The Price of Power: The Valuation of Power and Weather Derivatives

Craig Pirrong
University of Houston
Houston, TX 77204
713-743-4466
cpirrong@uh.edu
Martin Jermakyan
ElectraPartners.com

December 7, 2005

Abstract. Pricing contingent claims on power presents numerous challenges due to (1) the nonlinearity of power price processes, and (2) time-dependent variations in prices. We propose and implement a model in which the spot price of power is a function of two state variables: demand (load or temperature) and fuel price. In this model, any power derivative price must satisfy a PDE with boundary conditions that reflect capacity limits and the non-linear relation between load and the spot price of power. Moreover, since power is non-storable and demand is not a traded asset, the power derivative price embeds a market price of risk. Using inverse problem techniques and power forward prices from the PJM market, we solve for this market price of risk function. During 1999-2001, the upward bias in the forward price was as large as $50/MWh for some days in July. By 2005, the largest estimated upward bias had fallen to $19/MWh. These large biases are plausibly due to the extreme right skewness of power prices; this induces left skewness in the payoff to short forward positions, and a large risk premium is required to induce traders to sell power forwards. This risk premium suggests that the power market is not fully integrated with the broader financial markets.

1 Introduction

Pricing contingent claims on power presents numerous difficulties. The price process for power is highly non-standard, and is not well captured by price process models commonly employed to price interest rate or equity derivatives. Electricity "spot" prices exhibit extreme non-linearities. The volatility of power prices displays extreme variations over relatively short time periods. Furthermore, power prices exhibit substantial mean reversion and seasonality. No reduced form, low-dimension price process model can readily capture these features. Finally, and perhaps most important, the non-storability of power creates non-hedgeable risks. Thus, preference free pricing in the style of Black-Scholes is not possible for power.

To address these problems, this article presents an equilibrium model to price power contingent claims. This model utilizes an underlying demand variable a fuel price as the state variables. The demand variable can be output (referred to as "load") or temperature. The price of power at the maturity of the contingent claim is related to the state variables through a terminal pricing function. This pricing function establishes the payoff of the contingent claim, and thus provides one of the boundary conditions required to value it. Given a specification of the dynamics of the state variables and the relevant boundary conditions, conventional PDE solution methods can be used to value the contingent claim.

Since the risks associated with the demand state variable are not hedgeable, any valuation depends on the market price of risk associated with this variable. We allow the market price of risk to be a function of load. Given this function, it is relatively straightforward to solve the "direct" problem of valuing power forwards and options. However, since the market price of risk function is not known, it must be inferred from market prices (analogously to determining an implied volatility or volatility surface). We use inverse problem methods to infer this function from observed forward prices. This solution for the market price of risk function can then be used to price any other power contingent claim not used to calibrate the risk price.

We implement this methodology to value power forward prices in the Pennsylvania-New Jersey-Maryland ("PJM") market. The results of this analysis are striking. First, given terminal pricing function derived from either generators' bids into PJM or econometric estimates, we find that the market price of risk for delivery during the summers of 1999-2005 is large, and represents a substantial fraction of the quoted forward price of power. In particular, this risk premium was as large as \$50/MWh for delivery in July 2000 (representing as much as 50 percent of the forward price), and remained as high as \$19/MWh (or nearly 30 percent of the forward price) for delivery in July 2005. Second, this market price of risk function exhibits large seasonalities. The market price of risk peaks in July and August, and is substantially smaller during the remainder of the year.[1]

These results imply that the market price of risk function is key to pricing power derivatives. Demand and cost fundamentals influence forward and option prices, but the market price of risk is quantitatively very important

[1]Indeed, in some years there is downward bias in forward prices for deliveries during shoulder months. Bessembinder-Lemon (2002) present a model in which prices can be upward biased for deliveries in high demand periods and downward biased in low demand periods.

in determining the forward price of power, at least in the current immature state of the wholesale power market. Ignoring this risk premium will have serious effects when attempting to value power contingent claims, including investments in power generation and transmission capacity.

In addition to pricing power derivatives, the approach advanced in this article can be readily extended to price claims with payoffs that depend on power volume (i.e., load sensitive claims) and weather. Indeed, the equilibrium approach provides a natural way of valuing and hedging power price, load, and weather sensitive claims in single unified framework. More traditional approaches to derivative valuation cannot readily do so.

The remainder of this article is organized as follows. Section 2 presents an equilibrium model of power derivatives pricing. We implement this model for the PJM market; an appendix briefly describes the operation of this market. Section 3 presents a method for estimating the seasonally time-varying mean of the demand process required to solve the valuation PDE, and implements it using PJM data. Section 4 analyzes the methods for estimating the terminal pricing functions required to estimate boundary conditions used in solution of the PDE. Section 5 employs inverse methods to solve for the market price of risk function and presents evidence on the size of the market price of risk for PJM. This section also discusses the implications of these findings. Section 6 shows how to integrate valuation of weather and power derivatives. Section 7 summarizes the article.

2 An Equilibrium Pricing Approach

The traditional approach in derivatives pricing is to write down a stochastic process for the price of the asset or commodity underlying the contingent claim. This approach poses difficulties in the power market because of the extreme non-linearities and seasonalities in the price of power. These features make it impractical to write down a "reduced form" power price process that is tractible and which captures the salient features of power price dynamics.

Figure 1 depicts hourly power prices for the PJM market for 2001-2003. An examination of this figure illustrates the characteristics that any power price dynamics model must solve. Linear diffusion models of the type underlying the Black-Scholes model clearly cannot capture the behavior depicted in Figure 1; there is no tendency of prices to wander as a traditional random walk model implies. Prices tend to vibrate around a particular level (approximately \$20 per megawatt hour) but sometimes jump upwards, at times reaching levels of \$1000/MWh.

To address the inherent non-linearities in power prices illustrated in Figure 1, some researchers have proposed models that include a jump component in power prices. This presents other difficulties. For example, a simple jump model like that proposed by Merton (1973) is inadequate because in that model the effect of a jump is permanent, whereas Figure 1 shows that jumps in electricity prices reverse themselves rapidly.

Moreover, the traditional jump model implies that prices can either jump up or down, whereas in electricity markets prices jump up and then decline soon after. Barz and Johnson (1999) incorporate mean reversion and exponentially distributed (and hence positive) jumps to address these difficulties. However, this model presumes that big shocks to power prices damp out at the same rate as small price moves. This is implausible in some power markets. Geman and Roncoroni (2006) present a model that eases this constraint, but in which, conditional on the price spiking upward beyond a threshold level, (a) the magnitude of the succeeding down jump is independent of the magnitude of the preceding up jump, and (b) the next jump is necessarily a down jump (i.e., successive up jumps are precluded once the price breaches the threshold). Moreover, in this model the intensity of the jump process does not depend on whether a jump has recently occurred. These are all problematic features. Barone-Adesi and Gigli (2002) address the problem through a regime shifting model. However, this model does not permit successive up jumps, and constraining down jumps to follow up jumps makes the model non-Markovian. Villaplana (2004) eases the constraint by specifying a price process that is the sum of two processes, one continuous, the other with jumps, that exhibit different speeds of mean reversion. The resulting price process is non-Markovian, which makes it difficult to use for contingent claim valuation.

Estimation of jump-type models also poses difficulties. In particular, a reasonable jump model should allow for seasonality in prices and a jump intensity and magnitude that are also seasonal with large jumps more likely when demand is high than when demand is low. Given the nature of demand in the US, this implies that large jumps are most likely to occur during the summer months. Moreover, changes in capacity and demand growth will affect the jump intensity and magnitude. Estimating such a model on the limited time series data available presents extreme challenges. Geman and Roncoroni (2006) allows such a feature, but most other models do not; furthermore, due to the computational intensity of the problem, even Geman-Roncoroni must specify the parameters of the non-homogeneous jump intensity function based on *a priori* considerations instead estimating it from the data. Fitting regime shifting models is also problematic, especially if they are non-Markovian as is necessary to make them a realistic characterization of power prices (Geman, 2005).

Even if jump models can accurately characterize the behavior of electricity prices under the "true measure," they pose acute difficulties as the basis for the valuation of power contingent claims. Jump risk is not hedgeable, and hence the power market is incomplete.[2] A realistic jump model that allows for multiple jump magnitudes (and preferably a continuum of jump sizes) requires multiple market risk prices for valuation purposes; a continuum of jump sizes necessitates a continuum of risk price functions to determine the equivalent measure that is relevant for valuation purposes. Moreover, these functions may be time varying. The high dimensionality of the resulting valuation problem vastly complicates the pricing of power contingent claims. Indeed, the more sophisticated the spot price model (with Geman-Roncoroni being the richest), the more complicated the task of determining the market price of risk functions.

There are also difficulties in applying jump models to the valuation of volumetric sensitive claims. For example, a utility that wants to hedge its

[2]The market would be incomplete even if power prices were continuous (as is possible in the model presented below) because power is non-storable. Non-storability makes it impossible to hold a hedging "position" in spot power.

revenues must model both the price process and the volume process. There must be some linkage between these two processes. Grafting a volume process to an already complex price process is problematic, especially when one recognizes that there is likely to be a complex pattern of correlation between load, jump intensity, and jump magnitude.

Relatedly, the relation between fuel prices and power prices is of particular interest to practitioners. For instance, the "spark spread" between power and fuel prices determines the profitability of operating a power plant. The relation between fuel and power prices is governed by the process of transforming fuel inputs into power outputs. This process can generate state-dependent correlations between input and output prices that is very difficult to capture using exogenously specified power and fuel price processes.

To address these limitations of traditional derivative pricing approaches in power market valuation, we propose instead an approach based on the economics of power production and consumption. In this approach, power prices are a function of two state variables. These two state variables capture the major drivers of electricity prices, are readily observed due to the transparency of fundamentals in the power market, and result in a model of sufficiently low dimension to be tractible.

The first state variable is a demand variable. To operationalize it, we employ two alternative definitions. The first measure of the demand state is load. The second is temperature. Since load and temperature are so closely related, these interpretations are essentially equivalent. To simplify the discussion, in what follows we use load as the demand variable. Later on we discuss how use of weather as the state variable permits unified valuation and hedging of power price, power volume, and weather sensitive claims.

An analysis of the dynamics of load from many markets reveals that this variable is very well behaved. Load is seasonal, with peaks in the summer and winter for Eastern, Midwestern, and Southern power markets. Moreover, load for each of the various National Electricity Reliability Council ("NERC") regions is nearly homoskedastic. There is little evidence of GARCH-type behavior in load. Finally, load exhibits strong mean reversion. That is, deviations of load from its seasonally-varying mean tend to reverse fairly rapidly.

We treat load as a controlled process. Defining load as q_t, note that $q_t \leq X$, where X is physical capacity of the generating and transmission system.[3] If load exceeds this system capacity, the system may fail, imposing substantial costs on power users. The operators of electric power systems (such as the independent system operator in the PJM region we discuss later) monitor load and intervene to reduce power usage when load approaches levels that threaten the physical reliability of the system.[4] Under certain techni-

[3]This characterization implicitly assumes that physical capacity is constant. Investment in new capacity, planned maintenance, and random generation and transmission outages cause variations in capacity. This framework is readily adapted to address this issue by interpreting q_t as capacity utilization and setting $X = 1$. Capacity utilization can vary in response to changes in load and changes in capacity. This approach incorporates the effect of outages, demand changes, and secular capacity growth on prices. The only obstacle to implementation of this approach is that data on capacity availability is not readily accessible. In ongoing research we are investigating treating capacity as a latent process, and using Bayesian econometric techniques to extract information about the capacity process from observed real time prices and load. The analysis of price-load relations in section 3 implies that load variations explain most peak load price variations in PJM prices, which suggests that at least over the short run ignoring capacity variation in this market is not critical. This may not be true for all markets.

[4]See various PJM operating manuals available at www.pjm.com for information on emergency procedures in PJM.

cal conditions (which are assumed to hold herein), the arguments of Harrison and Taksar (1983) imply that under these circumstances the controlled load process will be a reflected Brownian motion.[5] Formally, the load will solve the following SDE:

$$dq_t = \alpha_q(q_t, t)q_t dt + \sigma_q q_t du_t - dL_t^u \tag{1}$$

where L_t^u is the so-called "local time" of the load on the capacity boundary.[6] The process L_t^u is increasing (i.e., $dL_t^u > 0$) if and only if $q_t = X$, with $dL_t = 0$ otherwise. That is, q_t is reflected at X.

The dependence of the drift term $\alpha_q(q_t, t)$ on calendar time t reflects the fact that output drift varies systematically both seasonally and within the day. Moreover, the dependence of the drift on q_t allows for mean reversion. One specification that captures these features is:

$$\alpha_q(q_t, t) = \mu(t) + k[\ln q_t - \theta_q(t)] \tag{2}$$

In this expression, $\ln q_t$ reverts to a time-varying mean $\theta_q(t)$. $\theta_q(t)$ can be specified as a sum of sine terms to reflect seasonal, predictable variations in electricity output. Alternatively, it can be represented as a function of calendar time fitted using non-parametric econometric techniques. The parameter $k \leq 0$ measures the speed of mean reversion; the larger $|k|$, the more rapid the reversal of load shocks. The function $\mu(t) = d\theta_q(t)/dt$ represents the portion of load drift that depends only on time (particularly time of day).

