Supply, Demand, and Risk Premiums in Electricity Markets*

Kris Jacobs Yu Li Craig Pirrong

University of Houston

November 8, 2017

Abstract

We model the impact of supply and demand on risk premiums in electricity futures, using daily data for 2003-2014. The model provides a satisfactory fit and allows for unspanned economic risk not embedded in the futures price. The spot risk premium and forward bias implied by the model are on average large and negative but highly time-varying. Risk premiums display strong seasonal patterns, are related to the variance and skewness of the electricity spot price, and help predict future returns. The risk premium associated with supply constitutes the largest component of the total risk premium embedded in electricity futures.

JEL Classification: G12, G13, Q02

Keywords: Electricity Futures; Economic Determinants; Supply; Demand; Risk Premium; Unspanned Risk.

*Correspondence to Craig Pirrong, Bauer College of Business, University of Houston, craigp@bauer.uh.edu.
We would like to thank our discussant Pradeep Yadav for helpful discussions and comments.

1 Introduction

The modeling of electricity prices and risk premiums in electricity markets is a long-standing research question, but the existing literature is relatively limited. Existing studies (Pirrong and Jermakyan (2008); Cartea and Villaplana (2008)) provide empirical evidence that risk premia depend on demand and supply variables. Another strand of the literature, commencing with Lucia and Schwartz (2002), applies no-arbitrage techniques in models with latent variables to price electricity futures. This approach explicitly distinguishes between the physical and risk-neutral model dynamics, and therefore allows for the estimation of risk premia. Heretofore, these literatures have developed independently.

This paper contributes to the understanding of the pricing of electricity derivatives, and hence of electricity risk premia, by integrating these two approaches. Specifically, we estimate a no-arbitrage model that provides a good fit to electricity futures prices, while also quantifying the impact of supply and demand variables on these prices. The model also allows for unspanned economic risk, which is risk captured by supply and demand variables but not identified by the futures prices. We use this model to estimate risk premiums embedded in electricity futures and study their characteristics and implications. The model allows a decomposition of risk premiums into several components, including the components due to supply and demand variables. Our empirical analysis reveals several new findings.

We first document that economic variables contain useful information about the risk premiums in electricity futures. After controlling for the information in the electricity futures curve, economic variables such as the natural gas price, load, and temperature have incremental forecasting power for future spot rates and returns on electricity futures. Second, while the supply and demand variables contain additional information on electricity risk premiums, the principal components of the futures curve summarize the majority of the information on futures prices. Specifically, we show that a model based on the first two principal components of the futures curve provides a good fit of the entire electricity futures curve.

Third, the estimated spot risk premium in the unspanned model is negative and very large. It is on average -0.84 percent per day, but it is highly time-varying and exhibits very large negative and positive outliers. For instance, during the 2014 polar vortex the spot risk premium for the unspanned model fluctuates between -50 percent and 170 percent per day.

Fourth, the spot risk premium implied by the model displays strong seasonal patterns. It is much larger (more negative) in the peak demand seasons of winter and summer. The spot risk premium is positively correlated with the volatility of the spot price and negatively correlated with the skewness of the spot price, consistent with the model of Bessembinder and Lemmon (2002). These results are consistent with the insight of Pirrong and Jermakyan (2008) that electricity futures prices incorporate a premium to compensate for the risk of price spikes that are more likely during peak demand periods, or when costs spike due to fuel price shocks.

Fifth, we find that unspanned economic risk associated with supply is the most important component of the spot risk premium on electricity futures. The estimated spot risk premium in the unspanned model is very different from the one implied by the spanned model and it provides better forecasts of future spot prices and returns on electricity futures compared to the risk premium of models that ignore this unspanned risk.

Sixth, the forward bias is also negative on average, implying that forward prices exceed expected spot prices. The average forward bias ranges from -$4 for the one-month maturity to -$7 for the twelve-month maturity, and is highly time-varying regardless of the maturity of the contract, but with larger fluctuations and outliers for shorter-maturity forwards. For instance, the day-ahead forward bias reaches a maximum of $340 and a minimum of -$80 during the polar vortex period. For longer maturities, the forward bias is much larger than the sample average for an extended period between 2006 and mid-2008.

What is the economic meaning of the large risk premia we find in these markets? The most likely explanation is that the risk premia are caused by barriers to the entry of risk bearing capital into these markets (Hirshleifer (1988); Bessembinder and Lemmon (2002)). The finding that the spot premium depends on the variance and skewness of the spot price, as predicted by the model of Bessembinder and Lemmon, is also consistent with such restrictions. This may also suggest that electricity markets are not fully integrated with the broader financial markets.

This paper is related to several strands of literature. An important literature uses reduced-form no-arbitrage models with latent variables to price electricity futures (see, for example, Lucia and Schwartz (2002); Cartea and Figueroa (2005); Deng and Oren (2006); Geman and Roncoroni (2006); Benth, Cartea, and Kiesel (2008); and Geman (2009)). Our proposed model nests this class of models but augments them with economic supply and demand variables. We find that the economic variables are important in explaining the risk premium associated with the electricity futures.

Another literature uses a more structural approach to price electricity futures. These papers use a bottom-up approach by first specifying the dynamics for supply and demand variables and then derive the spot price as a function of those variables. This approach is more intuitively appealing because it exploits the information contained in supply and demand variables suggested by economic theory (see, for example, Pirrong and Jermakyan (2008); Cartea and Villaplana (2008); and Pirrong (2011)). We show that while this approach is economically appealing and while the economic variables are important for explaining the risk premium, latent factors significantly improve model fit. We also demonstrate that it is critical to model the the supply and demand variables as unspanned.

Because our model contains supply and demand variables, it is also related to the literature which develops equilibrium models to study the determinants of the risk premium of electricity futures (Bessembinder and Lemmon (2002); Longstaff and Wang (2004); Dong and Liu (2007); Douglas and Popova (2008); Bunn and Chen (2013)). Finally, several related papers emphasize the importance of economic variables for modeling the risk premium of commodity futures (see, for example, Khan, Khokher, and Simin (2016); and Heath (2016)).

The remainder of the paper proceeds as follows. Section 2 describes the data and provides a discussion of the economics of electricity markets. Section 3 outlines the model specification and estimation. Section 4 discusses the estimation results and Section 5 discusses the model's implications for risk premiums. Section 6 concludes.

2 Electricity Markets

We estimate the model using electricity data for the PJM (Pennsylvania-New Jersey-Maryland) Western Hub market. We now discuss the institutional features of the PJM market, the electricity futures prices and returns we use in the empirical analysis, and the economic demand and supply variables used to explain these prices.

PJM is a "Regional Transmission Organization" that operates centralized day-ahead and real time markets for electricity. Operators of generation assets submit offers to the RTO that indicate the amount of power they are willing to generate as a function of price the day prior to the operating day. Consumers of electricity ("load") submit bids to purchase electricity, where bids can vary by time of day. The RTO aggregates the generation offers to construct a supply curve, and uses the bids to construct a demand curve. For each hour of the operating day, the RTO sets the *day-ahead forward price* equal to that which clears the market, i.e., sets quantity supplied equal to quantity demanded.

In reality, things are somewhat more complicated due to the fact that production and consumption of electricity are spatially dispersed, and there can be rather complex constraints on transmitting power over distance to move from generators to load. Based on the generation offers and load bids, the PJM RTO solves a constrained optimization program that maximizes the sum of consumer and producer surpluses, subject to the transmission constraints. The RTO sets the day-ahead forward prices for each transmission constraint location in the network equal to the shadow prices associated with that constraint produced by the solution to this optimization problem.

In real time, electricity demand can vary randomly, and differ from the amount forecast the day before, which is represented by the bids. Operation in real time requires exact balancing of generation and load, and must respect transmission constraints. As load varies over time and across the PJM region, the RTO dispatches generation to ensure the system remains in balance. To optimize dispatch, the RTO solves the surplus maximization constrained optimization problem, and sets market clearing *spot price* equal to the relevant shadow price in this optimization problem.

In addition to the day-ahead and real-time markets for physical energy, there are derivatives markets on PJM electricity. In particular, there are cash-settled *futures contracts* on PJM electricity. One such contract is the Peak PJM Western Hub Real Time contract. This contract has a payoff based on the arithmetic average of the PJM Western Hub market clearing real time price for each peak hour (8AM-11PM) of the contract calendar month. The notional quantity in this contract is 2.5 megawatts (MW). This contract is traded on the CME.

In our empirical analysis, we use the average real-time peak hour spot and day-ahead average peak hour prices in the PJM Western Hub market, and the prices of PJM Western Hub real-time peak calendar-month 2.5 MW futures. The real-time and day-ahead price are downloaded from the PJM website.[1] We model the day-ahead price as a short-term futures contract which matures in one day. Data on the PJM Western Hub real-time peak calendar-month 2.5 MW futures contracts are obtained from the CME. We include futures contracts with maturities of 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, and 12 months. The data frequency is daily. Each day, the sample therefore consists of thirteen futures prices. The sample period is from May 1, 2003 to May 30, 2014.

[1] See `https://dataminer.pjm.com/dataminerui/pages/public/lmp.jsf`.

PJM futures markets are quite liquid. Total open interest for PJM futures contracts as of the end of our sample period was 8,358,662 contracts, on three different exchanges, of which 76 percent was held by long commercials and 83 percent by short commercials.

The challenges in modeling these markets are apparent from Figure 1 and Table 1. Panel A of Figure 1 plots the time series of the daily spot prices. Panel A of Table 1 indicates that the average spot price over the sample is $58.16, but the fluctuations around this mean are enormous, with a minimum price of $21.21 and a maximum price of $757.42 over the sample period. The first row of Panel B of Table 1 reports descriptive statistics for log returns. The standard deviation of the daily log return is 30 percent, close to the 34 percent reported by Bessembinder and Lemmon (2002) for percentage returns in the 1997-2000 sample period. For comparison, the standard deviation for daily log returns on the $S\&P500$ over our sample is half a percent.

Panel A of Table 1 indicates that for futures with a maturity of more than one month, the futures price on average exceeds the spot price. The price of the twelve-month futures contract is $61.24 on average, or on average $3.08 higher than the spot price. The differences in higher moments are larger. Compare the time series of the daily twelve-month futures price in Panel B of Figure 1 with the spot price in Panel A. The future price also fluctuates considerably, but these fluctuations are much smaller, resulting in a smaller standard deviation. An even more important difference is in the fourth moment. The much lower kurtosis of the twelve-month contract is clearly visible in Figure 1. The spot price in Panel A is very jagged and the futures price in Panel B is much smoother.

The spot price in Panel A of Figure 1 is characterized by very sharp spikes, with a maximum of $757.42 during the time of the polar vortex, and other large spikes in 2005, 2006, and 2008. In contrast, the maximum value of the twelve-month futures contract is $138.39, which occurs in 2008. The maximums for the spot and twelve-month futures prices therefore occur at different times. Panel C of Figure 1 plots the difference between the spot price and the twelve-month futures price. We report monthly averages in Panel C, because for daily differences the extreme observations completely dominate the figure, and as a result it is not informative. As indicated in Table 1, the difference is negative on average, but Panel C indicates that it is also often positive. Similar observations apply to the other futures contracts.

Panel A of Table 1 presents descriptive statistics on prices, but when we report on risk premiums we are effectively using (log) returns. Panel B of Table 1 therefore reports descriptive statistics for the relevant log returns, which have very different statistical properties. We report on the log return on the spot contract, the log return on the day-ahead contract, and the natural logarithm of the ratio of the spot price at t+1 and the day-ahead contract at t. Note that in our sample, the electricity spot price at the end of the sample is lower than at the start of the sample, which gives a negative average log spot return. The most important observation is that the spot price as well as the day-ahead price are characterized by large positive kurtosis, but for the log returns the kurtosis is much smaller. Also, instead of the large positive skewness in prices, the skewness in returns is small. These results are partly due to the difference between returns and prices, and partly due to the use of log returns, because the logarithms effectively reduce the impact of outliers.

Seasonalities are extremely important in electricity markets. We follow the existing literature and de-seasonalize the electricity prices as well as the economic variables. The de-seasonalization method is discussed in Section 3 below. Panels A-D of Figure 2 plot the raw price, the seasonal component, and the de-seasonalized price for the spot price, the day-ahead contract and the 6-month and 12-month futures contracts. In order to better highlight the seasonalities we plot monthly averages rather than daily prices, which contain much high-frequency variation that is irrelevant for illustrating seasonalities. The seasonal patterns in the price data are readily evident from Figure 2.

The economic data include demand and supply variables. Following Pirrong and Jermakyan (2008), we use the natural gas price (PX) as the supply variable. We utilize the price of natural gas as the supply variable because gas-fired generating units usually produce the marginal megawatt, and hence the price of gas is a primary determinant of the marginal cost of production, given that capacity is fixed in the short run. We obtain daily spot natural gas middle prices for Columbia Gas and Texas Eastern Pipeline zone M-3 from Bloomberg. The demand variable is either a load or a temperature variable. For load, we obtain both the daily average and the maximum load for the PJM Western Hub market from the PJM website.[2] The temperature data is from the National Climatic Data Center (NCDC). We first get the daily average, maximum, and minimum temperature for Washington, D.C. and Pittsburgh. Then we calculate the average temperature of the two cities and use it as a temperature proxy for the PJM Western Hub market. We also compute the cooling degree days (CDD) and heating degree days (HDD) according to the weather derivatives literature (see for example Alaton, Djehiche, and Stillberger (2002) and Jewson and Brix (2005)).[3] Table 2 reports summary statistics for the supply and demand variables. Figure 3 plots the time series of the natural gas price, the maximum load, and the CDD. We plot the raw as well as the de-seasonalized series. Again the seasonal patterns in the economic data are readily evident from Figure 3.

In the model, we always use one supply variable and one demand variable. By combining the single supply variable with the seven different demand variables (average load, maximum load, maximum temperature, minimum temperature, average temperature, CDD, and HDD), we obtain seven different combinations of supply and demand variables. Empirical results for these seven models are very similar. In the empirical section below, we report on the model with the natural gas price and CDD as the benchmark model. When there are no space constraints, we report on two models: the first one uses the natural gas price and CDD and the second one uses the natural gas price and maximum load. Other results for the natural gas price and maximum load are relegated to the Online Appendix, and the Online Appendix also reports some summary results for the other five demand variables. Note that we choose a temperature variable as our benchmark demand variable instead of load because the time series for maximum load contains a structural break in

[2] See http://www.pjm.com/markets-and-operations/ops-analysis/.

[3] CDD is defined as max(Average Temperature - 18, 0). HDD is defined as max(18 - Average Temperature, 0). Note that we use 18 as the reference temperature because the temperature is expressed in degrees Celsius.

2004, which is due to a geographical enlargement of the PJM market. In spite of this, the empirical results are similar when we use the load variable instead of the temperature variable.

It is instructive to compare the patterns in the spot and futures data in Figure 2 with those in the economic variables in Figure 3. For the supply variable in Panel A of Figure 3, the natural gas price, it can be clearly seen that the spike in the natural gas price at the start of 2014 is accompanied by a large spike in the spot and day-ahead prices, but a much smaller increase in the twelve-month futures price. On the other hand, increases in the natural gas price in 2005 and 2008 are accompanied by increases in spot as well as futures prices in Figure 2. It is less obvious to detect relations between the demand variables, load and temperature, and the price data, partly because the raw data contain such strong seasonalities. Both for the load variable in Panel B and the CDD variable in Panel C, the deseasonalized data contain small spikes in 2011 and 2012, but these are not accompanied by large price increases in Figure 2.

3 Models For Electricity Futures

This section presents three different models of electricity futures prices. We first outline a general affine framework which nests these three models. We then discuss the unspanned model, the spanned model with latent variables, and the spanned model with economic variables.

