^x}\mu_{P_W}, && \text{(B.4)} \end{aligned}$$

where $\mu_{P_W^x} \equiv E(P_W^x)$. Similarly,

$$\operatorname{cov}(P_W^{x+1}, P_W) \;=\; \sum_{k=1}^{3} \phi_k \psi_k^{x+2} - \mu_{P_W^{x+1}}\mu_{P_W}, \quad \text{(B.5)}$$

where $\mu_{P_W^{x+1}} \equiv E(P_W^{x+1})$.

Substitution of equations (B.4) and (B.5) into equation (1.17) results in equation (1.30) in the text. □

BIBLIOGRAPHY

[1] Aït-Sahalia, Y., 1996, Testing continuous-time models of the spot interest rate, *Review of Financial Studies* 9: 385-426

[2] Aït-Sahalia, Y., 1999, Transition densities for interest rates and other nonlinear diffusions, *Journal of Finance* 54: 1361-1395.

[3] Aït-Sahalia, Y., 2002, Maximum likelihood estimation of discretely sampled diffusions: A closed-form approximation approach, *Econometrica* 70: 223-262.

[4] Andersen, T., T. Bollerslev, F. Diebold, H. Ebens, 2001. The distribution of realized stock return volatility. *Journal of Financial Economics* 61, 43-76.

[5] Andersen, T., T. Bollerslev, F. Diebold, P. Labys, 2003. Modeling and forecasting realized volatility. *Econometrica* 71, No. 2, 579-625.

[6] Andersen, T., T. Bollerslev, F. Diebold, P. Labys, 2001. The distribution of realized exchange rate volatility. *Journal of the American Statistical Association* 96, No. 453, 42-55.

[7] Anderson, Ronald, and Jean-Pierre Danthine, 1980, Hedging and joint production: Theory and illustrations, *Journal of Finance* 35: 487-501.

[8] Bakshi, G., P. Carr, and L. Wu, 2005, Stochastic risk premiums, stochastic skewness in currency options, and stochastic discount factors in international economies, Working paper (University of Maryland, College Park, MD).

[9] Bakshi, G., N. Ju, 2005, A refinement to Aït-Sahalia's (2002) maximum likelihood estimation of discretely sampled diffusions, A closed form approximation approach, *Journal of Business* 78, No. 5.

[10] Bakshi, G., N. Ju, and H. Ou-Yang, 2005, Estimation of continuous-time models with an application to equity volatility, *Journal of Financial Economics* (forthcoming).

[11] Bakshi, G., N. Kapadia, and D. Madan, 2003, Stock return characteristics, skew laws, and the differential pricing of individual equity options, *Review of Financial Studies* 16, 101-143.

[12] Bakshi, G., and D. Madan, 2000, Spanning and derivative security valuation, *Journal of Financial Economics* 55: 205-238.

[13] Banerjee, Suman, and Tom Noe, 2002, Exotics and electrocns: Electric power crises and financial risk management, Working paper, Tulane University.

[14] Barndorff-Nielsen, O., and N. Shephard, 2004, Econometric analysis of realized covariation: high frequency based covariance, regression, and correlation in financial economics, *Econometrica* 72: 885-925.

[15] Bates, D., 1996, Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche mark options, *Review of Financial Studies* 9: 69-107.

[16] Bates, D., 2004, Maximum likelihood estimation of latent affine processes, *Review of Financial Studies* (forthcoming).

[17] Bessembinder, Hendrik, and Michael L. Lemmon, 2002, Equilibrium pricing and optimal hedging in electricity forward markets, *Journal of Finance* 57: 1347-1382.

[18] Bliss, R., N. Panigirtzoglou, 2004. Option-implied risk aversion estimates. *Journal of Finance* 59 (1), 407-446.

[19] Bollerslev, T., and H. Zhou, 2002, Estimating stochastic volatility diffusion using conditional moments of integrated volatility, *Journal of Econometrics* 109: 33-65.

[20] Brandt, M., and P. Santa-Clara, 2002, Simulated likelihood estimation of diffusions with an application to exchange rate dynamics in incomplete markets, *Journal of Financial Economics* 63: 161-210.

[21] Britten-Jones, M., A. Neuberger, 2000. Option prices, implied price processes, and stochastic volatility. *Journal of Finance* 55, 839-866.

[22] Carr, P., and D. Madan, 2001, Optimal positioning in derivative securities, *Quantitative Finance* 1: 19-37.