[5] The conditions are (1) there exists a "penalty function" $h(q)$ that is convex in some interval, but is infinite outside the interval, and (2) in the absence of any control, q would evolve as the solution to $dq = \mu dt + \sigma dW$. The penalty function can be interpreted as the cost associated with large loads. If $q > X$, the system may fail, resulting in huge costs. We thank Heber Farnsworth for making us aware of the Harrison-Taksar approach.

[6] This is an example of a Skorokhod Equation.

For instance, given $\ln q_t - \theta_q(t)$, load tends to rise from around 3AM to 5PM and then fall from 5PM to 3AM on summer days.

The load volatility σ_q in (1) is represented as a constant, but it can depend on q_t and t. There is some empirical evidence of slight seasonality in the variance of q_t.

The second state variable is a fuel price. For some regions of the country, natural gas is the marginal fuel. In other regions, coal is the marginal fuel. In some regions, natural gas is the marginal fuel sometimes and coal is the marginal fuel at others. We abstract from these complications and specify the process for the marginal fuel price. The process for the forward price of the marginal fuel is:

$$\frac{df_{t,T}}{f_{t,T}} = \alpha_f(f_{t,T}, t) + \sigma_f(f_{t,T}, t)dz_t \tag{3}$$

where $f_{t,T}$ is the price of fuel for delivery on date T as of t and dz is a standard Brownian motion. Note that $f_{T,T}$ is the spot price of fuel on date T.

The processes $\{q_t, f_{t,T}, t \geq 0\}$ solve (1) and (3) under the "true" probability measure $\mathcal{P}$. To price power contingent claims, we need to find an equivalent measure $\mathcal{Q}$ under which deflated prices for claims with payoffs that depend on q_t and $f_{t,T}$ are martingales. Since $\mathcal{P}$ and $\mathcal{Q}$ must share sets of measure 0, q_t must reflect at X under $\mathcal{Q}$ as it does under $\mathcal{P}$. Therefore, under $\mathcal{Q}$, q_t solves the SDE:

$$dq_t = [\alpha_q(q_t, t) - \sigma_q \lambda(q_t, t)]q_t dt + \sigma_q q_t du_t^* - dL_t^u$$

In this expression $\lambda(q_t, t)$ is the market price of risk function and du_t^* is a $\mathcal{Q}$ martingale. Since fuel is a traded asset, under the equivalent measure

$df_{t,T}/f_{t,T} = \sigma_f dz_t^*$, where dz_t^* is a $\mathcal{Q}$ martingale. The change in the drift functions is due to the change in measure.

Define the discount factor $Y_t = \exp(-\int_0^t r_s ds)$ where r_s is the (assumed deterministic) interest rate at time s. (Later we assume that the interest rate is a constant r.) Under $\mathcal{Q}$, the evolution of a deflated power price contingent claim C is:

$$Y_t C_t = Y_0 C_0 + \int_0^t C_s dY_s + \int_0^t Y_s dC_s$$

In this expression, C_s indicates the value of the derivative at time s and Y_s denotes the value of one dollar received at time s as of time 0. Using Ito's lemma, this can be rewritten as:

$$Y_t C_t = C_0 + \int_0^t Y_s(\mathcal{A}C + \frac{\partial C}{\partial s} - r_s C_s)ds + \int_0^t [\frac{\partial C}{\partial q} du_s^* + \frac{\partial C}{\partial f} dz_s^*] - \int_0^t Y_s \frac{\partial C}{\partial q} dL_s^u$$

where $\mathcal{A}$ is an operator such that:

$$\begin{aligned} \mathcal{A}C &= \frac{\partial C}{\partial q_t}[\alpha_q(q_t, t) - \sigma_q \lambda(q_t, t)]q_t \\ &+ .5\frac{\partial^2 C}{\partial q_t^2}\sigma_q^2 q_t^2 + .5\frac{\partial^2 C}{\partial f_{t,T}^2}\sigma_f^2 f_{t,T}^2 + \frac{\partial^2 C}{\partial q_t \partial f_{t,T}}\sigma_f \sigma_q \rho_{qf} q_t f_{t,T}. \end{aligned} \quad (4)$$

For the deflated price of the power contingent claim to be a $\mathcal{Q}$ martingale, it must be the case that:

$$E[\int_0^t Y_s(\mathcal{A}C + \frac{\partial C}{\partial s} - r_s C_s)ds] = 0$$

and

$$E[\int_0^t Y_s \frac{\partial C}{\partial q} dL_s^u] = 0$$

for all t. Since (1) $Y_t > 0$, and (2) $dL_t^u > 0$ only when $q_t = X$, with a constant interest rate r, we can rewrite these conditions as:

$$\mathcal{A}C + \frac{\partial C}{\partial t} - rC = 0 \quad (5)$$

and

$$\frac{\partial C}{\partial q} = 0 \text{ when } q_t = X \tag{6}$$

It is obvious that (5) and (6) are sufficient to ensure that C is a martingale under $\mathcal{Q}$; it is possible to show that these conditions are necessary as well.

Expression (5) can be rewritten as the fundamental valuation PDE:[7]

$$\begin{aligned} rC &= \frac{\partial C}{\partial t} + \frac{\partial C}{\partial q_t}[\alpha_q(q_t,t) - \sigma_q\lambda(q_t,t)]q_t \\ &+ .5\frac{\partial^2 C}{\partial q_t^2}\sigma_q^2 q_t^2 + .5\frac{\partial^2 C}{\partial f_{t,T}^2}\sigma_f^2 f_{t,T}^2 + \frac{\partial^2 C}{\partial q_t \partial f_{t,T}}\sigma_f\sigma_q\rho_{qf}q_t f_{t,T} \end{aligned} \tag{7}$$

For a forward contract, after changing the time variable to $\tau = T - t$, the relevant PDE is:

$$\begin{aligned} \frac{\partial F_{t,T}}{\partial \tau} &= \frac{\partial F_{t,T}}{\partial q_t}[\alpha_q(q_t,t) - \sigma_q\lambda(q_t,t)]q_t \\ &+ .5\frac{\partial^2_{t,T}}{\partial q_t^2}q_t^2\sigma_q^2 + .5\frac{\partial^2_{t,T}}{\partial f_{t,T}^2}\sigma_f^2 f_{t,T}^2 + \frac{\partial^2_{t,T}}{\partial q_t \partial f_{t,T}}q_t f_{t,T}\sigma_f\sigma_q\rho_{qf} \end{aligned} \tag{8}$$

where $F_{t,T}$ is the price at t for delivery of one unit of power at $T > t$.

Expression (6) is a boundary condition of the Neumann type. This boundary condition is due to the reflecting barrier that is inherent in the physical capacity constraints in the power market.[8] The condition has an intuitive interpretation. If load is at the upper boundary, it will fall almost certainly. If the derivative of the contingent claim with respect to load is non-zero at the boundary, arbitrage is possible. For instance, if the partial derivative is positive, selling the contingent claim cannot generate a loss and almost certainly generates a profit.

[7] Through a change of variables (to natural logarithms of the state variables) this equation can be transformed to one with constant coefficients on the second-order terms.

[8] If there is a lower bound on load (a minimum load constraint) there exists another local time process and another Neumann-type boundary condition.

In (7)-(8), there is a market price of risk function $\lambda(q_t, t)$. The valuation PDE *must* contain a market price of risk because load is not a traded claim and hence load risk is not hedgeable. Accurate valuation of a power contingent claim therefore depends on accurate specification and estimation of the $\lambda(q_t, t)$ function.

Valuation of a power contingent claim ("PCC") also requires specification of initial boundary conditions that link the state variables (load and the fuel price) and power prices at the expiration of a PCC. In most cases, the buyer of a PCC obtains the obligation to purchase a fixed amount of power (e.g., 25 megawatts) over some period, such as every peak hour of a particular business day or every peak hour during a particular month. Similarly, the seller of a PCC is obligated to deliver a fixed amount of power over some time period. Therefore, the payoff to a forward contract at expiration is:

$$F(0) = \int_{t'}^{t''} \delta(s) P^*(q(s), f(s), s) ds \tag{9}$$

where F is the forward price, $q(s)$ is load at time s, $f(s)$ is the fuel spot price at s, $\delta(s)$ is a function that equals 1 if the forward contract requires delivery of power at s and 0 otherwise, $P^*(.)$ is a function that gives the instantaneous price of power as a function of load and fuel price, t' is the beginning of the delivery period under the forward contract, and t'' is the end of the delivery period. In words (9) states that the payoff to the forward equals the value of the power, measured by the spot price, received over the delivery period. For instance, if the forward is a monthly forward contract for the delivery of 1 megawatt of power during each peak hour in the month, $\delta(s)$ will equal 1 if s falls between 6 AM and 10 PM on a weekday during that month, and will equal 0 otherwise.

Economic considerations suggest that the price function $P^*(.)$ is increasing and convex in q; section 4 provides evidence in support of this conjecture. As load increases, producers must employ progressively less efficient generating units to service it. The spot price function should also be a function of calendar time, with higher prices (given load) in spring and fall months than in summer months due to the fact that utilities schedule their routine maintenance to coincide with the seasonal demand "shoulders."

This pricing function determines the dynamics of the instantaneous power price. Using Ito's lemma,

$$dP^* = \Phi(q_t, f_{t,f}, t)dt + P_q^* \sigma_q q_t du_t + P_f^* \sigma_f f_{t,f} dz_t \tag{10}$$

with

$$\begin{aligned}\Phi(q_t, f_{t,f}, t) = P_q^* \alpha_q(q_t, t) q_t \quad &+ \quad P_f^* \alpha_f(f_{t,t}, t) f_{t,t} \\ &+ \quad .5 P_{qq}^* \sigma_q^2 q_t^2 + .5 P_{ff}^* \sigma_f^2 f_{t,f}^2 + P_{qf}^* q_t f_{t,f} \sigma_q \sigma_f \rho_{qf}\end{aligned}$$

where ρ_{qf} is the correlation between q_t and $f_{t,T}$; this correlation may depend on q_t, $f_{t,T}$, and t.[9] The volatility of the instantaneous price in this setup is time varying because P^* is a convex, increasing function of q. Specifically, the variance is

$$\sigma_P^2(q_t, f_{t,t}, t) = P_q^{*2} \sigma_q^2 q_t^2 + P_f^{*2} f_{t,t}^2 \sigma_f^2 + 2 P_f^* P_q^* q_t f_{t,t} \rho_{qf} \sigma_q \sigma_f. \tag{11}$$

[9] The spot price process is continuous if P^* has continuous first derivatives. Nonetheless, the market is still incomplete since q_t is not traded. Moreover, when output nears capacity and hence P_q^* becomes very large, the price can appear to exhibit large jumps even if prices are observed at high frequency (e.g., hourly). The spot price process is also likely to be dis-continuous due to discontinuities in generators' bids to sell power. These bids are step functions.

Since P_q^* is increasing with q, demand shocks have a bigger impact on the instantaneous price when load is high (i.e., demand is near capacity) than when it is low. In particular, if the price function becomes nearly vertical when demand approaches capacity, small movements in load can cause extreme movements in the instantaneous price. Moreover, given the speed of load mean reversion, the convexity of P^* implies that the speed of price mean reversion is state dependent; prices revert more rapidly when load (and prices) are high than when they are low. These non-linearities are a fundamental feature of electricity price dynamics, and explain many salient and well-known features of power prices, most notably the "spikes" in prices when demand approaches capacity and the variability of power price volatility.

The model also implies that the correlation between the fuel price and the power price will vary. Assuming that $\rho_{qf} = 0$ (which is approximately correct in most markets), then

$$corr(dP^*, df) = \frac{P_f^* \sigma_f f_{t,T}}{\sqrt{P_q^{*2} q_t^2 \sigma_q^2 + P_f^{*2} f_{t,T}^2 \sigma_f^2}}$$

Note that when load is small, $P_q^* \approx 0$, in which case $corr(dP^*, df) = 1$. Moreover, when load is large, $P_q^* \approx \infty$, in which case $corr(dP^*, df) = 0$. It is also straightforward to show that the correlation declines monotonically with q_t because P_q^* increases monotonically with q_t.[10] Thus, the model can generate rich patterns of correlation between power and fuel prices, and commensurately rich patterns of spark spread behavior.

The following sections discuss implementation of this model and describe

[10]This result can be generalized to $\rho_{qf} \neq 0$. The same basic results hold; power price-fuel price correlations are high when load is small, and the correlations are small when load is high.

some of its implications.

3 Estimating the Demand Process

The drift process for load given by (1) and (2) is complex because the change in load (conditional on the deviation between load and its mean) and the mean load both vary systematically by time of day, day of the week, and time of the year. To capture these various effects we utilize nonparametric techniques.