3.1 A Class of Affine Models

We outline a class of affine models, which nests the models that we investigate in our empirical work. Suppose that there are N state variables that fully determine the state of the electricity market. These variables can be latent variables or economic (demand and supply) variables. Generally denote this vector of state variables by X. We assume

X follows a Gaussian VAR under the P measure, where the P-dynamic of X is denoted as follows:

$$X_{t+1} = \text{Seas}_{X,t+1} + K_0^P + K_1^P \times X_t + \Sigma^P \times \epsilon_{t+1}^P \tag{1}$$

where X_t is the state vector at time t, $\text{Seas}_{X,t}$ is an N by 1 vector denoting the seasonal component of the state variables, K_0^P is an N by 1 vector, K_1^P is an N by N matrix, Σ^P is an N by N upper triangular matrix, and ϵ_t^P is an N by 1 vector of independent Brownian motions.

The stochastic discount factor (SDF) is assumed to be of the following form:

$$\text{SDF}_{t+1} = e^{(\Lambda_0 + \Lambda_1 \times X_t)' \times \epsilon_{t+1}} \tag{2}$$

where Λ_0 is an N by 1 vector and Λ_1 is an N by N matrix.

Given these assumptions, we have the following dynamic of the state variables under the risk-neutral measure Q:

$$X_{t+1} = \text{Seas}_{X,t+1} + K_0^Q + K_1^Q \times X_t + \Sigma^Q \times \epsilon_{t+1}^Q \tag{3}$$

where K_0^Q is an N by 1 vector, K_1^Q is an N by N matrix, Σ^Q is the upper left N by N matrix of Σ^P, and ϵ_t^Q is an N^Q by 1 vector of independent Brownian motions.

As in the log price model in Lucia and Schwartz (2002), we assume that the natural logarithm of the electricity spot price is a linear function of the state variables. Denoting the natural logarithm of the spot price S_t at time t as s_t, this gives:

$$\text{s}_t = \text{Seas}_{s,t} + \rho_0 + \rho_1 \times X_t \tag{4}$$

where ρ_0 is a scalar, ρ_1 is an 1 by N matrix, and $\text{Seas}_{s,t}$ is a scalar denoting the seasonal component of the log spot rate. This seasonal component is a scaled version of the seasonal component of the state vector.

Based on equation (4), futures prices can be derived recursively. Denoting the log price of the futures contract with maturity j at time t as f_t^j, we can show that f_t^j is given by

$$f_t^j = \text{Seas}_{f,t+j} + A_j + B_j \times X_t \tag{5}$$

where $\text{Seas}_{f,t+j}$ denotes the seasonal component of the forward contract with maturity $t+j$ and

$$A_j = A_{j-1} + B_{j-1}K_0^Q + \frac{1}{2}B_{j-1}\Sigma^Q\Sigma'^Q B'_{j-1} \tag{6}$$

$$B_j = B_{j-1}(I_{N^Q} + K_1^Q) \tag{7}$$

$$A_0 = \rho_0 \quad \text{and} \quad B_0 = \rho_1 \tag{8}$$

3.2 The Unspanned Model

We now assume that there are N^S state variables that fully determine the price of the electricity futures. Denote the vector of those state variables as X^S. The unspanned model assumes that the information in the futures price can only span part of the information in the economy. Denote the part that cannot be explained, or the unspanned part, by US_t, and rewrite X_t as follows.

$$X_t = \begin{bmatrix} X_t^S \\ \text{US}_t \end{bmatrix} \tag{9}$$

where $X_t^S \cup \text{US}_t = X_t$ and $X_t^S \cap \text{US}_t = \varnothing$. Substituting equation (9) into equation (1), we get the P-dynamic of the unspanned model.

$$\begin{bmatrix} X_{t+1}^S \\ \text{US}_{t+1} \end{bmatrix} = \text{Seas}_{X,t+1} + K_0^P + K_1^P \times \begin{bmatrix} X_t^S \\ \text{US}_t \end{bmatrix} + \Sigma^P \times \epsilon_{t+1}^P \tag{10}$$

In these models, the variables can be rotated, which means that we can re-define an equivalent model that is written in terms of different variables. In our empirical work, we rotate the unspanned part of economic variables to the economic variables EC_t themselves in order to provide a more intuitive interpretation of the estimated coefficients. We therefore estimate the following version of the unspanned model:

$$X_t = \begin{bmatrix} X_t^S \\ \mathrm{EC}_t \end{bmatrix} \tag{11}$$

$$\begin{bmatrix} X_{t+1}^S \\ \mathrm{EC}_{t+1} \end{bmatrix} = \mathrm{Seas}_{X,t+1} + K_0^P + K_1^P \times \begin{bmatrix} X_t^S \\ \mathrm{EC}_t \end{bmatrix} + \Sigma^P \times \epsilon_{t+1}^P \tag{12}$$

$$X_{t+1}^S = \mathrm{Seas}_{X^S,t+1} + K_0^Q + K_1^Q \times X_t^S + \Sigma^Q \times \epsilon_{t+1}^Q \tag{13}$$

$$\mathrm{SDF}_{t+1} = e^{(\Lambda_0+\Lambda_1\times X_t)'\times\epsilon_{t+1}} \tag{14}$$

$$\mathrm{s}_t = \mathrm{Seas}_{s,t} + \rho_0 + \rho_1 \times X_t^S \tag{15}$$

Joslin, Priebsch, and Singleton (2014) show that under certain assumptions, one can use principal components (PCs) of the futures data to estimate the unspanned model. Moreover, they show that it is possible to obtain consistent estimates of the P- and Q-parameters by breaking up the estimation problem in two parts. We follow Joslin, Priebsch, and Singleton (2014) and use the PCs of the electricity futures prices as the state variables under the risk neutral measure Q. We augment the PCs with economic variables to get the state vector under the physical measure P. Because the PCs and the economic variables are both observed, we can use a vector autoregressive approach to estimate the physical dynamic given in equation (12). Subsequently, we estimate the Q parameters in equation (13) by minimizing the root mean squared error based on the difference between observed futures prices and model prices.

3.3 The Spanned Model with Latent Factors

To highlight the importance of the unspanned relation between the demand and supply variables and the latent variables, we consider two alternative models which remove the unspanned economic component. The first model removes the unspanned economic component by dropping the economic variables. We refer to this model as the spanned model with latent factors. In this model with latent factors, the state variables under both the P- and Q-dynamics are equal to X^S. The dynamics for this model are:

$$X^S_{t+1} = \text{Seas}_{X^S,t+1} + K^P_0 + K^P_1 \times X^S_t + \Sigma^P \times \epsilon^P_{t+1} \tag{16}$$

$$X^S_{t+1} = \text{Seas}_{X^S,t+1} + K^Q_0 + K^Q_1 \times X^S_t + \Sigma^Q \times \epsilon^Q_{t+1} \tag{17}$$

$$\text{SDF}_{t+1} = e^{(\Lambda_0+\Lambda_1 \times X^S_t)' \times \epsilon_{t+1}} \tag{18}$$

$$\text{s}_t = \text{Seas}_{s,t} + \rho_0 + \rho_1 \times X^S_t \tag{19}$$

This model belongs to a class of reduced-form models that only use latent factors to price futures. In this class of models, Lucia and Schwartz (2002) propose two models that are based on the log power price. Either of these models can be seen as a special case of the model in this section. To benchmark the performance of our models, we therefore also estimate the two-factor model of Lucia and Schwartz (2002). Details on the specification and estimation of this model are given in the Appendix.

We estimate the spanned model with latent factors using a method very similar to the one used for the unspanned model. First, use the PCs of the futures curve to estimate the P-parameters in equation (16) using a vector autoregressive approach. Subsequently the Q parameters in equation (17) are estimated by minimizing the root mean squared error.

3.4 The Spanned Model with Economic Variables

Another special case of the unspanned model is a model without latent variables. In this case, the state variables under the physical measure only consist of economic variables.[4] The same economic variables are also the state variables under the risk-neutral measure Q and thus fully determine the prices of electricity futures. The futures prices fully span the economy, and conversely the economic variables are fully spanned by the electricity futures. We refer to this model as the spanned model with economic variables. The resulting model is related to the framework of Pirrong and Jermakyan (2008) and Pirrong (2011), who exclusively use demand and supply variables to price electricity futures. In summary, this model is given by:

$$\mathrm{EC}_{t+1} = \mathrm{Seas}_{EC,t+1} + K_0^P + K_1^P \times \mathrm{EC}_t + \Sigma^P \times \epsilon_{t+1}^P \tag{20}$$

$$\mathrm{EC}_{t+1} = \mathrm{Seas}_{EC,t+1} + K_0^Q + K_1^Q \times \mathrm{EC}_t + \Sigma^Q \times \epsilon_{t+1}^Q \tag{21}$$

$$\mathrm{SDF}_{t+1} = e^{(\Lambda_0 + \Lambda_1 \times \mathrm{EC}_t)' \times \epsilon_{t+1}} \tag{22}$$

$$\mathrm{s}_t = \mathrm{Seas}_{s,t} + \rho_0 + \rho_1 \times \mathrm{EC}_t \tag{23}$$

The economic variables are observed and thus we can estimate the P dynamic in equation (20) using a vector autoregressive approach. Then, we use the economic variables as the state variables under Q and we estimate the Q dynamic in equation (21) by minimizing the dollar root mean squared errors.

[4] Strictly speaking, it is incorrect to refer to this model as being nested by the unspanned model. In the unspanned model, the state variables under P consist of the unspanned part of the economic variables, whereas in a spanned model with economic variables, the state variables are the economic variables themselves. We can refer to the unspanned part of the economic variables as the economic variables due to the rotation.

3.5 Modeling the Seasonal Component

We specify the seasonal component of the log of the electricity spot price and the economic variables following Lucia and Schwartz (2002). For instance, for the log spot rate:

$$\text{Seas}_{s,t} = \beta_1 \times M_1(t) + \beta_2 \times M_2(t) + ... + \beta_{12} \times M_{12}(t) \tag{24}$$

where $M_i(t), i = 1, 2, ..., 12$ are monthly dummies. For example, $M_1(t)$ is defined as follows.

$$M_1(t) = \begin{cases} 1, & \text{if t is in January} \\ 0, & \text{otherwise} \end{cases} \tag{25}$$

The other $M_i(t), i = 2, 3, ..., 12$ are defined similarly.

Following Lucia and Schwartz (2002), we first use OLS to estimate the following regression to get the seasonal component of both the log spot price and the economic variables.

$$\begin{aligned} \text{s}_t &= \beta_1 + \beta_2 \times M_2(t) + \beta_3 \times M_3(t) + ... + \beta_{12} \times M_{12}(t) + \epsilon_t \\ \text{EC}_t &= \gamma_1 + \gamma_2 \times M_2(t) + \gamma_3 \times M_3(t) + ... + \gamma_{12} \times M_{12}(t) + \epsilon_t \end{aligned} \tag{26}$$

Then we de-seasonalize the log spot price and the economic variables and obtain the corresponding de-seasonalized series.

$$\begin{aligned} \text{De-Seasonalized s}_t &= \text{s}_t - (\hat{\beta}_1 + \hat{\beta}_2 \times M_2(t) + ... + \hat{\beta}_{12} \times M_{12}(t)) \\ \text{De-Seasonalized EC}_t &= \text{EC}_t - (\hat{\gamma}_1 + \hat{\gamma}_2 \times M_2(t) + ... + \hat{\gamma}_{12} \times M_{12}(t)) \end{aligned} \tag{27}$$

The de-seasonalized log futures prices are obtained by adjusting the raw log futures prices with the value of the seasonal component of the spot rate in the month when the futures mature. The definition of de-seasonalized futures price is thus as follows.

$$\text{De-Seasonalized } f_t^j = f_t^j - (\hat{\beta}_1 + \hat{\beta}_2 \times M_2(t+j) + ... + \hat{\beta}_{12} \times M_{12}(t+j)) \tag{28}$$

We use the de-seasonalized series in equations (27) and (28) to estimate the model parameters.

This de-seasonalization approach deserves some comment. Theory suggests that the risk premiums in futures prices, and hence futures prices themselves, depend on the likelihood and magnitude of price spikes (Bessembinder and Lemmon (2002); Pirrong and Jermakyan (2008)). Furthermore, the likelihood of price spikes is seasonal because spikes are more likely to occur when capacity utilization is high, which is most likely during seasonal demand peaks that occur in the summer and winter in the United States. Thus, risk premia are likely to be seasonal.

Deseasonalizing futures prices themselves using standard techniques would make it impossible to detect any such seasonality in risk premia. The approach we implement quantifies the seasonality in the expectation of the spot price under the physical measure, and by removing this seasonal component we can identify seasonalities in the risk premium.

4 Model Estimates

We first establish that electricity futures prices can be adequately summarized by their first two principal components (PCs). Then we show that the demand and supply variables contain additional information beyond the PCs, suggesting that they are unspanned by the electricity futures. We then estimate the model, discuss the fit and economic implications of the unspanned model, and compare it with other models. We also discuss the implications of the spanning assumption.

4.1 The Information in the Principal Components and the Economic Variables

We investigate if the supply and demand variables are spanned by the electricity futures. To this end we first need to parsimoniously represent the information in the electricity futures. We use principal component analysis to analyze the electricity futures curve. Figure 4 shows the loadings of the first two principal components (PCs) and the fraction of total variance explained by each PC. The first two PCs explain more than 92% of the total variation of the price of electricity futures. We therefore conclude that most information in the electricity futures curve can be adequately summarized by the first two PCs.

The interpretation of these two PCs is similar to that of yield curve PCs. The loading on the first PC is virtually identical for all maturities from one day to 12 months as well as the spot, meaning that this first component affects the prices of all maturities similarly, and therefore causes parallel shifts in the forward curve; this is a level effect. The loading of PC^2 is large and positive for short maturities, and negative and relatively small (in absolute value) for longer maturities. Thus, this PC basically drives the slope of the forward curve. The time series of PC^1 is therefore very similar to Panel B of Figure 1. The time series of PC^2 is highly correlated with Panel C of Figure 1, but note that Panel C of Figure 1 reports monthly averages. The time series of PC^2 therefore contains much more short-term variation.

We next verify if the PCs can span the supply and demand variables. We run the following regression:

$$\mathrm{EC}_t = \gamma_0 + \gamma_{\mathrm{pc}} \mathrm{PC}_t^{1-5} + \mathrm{unEC}_t \tag{29}$$

Equation (29) projects the demand and supply variables on the first five PCs of the electricity futures curve. If the economic variable is spanned, we expect a high explanatory power of PCs for the demand and supply variables, and therefore a high adjusted R^2.

Panel A of Table 3 reports the regression results for the natural gas price (PX), the maximum load, and the temperature. The adjusted R^2 is approximately 44% for PX, 11% for maximum load, and 29% for CDD. Thus, the first five PCs at most explain half of the variation of the economic variables. Note that this is not due to the fact that the PCs do a poor job of summarizing the information in the electricity futures; instead the explanation is that the demand and supply variables contain additional information.

The residuals from the regression, unEC, represent the unspanned component of the economic variables. We now proceed to show that they are important factors that affect the risk premium in electricity futures, rather than random noises. To determine if the unspanned component of the demand and supply variables (unEC) affects the risk premium in electricity futures, we use the unspanned component to forecast changes in the first two PCs. The regression is specified as follows:

$$\Delta \mathrm{PC}^{1-2}_{t\to t+1} = \mathrm{Const.} + \beta_{\mathrm{pc}}\mathrm{PC}^{1-5}_t + \beta_{\mathrm{unEC}}\mathrm{unEC}_t + \epsilon_t \tag{30}$$

If the unspanned components are unspanned by the electricity futures, the loading on unEC in this forecasting regression should be statistically significant and the adjusted R^2 should increase when adding the unspanned components to the regression.

The results in Panel B of Table 3 indicate that this is indeed the case. The loading on the unspanned component of the natural gas price (unPX) and the maximum load (unMax Load) are significant and positive for both PC^1 and PC^2, suggesting that the unspanned components of the economic variables impact the future realized changes in the PCs. Moreover, after including the unspanned components, the adjusted R^2 substantially increase. Higher adjusted R^2s mean that the unspanned components of the demand and supply variables contain useful information about future changes in electricity prices, which provides strong support for the hypothesis that the demand and supply variables are unspanned by the electricity futures.