[23] Carr, P., L. Wu, 2004. Variance risk premia. *Mimeo*, Baruch College.

[24] Chan, K.C., G.A. Karolyi, F.A. Longstaff, and A.B. Sanders, 1992, An empirical investigation of alternative models of the short interest rate, *Journal of Finance* 47: 1209-1228.

[25] Deng, Shi-Jie, 2000, Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes, Working paper, Georgia Institute of Technology.

[26] Deng, Shi-Jie, Blake Johnson, and Aram Sogomonian, 2001, Exotic electricity options and the valuation of electricity generation and transmission assets, *Decision Support Systems* 30, 383-392.

[27] Deng, Shi-Jie, and S.S. Oren, 2003, Incorporating operational characteristics and startup costs in option-based valuation of power generation capacity, *Probability in the Engineering and Information Sciences* 17, 155-181.

[28] Durham, G., 2003, Likelihood-based specification analysis of continuous-time models of the short-term interest rate, *Journal of Financial Economics* 70: 463-487.

[29] Durham, G., and R. Gallant, 2002, Numerical techniques for maximum likelihood estimation of continuous-tiem diffusion processes, *Journal of Business and Economic Statistics* 20: 297-316.

[30] Engle, R., 2004, Risk and volatility: econometric models and financial practice. *American Economic Review* 94, No. 3, 405-420.

[31] Eraker, B., 2001, MCMC analysis of diffusion models with application to finance, *Journal of Business and Economic Statistics* 19: 177-191.

[32] Eydeland, Alexander, and Krzysztof Wolyniec, 2003, Energy and Power Risk Management: New Developments in Modeling, Pricing, and Hedging, John Wiley and Sons, Hoboken, New Jersey.

[33] Eydeland, Alexander, and Helyette Geman, 1999, Pricing power derivatives, *Risk*, 71-74.

[34] Gallant, A.R., and G. Tauchen, 1996, Which moments to match?, *Econometric Theory* 12: 657-681.

[35] Garman, M., and S. Kohlhagan, 1983, Foreign currency option values, *Journal of International Money and Finance* 2: 231-237.

[36] Geman, Helyette, and Andrea Roncoroni, 2004, Understanding the fine structure of electricity prices, *Journal of Business* (forthcoming).

[37] Gouriéroux, C., A. Monfort, and E. Renault, 1993, Indirect inference, *Journal of Applied Econometrics* 8: s85-s118.

[38] Hansen, L.P., and J.A. Scheinkman, 1995, Back to the future: Generating moment implications for continuous time Markov processes, *Econometrica* 63: 767-804.

[39] Harvey, A.C., and N. Shephard, 1994, Estimation of an asymmetric stochastic volatility model for asset returns, *Journal of Business and Economic Statstics* 14: 429-434.

[40] Heston, S., 1993, A closed form solution for options with stochastic volatility with applications to bonds and currency options, *Review of Financial Studies* 6 (2): 327-343.

[41] Hirshleifer, David, and Avanidhar Subramanyam, 1993, Futures versus share contracting as a means of diversifying output risk, *The Economic Journal* 103: 620-638.

[42] Jacquier, E., N.G. Polson, and P.E. Rossi, 1994, Bayesian analysis of stochastic volatility models, *Journal of Business and Economic Statstics* 12: 371-389.

[43] Jiang, G., Y. Tian, 2005, The model-free implied volatility and its information content. *Review of Financial Studies* 18, 1305-1342.

[44] Kessler, M., and M. Sorensen, 1999, Estimating equations based on eigenfunctions for a discretely observed diffusion, *Bernoulli* 5: 299-314.

[45] Longstaff, Francis A., and Ashley W. Wang, 2004, Electricity forward prices: A high-frequency empirical analysis, *Journal of Finance* 59: 1877-1900.

[46] Pirrong, Craig, and Martin Jermakyan, 1999, Valuing power and weather derivatives on a mesh using finite difference methods, in *Energy Modelling and the Management of Uncertainty*, (Risk Books, London).

[47] Routledge, Brian,Duane Seppi, and Chester Spatt, 1999, The spark spread: An equilibrium model of cross-commodity price relationships in electricity, Working paper, Carnegie Mellon University.

[48] Routledge, Brian,Duane Seppi, and Chester Spatt, 2000, Euilibrium forward curves for commodities, *Journal of Finance* 55: 1297-1338.

[49] Ruiz, E., 1994, Quasi-maximum likelihood estimation of stochastic volatility models, *Journal of Econometrics* 63: 289-306.