It is necessary to estimate $\mu(t)$, $\theta_q(t)$, and k. To see how this is done, consider a discrete version of (1) and (2) that ignores the local time:

$$\frac{\Delta q_t}{q_t} = \mu(t)\Delta t + k[\ln q_t - \theta_q(t)]\Delta t + \sigma_q \sqrt{\Delta t}\epsilon_t$$

where ϵ_t is an i.i.d. standard normal variate. We have hourly load data, so Δt is one hour. Note that:

$$\frac{\Delta q_t}{q_t} - E[\frac{\Delta q_t}{q_t}|t] = k\{\ln q_t - E[\ln q_t|t]\}\Delta t + \sigma_q \sqrt{\Delta t}\epsilon_t \tag{12}$$

Simple algebra demonstrates that $\mu(t) = E[\Delta q_t/q_t|t]$ and $\theta_q(t) = E[\ln q_t|t]$. Once these conditional expectations are known, k can be estimated by OLS.

Two nonparametric approaches were utilized to estimate how the log of expected load depends on time of day, the day of the year, and the day of the week. Both approaches give virtually identical results, so for brevity we describe only one method.[11] In this approach, we first create a 53 by 24 by 7 grid. The first dimension measures day of the year d, which runs between 1

[11]The other approach is also nonparametric, but estimates the day of the week effects using day-of-week dummies in a kernel regression. We have also modeled the mean load as a sum of sine functions. All methods give similar pricing results.

and 365 in increments of 7 days. The second dimension measures hour of the day h, and the third measures day of the week w, with 1 corresponding to Monday, 2 to Tuesday, and so on. Given this grid, we then estimate expected log load as a function of the three time variables using hourly load data from PJM for 1 January, 1992 to 28 February, 2003. The data are first detrended by assuming a 2 percent annual load growth rate.

At each point of the grid, we estimate two sets of local linear regressions; each set entails estimation of $53 \cdot 24 \cdot 7 = 8904$ sets of weighted least squares regressions.

In the first set, at each point of the grid we estimate a local linear regression with $\Delta q_t/q_t$ as the dependent variable and $d^*_{t,i}$ and a constant as the independent variables. The variable $d^*_{t,i}$ is the number of days between t and day i; this number is less than or equal 182 (except in a leap year, when it is less than or equal to 183). Note that $d^*_{t,i} = d_t - d_i$ if $|d_t - d_i| < 365 - |d_t - d_i|$, $d^*_{t,i} = d_t - d_i - 365$ if $|d_t - d_i| \geq 365 - |d_t - d_i|$ and $d_i < d_t$, and $d^*_{t,i} = d_i + d_t - 365$ if $|d_t - d_i| \geq 365 - |d_t - d_i|$ and $d_i \geq d_t$, where d_t is the day of the year (a number between 1 and 365) corresponding to time t, and d_i is the day of the year corresponding to point i in the day of the year dimension of the grid. Although we do not include these variables as regressors, we also define the hour distance (which is less than or equal to 12) as $h^*_{t,j} = h_t - h_j$ if $|h_t - h_j| < 24 - |h_t - h_j|$, $h^*_{t,j} = h_t - h_j - 24$ if $|h_t - h_j| \geq 24 - |h_t - h_j|$ and $h_j < h_t$, and $h^*_{t,j} = h_j + h_t - 24$ if $|h_t - h_j| \geq 24 - |h_t - h_j|$ and $h_j \geq h_t$, where h_t is the hour of the day corresponding to time t, h_j is the hour corresponding to point j on the hour dimension of the grid, and the day of the week distance (which is less than or equal to 3) as $d^*_{t,k} = d_t - d_k$

if $|d_t - d_k| < 7 - |d_t - d_k|$, $d^*_{t,k} = d_t - d_k - 7$ if $|d_t - d_j| \geq 7 - |d_t - d_k|$ and $d_k < d_t$, and $d^*_{t,k} = d_k + d_t - 7$ if $|d_t - d_k| \geq 7 - |d_t - d_k|$ and $d_k \geq d_t$, where w_t is the day of the week corresponding to t, and w_k is the k'th point on the day of the week dimension. The hour and week day variables are relevant in determining the kernel weights.

Each observation is weighted by the square root of the multiplicative Gaussian kernel:

$$K(d_t, h_t, w_t | d_i, h_j, w_k) = \frac{1}{b_d b_h b_w} n(\frac{d^*_{t,i}}{b_d}) n(\frac{h^*_{t,j}}{b_h}) n(\frac{w^*_{t,k}}{b_w})$$

where $n(.)$ is the standard normal density, b_d is the day of year bandwidth (in days), b_h is the hour of day bandwidth (in hours), and b_w is the day of week bandwith (in days). After some experimentation, we chose bandwiths $b_d = 14$, $b_h = 2$ and $b_w = 1$; results are not very sensitive to this choice. The constant from the regression at the point on the grid corresponding to day of the year d_i, hour of the day h_j, and day of the week w_k is the value of $\mu(t)$ for this specific time. For other t not corresponding to grid points, $\mu(t)$ is estimated through interpolation. We find that $\mu(t)$ is important; this function explains approximately 35 percent of the variation in load changes.

In the second set of regressions, we estimate local linear regressions with $\ln q_t$ as the dependent variable and $d^*_{t,i}$ and a constant as independent variables. The observations are weighted by $\sqrt{K(.)}$. The value of $\theta_q(t)$ on d_i, h_j, and w_k is given by the constant from the regression estimated for this point on the grid; $\theta_q(t)$ for other t is determined through interpolation. Figure 2 illustrates this function for the PJM market. Note the two seasonal peaks, one corresponding to the summer cooling season and the other corresponding to the winter heating season, with seasonal troughs in the spring and fall. The summer peak is larger than the winter one. Moreover, there is a clear variation in load by time of day, with low loads during the night and higher loads during the day. The intraday variation is more pronounced in the summer than the winter.

Given the estimates of $\mu(t)$ and $\theta_q(t)$ we estimate the speed of mean reversion k from (12) using OLS and the 1992-2003 PJM hourly load data. Using hourly data, $\hat{k} = -.0614$; the annualized value of k is therefore -537.64 indicating extremely rapid reversal of load shocks. Indeed, the half-life of a load shock is only 11.3 hours. The sample variance of the error term is $\hat{\sigma}_q^2 = .00091$. This variance and $\hat{k}$ imply that the unconditional variance of $\ln q_t - \theta_q(t)$ is $-\hat{\sigma}_q^2/2\hat{k} = .0074$. Due to the rapid mean reversion, the conditional variance of this difference converges to its unconditional value quite quickly.

4 Determining the Terminal Pricing Function

Valuation of a PCC using the equilibrium model requires estimation of the payoff to a forward contract (or other derivative) to serve as the initial boundary condition, where this payoff usually has the form given in (9) above. Determining this payoff function poses several challenges.

First, solution of the valuation PDE (8) using finite difference methods requires discretization of time and load steps, so it will be necessary to create a discretized approximation of (9). Relatedly, computational considerations call for using a relatively course time grid, so even an approximation of the integral in (9) with a sum of hourly prices during the delivery period is not practical; this is especially true when solving the inverse problem (as in the following section) as this requires solving PDEs for each and every maturity from the present to the last day of the most distant contract's delivery period. A daily time step is reasonable, so it is necessary to approximate the integral in (9) with a function of load at a single time on a given day.

Second, it is necessary to understand the relation betweeen price and load. There are three basic approaches that one can employ to do so. The first is to assume that the power market is competitive and utilize data on marginal generation costs as a function of load and fuel prices to determine the terminal power price as a function of these state variables. This is the approach advanced by Eydeland and Geman (1999). The second is an econometric approach that does not assume perfect competition. The third utilizes generators energy bids in centralized spot markets (where available) to construct a bid stack; this approach does not assume competition because generators' bids reflect any market power they possess.

There are numerous studies that document market power in pooled markets with generation bidding such as PJM. Examples include Rudkevich and Duckworth (1998); Green and Newbery (1992); Newbery (1995); Wolak and Patrick (1997), Wolfram (1999), and Hortaçsu and Puller (2005).[12] Thus the first alternative is problematic.

Since prices reflect market power, the second and third approaches are preferable to the first. Where the relevant data are available (as it is for PJM and some other markets with a system operator), the bid based approach has

[12]Since bidders in most pools submit supply schedules and most pools utilize second price rather than first price mechanisms, the analysis of Back and Zender (1992) implies that non-competitive outcomes are plausible in power pools.

several virtues. Most notably, system operators actually use bids to set spot prices, so there is a direct relation between bids and realized spot prices. Implementing this method does pose some challenges, however.

First, the economics of generation are actually quite complex. Startup, shutdown and no load costs imply that optimal dispatch requires solution of a rather complex (and non-convex) dynamic programming problem. Moreover, due to outages (planned and forced) the set of generating assets available varies over time. The spatial pattern of load can also vary, and in the presence of transmission constraints such varations can cause prices at a given point to vary even if aggregate load in a region is unchanged; generation may be dispatched out of merit order due to transmission constraints, which can cause price fluctuations even in the absence of fluctuations in aggregate load. These factors, in turn, imply that the marginal cost of generation at any instant is a function of past loads and operating decisions, the set of available generators, the spatial pattern of load, and the existence of transmission constraints; due to these factors there is no unique mapping between load and the marginal bid/market price in the market. However, taking these complexities into account would greatly increase the dimensionality of the problem, making it computationally intractible.

Second, bids may differ by hour, and loads certainly do in a systematic way. Thus, calculation of a peak load price requires the knowledge of loads for every instant of the on-peak period. Due to the necessity of discretizing, it is necessary to find some way to calculate a peak load price based on loads from only a subset of the peak hours, and perhaps load from only a single hour.

Third, most bid data is publically available only with a lag (six months for PJM, for instance), and bids may change over time with entry of new generating or transmission capacity, or changes in fuel prices. Thus, whereas the boundary conditions for the PDE should be forward looking, available data is backward looking.

Despite these challenges, the bid data have a crucial and desirable feature: they reflect market participants' intimate knowledge of the characteristics of generating assets and the competitiveness of the market. Hence, I utilize PJM bid data to estimate the terminal price function.[13]

Numerical and economic considerations also suggest imposing some additional structure on the problem. Specifically, as just noted, market participants' bids are for power only, but should vary systematically with fuel prices. It is therefore necessary to adjust historical bid information to reflect the impact of variations in fuel prices on future payoffs. This can be done by utilizing the concept of a "heat rate" that represents the number of BTU of fuel needed to generate the marginal megawatt of power, where heat rate is an increasing function of load to reflect the use of progressively less efficient generating units to serve larger loads. The heat rate is the conventional way for practitioners to analyze the impact of fuel price changes on power prices.

If the market heat rate fuction is $\phi(q_t)$, with $\phi'(.) > 0$,

$$P^*(q_t, f_{t,t}) = f_{t,t}\phi(q_t)$$

[13] Earlier versions of this paper utilize an econometric technique. This technique is described in an Appendix. The basic conclusions presented herein obtain under both terminal pricing function methods. Sepcifically, both methods imply that forward prices are significantly upward-biased.

More generally, one can consider a function of the form

$$P^*(q_t, f_{t,t}) = f_{t,t}^{\gamma}\phi(q_t)$$

The first specification imposes the restriction that the elasticity of the power price with respect to the fuel price equals one. The second specification does not.

In addition to being an economically sensible way of adjusting bids to reflect fuel price variations, this approach has computational benefits. Specficially, it permits the reduction of the dimensionality of the problem. Posit that the forward price function is of the form $F(q_t, f_{t,T}, \tau) = f_{t,T}^{\gamma} V(q_t, \tau)$. Then, making the appropriate substitutions into (9) produces the new one dimensional PDE:

$$\frac{\partial V}{\partial \tau} = .5\sigma_q^2 q^2 \frac{\partial^2 V}{\partial q^2} + [\gamma\rho\sigma_f q + a]\frac{\partial V}{\partial q} + .5\sigma_f^2\gamma(\gamma - 1)V \qquad (13)$$

where $a = (\alpha_q(\tau, q) - \sigma_q\lambda(\tau, q))q$, and where the equation must be solved subject to the von Neuman boundary condition $\partial V(X, \tau)/\partial q = 0$ and the initial condition $V(q_T, 0) = \phi(q_T)$. Once this function is solved for, the power forward price is obtained by multiplying $V(.)$ by $f_{t,T}^{\gamma}$. The reduction in dimensionality is especially welcome when solving the computationally intense inverse problem for $\lambda(q, t)$ as described in the next section.[14]

Given these preliminaries, the analysis proceeds as follows. Assume that it is 1 June, 2005, and that the objective is to determine the boundary

[14]This reduction is not feasible for all power contingent claims, particularly options. Even if the spot price function is multiplicatively separable in fuel price and load, the option payoff ($[f_{T,T}^{\gamma}\phi(q_T) - K]^+$ for a call at strike K) is not multiplicatively separable, so the decomposition $f_{t,T}^{\gamma}V(q, \tau)$ is not appropriate for such an option. If $\gamma = 1$, the decomposition works for a spark spread option that has payoff $[f_{T,T}\phi(q_T) - f_{T,T}H^*]^+$ where H^* is the heat rate strike specified in the contract.

condition for a foward contract maturing on 15 July, 2005; the analysis for other maturity dates is identical. Then:

1. Collect PJM bid data for 15 July, 2004 (the most recent July for which such data is available). For each generating unit, this data reports a set of price-quantity pairs, with a maximum of 10 different pairs per unit. The quantity element of the pair represents the amount the bidder is willing to generate at the price element of the pair. Sort all such pairs by price, and to determine the amount supplied at a given price P sum all of the quantities bid at prices P or lower. This is the "bid stack." The bid stack characterizes price as a function of instantaneous load–i.e., P^*. Figure 3 depicts a bid stack from 15 July, 2004.