4.2 The Dynamics of the Unspanned Model

Now that we have established that the demand and supply variables are unspanned by the electricity futures, we proceed to estimate the physical (P) and risk neutral (Q) dynamics of the unspanned model.

Panel A of Table 4 reports the risk-neutral model estimates. The upper left entry of the K_1^Q matrix is very close to one and highly statistically significant. This parameter captures the persistence of the model implied spot price under the risk neutral measure. The spot price is close to a unit root process under this measure. This of course reflects not only the dynamic of the spot price under the physical measure, but also the risk premium.

The loading of PC^1 on PC^2 is negative and statistically significant. A larger PC^2 indicates a flatter slope, thus the negative sign indicates that the level of the electricity price will decrease when the slope flattens. The bottom left entry of K_1^Q is insignificant, indicating that the level of the futures curve does not predict its slope. Finally, the bottom right entry of K_1^Q indicates that the futures slope is mean reverting.

Panel B of Table 4 reports the estimated P dynamic of the unspanned model. Not surprisingly, the first PC, which captures the level of the futures prices, is much more persistent than the second PC, which captures the slope. Nevertheless, the loading of PC_{t+1}^1 on PC_t^1 is 0.96, which is not very high given that the data are daily and we are investigating the pricing implications of this factor one month or one year ahead. The supply and demand variables are stationary: indeed, temperature and load are rapidly mean reverting. This means that shocks to the supply and demand variables do not persist, and have a bigger impact in the short term than over longer horizons. Put differently, shocks to the supply and demand variables may be informative about short-term movements in prices, but have little information about longer term movements.

The estimates in Panel B of Table 4 also capture the interaction between the electricity prices and the demand and supply variables under the physical measure. The estimates reflect that natural gas is the marginal fuel for electricity production, and consequently the natural gas price is closely tied to electricity prices. The loading of the level of the electricity price PC^1 on PX_t is positive and highly statistically significant. The positive sign reflects the economic relation between production cost and electricity price.

The natural gas price not only affects the spot price but also the slope of the electricity futures curve, as the loading of PC^2_{t+1} on PX_t is also positive and significant. A smaller PC^2 indicates a steeper futures slope, so a higher natural gas price predicts a flatter slope. This reflects that the natural gas price mainly affects the short-term electricity price.

The model also indicates that the electricity price affects the natural gas price. The loadings of PX_{t+1} on PC^1_t are positive and significant. These positive signs indicate that higher electricity spot prices lead to higher natural gas prices. The high electricity price might result from high demand for electricity, which in turn leads to a higher usage of natural gas and thus higher prices.

Temperature is a proxy for electricity demand. The results in Table 4 are therefore consistent with economic intuition. First, CDD positively affects the PC^1. This reflects the fact that higher temperature generally leads to higher electricity demand. Second, because a higher PC^2_{t+1} implies a flatter futures curve, the positive impact of CDD on PC^2_{t+1} implies that CDD negatively affects the slope of the futures curve. Third, we find that temperature is autoregressive but that none of the variables (except for the PC^2) can predict temperature. This is consistent with the fact that changes in temperature are very difficult to predict.

4.3 Model Fit

Table 5 reports the fit of the unspanned model and compares its performance with that of the spanned model with economic variables. We also compare the model's fit with the two-factor log price model of Lucia and Schwartz (2002), which is a benchmark model in the literature. We do not report on the fit of the spanned model, because by definition it is identical to the fit of the unspanned model. For each of these three models, Table 5 reports the root mean squared error (RMSE) and the relative root mean squared error (RRMSE) for the spot and each futures contract, as well as the overall RMSE and RRMSE. Figure 5 graphically illustrates the fit of the unspanned model for the spot, the day-ahead price, the 6-month futures contract, and the 12-month futures contract.

The unspanned model has the smallest RMSE and RRMSE among the three models. The overall RMSE (RRMSE) is 5.57 (0.0649) for the unspanned model, compared to 52.83 (0.7487) for the spanned model and 6.10 (0.0800) for the Lucia and Schwartz model. The poor fit of the spanned model with economic variables is not surprising: the spanning assumption forces all the information in the economic variables to enter the futures prices, which results in a poor fit. The main objective of a model with economic variables only is not to provide the best possible fit, but rather to provide the best possible economic intuition. The Lucia and Schwartz model overall results in a good fit, but it is outperformed by the unspanned model.[5] This is also not surprising: the Lucia and Schwartz model imposes constraints on the dynamics of the state variables, while the unspanned model does not impose such constraints.

5 Analyzing Risk Premiums

The main conclusion from Table 5 is that the unspanned model provides a good fit for the futures price. It must be emphasized that this fit is identical to the fit of the spanned model. The difference between the two modeling approaches emerges when studying risk premiums. This highlights the fact that the demand and supply variables contain additional information that is relevant for the futures prices under the physical measure, and hence the risk premiums. It is important to separate this information, and it is to this task that we now turn. We first discuss the estimates and properties of the spot premium. We then discuss the forward bias.

[5] To benchmark the models' performance, note that Lucia and Schwartz (2002) report a RRMSE of more than 0.10 for this model. The fit in our application is somewhat better, presumably due to the use of a different and longer sample period.

5.1 The Spot Premium in the Unspanned Model

We analyze the spot premium implied by the unspanned model and compare it with spot premiums implied by alternative models. The spot premium measures the compensation required by investors for investing in electricity futures. It is defined as the expectation under the physical measure of the difference between the log spot price and the log day-ahead price.

$$\text{Spot Premium}_t = E_t^P[\text{Log}(S_{t+1}) - \text{Log}(F_t^{DA})] \tag{31}$$

Table 6 reports the average spot premium for the unspanned model. We report results for the entire sample period as well as by season. The average estimated spot premium is approximately -0.84 percent per day. It is negative on average in every season, but it is larger (more negative) in the winter and the summer. This is consistent with the intuition that there is a greater risk of power price spikes in the peak seasons (summer and winter). Those who are short power (i.e., distribution companies that must buy power at the market price to sell to customers at fixed rates) are at risk to these price spikes, which can impose large losses on them. Risk averse physical shorts can hedge these risks by purchasing futures, thereby creating hedging pressure on prices: this pressure tends to cause upward biased futures prices, which in the context of the model means a negative risk premium.[6] In the non-peak seasons, price spikes are less likely, and the need to hedge is commensurately less. The lower hedging pressure from power consumers reduces the upward bias in futures prices. Indeed, since there can be short hedging pressure from generation operators looking to hedge electricity price risk, prices can actually be biased downwards, especially in the low-demand “shoulder" months of the spring and fall.

[6]See Keynes (1923), Hirshleifer (1988), or Hirshleifer (1990) for models of commodity markets in which hedging pressure is a determinant of price bias and risk premia. Upward bias is associated with a negative risk premium because a negative risk premium means that spot prices drift up more (down less) in the equivalent (pricing) measure than the physical measure.

Figure 6 highlights the differences between the spot risk premiums for the different models. It plots the time series of the spot premium for the unspanned model (Panel A), the spanned model (Panel B), and the model with economic variables (Panel C). Panel D compares the three models; here we report weekly averages because the three daily plots are too noisy in one panel. The properties of the spot premium for the unspanned model are quite different from the other two models. Figure 6 clearly indicates that the risk premium in the unspanned model is most variable, followed by the spanned model, and the model with economic variables. More importantly, Figure 6 shows that the unspanned model is capable of generating occasional large positive spikes in the spot risk premium, most notably in 2014, at the time of the polar vortex. The two other models generate large negative risk premiums on that occasion. The unspanned model is able to capture this spike due to the spike in the natural gas price, evident from Figure 3. While the model with economic variables of course also includes the natural gas price, it is constrained because it does not allow for pricing factors other than the economic variables, which restricts its flexibility to capture atypical patterns in risk premiums in the polar vortex period.

Figure 6 illustrates that the spot risk premium in the unspanned model is highly time-varying. Figure 7 provides more perspective on these fluctuations and the differences with other models by plotting the spot premiums of the different models during the 2014 polar vortex period. The models with economic variables (the unspanned model and the spanned model with economic variables) exhibit dramatic changes in risk premiums during this period, whereas models without economic variables (the spanned model with latent variables) cannot. While the estimated spot risk premium in the unspanned model is on average -0.84 % per day, during the polar vortex period it fluctuates between approximately -50 percent and 175 percent per day.

Finally, to further investigate the dynamics of the spot premium, we decompose the spot premium into four components plus a constant. The four components represent the component associated with the level of the electricity futures curve, the component associated with the slope of the electricity futures curve, the component associated with the natural gas price, and the component associated with the temperature. The Appendix provides details on this decomposition.

Figure 8 depicts the time-series of these four components. Several conclusions obtain. First, the components associated with the electricity level and the slope are time-varying. They are sometimes positive and sometimes negative, indicating that investors sometimes require compensation to bear this risk while at times paying to hedge this risk. Second, the risk premiums associated with the natural gas price are much larger than the ones associated with the CDD. This suggests that the impact of economic variables on risk premiums mainly originates on the supply side rather than the demand side, confirming the evidence in Figures 2 and 3 .

5.2 The Distribution of Spot Prices and the Spot Premium

The model of Bessembinder and Lemmon (2002) implies that the spot premium should be related to the statistical properties of spot prices. Specifically, they predict that the spot premium should be negatively correlated with the variance of spot prices and positively correlated with the skewness of spot prices. To verify if the unspanned model can capture this stylized fact, we regress the estimated spot premium for the unspanned model against the variance and skewness of spot prices. The regression specification is as follows:

$$\text{Spot Premium}_t = \alpha + \beta_{\text{Variance}} \times \text{Variance}_t + \beta_{\text{Skewness}} \times \text{Skewness}_t + \epsilon_t \tag{32}$$

where Spot Premium_t is the average daily spot premium of the unspanned model in period t. Variance_t is the variance of daily real-time electricity prices in period t, and Skewness_t is the skewness of real-time electricity prices in period t. Note that our definition of the spot premium implies that these signs are the opposite of the signs in Bessembinder and Lemmon (2002). We expect a positive β_{Variance} and a negative β_{Skewness}.

Table 7 reports the estimates of equation (32). We report on three different implementations of this regression, where a period is defined to be either a month, a season, or a year. The results are consistent with the model of Bessembinder and Lemmon (2002). For all estimates of the spot premiums, the spot premium is positively related with variance and negatively related with skewness. Of course, when using years as the observation period, the estimates are imprecise because our sample is very small.

This again reflects the effects of hedging pressures. As Bessembinder and Lemmon (2002) show, long hedgers (i.e., those with commitments to sell power at fixed prices who must buy at spot prices to cover those commitments) are primarily at risk to price spikes, and greater frequency and intensity of price spikes cause greater skewness in prices. Long hedging pressure increases upward bias in futures prices (reduces the risk premium/makes the risk premium more negative). Conversely, power generators who sell at spot prices benefit from price spikes, but incur greater risk when price variances are larger. Thus, short hedging pressure depends primarily on variance, meaning that higher variance increases short hedging pressure, thereby causing the bias in futures prices to fall and the risk premium to rise. Thus, the empirical estimates we present here are consistent with the economics of hedging pressure in the electricity market.

5.3 Predicting Returns with the Estimated Spot Premium

Our results indicate that the spot premiums from different models have different properties. While the spot premium from the unspanned model seems to have some plausible properties, strictly speaking this does prove it is a superior measure of the risk preferences of investors in electricity markets. We therefore conduct an out-of-sample exercise in which we use the estimated spot premium to predict the future realized changes of the log spot price. The regression specification is as follows:

$$s_t - s_{t-1} = \text{Const.} + \beta \times \text{Spot Premium}_{t-1} + \epsilon_t \tag{33}$$

If the spot premium of the unspanned model is a better indicator of the risk premium people pay for electricity futures, then the prediction of the unspanned model should be statistically significant and more importantly, the adjusted R^2 should be higher compared to the other models.

Table 8 shows results for the regression (33) for the entire sample period and two sub-samples. We also control for lagged returns. Among the three models we consider, the unspanned model has the highest adjusted R^2. This results holds when we discuss the forecasting power of the estimated spot premiums separately as well as jointly. The same result holds for the sub-samples, and when including control variables. These findings suggest that the spot premium of the unspanned model better captures the risk preferences of investors in electricity markets. The Online Appendix contains results for an additional out-of-sample exercise where we predict the future log day-ahead price instead of the future log spot price. Once again the unspanned model outperforms the other models.

5.4 Measuring the Forward Bias

There are several ways to characterize the risk premiums in forward contracts. We follow the existing literature on electricity markets and define the forward bias as the difference between the expected average spot price and the futures price. More precisely, it is defined as follows:

$$\text{Forward Bias}_t^i = E_t^P[\bar{S}_{t+i}] - F_t^i \tag{34}$$

where F_t^i denotes the price of the i-th month futures contract at time t and $E_t^P[\bar{S}_{t+i}]$ denotes the time-t average expected spot price in the maturity month t+i, where the expectation is under the P measure.

Figure 9 plots time series of the forward bias. Panel A plots the forward bias for the day-ahead contract, Panel B for the 1-month futures contract, and Panel C for the 12-month futures contract. Note that the forward bias is expressed in dollar terms, whereas the spot premium in Figure 6 is expressed in daily percentage returns. We report weekly averages, because for daily differences the extreme observations completely dominate the figure in Panel A, and as a result it is not informative.

Figure 9 reports on the unspanned model. The differences between the spanned model and the unspanned model with economic variables are more pronounced for short maturities than for long maturities. The forward bias in Panels B and C is therefore highly correlated with the forward bias in the spanned model. For the forward bias in the day ahead contract, the correlation with the forward bias in the spanned model is only 35%. These findings are due to the fact that the supply and demand variables are rapidly mean reverting. Shocks to the supply and demand variables do not persist, and have a bigger impact in the short term than over longer horizons. Put differently, shocks to the supply and demand variables are informative about short-term risk premia but have little information about longer term risk premia. To some extent these findings may also be due to the fact that we estimate the model using daily data. Estimating the same models using mildly autoregressive economic variables at the weekly or monthly frequency may yield different implications for longer horizons.

The forward biases for the 1-month futures contract in Panel B and the 12-month futures contract in Panel C are positively correlated, at 32%. The correlations between the time series of the forward bias for contracts with a maturity of one month and more (not reported) are all positive, and not surprisingly the correlations are high for contracts with similar maturities. This is not the case for the forward bias in the day-ahead contract. The correlation between the forward bias in Panels A and C is 1%, and the correlation between Panels A and B is even lower. This result once again follows from the fact that the unspanned feature of the model mainly has implications for risk premiums at shorter maturities.

Although the one-month forward bias and the twelve-month forward bias are positively correlated, Panels B and C highlight some important differences. There is a large positive spike in the one-month forward bias at the time of the polar vortex in 2014, which is much less pronounced for the twelve-month contract. The twelve-month forward bias on the other hand is positive and large for an extended period of time during 2005-2009, and this is not the case for on one-month forward bias. As expected from the small correlations, the patterns in the one-day ahead forward bias in Panel A are entirely different from those in Panels B and C. In 2014, we first have a large negative spike in the day-ahead forward bias, followed by a smaller positive one.

The most important difference between the forward bias in the day-ahead and 1-month contract on the one hand, and the twelve-month contract on the other hand, is that for the forward bias for the twelve-month contract in Panel C, reversion to the mean is a lot slower. The forward bias first increases (becomes more negative) between 2004 and 2006, stays at this heightened level between 2006 and mid-2008, and then decreases in absolute value, eventually turning positive. No such patterns are evident in Panels A and B. It is also worth pointing out that no such patterns are evident from the slope of the futures curve in Panel C of Figure 1. The time series of the forward bias for contracts with maturities between one month and one year (not reported) show that as the maturity of the contract gets longer, the pattern in the time series of the forward bias increasingly resembles the pattern in Panel C

Figure 10 plots the average difference between the expected average spot rate and the forward rate in each season (Panels A-D) and for the entire sample (Panel E). In each panel, we plot this difference as a function of the time to maturity of the forward contract, which ranges up to twelve months. Note that the horizontal axis refers to the contract number, with 1 denoting the spot rate and 13 the twelve-month forward. This analysis follows Pirrong and Jermakyan (2008), who conduct a similar exercise for a single day. Figure 10 indicates that this premium increases (becomes more negative) as a function of maturity. When considering the entire sample in Panel E, this increase is almost monotonic, but not surprisingly this is not the case in a given season. This finding is consistent with Pirrong and Jermakyan (2008) and with our findings on the seasonality of risk premiums in Table 6.

Finally, the forward bias is very large and economically significant. Recall from Table 1 that over our sample period the average spot and futures price is approximately \$60. Panel E of Figure 10 indicates that on average the forward bias is -\$4 for the one-month maturity and -\$7 for the twelve-month maturity, which in both cases represents a significant percentage of the price. However, Figure 9 shows that the forward bias dramatically fluctuates around the sample average, with larger fluctuations and outliers for shorter-maturity contracts. For instance, the day-ahead forward bias contains a positive outlier of \$340 and a negative outlier of -\$80 during the polar vortex period. The one-month forward bias has negative outliers of approximately -\$60 in 2014 and -\$40 in 2005 and 2008. Note that these are weekly averages and therefore the daily outliers are even larger.

6 Concluding Remarks

We model the impact of supply and demand on the price of electricity futures in a no-arbitrage model, using daily data between 2003 and 2014. By design, the model fits electricity futures as well as a fully latent model. Additionally, the model allows for unspanned economic risk which is captured by the supply and demand variables but not identified by the futures prices.

We find that demand and supply variables contain valuable information about the spot risk premiums embedded in electricity futures prices. The spot risk premiums implied by the unspanned model are on average negative, very different from the risk premiums implied by the spanned model, display strong seasonal patterns, are related to the variance and skewness of the electricity spot price, and better predict future changes in spot prices. We decompose the spot risk premium into components associated with demand and supply. The unspanned risk premium associated with supply is economically large and highly time-varying and constitutes the most important component of the total spot risk premium embedded in electricity futures. The forward bias implied by the model is also on average negative and large. It is highly time-varying and increases as a function of maturity.

Our findings thus suggest that the use of latent variables to model electricity prices provides a good fit to the data, but additionally including unspanned economic variables generates more plausible economic implications.

Several extensions of our approach may prove interesting. Our out-of-sample analysis uses the estimated spot premium to predict the future realized changes of the log spot price and the log day-ahead price, but this analysis uses estimates obtained using the entire sample. It may prove interesting to repeat this exercise using a recursive estimation. We could also estimate a quadratic model to investigate the robustness of our results on spanning and the measurement of risk premiums to the assumption that the log spot price is linear in the state variables. Finally, as discussed above, our use of daily data makes it likely that the model implications for the forward bias are similar to those of the unspanned model for longer horizons. Implementing the model with weekly or monthly data may yield different results at longer horizons.

References

Alaton, P., B. Djehiche, and D. Stillberger. 2002. On modelling and pricing weather derivatives. *Applied Mathematical Finance* 9:1–20.

Benth, F. E., Á. Cartea, and R. Kiesel. 2008. Pricing forward contracts in power markets by the certainty equivalence principle: explaining the sign of the market risk premium. *Journal of Banking & Finance* 32:2006–21.

Bessembinder, H., and M. L. Lemmon. 2002. Equilibrium pricing and optimal hedging in electricity forward markets. *The Journal of Finance* 57:1347–82.

Bunn, D. W., and D. Chen. 2013. The forward premium in electricity futures. *Journal of Empirical Finance* 23:173–86.

Cartea, A., and M. G. Figueroa. 2005. Pricing in electricity markets: a mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance* 12:313–35.

Cartea, Á., and P. Villaplana. 2008. Spot price modeling and the valuation of electricity forward contracts: the role of demand and capacity. *Journal of Banking & Finance* 32:2502–19.

Deng, S.-J., and S. S. Oren. 2006. Electricity derivatives and risk management. *Energy* 31:940–53.

Dong, L., and H. Liu. 2007. Equilibrium forward contracts on nonstorable commodities in the presence of market power. *Operations Research* 55:128–45.

Douglas, S., and J. Popova. 2008. Storage and the electricity forward premium. *Energy Economics* 30:1712–27.

Geman, H. 2009. *Commodities and commodity derivatives: modeling and pricing for agriculturals, metals and energy*. John Wiley & Sons.

Geman, H., and A. Roncoroni. 2006. Understanding the fine structure of electricity prices. *The Journal of Business* 79:1225–61.

Heath, D. 2016. Unspanned macroeconomic risks in oil futures. Working Paper, University of Utah.

Hirshleifer, D. 1988. Residual risk, trading costs, and commodity futures risk premia. *The Review of Financial Studies* 1:173–93.

———. 1990. Hedging pressure and futures price movements in a general equilibrium model. *Econometrica: Journal of the Econometric Society* 411–28.

Jewson, S., and A. Brix. 2005. *Weather derivative valuation: the meteorological, statistical, financial and mathematical foundations*. Cambridge University Press.

Joslin, S., M. Priebsch, and K. J. Singleton. 2014. Risk premiums in dynamic term structure models with unspanned macro risks. *The Journal of Finance* 69:1197–233.

Keynes, J. M. 1923. Some aspects of commodity markets. *Manchester Guardian Commercial: European Reconstruction Series* 13:784–6.

Khan, S., Z. Khokher, and T. Simin. 2016. The information content of a nonlinear macro-finance model for commodity prices. *The Review of Financial Studies* 30:2818–50.

Longstaff, F. A., and A. W. Wang. 2004. Electricity forward prices: a high-frequency empirical analysis. *The Journal of Finance* 59:1877–900.

Lucia, J. J., and E. S. Schwartz. 2002. Electricity prices and power derivatives: Evidence from the nordic power exchange. *Review of Derivatives Research* 5:5–50.

Pirrong, C. 2011. *Commodity price dynamics: A structural approach*. Cambridge University Press.

Pirrong, C., and M. Jermakyan. 2008. The price of power: The valuation of power and weather derivatives. *Journal of Banking & Finance* 32:2520–9.

Table 1: Descriptive Statistics. Electricity Prices

We report descriptive statistics for electricity prices. Panel A reports the number of observations, the sample mean, standard deviation, skewness, kurtosis, the minimum, maximum, and autocorrelation coefficient of order 1 for the spot and futures prices. We report on PJM Western Hub real-time spot, day-ahead futures (F^{DA}), and PJM Western Hub real-time peak calendar-month 2.5 MW futures with maturities between 1 month and 12 months (F^1 to F^{12}). The unit of the price is Dollar per MWh. Panel B reports on log returns in daily percentage returns, where s_t denotes the log of the spot price at time t and f_t^{DA} denotes the log of the day-ahead price at time t. The sample period is from May 2003 to May 2014.

Panel A: Prices

	Nobs	Mean	Std	Skew	Kurt	Min	Max	AC1
Spot	2767	58.1649	32.4250	6.5112	101.9091	21.2089	757.4217	0.6136
F^{DA}	2767	58.0342	29.8561	6.8500	105.4550	20.6838	683.1615	0.7566
F^1	2767	58.0873	18.6850	1.7388	7.3168	31.3500	164.7500	0.9866
F^2	2767	58.9385	19.3147	1.7136	6.9351	33.0600	154.3300	0.9932
F^3	2767	59.5083	19.1678	1.5195	5.5145	35.1700	140.0000	0.9937
F^4	2767	59.8302	18.9369	1.3944	4.8339	34.1700	139.2300	0.9937
F^5	2767	59.9655	18.2606	1.1540	3.8122	35.0400	128.2900	0.9932
F^6	2767	60.1071	17.9742	1.0423	3.3712	34.1100	128.4200	0.9928
F^7	2767	60.4287	18.1264	1.0697	3.7185	34.3200	128.4200	0.9932
F^8	2767	61.0319	18.5374	1.0138	3.5116	35.3600	127.0800	0.9939
F^9	2767	61.3238	18.5972	0.9314	3.0770	35.3200	116.2500	0.9943
F^{10}	2767	61.3462	18.1073	0.8598	2.8350	35.0000	114.3300	0.9935
F^{11}	2767	61.2828	17.7731	0.8668	2.8897	34.3000	119.5000	0.9930
F^{12}	2767	61.2449	17.9752	1.0133	3.6591	34.2600	138.4900	0.9932

Panel B: Log Returns

	Nobs	Mean	Std	Skew	Kurt	Min	Max	AC1
s	2766	-0.0156%	0.2988	0.0551	6.3799	-1.7030	1.8596	-0.2873
f^{DA}	2766	-0.0028%	0.1845	-0.1748	11.2055	-1.4306	1.2689	-0.0552
$s_{t+1} - f_t^{DA}$	2766	-1.3691%	0.2738	0.1567	6.6494	-1.6341	1.8042	0.1277

Table 2: Descriptive Statistics. Demand and Supply Variables

We report descriptive statistics for the raw economic variables. The economic variables are the natural gas price (PX), the average load (Avg Load), the maximum load (Max Load), the maximum temperature (Max T), the minimum temperature (Min T), the average temperature (Avg T), the cooling degree days (CDD), and the heating degree days (HDD). The sample period is from May 2003 to May 2014.

	Nobs	Mean	Std	Skew	Kurt	Min	Max	AC1
PX	2767	6.3659	4.1495	10.6412	211.3608	1.9900	99.6600	0.7445
Avg Load	2767	39.9420	13.4143	-1.1067	4.3669	5.4614	74.0745	0.9820
Max Load	2767	42.4714	14.5917	-0.9495	4.1839	5.6450	80.1791	0.9772
Max T	2767	17.8765	10.3798	-0.3876	2.0988	-11.6000	38.1000	0.9008
Min T	2767	7.0997	9.3482	-0.2546	2.1015	-19.9000	24.4500	0.9105
Avg T	2767	12.4881	9.6920	-0.3302	2.0892	-15.7500	30.7000	0.9258
CDD	2767	1.6956	2.7615	1.4701	3.9566	0.0000	12.7000	0.8800
HDD	2767	7.2075	7.8649	0.7892	2.4767	0.0000	33.7500	0.9077

Table 3: Spanning Regressions

In Panel A, we report the results if projecting the economic variables EC_t on the first five principal components (PC_t^{1-5}) of the electricity futures. The regression is specified as follows:

$$\text{EC}_t = \text{Const.} + \gamma_{pc}\text{PC}_t^{1-5} + \text{unEC}_t$$

In Panel B, we forecast changes in the first two principal components using the unspanned component of demand and supply variables, controlling for lagged values of the PCs. The regression is specified as follows:

$$\Delta\text{PC}_{t\to t+1}^{1-2} = \text{Const.} + \beta_{pc}\text{PC}_t^{1-5} + \beta_{unEC}\text{unEC}_t + \epsilon_t$$

For both regressions, standard errors are reported in parentheses. The sample period is from May 2003 to May 2014.

Panel A: Projecting ECs on PCs

	Const.	PC1_t	PC2_t	PC3_t	PC4_t	PC5_t	Adj. R^2	N.Obs.
PX_t	-0.1916	2.4991	2.6934	4.9415	-1.0160	2.4471	0.4410	2767
	(0.0653)	(0.0605)	(0.1612)	(0.3730)	(0.4307)	(0.5168)		
Max Load_t	-0.3187	1.1686	8.3349	-11.5129	-17.0962	-16.5058	0.1072	2767
	(0.2891)	(0.2678)	(0.7131)	(1.6503)	(1.9055)	(2.2864)		
CDD_t	0.1654	0.1157	2.3044	-1.0969	-1.0573	-1.8853	0.2878	2767
	(0.0300)	(0.0278)	(0.0739)	(0.1710)	(0.1975)	(0.2370)		

Panel B: Forecasting Changes in PCs

	$\Delta\text{PC1}_{t\to t+1}$	$\Delta\text{PC1}_{t\to t+1}$	$\Delta\text{PC1}_{t\to t+1}$	$\Delta\text{PC2}_{t\to t+1}$	$\Delta\text{PC2}_{t\to t+1}$	$\Delta\text{PC2}_{t\to t+1}$
unPX_t		1.1555	1.1509		2.4028	2.3770
		(0.0755)	(0.0768)		0.1495	0.1522
unMax Load_t		0.0975			0.2232	
		(0.0171)			0.0338	
unCDD_t			0.6041			1.2273
			(0.1674)			0.3318
PC_t^{1-5}	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R^2	0.1357	0.2059	0.2003	0.1914	0.2646	0.2566
N.Obs.	2766	2766	2766	2766	2766	2766

Table 4: Estimates of the Unspanned Model

We report the estimated Q- and P-dynamics of the unspanned model. The Q- and P-dynamics are specified as follows:

$$\begin{aligned} \text{States in } \mathrm{Q}_{t+1} &= K_0^Q + K_1^Q \times \text{States in } \mathrm{Q}_t + \Sigma^Q \times \epsilon_{t+1}^Q \\ \text{States in } \mathrm{P}_{t+1} &= K_0^P + K_1^P \times \text{States in } \mathrm{P}_t + \Sigma^P \times \epsilon_{t+1}^P \end{aligned}$$

where ϵ_t^Q and ϵ_t^P are standard Brownian motions, and Σ^Q is the left upper 2 by 2 sub-matrix of Σ^P. The state variables in the Q-dynamic are the PC1 and the PC2. The state variables in the P-dynamic are the PC1, the PC2, the natural gas price, and the CDD. Standard errors are reported in parentheses. The sample period is from May 2003 to May 2014.

Panel A: Q-Dynamic

	K_0^Q	K_1^Q	
		$\mathrm{PC1}_t$	$\mathrm{PC2}_t$
$\mathrm{PC1}_{t+1}$	-0.0577	0.9991	-0.0269
	(0.0094)	(0.0004)	(0.0006)
$\mathrm{PC2}_{t+1}$	0.0152	-0.0016	0.9475
	(0.0090)	(0.0014)	(0.0017)

Panel B: P-Dynamic

	K_0^P	K_1^P			
		$\mathrm{PC1}_t$	$\mathrm{PC2}_t$	PX_t	CDD_t
$\mathrm{PC1}_{t+1}$	-0.0096	0.9588	-0.1713	0.0123	0.0054
	(0.0025)	(0.0031)	(0.0079)	(0.0007)	(0.0016)
$\mathrm{PC2}_{t+1}$	-0.0242	-0.0871	0.6078	0.0268	0.0110
	(0.0050)	(0.0061)	(0.0158)	(0.0015)	(0.0033)
PX_{t+1}	-0.2439	0.9379	0.1182	0.5965	-0.0606
	(0.0566)	(0.0689)	(0.1777)	(0.0168)	(0.0370)
CDD_{t+1}	0.0250	0.0260	0.2678	-0.0097	0.6071
	(0.0267)	(0.0325)	(0.0838)	(0.0079)	(0.0175)

Panel C: Estimates of $\Sigma^P(\Sigma^P)'$

	$\mathrm{PC1}_t$	$\mathrm{PC2}_t$	PX_t	CDD_t
$\mathrm{PC1}_t$	0.0141	0.0250	0.0481	0.0684
$\mathrm{PC2}_t$		0.0566	0.0915	0.1523
PX_t			7.1277	0.0157
CDD_t				1.5859

Table 5: Model Fit

We report the fit of the unspanned model (Panel A), the spanned model with economic variables (Panel B), and the Lucia and Schwartz model (Panel C) for futures prices. The economic variables are the natural gas price and CDD. For each model, we report both the root mean squared error (RMSE) and relative root mean squared error (RRMSE), defined as follows:

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(\hat{\text{F}}^{i} - \text{F}^{i})^2} \qquad \text{RRMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\frac{(\hat{\text{F}}^{i} - \text{F}^{i})^2}{(\text{F}^{i})^2}}$$

where $\hat{\text{F}}^{i}$ is the fitted futures price and F^{i} is the realized futures price. F^{DA} denotes the day-ahead futures price. F^{1} to F^{12} denote the futures maturing in 1 month to 12 months respectively. Overall RMSE and RRMSE are calculated as simple averages of the RMSE and RRMSE over all futures contracts. The sample period is from May 2003 to May 2014.