2. Convert the bid stack into a heat rate stack by dividing the bid stack by the price of fuel on 15 July, 2004. We use daily data on day-ahead natural gas prices for Columbia Gas and Texas Eastern Pipeline zone M-3 obtained from Bloomberg. Columbia pipeline and Texas Eastern pipeline zone M-3 serve plants in the PJM territory. This is referred to as the "market heat rate stack" because the heat rate reflects market bids divided by the price of fuel. It may differ from the true heat rate stack (i.e., from true marginal costs) because bids may differ from marginal costs.

3. Define a vector of market loads that may be observed at 4 PM on 15 July, 2005. The number of load points is given by the number of such points in the finite difference valuation grid to be used to solve the

valuation PDE.[15]

4. Utilize the load surface described in section 3 above that relates expected load to time of day, day of week, and day of the year to determine a "load shape" for 15 July, 2005. This load shape indicates the expected load for each hour of this date. Divide each of the 16 expected loads for peak hours given by the load shape by the 4 PM expected load.

5. Multiply each load value in the load vector by the load shape. Each product gives 16 hourly loads corresponding to that value of 4 PM load. Since there are multiple values of 4 PM load in the valuation grid, there are multiple 16 hour load vectors–one for each load step in the valuation grid.

6. Using each of the 16 hour load vectors, input the load for each of the peak hours into the heat rate stack function that relates price to load. This produces 16 different market heat rates (one for each peak hour) corresponding to a particular value of 4 PM load from the valuation grid. Average these 16 heat rates to produce a peak heat rate on July 15, 2005 conditional on 4 PM load for that date. This represents the payoff (boundary condition) for that value of the 4 PM corresponding to a particular point on the valuation grid. That is, this is an approximation of the integral in (9) that takes into account variations in load over time and systematic intra-day patterns in load. There is a peak heat rate estimate for each load value in the valuation grid.

[15] See the next section for a detailed description of the PDE solution technique.

This approach implicitly assumes that generator bidding strategies are relatively stable over time, that generators effectively bid heat rates multiplied by fuel prices, and that generators are price takers in the fuel market. The first assumption is plausible when valuing relatively short tenor forward prices, but is problematic over longer time periods when entry, exit, demand growth, and changes in market rules may affect the market power of generators. The second assumption comports with conventional analyses of generating economics, but should be validated empirically; choice of a $\gamma \neq 1$ could be used if generators do not adjust power bids by x percent in response to an x percent change in fuel price.

Moreover, due to the fact that load does not map one-to-one into a heat rate (due to the complexities of generation), even conditional on load there will be a difference between realized power prices and those implied by the heat rate function. The existence of such a noise term does not impact the estimation of the market price of risk function under certain simplifying assumptions. Specifically, if $P_T = f_{T,T}^{\gamma}\phi(q_T) + \epsilon_T$, where ϵ_T, the divergence between observed spot prices and the heat rate function that results from the factors discussed earlier, is unpriced in equilibrium, then the forward price is $F_{t,T} = \tilde{E}_t f_{T,T}^{\gamma}\phi(q_T) + \tilde{E}_t \epsilon_T$ where $\tilde{E}_t$ indicates the time-t expectation under $\mathcal{Q}$. If ϵ_t mean reverts very rapidly (with a half-life measured in hours, for instance, as is plausible for price movements caused by forced outages or variations in the spatial variation in load), and this risk is not priced, then if $T - t$ is as little as a day then $\tilde{E}_t \epsilon_T$ (which is conditional on ϵ_t) is very nearly zero. Hence, $F_{t,T}$ can reasonably be considered a function of q_t and $f_{t,T}$ alone. The Feynman-Kac Theorem implies that $\tilde{E}_t f_{T,T}^{\gamma}\phi(q_T)$ is given by the solution to (8) with initial condition $F_{T,T} = f^{\gamma}_{T,T}\phi(q_T)$.

5 Power Forward Prices and Expected Spot Prices: The Market Price of Risk

As noted earlier, it is essential to incorporate the market price of risk in any power derivative pricing exercise. The market price of risk is inevitably present in any valuation problem due to the fundamental nature of electricity. Moreover, the data make it clear that ignoring the market price of risk is likely to lead to serious pricing errors because it is large.

Data from PJM illustrate this point clearly. If the market price of risk is nonzero, the forward price will differ from the expected spot price. Therefore, systematic differences between forward prices and realized spot prices are evidence of a market price of risk. For the PJM West Hub, on average there are systematic differences between one-day forward prices and realized spot prices over the 1997-1999 period (where the day ahead prices are bilateral trade prices reported on Bloomberg) and the 2000-2003 period (where the forward prices are from the PJM day ahead market). Over the 1997-1999 period, the forward price for peak power delivered on the following day exceeded the average realized peak hourly price of power in PJM West on the following day by an average of \$.92/MWh. Moreover, the median of the difference between the one day forward price and the realized average peak hourly price on the following day was \$1.36/MWh. This large median indicates that the difference between the forward price and the realized spot price is not due to a few outliers. Furthermore, the forward price exceeded the realized spot price on 311 of the 503 days in the sample. The forward price is also a biased predictor of the next day's realized spot price. The intercept in a regression of the day t average peak spot price against the day $t-1$ one day forward price is 8.75 and the slope coefficient is .6545. The standard error on the intercept is 1.5, while that on the slope coefficient is .047. One can therefore reject the null that the intercept is zero and the slope is one at any conventional significance level.

Similar results obtain for the 2000-2003 period. During this period, the day ahead price for PJM West averaged \$1.03 more than the realized real time price. The median difference between the day ahead and real time prices was \$1.12. The constant in a regression of the real time on the forward was 2.26, and the slope .89. One can again reject the null that the forward is an unbiased predictor of the real time price.

These data make it clear that the forward price is not an unbiased predictor of realized spot prices even one day hence.[16] Indeed, the bias is large. This indicates that even over a horizon as short of a day there is a risk premium embedded in power forward prices.[17]

[16]There is evidence of such bias in other markets. Averaged across all PJM pricing points, the median difference between the day ahead and real time prices was \$1.73 in 2002. In 2000, the day ahead price was \$4.00 greater on average than the real time price. Borenstein *et al* also document large disparities between day ahead and realized real time prices in the California market. They attribute these disparities to market inefficiencies and market power, and rule out risk premia as an explanation on largely *a priori* grounds. In particular, they invoke the CAPM to argue that the low correlation between power prices and the overall market (proxied by the S& P 500) implies that risk premia should be small. As Bessimbinder and Lemon (2002) note, however, (a) this presumes that the power market is integrated with the broader financial market, and (b) there is considerable reason to believe that in fact the power market is not so integrated. Moreover, the fact that forward prices are biased predictors of spot prices in markets other than California casts doubt on the view that the Borenstein *et al* results reflect only dysfunction in California.

[17]This finding is remarkable given the difficulty of detecting risk premia in other commodities. Economists since Telser (1956) have used increasingly sophisticated methods to attempt to find risk premia in commodities, with mixed results.

Since the market price of risk is potentially large, it is imperative to take it into account when valuing derivatives. Unfortunately, the market price of risk is not observable directly. However, it can be inferred from prices of traded instruments. In particular, given a set of quoted forward prices, inverse problem methods can be applied to (8) to generate an implied market price of risk.

At any time there are only a finite number of forward prices quoted in the marketplace. Call the set of available forward quotes $\mathcal{D}$. There is an arbitrary number of functions λ that could equate exactly the solutions of (8) for a finite number of quotes in $\mathcal{D}$. Thus, the problem of determining the market price of risk is ill-posed. If a problem is ill-posed, small changes in the data input (i.e., in the forward quotes) can lead to large changes in the estimates of the λ function. These problems are quite common in a variety of physics and engineering contexts, and methods have been developed to solve them. These involve use of regularization techniques.[18]

The solution technique involves choosing a function λ that minimizes the sum of squared deviations between the forward prices implied by (8) for a given set of delivery dates and the prices quoted for these dates, subject to some regularization constraint. To make the problem tractible, λ is a function of load only. We use the H^2 norm as our regularizer. Formally, the regularizer is:

$$R(\lambda) = \int_0^X \left[\lambda^2 + \left(\frac{\partial \lambda}{\partial q} \right)^2 \right] dq.$$

where as before X is the maximum load (given by the physical capacity of

[18]See Tikhonov and Arsenin (1977).

the power system). We choose a function λ to minimize:

$$\sum_{i \in \mathcal{D}} [F_i(q_t, f_{t,T}|\lambda) - \mathcal{F}_i]^2 + \kappa R(\lambda)$$

In this expression, $F_i(q_t, f_{t,T}|\lambda)$ is the solution to (8) for a forward contract corresponding to forward price quote $i \in \mathcal{D}$ given current load q_t and current fuel price $f_{t,T}$; note that this forward price depends on the λ function. Moreover, $\mathcal{F}_i$ is the quoted ith forward price. Finally, κ is the regularization parameter.

The regularization technique in essence penalizes overfitting. In the regularized problem, there is a trade-off between the precision with which the forward quotes are fit and the smoothness of the λ function. Note that $R(\lambda)$ is large when $|\partial\lambda/\partial q|$ is large. Thus, $R(\lambda)$ is large (small) when λ is very jagged (smooth). Choice of a regularization parameter κ determines the smoothness of the resulting fit; the bigger the value of this parameter, the greater penalty for non-smoothness, and the smoother the resulting solution.

This problem is solved using finite difference techniques. We create a valuation grid in q and τ with increments Δq and $\Delta\tau$; there are N total time steps and M total q values. An initial guess for λ is made. In the valuation grid, λ is represented as a vector, where the length of the vector is equal to M.

The estimation then proceeds by the Method of Small Parameter. Specifically, letting λ_0 denote the initial guess for the λ function, define

$$\lambda(q) = \lambda_0(q) + \sum_{j=1}^{\infty} \lambda_k(q)$$

where the λ_k are improvements to the initial guess. Similarly, define V_0 as the solution to (13) based on λ_0, and let

$$V = V_0 + \sum_{j=1}^{\infty} V_k$$

where the V_k are improvements to the initial guess. Note that the use of the $V(.)$ function exploits the dimensionality reduction.

Consider an ε-family of market price of risk functions:

$$\lambda(q) = \lambda_0 + \sum_{j=1}^{\infty} \varepsilon^k \lambda_k(q)$$

with a corresponding representation of V:

$$V(\tau, q) = V_0(\tau, q) + \sum_{j=1}^{\infty} \varepsilon^k V_k(\tau, q).$$

Plugging these various equations into the valuation PDE and equating powers of ε implies:

$$\frac{\partial V_0}{\partial \tau} = c_0 \frac{\partial^2 V_0}{\partial q^2} + c_1[\lambda_0] \frac{\partial V_0}{\partial q} + c_0 V_0 \quad (14)$$

$$\frac{\partial V_1}{\partial \tau} = c_0 \frac{\partial^2 V_1}{\partial q^2} + c_1[\lambda_0] \frac{\partial V_1}{\partial q} + c_0 V_1 - \sigma_q q \lambda_1 \frac{\partial V_0}{\partial q} \quad (15)$$

$$\frac{\partial V_k}{\partial \tau} = c_0 \frac{\partial^2 V_k}{\partial q^2} + c_1[\lambda_0] \frac{\partial V_k}{\partial q} + c_0 V_k - \sum_{j=0}^{k} \sigma_q q \lambda_j \frac{\partial V_{k-j}}{\partial q} \quad (16)$$

for $k = 2, 3, \ldots$ where $c_0 = .5\sigma_f^2 \gamma(\gamma - 1)$, $c_1[\lambda_0] = q[\gamma\rho\sigma_q\sigma_f + \alpha_q(\tau, q) - \sigma_q \lambda_0(\tau, q)]$, and $c_2 = .5\sigma_q^2 q^2$. There is a set of equations (14)-(16) for each maturity date included in the analysis.

The PDE (14) is solved implicitly by switching between forward and backward difference approximations for the first partial derivative term depending on the sign of the drift term. Now note that it can be shown from (15) that in the discretized q and τ grid, a recursion relationship holds:

$$A_{n+1} V_1^{(i)}[n+1] = V_1^{(i)}[n] + G_{n+1}^{(i)} \lambda_1 \quad (17)$$

where $n+1$ indicates the time step, A_{n+1} is a tri-diagonal $(M-2)\times(M-2)$ matrix, $V_1^{(i)}[n]$ is an $M-2$ vector giving the value of V_1 at each interior load point for maturity $i \in \mathcal{D}$ at time step n, and

$$G_{n+1}^{(i)} = -.5\frac{\Delta\tau}{\Delta q}\sigma_q q \frac{\partial V_0^{(i)}(q,(n+1)\Delta\tau)}{\partial q}$$

where the partial derivative is estimated using a central finite difference. Furthermore, $G_n^{(i)} = 0$ if $n\Delta t$ is greater than the time to maturity of forward contract $i \in \mathcal{D}$. Completing this recursion implies:

$$V_1^{(i)}[N] = \sum_{j=1}^{N-1} (\Pi_{k=0}^{j} A_{N-j}) G_{N-j}^{(i)} \lambda_1 \equiv \mathcal{B}^{(i)}\lambda_1 \tag{18}$$

Now note that the regularized objective function becomes:

$$\sum_{i\in\mathcal{D}} [V_0^{(i)} + \mathcal{B}^{(i)}\lambda_1 - \frac{\mathcal{F}_i}{f_i^{\gamma}}]^2 + \kappa R(\lambda_1) \tag{19}$$

where f_i is the fuel forward price with the same maturity as power forward contract i. One can show that if one uses the trapezoidal rule to approximate the integral in the expression for the regularizer, $R(\lambda_1)$ is quadratic in λ_1. Therefore, minimizing (19) with respect to λ_1 produces a set of linear equations that can be solved for λ_1.

Given this improvement, a similar method can be used to solve for λ_k, $k > 1$, based on (16); the main difference in the solution technique for $k = 2, 3, \ldots$ as opposed to $k = 1$ is the presence of additional forcing terms that depend on the q-derivatives of V improvements $k-1, k-2, \ldots, 2$. The user can choose the total number of improvements to implement, with more improvements involving greater computational cost (particularly storage).

Using this method, we evaluate the market price of risk function for PJM prices for a variety of dates from 1999-2005. We present detailed results for 7

June, 2005. For this date, we obtained prices for monthly PJM forwards for each delivery month July, 2005 through December, 2005. These prices were for PJM monthly futures transactions executed on the NYMEX ClearPort system for delivery dates in July, 2005 through December, 2005. These contracts call for delivery of power during each peak hour of the delivery month.

Figure 4 depicts the λ function fitted to these seven forward prices based on four improvements. Note that the market price of risk function is uniformly negative. This implies that under the equivalent measure load drifts up more rapidly than under the physical measure. Given that the power price is monotonically increasing in load, this in turn implies that forward prices are upward biased; that is, the expectation under the equivalent measure exceeds the expectation under the physical measure. Note also that the absolute value of the λ function is increasing in load, which implies that upward bias should be more extreme for forwards expiring during high demand periods.

The $\lambda(q)$ function thus determined can then be used in expression (8) to solve the direct problem for daily forward contracts for each possible delivery date during the 8 June, 2005-31 December, 2005 period. The solutions to these direct problems give theoretical forward prices for each such delivery date that are consistent with the transaction prices for the monthly forward contracts. This is similar to using volatilities implied from traded options to value other options.

The shape of the λ function generated by the model and the solution of the inverse problem implies that the buyer of a July or August forward position must pay a substantial risk premium to the seller. This is illustrated by Figure 5, which depicts (1) the forward prices for each possible daily delivery date determined as just described, and (2) the forward prices for each such delivery date implied by the model assuming that $\lambda = 0$. The forward price with $\lambda = 0$ is the expected spot price implied by the dynamics of the load and fuel prices and by the terminal pricing function. The jagged nature of the lines is due to the systematic intra-weekly variations in load, which affect both the expected spot price and the likelihood of price spikes.[19]

Note that the difference between these two functions exhibits seasonalities, as illustrated in Figure 6. Although the bias is positive throughout the year, for the summer peak delivery dates the difference between the forward price curve derived from the solution of the inverse problem and the forward curve assuming no market price of risk is particularly large. It reaches its maximum in mid-summer when the difference between the forward price and the expected spot price reaches a maximum of almost \$19/MWh, which is about 24 percent of the calibrated forward price for the relevant day. The average difference between the forward price and expected price curves is \$17.71/MWh for July delivery. In contrast, the upward bias in Octoberr is only about \$6.64/MWh.

Although this upward bias is large, it is actually smaller than observed in other years. We have estimated the λ function for various dates in the 1998-2005 time period.[20] In all cases, forward prices are upward biased for July

[19] The vertical displacements in the figure reflect the use of monthly fuel prices that are substantially different month-to-month for the $f_{t,T}$.

[20] For years prior to the launch of PJM futures on ClearPort, we utilize OTC forward price quotes obtained from Bloomberg's Volt page. We also used the econometric method to estimate the payoff. For 2005, the econometric method gives similar values to those

and August deliveries. In general, the degree of upward bias has declined over time. In May, 2000, for instance, the average upward bias for July-August deliveries averaged 35 percent of the forward price, and reached a maximum of $50/MWh–or about 50 percent of the calibrated forward price. In 2003, the maximum upward bias was substantially smaller–about $35/MWh. As just noted, it is smaller still in 2005.

The existence of a substantial upward bias is consistent with the existence of a "skewness premium" or "spike-a-phobia" in power prices. During the summer months in particular, power prices can spike up, generating substantial right skewness in prices. A price spike of this sort can impose a large loss on a seller of power forwards; for example, during the summer of 1998 one large utility (Illinova) that was short power had an entire year's earnings wiped out by a single day's trading losses, and other utilities (including Pacificorp and others) lost heavily as well. Thus, the profit distribution of a short power forward position exhibits substantial left skewness. It is well known that those with consistent preferences exhibiting risk aversion dislike left skewness and therefore demand a risk premium to bear it. The above results suggest that this premium to bear skewness risk is large.[21]

Moreover, the results are consistent with the main theoretical analysis of

presented above, which were estimated based on the bid stack data.

[21] It is also somewhat variable. The variance in the daily changes in July and August forward prices is on the order of 5 times the variance of the change in daily changes in the corresponding expected spot prices. Due to the rapid mean reversion in load, there is little variation in expected spot prices until the time to delivery becomes very short, and what variation that exists is due primarily to fluctuations in fuel forward prices. Fluctuations in power transactions prices could reflect microstructure effects–the power forward market is relatively illiquid, and hence small transactions volumes can have big price impacts. Relatedly, they may also reflect fluctuations in price bias due to changes in hedging pressure.

risk premia in power markets. Bessimbinder and Lemmon (2002) show that power forward prices should be upward biased when prices are highly right skewed. This reflects hedging pressure. Right skewness induces long hedging pressure and variance induces short hedging pressure. The former effect dominates when prices are highly skewed (during high demand periods). Thus, Bessimbinder-Lemon predict substantial upward biases in seasonal peak periods. We document these effects here.

The decline in upward bias that has occurred in the recent past is also broadly consistent with this theoretical explanation. There was substantial entry of new generating capacity in PJM in the 2000-2002 period. *Ceteris paribus,* the entry of new generating capacity reduces the likelihood of price spikes, which in turn should reduce the degree of bias–this is exactly what we observe. Moreover, Bessembinder-Lemon argue that the lack of integration between power markets and the broader financial market impedes the shifting of risk to speculators who can diversify it away and therefore exacerbates the bias. During the past several years the degree of integration has increased. Several major financial institutions have entered power trading. This has been facilitated in part by the development of financially settled (as opposed to physical delivery-settled) derivatives and the development of other mechanisms (such as "virtual bidding" on PJM day ahead markets) that permit speculators to supply risk bearing services without having to navigate the complexities of making or taking delivery of physical power. Both new capacity and the entry of risk bearing capacity should reduce upward bias, and we have found such a reduction in PJM.

These preliminary empirical results strongly suggest that there is a substantial risk premium in PJM power forward prices for summer delivery. However, the magnitude of the upward bias is extraordinary, so alternative explanations should be considered. For instance, a so-called "peso problem" may be at work here. That is, the forward prices may incorporate expectations of events that did not occur in the historical data employed to determine the terminal boundary conditions and the probability distribution for load. In this case, we may underestimate the expected spot prices and therefore overestimate the market price of risk. However, the persistence of these results over seven years strengthens the case for the upward bias interpretation and undermines the plausibility of the peso problem explanation. Moreover, other investigators (e.g., Ronn-Dincerler (2001), Geman-Vasicek (2001), Bessembinder-Lemon (2002), and Longstaff-Wang (2004)), using a variety of methodologies have presented evidence of upward bias from a variety of power markets. This further suggests that the risk premium explanation is robust.

The finding of such a large risk premium in power forward prices suggests that the power forward market remains incompletely integrated with other financial markets despite the recent entry of speculative risk bearing capacity to the marketplace. Although shocks to power prices can be extremely large, they are plausibly completely idiosyncratic and hence diversifiable. If this were indeed the case, power forward prices would embed no risk premium if the power forward market were completely integrated with the financial markets.

Presumably the existence of large risk premia like those documented here will attract additional speculative capital to the power markets. This should lead to a continued decline in the risk premium. The risk premium may not disappear altogether even if "spike risk" is diversifiable if there are economies of scale in trading. Firms (including trading firms, hedge funds, and other pools of speculative capital) must invest in specialized expertise and systems to trade a particular class of claims (such as power). Scale economies induce some concentration in position size, which given agency costs will impose some deadweight costs on the firms trading power. Moreover, the potential for extreme price moves in power implies that those providing risk bearing capacity to the power market will incur higher moral hazard and monitoring costs than firms providing such capacity to other markets (e.g., the currency or fixed income markets). Thus, although the risk premium in power forwards is likely to decline as the power market becomes more closely integrated with the broader financial market, it is unlikely to disappear altogether.

6 Integrated Valuation of Power Price, Volume, and Weather-Sensitive Claims

The foregoing analysis utilizes load as the demand state variable. It is well known that load is largely determined by weather, especially by temperature. Therefore, it is possible to recast the model using weather as the state variable. This allows unified hedging of derivatives or assets with values that depend on power prices, loads, or weather, or all three.

Formally, define w_t as the value of the weather variable as of time t. The weather variable follows an Ito process:

$$dw_t = \alpha_w(w_t, t)dt + \sigma_w(w_t, t)dz_t \tag{20}$$

where dz_t is a standard Brownian motion. Whereas in our original analysis power spot prices depended on load, we now specify that they depend on this weather variable and calendar time:[22]

$$P_t = g(w_t, t) \tag{21}$$

Moreover, power output as t depends on w_t:

$$q_t = h(w_t, t) \tag{22}$$

Consider the value of any power derivative with a payoff that depends on the spot price of power at $T > t$. The value of this derivative at t can be written as $V(w_t, t, T)$. The value of a derivative with a positive value at t (such as an option) must satisfy the partial differential equation:

$$rV = \frac{\partial V}{\partial t} + \frac{\partial V}{\partial w_t}[\alpha_w - \sigma_w \lambda(w_t, t)] + .5\frac{\partial^2 V}{\partial w_t^2}\sigma_w^2 \tag{23}$$

The term $\lambda(w_t, t)$ is a market price of w_t risk. The expression must include a market price of risk because w_t is not a traded asset, and thus is not a hedgeable risk. For a forward contract, the PDE is:

$$0 = \frac{\partial F}{\partial t} + \frac{\partial F}{\partial w_t}[\alpha_w - \sigma_w \lambda(w_t, t)] + .5\frac{\partial^2 F}{\partial w_t^2}\sigma_w^2 \tag{24}$$

This PDE must be solved subject to boundary conditions. For a daily strike call option on a forward contract with strike price K, the payoff is $V(w_T, T, T) = \max[g(w_T, T) - K, 0]$. For a put option with strike price K, the boundary condition at expiration is $V(w_T, T, T) = \max[K - g(w_T, T), 0]$. For a forward contract, boundary condition is $F(w_T, T, T) = g(w_T, T)$.

[22]The dependence of power prices on the price of fuel is suppressed to simplify the notation.

Some claims may have payoffs that depend on both price and load. For example, the revenue of a utility is given by the product of price and load. Thus, a utility interested in hedging its revenue is interested in the value of a claim with a payoff of $g(w_T, T)h(w_T, T)$. As another example, the value of a power plant depends on both the price at which it sells power and the amount of power the plant generates. Denoting the output of the plant as a function of the weather variable as $x(w_T, T)$, the value of the right to operate the plant and sell power at T is given by the solution of (18) with a payoff at T of $\max[x(w_T, T)(g(w_T, T) - c) - a, 0]$ where c is the (assumed constant) marginal cost of operating the plant, and a is the avoidable, non-output dependent cost of operation (i.e., the no-load cost).[23]

Weather derivatives can be evaluated in an identical framework. In general, a weather derivative is a claim with a payoff that depends on the realization of some weather variable. The most common weather derivatives that have payoffs that depend on temperature over some time period. These include degree day swaps and degree day options. The payoffs to these claims depend on the average of deviations of temperatures at some location from 65° Farenheit over some time period. Since temperature and load are so closely related, some power producers and consumers use weather derivatives to hedge volumetric risks.

To see how to value weather derivatives, call $W(w_t, t, T)$ the price a weather contingent claim with a payoff that depends on w_t that expires at

[23] The analysis is also readily extended to the case of a risky fuel price and "spark spread."