Panel A: The Unspanned Model

	Spot	F^{DA}	F^1	F^2	F^3	F^4	F^5	F^6	F^7	F^8	F^9	F^{10}	F^{11}	F^{12}	Overall
RMSE	9.8509	11.0573	6.1483	5.4906	4.8749	4.5452	4.0922	3.9061	3.9283	4.4184	4.7793	4.7363	5.0133	5.1126	5.5681
RRMSE	0.0953	0.1125	0.0933	0.0785	0.0662	0.0588	0.0567	0.0567	0.0584	0.0642	0.0691	0.0670	0.0721	0.0772	0.0733

Panel B: The Spanned Model with Economic Variables

	Spot	F^{DA}	F^1	F^2	F^3	F^4	F^5	F^6	F^7	F^8	F^9	F^{10}	F^{11}	F^{12}	Overall
RMSE	161.9312	137.6671	49.5725	46.8987	42.5467	41.3189	39.2372	42.5299	37.9617	31.0323	28.8998	27.1105	26.1147	26.8088	52.8307
RRMSE	1.5846	0.7146	0.5713	0.8756	0.9035	0.8711	0.7722	0.6352	0.6485	0.6712	0.6757	0.6096	0.5198	0.4294	0.7487

Panel C: The Lucia and Schwartz Model

	Spot	F^{DA}	F^1	F^2	F^3	F^4	F^5	F^6	F^7	F^8	F^9	F^{10}	F^{11}	F^{12}	Overall
RMSE	14.0016	8.6296	7.8107	6.0312	6.3501	6.2814	5.2817	3.8809	2.3308	4.0228	5.5056	5.8228	5.3645	4.0535	6.0977
RRMSE	0.1388	0.0762	0.1219	0.0899	0.0873	0.0833	0.0722	0.0541	0.0337	0.0582	0.0795	0.0831	0.0790	0.0627	0.0800

Table 6: Estimated Spot Premiums

We report the daily average of the spot premiums of the unspanned model. The spot premium is defined as follows:

$$\text{Spot Premium}_t = E_t^P[\text{Log}(S_{t+1}) - \text{Log}(F_t^{DA})]$$

For each model, we report the average spot premium in each season as well as over the entire sample period. The definition of the seasons is as follows: Winter is defined as December, January, and February. Spring is defined as March, April, and May. Summer is defined as June, July, and August. Fall is defined as September, October, and November. The economic variables that are used to calculate the spot premium of the unspanned model are the natural gas prices and CDD. The reported numbers are raw daily log returns. The sample period is from May 2003 to May 2014.

	Nobs	Mean	Std	Skew	Kurt	Min	Max	AC1
Winter	657	-0.0158	0.1344	5.9916	74.5678	-0.5230	1.7777	0.3148
Spring	722	-0.0026	0.0661	-1.6762	11.9224	-0.5095	0.2103	0.4794
Summer	709	-0.0101	0.0789	-1.2472	5.7594	-0.3753	0.1812	0.6255
Fall	679	-0.0059	0.0573	-0.2633	4.3276	-0.2691	0.1787	0.5851
All Seasons	2767	-0.0084	0.0886	4.4316	94.1437	-0.5230	1.7777	0.4426

Table 7: Regressing Spot Premiums on Variance and Skewness of Real-Time Spot Prices

We report the estimates of regressing spot premiums on the variance and skewness of real-time prices. The regression specification is as follows:

$$\text{Spot Premium}_t = \text{Const.} + \beta_{\text{Variance}} \times \text{Variance}_t + \beta_{\text{Skewness}} \times \text{Skewness}_t + \epsilon_t$$

where Spot Premium_t is the average spot premium of the unspanned model in period t, Variance_t is the variance of the daily electricity real-time price in period t, and Skewness_t is the skewness of the daily electricity real-time price in period t. We consider three frequencies when calculating those numbers, i.e. a month, a season, and a year. Spot Premium_t, Variance_t, and Skewness_t are all standardized. The economic variables that are used in the unspanned model are natural gas price and CDD. Standard errors are reported in parentheses. The sample period is from May 2003 to May 2014.

	Const.	β_{Variance}	β_{Skewness}	Adj. R^2	N.Obs.
By Month	0.0000	0.2051	-0.1367	0.0442	133
	(0.0848)	(0.0852)	(0.0852)		
By Season	0.0000	0.1782	-0.2958	0.0889	46
	(0.1407)	(0.1430)	(0.1430)		
By Year	0.0000	0.0459	-0.4156	0.0032	12
	(0.2882)	(0.3110)	(0.3110)		

Table 8: Predicting the Log Spot Return with the Estimated Spot Premium

We compare the predictive power of the spot premium implied by different models for the log spot return. The predictive regression is specified as follows:

$$s_t - s_{t-1} = \text{Const.} + \beta \times \text{Spot Premium}_{t-1} + \epsilon_t$$

where s_t is the log of the spot price at time t. The economic variables that are used in the unspanned model and the spanned model with economic variables are the natural gas price and CDD. Standard errors are reported in parentheses. The sample period is from May 2003 to May 2014.

Panel A: Baseline Regression

Unspan	1.5176			1.1057
	(0.0574)			(0.1077)
Span with Latents		2.0789		1.0228
		0.0943		0.2020
Span with Econs			(0.2047)	-(0.0864)
			0.0967	0.1502
Const.	0.0127	0.0174	0.0015	0.0171
	(0.0051)	(0.0053)	(0.0057)	(0.0051)
Adj. R^2	0.2017	0.1493	0.0013	0.2192
N.Obs.	2766	2766	2766	2766

Panel B: Sub-Samples

	2003 - 2008				2009 - 2014			
Unspan	1.8557			1.1875	1.3970			1.1490
	(0.0980)			(0.1877)	(0.0728)			(0.1879)
Span with Latents		2.3210		1.4081		1.9396		0.8474
		(0.1424)		(0.3180)		(0.1277)		(0.3793)
Span with Econs			0.5441	-0.8086			0.0643	0.1067
			(0.1651)	(0.2059)			(0.1326)	(0.3221)
Const.	-0.0037	0.0111	0.0146	-0.0172	0.0269	0.0237	-0.0008	0.0313
	(0.0067)	(0.0070)	(0.0087)	(0.0081)	(0.0077)	(0.0081)	(0.0087)	(0.0090)
Adj. R^2	0.2025	0.1583	0.0070	0.2133	0.2134	0.1450	-0.0006	0.2382
N.Obs.	1409	1409	1409	1409	1356	1356	1356	1356

Panel C: Controlling for Lagged Returns

Unspan		1.3636			1.0918
		(0.0579)			(0.1061)
Span with Latents			1.7574		0.7228
			(0.0989)		(0.2016)
Span with Econs				-0.0294	-0.0691
				(0.0940)	(0.1480)
Lagged Returns	-0.2863	-0.1888	-0.1737	-0.2872	-0.1641
	(0.0182)	(0.0171)	(0.0184)	(0.0185)	(0.0176)
Const.	-0.0001	0.0114	0.0147	-0.0004	0.0146
	(0.0054)	(0.0050)	(0.0052)	(0.0055)	(0.0050)
Adj. R^2	0.0816	0.2349	0.1756	0.0813	0.2426
N.Obs.	2765	2765	2765	2765	2765

Figure 1: Electricity Spot Prices and Futures Prices

We plot the daily spot price (Panel A), the price of the 12-month futures contract (Panel B), and the difference between them (Panel C) over the entire sample period. The difference is calculated as the average of daily differences in each month. The sample period is from May 2003 to May 2014.

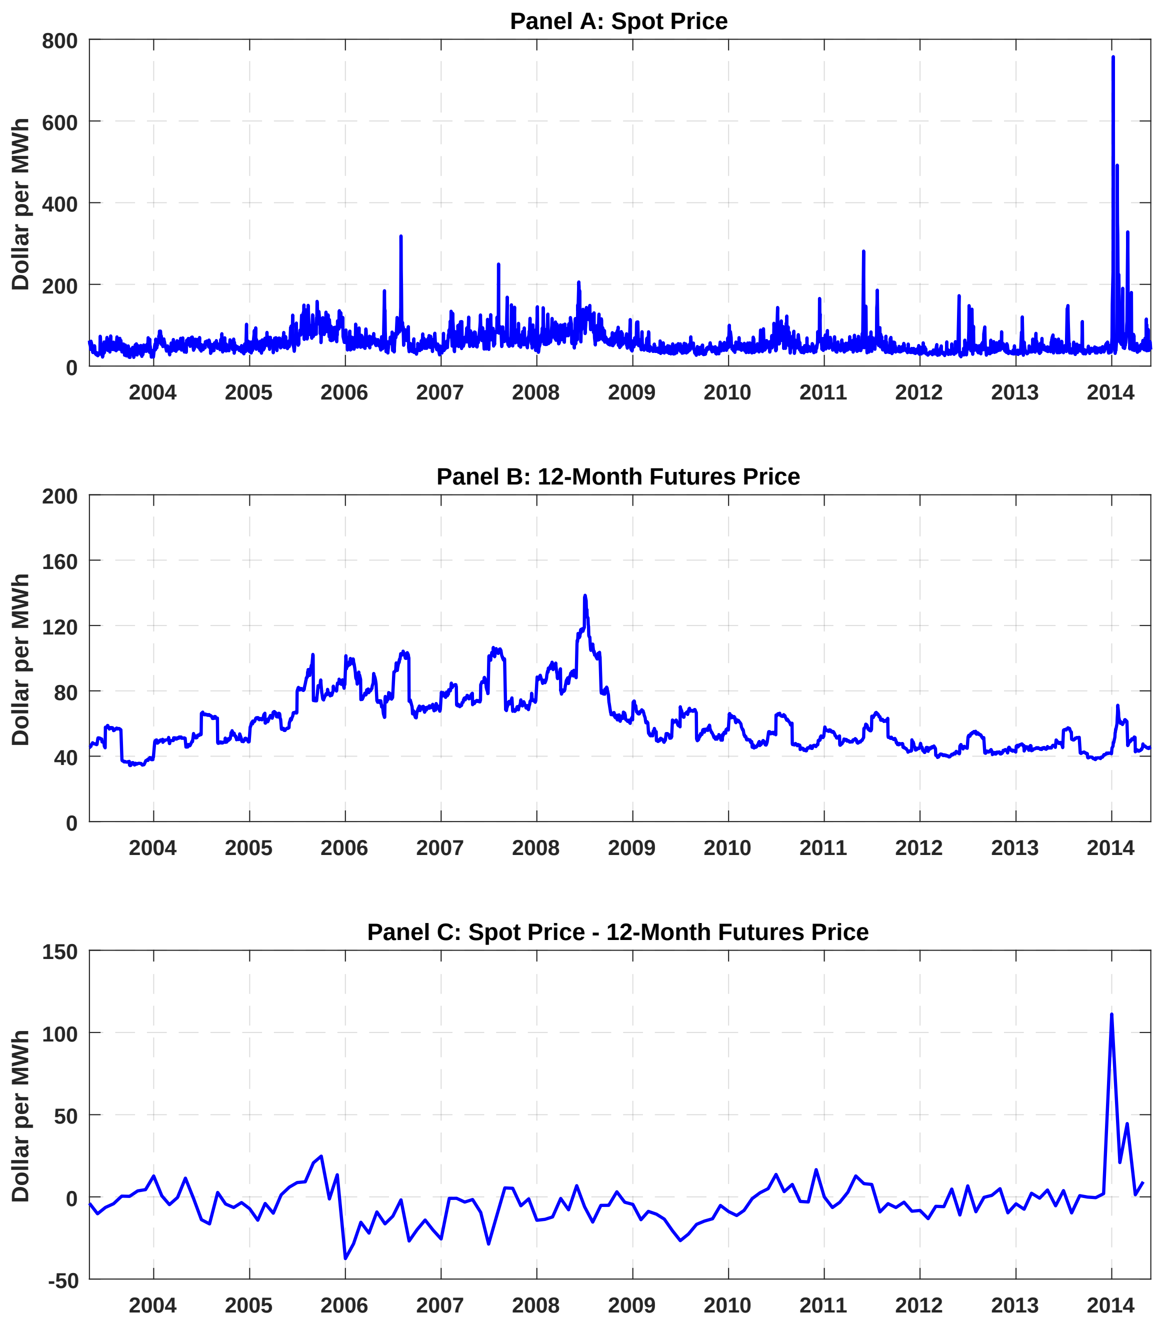

Figure 2: Electricity Prices

We plot the average spot price (Panel A), the day-ahead price (Panel B), the price of the 6-month futures contract (Panel C), and the price of the 12-month futures contract (Panel D) in each month of the sample period. In each panel, we plot the raw price (the blue solid line), the seasonal component (the dotted line), and the de-seasonalized price (the green solid line). The sample period is from May 2003 to May 2014.

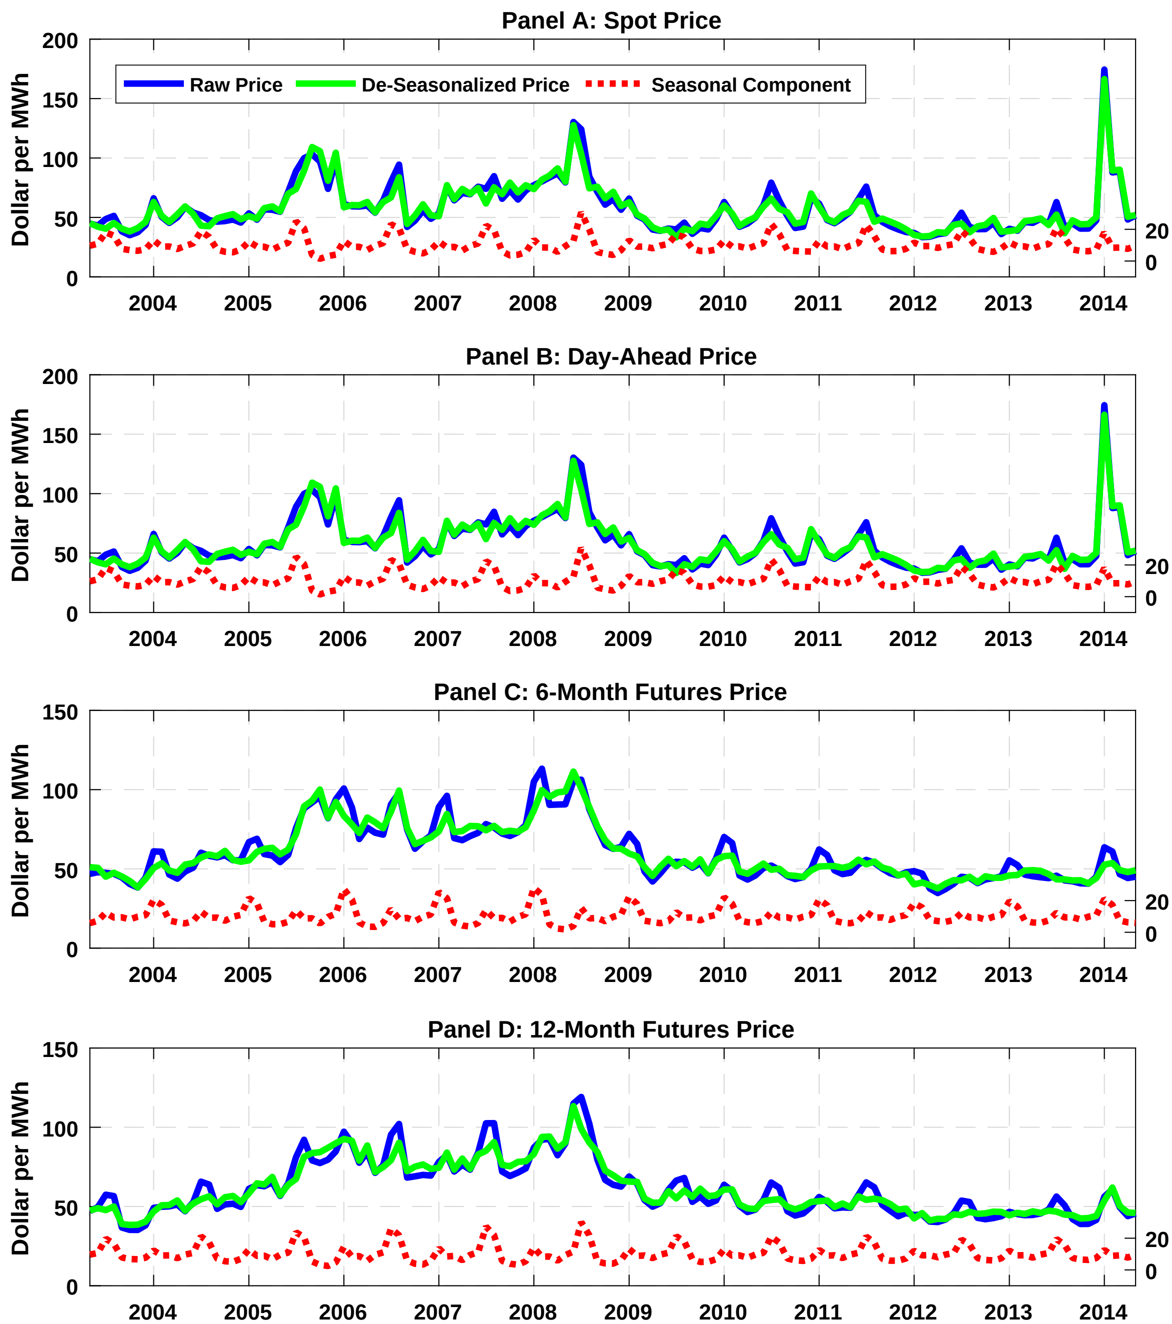

Figure 3: Demand and Supply Variables

We plot the average natural gas price (Panel A), the maximum load (Panel B), and the CDD (Panel C) in each month of the sample period. In each panel, we plot the raw series (the blue solid line), the seasonal component (the dotted line), and the de-seasonalized series (the green solid line). The sample period is from May 2003 to May 2014.

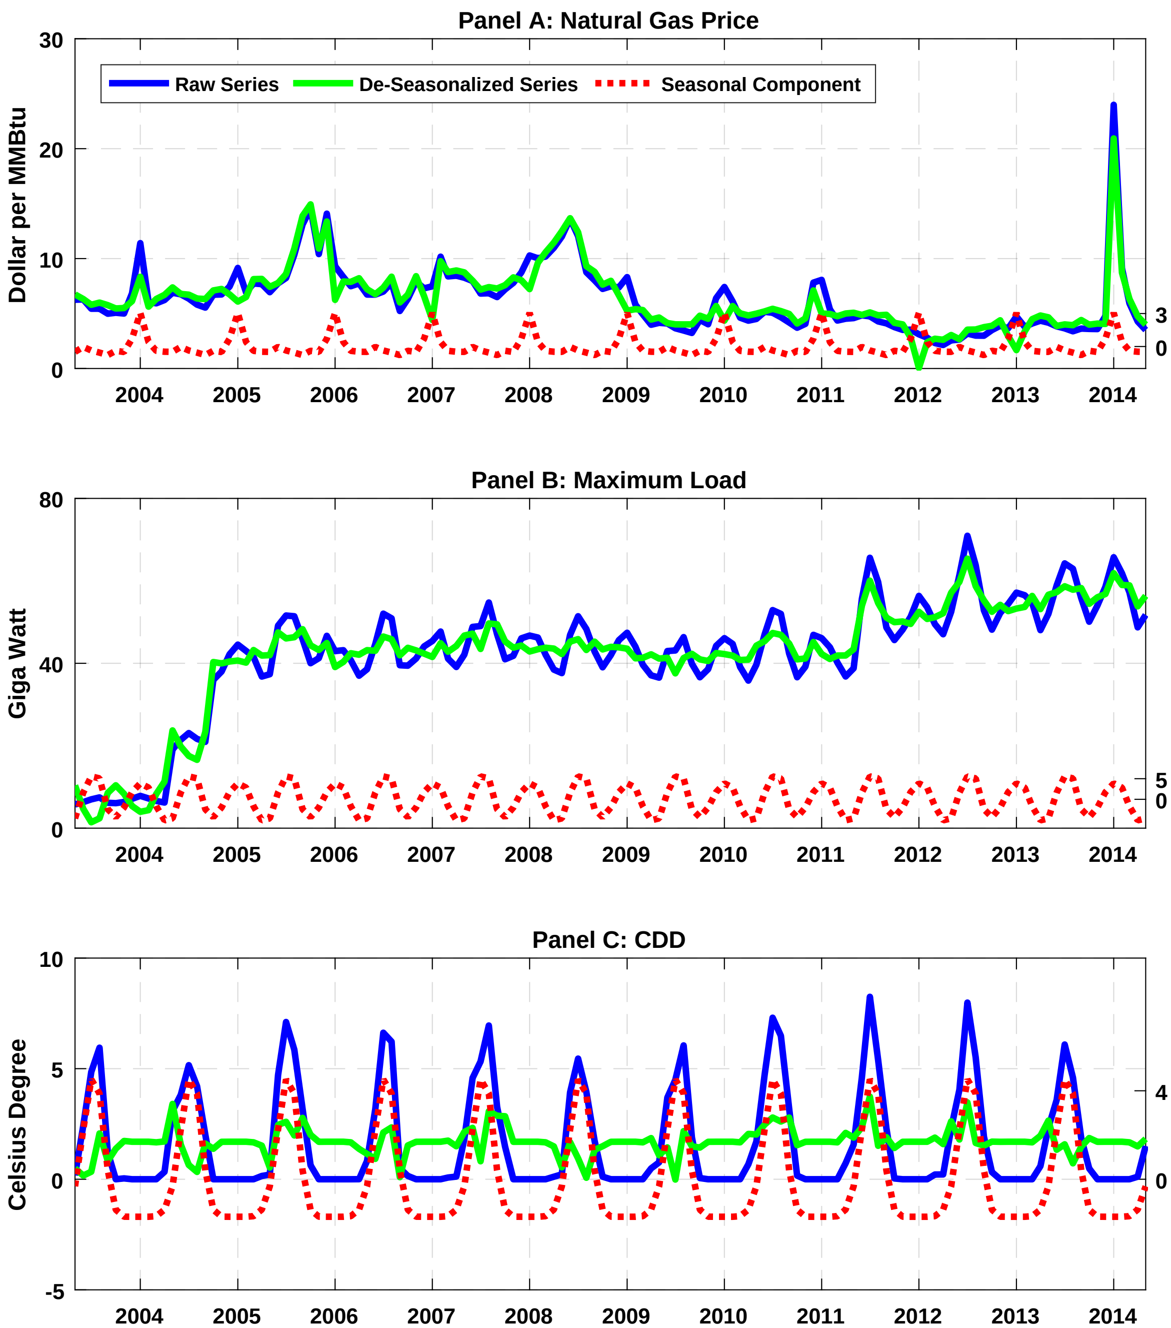

Figure 4: Principal Components of Electricity Prices

We plot the loadings of the first two principal components (PCs) of log electricity prices. The first PC is represented by the dashed line and the second PC is represented by the dotted line. The legend displays the fraction of the total variance explained by each of the principal components. The sample period is from May 2003 to May 2014.

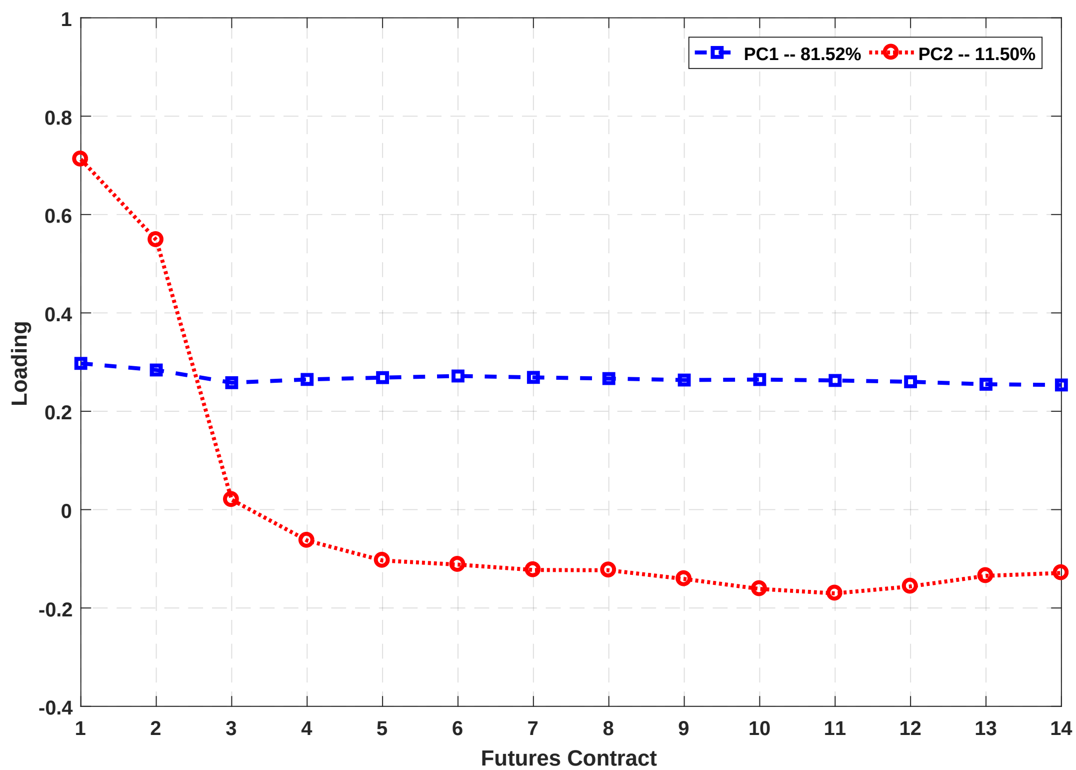

Figure 5: Model Fit

For each month of the sample period, we plot the raw data (the solid line), the model price (the dotted line), and the difference between the data with the model price (the dashed line) for the spot price (Panel A), the day-ahead price (Panel B), the price of 6-month futures (Panel C), and the price of 12-month futures (Panel D). The sample period is from May 2003 to May 2014.

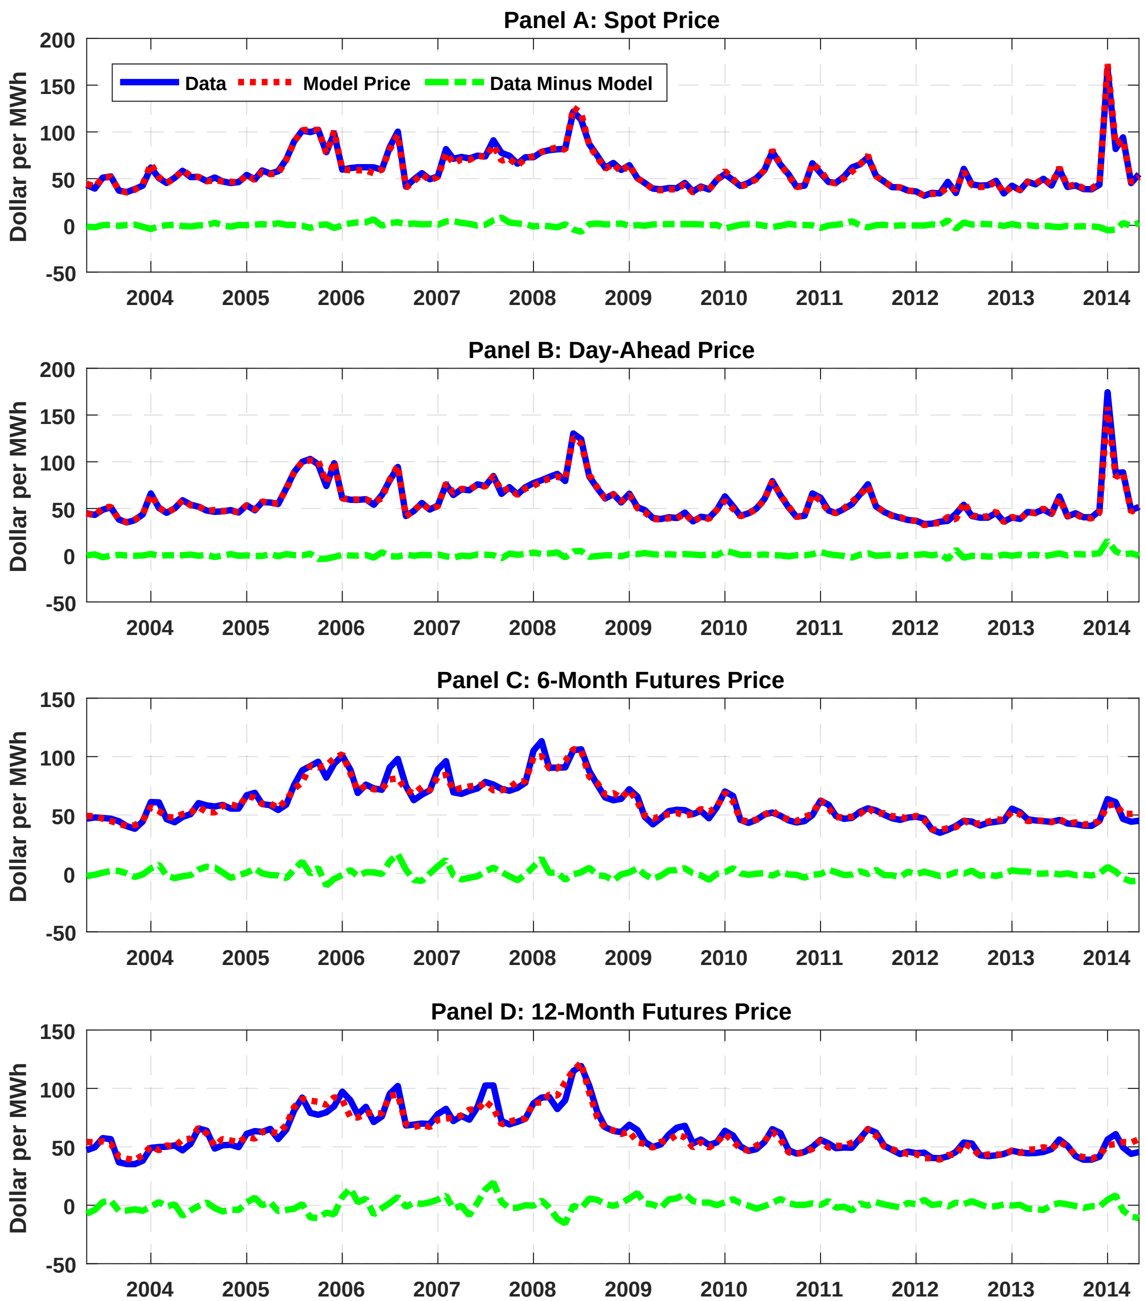

Figure 6: Spot Premiums. Various Models

We plot the model-implied spot premium for the unspanned model (Panel A), the model with latent variables (Panel B), the model with economic variables (Panel C), and their weekly averages (Panel D). The economic variables that are used to calculate the spot premium are the natural gas price and CDD. The sample period is from May 2003 to May 2014.