T. This contingent claim must satisfy the following PDE:

$$rW = \frac{\partial W}{\partial t} + \frac{\partial W}{\partial w_t}[\alpha_w - \sigma_w \lambda(w_t, t)] + .5\frac{\partial^2 W}{\partial w_t^2}\sigma_w^2 \tag{25}$$

subject to the appropriate boundary conditions implied by the payoff function for the weather derivative.

The solutions to the PDEs for a power contingent claim and a weather contingent claim generate information necessary to construct a hedge position. Specifically, consider a weather derivative with $\Delta_w = \partial W(w_t, t, T)/\partial w_t$ and a power derivative with $\Delta_v = \partial V(w_t, t, T)/\partial w_t$. A firm holding a long position in the power derivative can hedge this postion against w_t risk by trading $-\Delta_v/\Delta_w$ units of the weather derivative.

Finite difference techniques are especially useful in valuing weather derivatives due to the path-dependent nature of the latter; payoffs to the most common weather derivatives depend on the average temperatures across many days, rather than the temperature on a single day. For example, consider an option with a payoff that depends on the average temperature over some time period $[t_0, t_N]$. This average temperature is given by the integral:

$$I = \frac{1}{t_0 - t_N} \int_{t_0}^{t_N} w(\tau) d\tau \tag{26}$$

It is possible show that the average temperature option must satisfy the following PDE (see Wilmott *et al* 1993):

$$rW(w_t, I, t) = \frac{\partial W}{\partial t} + \frac{\partial W}{\partial w_t}[\alpha_w - \sigma_w \lambda(w_t, t)] + .5\frac{\partial^2 W}{\partial w_t^2}\sigma_w^2 + \delta(t) w_t \frac{\partial W}{\partial I} \tag{27}$$

where $\delta(t) = 1$ for $t \in [t_0, t_N]$, and $\delta(t) = 0$ otherwise. Capturing the state dependence therefore requires increasing the dimensionality of the problem.

However, although this makes the problem more computationally expensive, it is still quite tractible and can be solved using traditional approaches. Now the pricing mesh must have an I dimension in addition to the w and t dimensions. In addition, the solution must satisfy certain "jump" conditions.

7 Summary

The valuation of power contingent claims presents acute difficulties due to the non-storability of power its resultant implications for power price dynamics. Traditional valuation approaches based on specifying a power price process face difficulties due to the unique behavior of power prices. Moreover, since the spot power price is not properly a traded asset (and hence the market is incomplete) due to non-storability, this approach cannot side-step the necessity of estimating a market price of risk.

We take a different approach to valuing power derivatives. Exploiting the transparency of the fundamentals of the power market, we specify that the power price is a function of underlying demand and cost state variables. These state variables are very well behaved, are readily characterized using standard diffusion models, and for the PJM market we study, there is a clear and close relation between these variables and spot power prices. Applying traditional valuation techniques implies that any power contingent claim must obey a partial differential equation that can be solved using traditional finite difference methods. This approach also allows integrated valuation and risk management of power derivatives, weather derivatives, and claims with payoffs that depend on volume.

Implementation of this approach faces one key challenge: since one of the state variables is not a traded asset, it is necessary to take account of its market price of risk. Since the market price of risk is not observable, it is necessary to estimate it. We extract the market price of risk from the prices of traded power claims for the PJM market using inverse methods. The key finding of the article is that this market price of risk was large during 1998-2005 for the PJM market studied, but has declined over this period. For delivery dates during the peak of the 2000 cooling season (July and August) this market price of risk was approximately 50 percent of the forward price. By the 2005 cooling season, likely due to the entry of generating and risk bearing capacity to the market, the upward bias for July delivery had fallen to about 15 percent of the forward price (reaching a maximum of 24 percent of the forward price).[24]

These findings regarding the market price of risk have important implications for valuation problems in power markets. First, they suggest that at present the power markets are not fully integrated with the broader financial markets. Second, they imply that ignoring the market price of risk will lead to substantial errors in valuing any power contingent claim.

This last result is of particular interest in the power markets at present. The industry is currently undergoing a substantial restructuring as result of the deregulation process. As part of this restructuring, market participants are evaluating whether to buy or sell existing generating assets, and whether to invest in new assets. These decisions require valuation of generation ca-

[24]We emphasize again that due to the non-storability of electricity that it will be necessary to incorporate a market price of risk into more traditional valuation approaches based on a specification of the spot power price process due to market incompleteness. Thus, the market price of risk is not an artifact of the particular model we implment; it is an inherent feature of power derivatives pricing.

pacity, which can be viewed as contingent claims. The results of this article imply that those ignoring the market price of risk will make large valuation errors. Given that (1) the market price of risk is not observable, (2) real generating assets have very long lives (extending far beyond the visible portion of the forward curve), and (3) the market price of risk is likely to change over time due to the flow of speculative capital to power market trading, these valuation problems are non-trivial even for those who recognize the importance of power risk premia.

A Econometric Estimation of the Payoff Function

The main text uses generator bid information to determine forward contract payoffs (boundary conditions). For some markets, bid data is not readily available. Moreover, the bid stack approach does impose some economic structure on the problem that may be deemed problematic. An alternative is to utilize econometric techniques to estimate a payoff function. Econometric estimates impose no assumptions about competition in the power market. They only assume that there is a close and stable relation between the state variables and power prices.

As noted in section 4, since numerical solution of the PDE requires discretization of time and load steps. Computational considerations make it desirable to approximate the payoff for a daily forward contract as a function of the load at a single point in the delivery day. Moreover, due to the difficulties of specifying a parametric function that captures the rapid increase in prices that occurs when loads near generation capacity, we use semiparametric techniques to estimate the terminal pricing function using daily data on prices and loads. Therefore, we specify that the relation between average peak price and load is:

$$P(q_j^*, f_j, j) = f_j^{\gamma} \phi(\ln q_j^*) \tag{28}$$

where $P(.)$ is the average 16 hour on peak price on a given day, q_j^* is the load during the selected hour of that peak period, f_j is the spot fuel price at that time, and $\phi(.)$ is an unknown function.

This functional form also captures the traditional way of characterizing generation economics. As noted in section 4, it is conventional to measure the efficiency of generating assets by their "heat rate." The heat rate measures the amount of fuel required (measured in BTU) to generate one megawatt of power. Given a heat rate, the cost of generating a megawatt of electricity is the heat rate multiplied by the price of fuel per BTU. In (13), the $\phi(.)$ function can be interpreted as the market marginal heat rate. Also as noted in section 4, computational considerations also make the functional form of fuel price times some function of load desirable. It permits a reduction of the $2D$ PDE (7) or (8) to a $1D$ PDE, which greatly eases computation. This is of particular importance for the solution of the inverse problem.

Data on the average price from the real time spot market during each hour and average load during each hour are available from the PJM web site. For f_j we use daily data on day-ahead gas prices for Columbia Gas and Texas Eastern Pipeline zone M-3 obtained from Bloomberg. Since the forward price data is for contracts that specify delivery of a constant amount of power per hour during each peak hour of the day, we use the average PJM spot price during the peak period on a given day as the dependent variable in our regression.[25] We use load during the hour ending at 4PM Eastern time on that day as our output variable. The sample period for this analysis is 1 January, 2000-28 December, 2004.[26]

We use a two-step semiparametric procedure to estimate γ and $\phi(\ln q)$. First, we estimate $E(\ln f|\ln q)$ and $E(\ln P|\ln q)$ using a local quadratic regression in $\ln q$ with an Epanechnikov kernel. Note that from (13):

$$[\ln P - E(\ln P|\ln q)] = \gamma[\ln f - E(\ln f|\ln q)]$$

Therefore, to estimate γ, we regress $[\ln P - E(\ln P|\ln q)]$ against $[\ln f - E(\ln f|\ln q)]$. Given the coefficient estimate $\hat{\gamma}$ (which is consistent), we determine $\phi(\ln q)$ by estimating a local quadratic regression with Epanechnikov kernel weights, $P_j/f_j^{\hat{\gamma}}$ as the dependent variable, and $\ln q_j^*$, $(\ln q_j^*)^2$, and a constant as the regressors.[27] Figure 7 demonstrates the relation between 4PM load and peak price along with a scatter of both observed prices (marked with an 'o') and fitted prices (marked with an 'x') against observed loads. It is evident that the semiparametric estimation captures the salient features of the load-price relation. It especially reproduces the rapid increase in prices for small load changes as output nears capacity. The fit is not perfect (the R^2

[25] Peak hours are between 6 AM and 10 PM.

[26] We also estimated a function based on data from the 1997-2003 period. However, a review of this model suggests that large capacity investments in the 2000-2001 period caused the supply curve to shift rightwards, as the model for the 1997-2003 period underpredicts peak prices in the 1997-1999 period and overpredicts peak prices for the 2000-2003 period. Adjustments for installed capacity (as reported by PJM) reduce these problems, but do not eliminate them. Moreover, adjusting loads to reflect scheduled outages (again as reported by PJM) do not substantially increase explanatory power.

[27] The Epanechnikov kernel is $.75(1 - \psi_{t,i}^2)$ for $|\psi_{t,i}| \leq 1$ and 0 otherwise, where $\psi_{t,i} = (\ln q_t - \ln q_i^*)/b_q$. In this expression, $\ln q_i^*$ is the i'th point in the estimation grid and b_q is the bandwidth.

is .68), but this is to be expected because, as noted earlier, the economics of generation imply that there is no unique mapping between the state variables and prices. Nonetheless, simulated market heat rates based on simulations of load filtered through the fitted function behave quite similarly to observed PJM market heat rates.[28]

Using interpolation, it is possible to estimate P for arbitrary $\ln q$. Therefore, this procedure can be used to determine forward contract payoffs when solving the valuation PDE.

B The PJM Market

We apply the equilibrium model to study the pricing of power in the PJM market. It is worthwhile to describe briefly the operation of the PJM market as these details affect the data available to us and thus affect our methodology. The PJM is one of the first markets to implement centralized dispatch and real time market pricing of power. PJM is governed by an independent system operator ("ISO") that dispatches generation to meet power demand.

PJM operates a two settlement system. On the day before power is to be delivered, owners of generation assets submit bids specifying various costs that they must cover to operate (e.g., startup and no-load costs) and per unit charges to deliver energy into the day ahead market–a one-day forward market. Load serving entities specify demand curves. Given these bids from

[28]The existence of the noise term does not impact the estimation of the market price of risk function under certain simplifying assumptions. Specficially, if $P_t = f^{\gamma}_{t,T}\phi(q_t) + \epsilon_t$ (as is assumed in the estimation), and if ϵ_t is unpriced in equilibrium, then the forward price is $F = \tilde{E}_t(f^{\gamma}_{T,T}\phi(q_T))$ where $\tilde{E}_t$ is the expectation under the equivalent measure. The Feynman-Kac Equation implies that this expectation is given by the solution to (9) with initial condition $F(0) = f^{\gamma}_{T,T}\phi(q_T)$.

suppliers and demanders, the ISO chooses which generators will operate at each instant of the next day in order to maximize net surplus; the solution to this "unit commitment problem" ensures that the generation and load schedules do not violate any transmission constraints. The marginal cost of supplying power during each hour establishes the day ahead forward price; due to transmission constraints, there can be different prices at every bus in the PJM system as the marginal cost at each bus may differ. After the close of the day ahead market, generators not scheduled in the day ahead market submit additional multi-part supply bids. Based on its forecast of load, PJM then schedules additional resources to meet expected load and provide reserves for reliablity. This second commitment iteration schedules generation to minimize no load and startup costs.

During the operating day, the ISO optimizes dispatch (based on the bids made the day before) to minimize cost while respecting system reliability and transmission constraints. In essence, given the load at a particular time, the ISO allocates generation to serve this load to maximize total surplus, subject to transmission and system reliability constraints. During real time, the price of power is set equal to the marginal cost of generation implied by these bids. Again, real time prices can differ by location depending on transmission constraints.

In the two settlement system, commitments to buy and sell power established in the day ahead market are settled using the day ahead clearing price. Deviations from commitments established in the day ahead market are settled using the real time price.

In addition to the real time spot market, buyers and sellers of electricity can engage in bilateral transactions in forward markets. These bilateral transactions may be for delivery during the next hour, the next day, or months in the future. Power producers who engage in bilateral transactions "self-schedule" their facilities. That is, they inform the ISO of their dispatch plans so the ISO can utilize this information to respect transmission and reliability constraints when dispatching generation. About 50 percent of the power generated in PJM is bought and sold through the spot market, the remainder through bilateral transactions. Since we employ price data from both the spot and bilateral markets, the fact that each supports a high volume of trade provides some confidence in the reliability of this data.

As noted above, the ISO is also responsible for maintaining the reliability and safety of the generating and transmission system. In the event that reliability or safety is jeapordized due to large loads or the failure of some generating or transmission the ISO can intervene to ensure continued safe operation of the system. For instance, the ISO can curtail loads or order the activation of additional generation in order to achieve this objective.

PJM disseminates information about hourly load and prices (day ahead and real time) via its web site. We utilize this data to calibrate our $\theta_q(t)$, $\mu(t)$, and payoff functions. Sections 3 and 4 describe the calibration methods we employ.

References

[1] Barone-Adesi, G., and A. Gigli. "Electricity Derivatives." Working paper, Universita della Svizzera Italiana (2002).