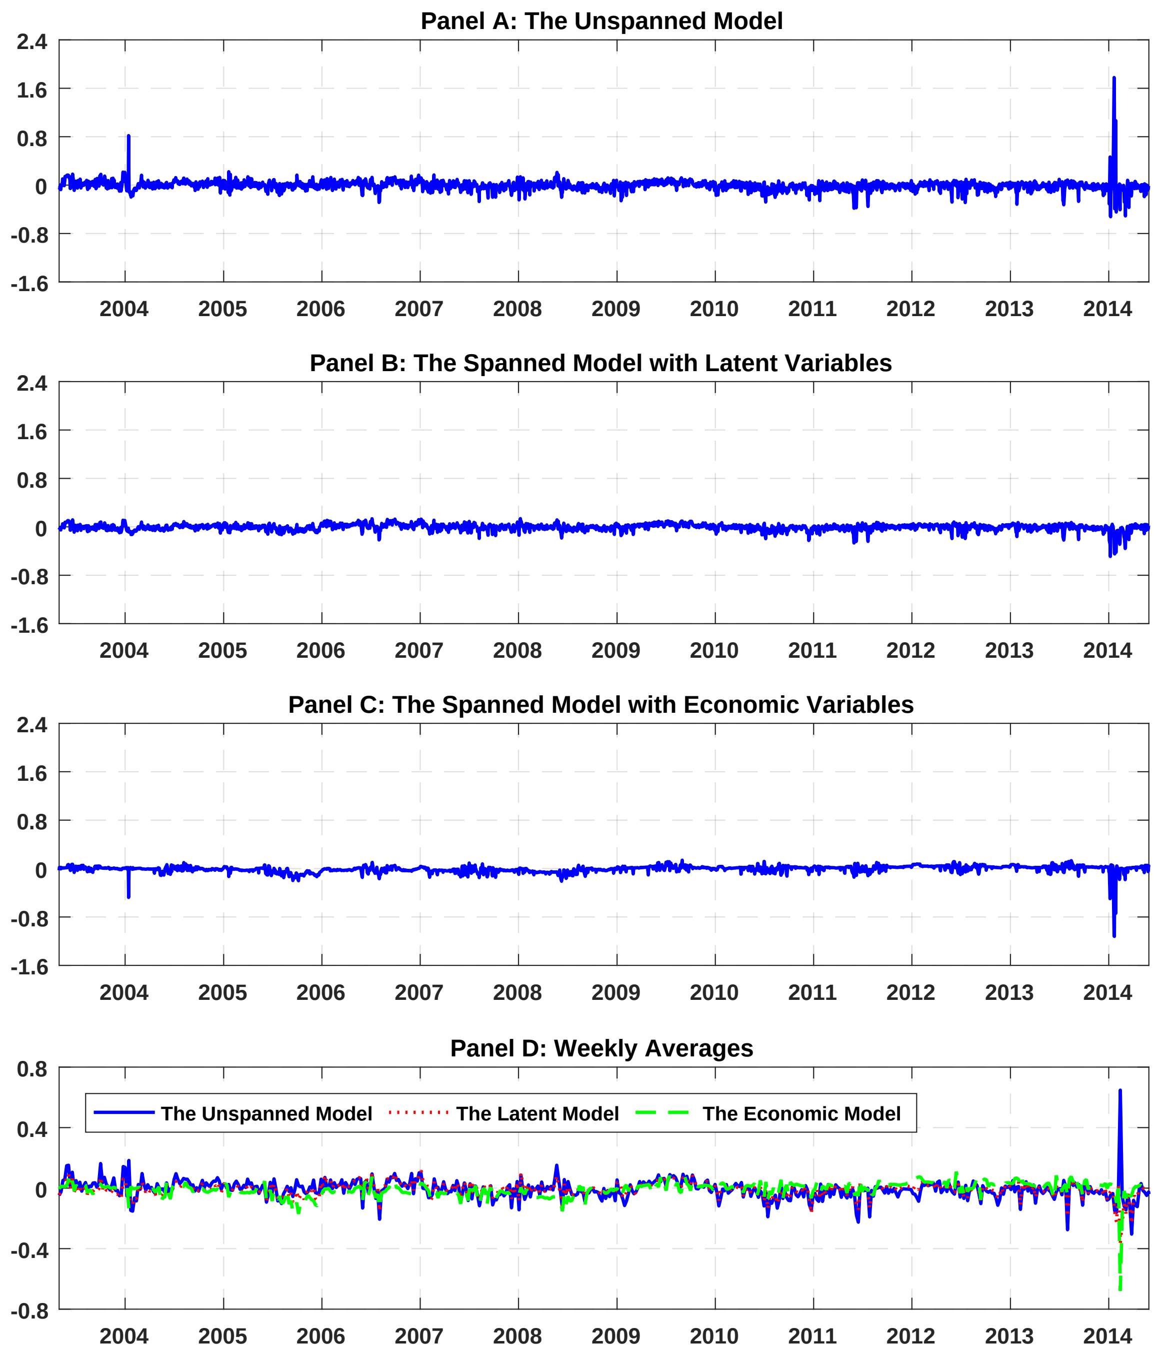

Figure 7: Spot Premiums During the 2014 Polar Vortex Period

We plot the estimated spot premium during the 2014 Polar Vortex period for various models. The economic variables that are used in the unspanned model and the spanned model with economic variables are the natural gas price and CDD. The sample period is from November 2013 to April 2014.

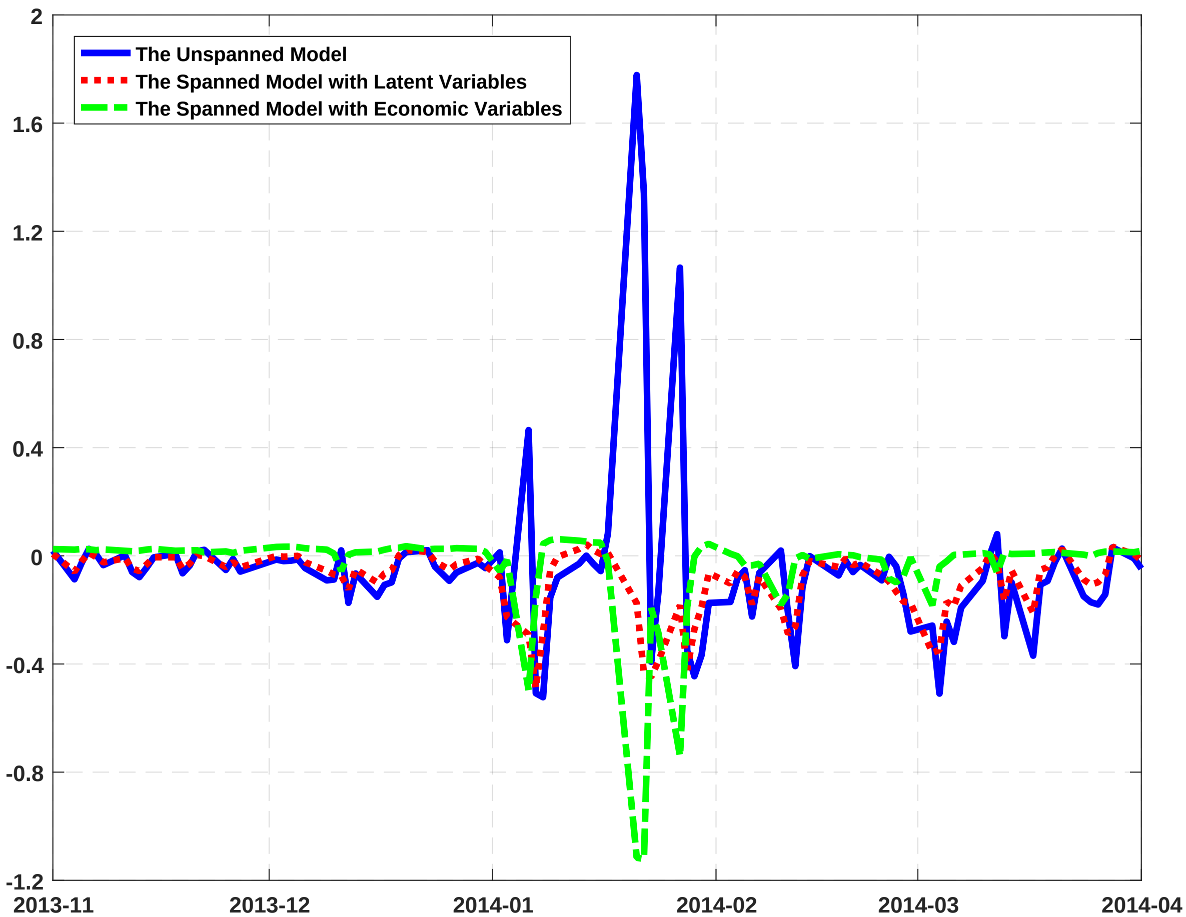

Figure 8: Decomposing the Spot Premium of the Unspanned Model

We decompose the spot premium of the unspanned model. Panel A plots the spot premium associated with PC^1, Panel B plots the spot premium associated with PC^2, Panel C plots the spot premium associated with the natural gas price, and Panel D plots the spot premium associated with the CDD. The details of the decomposition are given in the Appendix. The sample period is from May 2003 to May 2014.

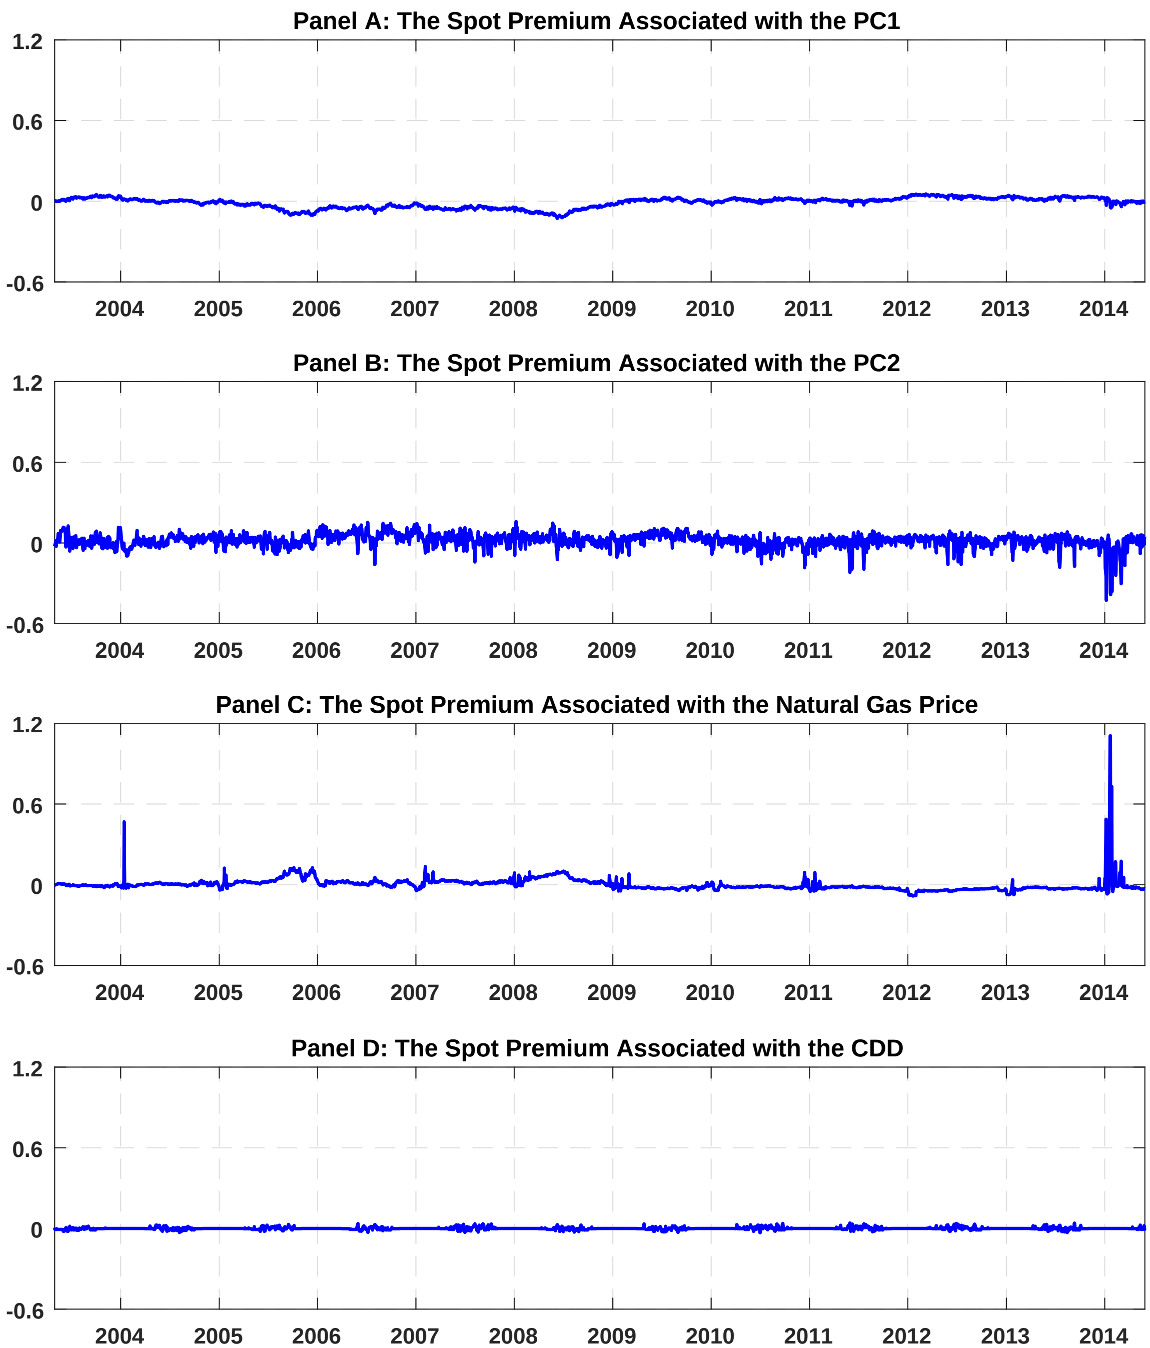

Figure 9: The Forward Bias. Various Contracts

We plot the forward bias in the unspanned model. For a futures contract maturing in month i, the forward bias is defined as follows:

$$\text{Forward Bias}_t^i = E_t^P[\bar{S}_{t+i}] - F_t^i$$

F_t^i denotes the price of the i-th month futures contract at time t and $E_t^P[\bar{S}_{t+i}]$ denotes the time-t average expected spot price in the maturity month (t+i), where the expectation is under the P measure. Panels A to C plot the weekly averages of the forward bias for contracts with maturities equal to 1 day, 1 month, and 12 months respectively. The economic variables that are used to calculate the expected spot price under P are the natural gas price and CDD. The sample period is from May 2003 to May 2014.