[2] Bessimbinder, H., and M. Lemmon. "Equilibrium Pricing and Optimal Hedging In Electricity Forward Markets. *Journal of Finance* 57 (2002).

[3] Borenstein, S., J. Bushnell, C. Knittel, and C. Wolfram. "Trading Inefficiencies in California's Electricity Markets." Working paper, University of California Energy Institute (2001).

[4] Eydeland, A. and H. Geman. "Fundamentals of Electricity Derivatives," in R. Jameson, ed., *Energy Modelling and the Management of Uncertainty.* London: Risk Publication, 1999.

[5] Geman, H. *Commodiies and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals, and Energy.* West Sussex: Wiley Finance, 2005.

[6] Geman, H., and A. Roncoroni. "Understanding the Fine Structure of Electricity Prices." *Journal of Business* 79 (2006).

[7] Geman, H., and O. Vasicek. "Forward and Futures Contracts on Non-Storable Commodities: The Case of Electricity." *Risk* (August, 2001).

[8] Green, R. J., and D. Newbery. "Competition in the British Electric Spot Market." *Journal of Political Economy* 100 (1992).

[9] Johnson, B., and G. Barz. “Selecting Stochastic Processes for Modelling Electricity Prices” in R. Jameson, ed., *Energy Modelling and the Management of Uncertainty.* London: Risk Publication, 1999.

[10] Harrison, M. J., and M. Taksar. “Instantaneous Control of Brownian Motion.” *Mathematics of Operations Research* (1983).

[11] Newbery, D. “Power Markets and Market Power.” *Energy Journal* 16 (1995).

[12] Ronn, E., and C. Dincerler. “Risk Premia and Price Dynamics in Power Markets.” Working paper, University of Texas.

[13] Rudkevich, A. and M. Duckworth. “Strategic Bidding in a Deregulated Generation Market: Implications for Electricity Prices, Asset Valuation, and Regulatory Response.” *Electricity Journal* (1998).

[14] Telser, L. G. “Futures Trading and the Storage of Cotton and Wheat.” *Journal of Political Economy* 66 (1958): 233-255.

[15] Tikhinov, A., and V. Arsenin. *Solution of Ill-posed Problems.* New York: Halsted Press, 1977.

[16] Villaplana, P. “A Two-state Variables Model for Electricity Prices.” *Third World Congress of the Bachlier Finance Society, Chicago* (2004).

[17] Wang, A., and F. Longstaff. “Electricity Forward Prices: A High Frequency Empirical Analysis.” *Journal of Finance* 59 2004.

[18] Wilmott, P., J. Dewynne, and S. Howison. *Option Pricing: Mathematical Models and Computation.* Oxford: Oxford University Press, 1993.

[19] Wolak, F., and R. Patrick. "The Impact of Market Rules and Market Structure on the Price Determination Process in the England and Wales Electricity Market." Working paper, Stanford University, 1997.

Figure 1

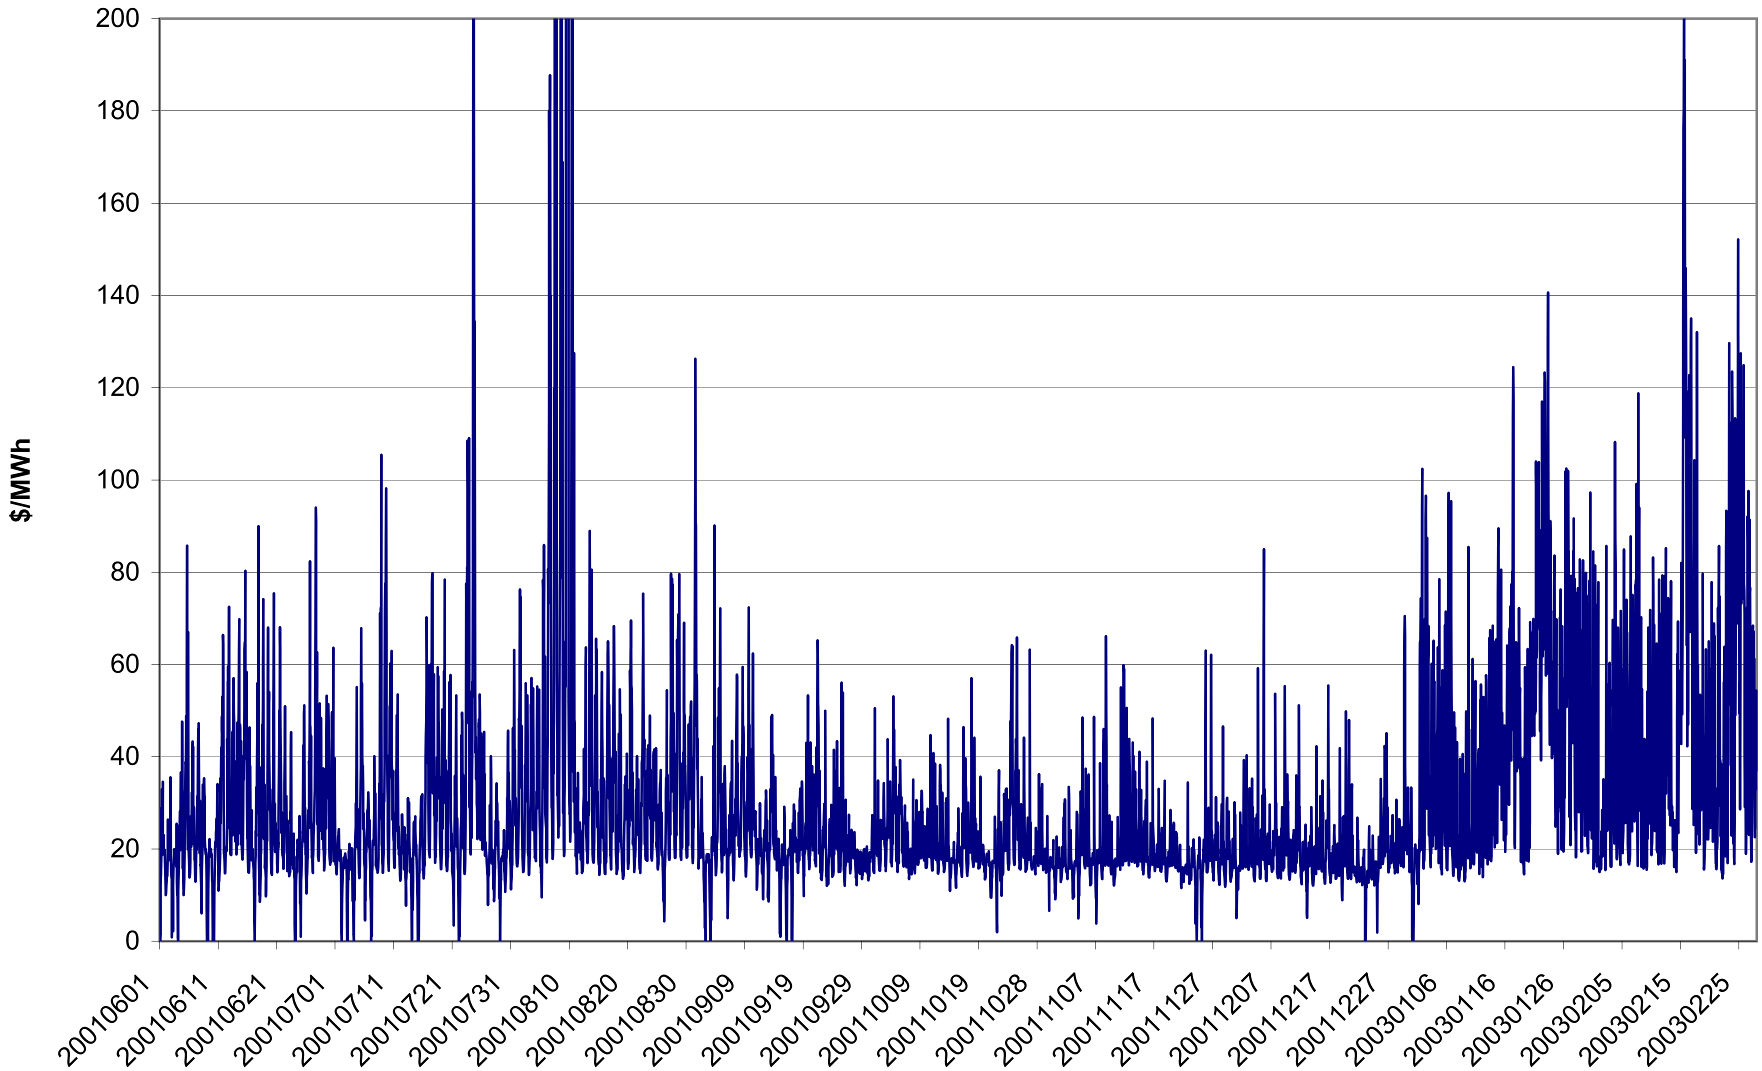

$/MWh
200
180
160
140
120
100
80
60
40
20
0
20010601
20010611
20010621
20010701
20010711
20010721
20010731
20010810
20010820
20010830
20010909
20010919
20010929
20011009
20011019
20011028
20011107
20011117
20011127
20011207
20011217
20011227
20030106
20030116
20030126
20030205
20030215
20030225

Figure 2

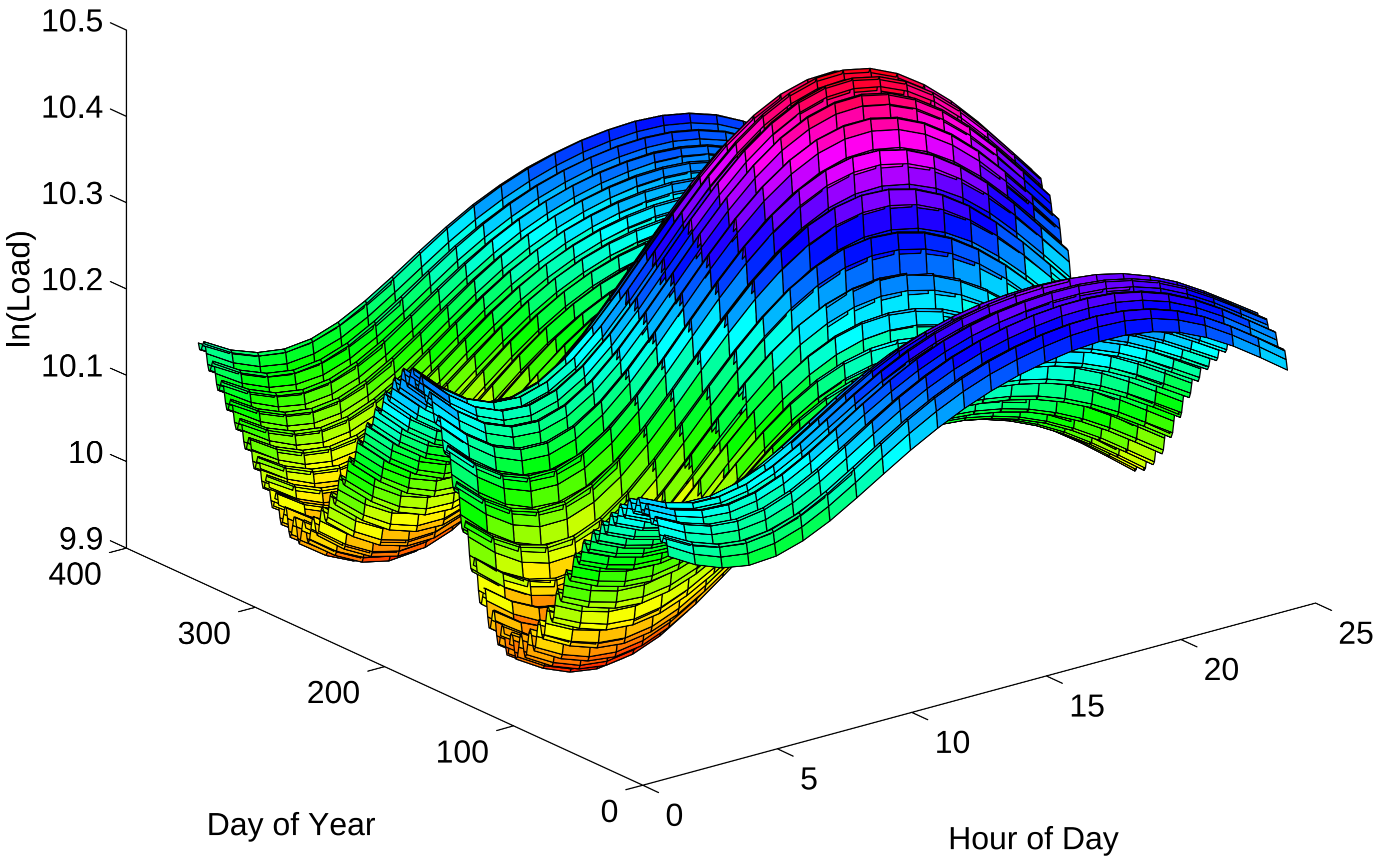

10.5
10.4
10.3
10.2
10.1
10
9.9
ln(Load)
400
300
200
100
0
Day of Year
0
5
10
15
20
25
Hour of Day

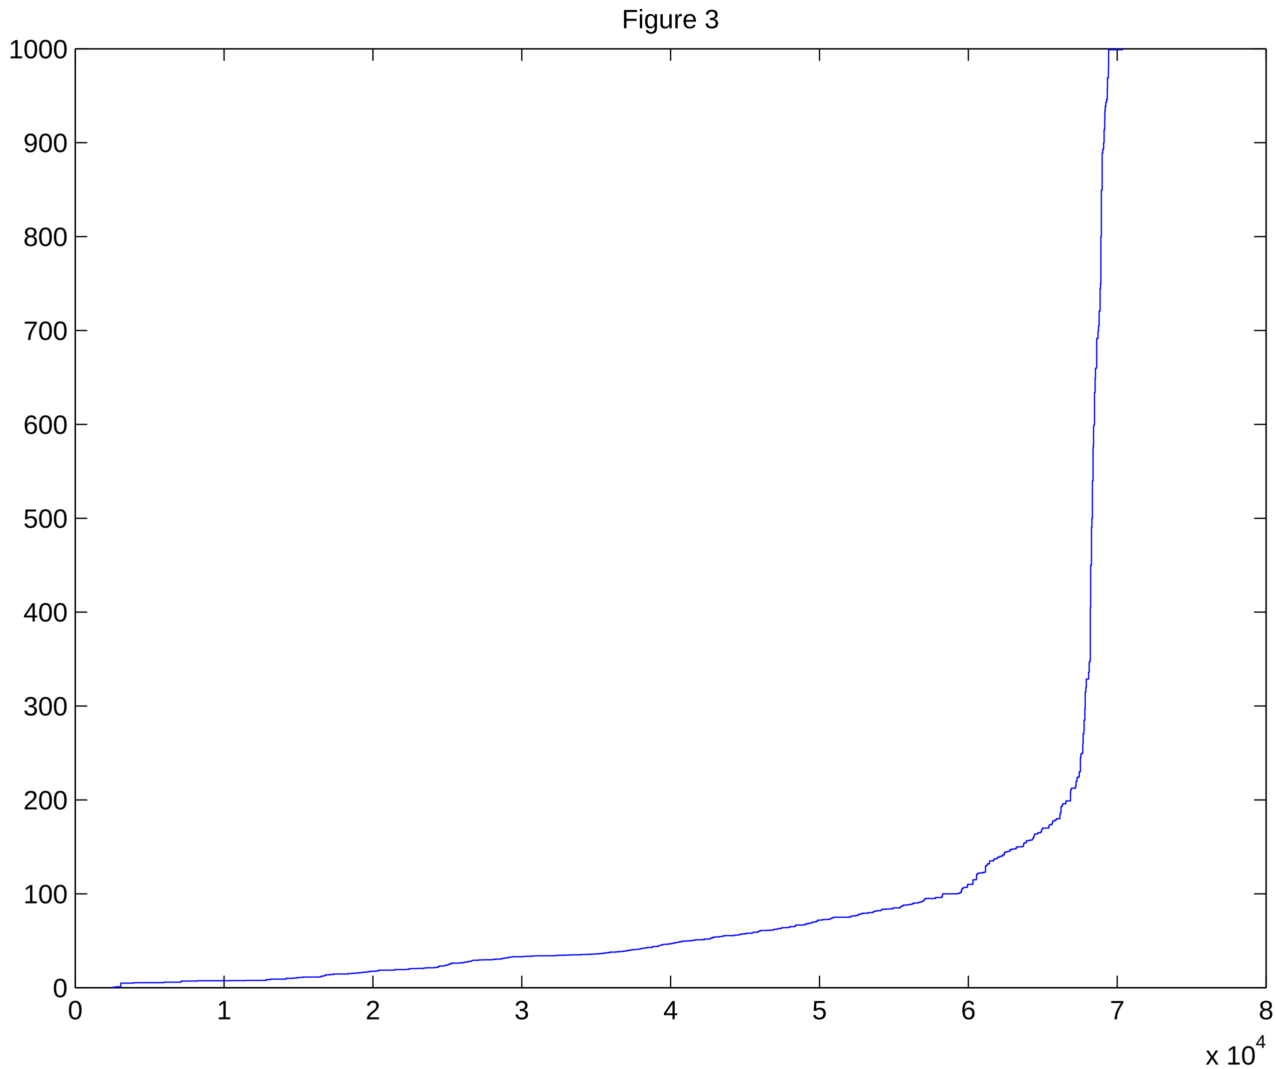
Figure 3
1000
900
800
700
600
500
400
300
200
100
0
0
1
2
3
4
5
6
7
8
x 10^4

Figure 4

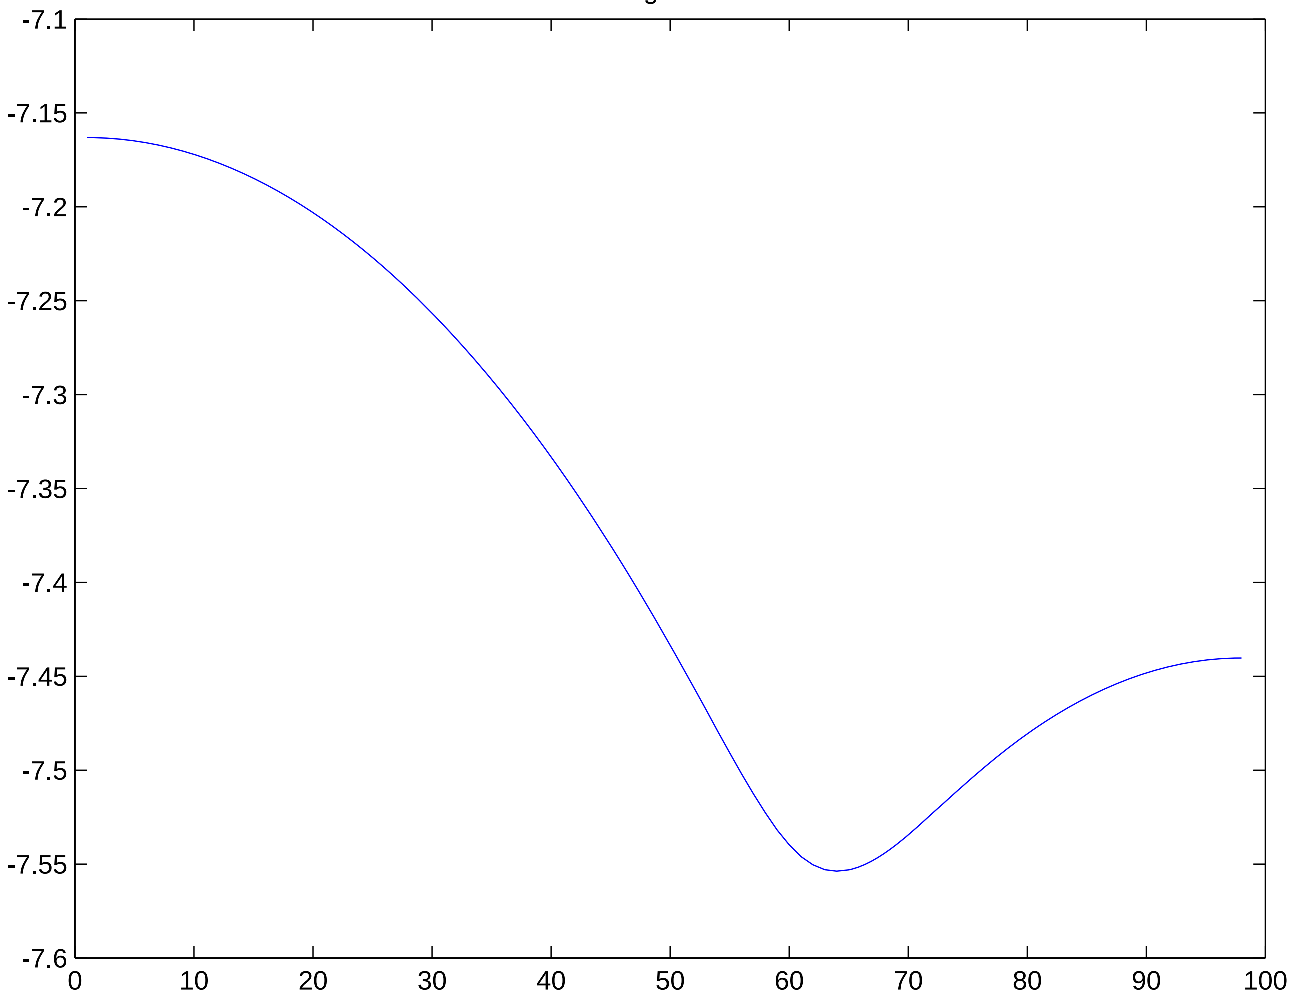

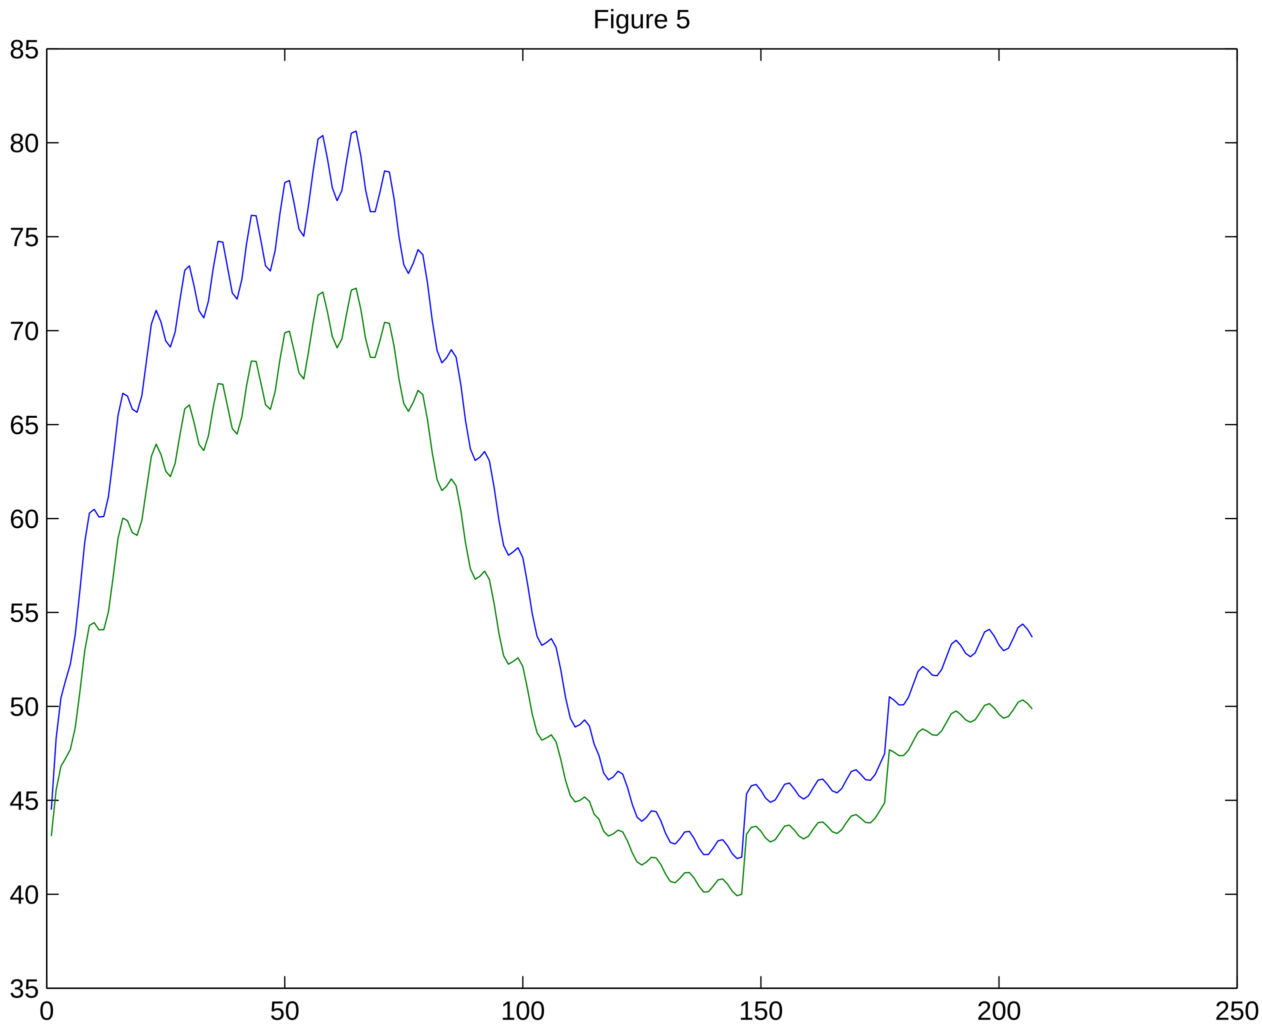
Figure 5
85
80
75
70
65
60
55
50
45
40
35
0
50
100
150
200
250

Figure 6