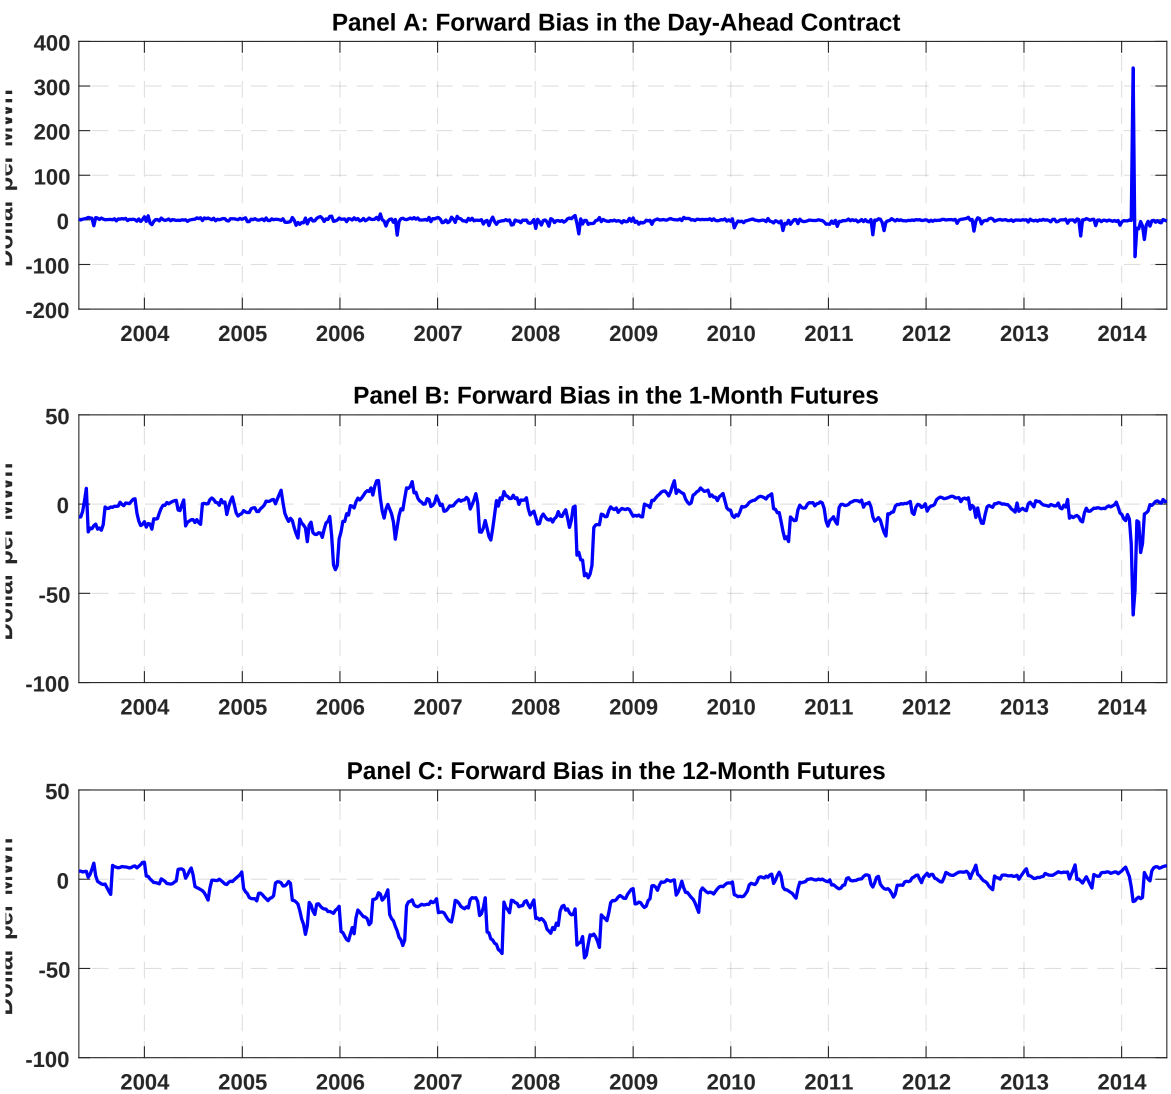

Figure 10: The Forward Bias in Different Seasons

We plot the average forward bias in winter (Panel A), spring (Panel B), summer (Panel C), and fall (Panel D). Panel E plots the average forward bias over the entire sample period. In each panel, the maturity ranges from 1 month (contract number = 2) up to 12 months (contract number = 13). For contract number 1, the maturity equals 1 day. The economic variables that are used to calculate the expected spot price under P are the natural gas price and CDD. The sample period is from May 2003 to May 2014.

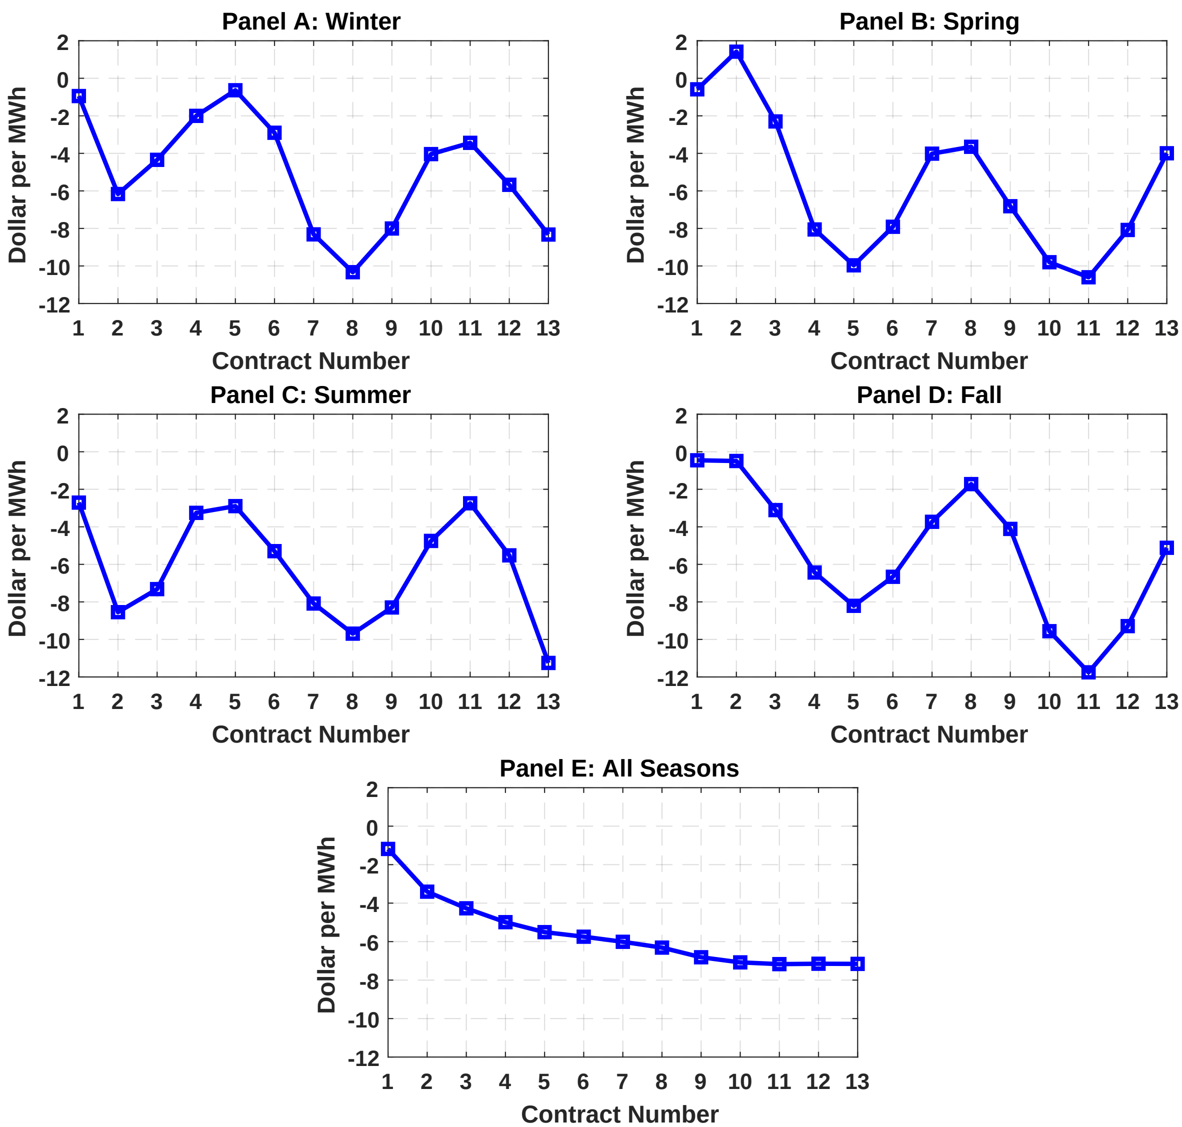

Appendix

A.1 Decomposing the Spot Premium of the Unspanned Model

We define the spot risk premium as the expected log return of holding a day-ahead contract. Then the spot premium is defined as follows:

$$\text{Spot Premium}_t = E_t^P[\text{Log}(\text{S}_{t+1}) - \text{Log}(\text{F}_t^{\text{DA}})] \tag{A.1.1}$$

where S_{t+1} denotes the electricity spot price at time t+1 and F_t^{DA} denotes the day-ahead price at time t. We can show that the spot premium is equal to the difference between the log expected spot under the P measure and the one under the Q measure plus a constant. The derivation is as follows:

$$\begin{aligned}
\text{Spot Premium}_t &= E_t^P[\text{Log}(\text{S}_{t+1}) - \text{Log}(\text{F}_t^{\text{DA}})] \\
&= E_t^P[\text{s}_{t+1} - f_t^{\text{DA}}] \\
&= E_t^P[\text{s}_{t+1}] - f_t^{\text{DA}} \\
&= E_t^P[\text{s}_{t+1}] - (E_t^Q[\text{s}_{t+1}] + \frac{1}{2}\sigma_\text{s}^2) \\
&= E_t^P[\text{s}_{t+1}] - E_t^Q[\text{s}_{t+1}] - \frac{1}{2}\sigma_\text{s}^2
\end{aligned} \tag{A.1.2}$$

where s_{t+1} denotes the log of the spot price at time t+1, f_t^{DA} denotes the log of the day-ahead price at time t, and σ_s denotes the volatility of the log spot price.

Suppose that we use [PC1, PC2, PX, CDD] as the state variables for the unspanned model, where PC1 denotes the first principal component of the log electricity prices (log spot price and log futures price), PC2 denotes the second principal component of the log electricity prices, PX denotes the natural gas price, and CDD denotes the cooling degree days. Then based on equation (12), the expected log spot price under P can be expressed as:

$$
\begin{aligned}
E_t^P[s_{t+1}] = {} & \text{Seas}_{s,t+1} + K_0^P(1) + K_1^P(1,\text{PC1}) \times \text{PC1}_t + K_1^P(1,\text{PC2}) \times \text{PC}_t^2 \\
& + K_1^P(1,\text{PX}) \times \text{PX}_t + K_1^P(1,\text{CDD}) \times \text{CDD}_t
\end{aligned}
\tag{A.1.3}
$$

where $K_0^P(1)$ denotes the first element of the column vector K_0^P, $K_1^P(1,\text{PC1})$ is the element at row one and column one of the matrix K_1^P, $K_1^P(1,\text{PC2})$ is the element at row one and column two of the matrix K_1^P, $K_1^P(1,\text{PX})$ is the element at row one and column three of the matrix K_1^P, and $K_1^P(1,\text{CDD})$ is the element at row one and column four of the matrix K_1^P.

Similarly, equation (13) implies that the expected price under Q can be expressed as:

$$
E_t^Q[s_{t+1}] = \text{Seas}_{s,t+1} + K_0^Q(1) + K_1^Q(1,\text{PC1}) \times \text{PC1}_t + K_1^Q(1,\text{PC2}) \times \text{PC}_t^2 \tag{A.1.4}
$$

where $K_0^Q(1)$ is the first element of the column vector K_0^Q, $K_1^Q(1,\text{PC1})$ and $K_1^Q(1,\text{PC2})$ are the first and second element of the first row of the matrix K_1^Q respectively.

Substituting equation (A.1.3) and (A.1.4) into equation (A.1.2) and, we can get

$$
\begin{aligned}
\text{Spot Premium}_t &= E_t^P[s_{t+1}] - E_t^Q[s_{t+1}] - \frac{1}{2}\sigma_s^2 \\
&= \text{Seas}_{s,t+1} + K_0^P(1) + K_1^P(1,\text{PC1}) \times \text{PC1}_t + K_1^P(1,\text{PC2}) \times \text{PC}_t^2 \\
&+ K_1^P(1,\text{PX}) \times \text{PX}_t + K_1^P(1,\text{CDD}) \times \text{CDD}_t - (\text{Seas}_{t+1} + K_0^Q(1) \\
&+ K_1^Q(1,\text{PC1}) \times \text{PC1}_t + K_1^Q(1,\text{PC2}) \times \text{PC}_t^2) - \frac{1}{2}\sigma_s^2 \\
&= \underbrace{(K_0^P(1) - K_0^Q(1)) - \frac{1}{2}\sigma_s^2}_{\text{Constant}} \\
&+ \underbrace{(K_1^P(1,\text{PC1}) - K_1^Q(1,\text{PC1})) \times \text{PC1}_t}_{\text{Spot premium associated with the PC1}} + \underbrace{(K_1^P(1,\text{PC2}) - K_1^Q(1,\text{PC2})) \times \text{PC2}_t}_{\text{Spot premium associated with the PC2}} \\
&+ \underbrace{K_1^P(1,\text{PX}) \times \text{PX}_t}_{\text{Spot premium associated with the PX}} + \underbrace{K_1^P(1,\text{CDD}) \times \text{CDD}_t}_{\text{Spot premium associated with the CDD}}
\end{aligned}
\tag{A.1.5}
$$

A.2 The Lucia and Schwartz Model

This section discusses the two-factor log price model that proposed by Lucia and Schwartz (2002). We first explain the model specification. Then we discuss its relationship with the spanned model with latent variables.

A.2.1 Model Specification

The log price is assumed to be consisted of a seasonal component with two latent factors. Denote the log spot price at time t as s_t, then s_t can be written as follows:

$$s_t = \text{Seas}_{s,t} + L_t + \epsilon_t \tag{A.2.6}$$

The seasonal component Seas_t is specified according to equation (24). L_t and ϵ_t are assumed to have the following P-dynamic:

$$dL_t = -\kappa L_t dt + \sigma_L dz_t^L \tag{A.2.7}$$

and

$$d\epsilon_t = \mu_\epsilon dt + \sigma_\epsilon dz_t^\epsilon \tag{A.2.8}$$

dz_t^L and dz_t^ϵ are two standard Brownian motions. The correlation between dz_t^L and dz_t^ϵ is equal to $\rho_{L\epsilon}$, i.e.

$$dz_t^L \times dz_t^\epsilon = \rho_{L\epsilon} dt \tag{A.2.9}$$

The market price of risk are assumed to be constant for both dz_t^L and dz_t^ϵ. Thus, the Q-dynamic can be written as follows:

$$dL_t = (-\kappa L_t - \lambda_L)dt + \sigma_L dz_t^L \tag{A.2.10}$$

$$d\epsilon_t = (\mu_\epsilon - \lambda_\epsilon)dt + \sigma_\epsilon dz_t^\epsilon \tag{A.2.11}$$

where λ_L and λ_ϵ are two scalars that represent the market price of risk of dz_t^L and dz_t^ϵ, respectively.

The log futures price can be expressed in closed form and the formula is

$$\begin{aligned} f_t^j &= E_t^Q[s_{t+j}] + \frac{1}{2}Var_t^Q[s_{t+j}] \\ &= \text{Seas}_{s,t+j} + e^{-\kappa\times j}L_t + \epsilon_t + \mu_\epsilon^* \times j - (1-e^{\kappa\times j})\frac{\lambda_L}{\kappa} \\ &+ \frac{1}{2}((1-e^{-2\kappa\times j})\frac{\sigma_L^2}{2\kappa} + \sigma_\epsilon^2 \times j + 2(1-e^{-\kappa\times j})\frac{\rho_{L\epsilon}\sigma_L\sigma_\epsilon}{\kappa}) \end{aligned} \tag{A.2.12}$$

where f_t^j denotes the log of the j-month futures price at time t. $\mu_\epsilon^* = \mu_\epsilon - \lambda_\epsilon$. We estimate the model using the Kalman filter.

A.2.2 Relationship with the Spanned Model with Latent Variables

The Lucia and Schwartz model can be written in the framework of the spanned model with the following matrices:

$$X_t^S = \begin{bmatrix} L_t \\ \epsilon_t \end{bmatrix}, \quad K_0^P = \begin{bmatrix} 0 \\ \mu_\epsilon \end{bmatrix}, \quad K_0^Q = \begin{bmatrix} -\lambda_L \\ \mu_\epsilon - \lambda_\epsilon \end{bmatrix} \tag{A.2.13}$$

$$K_1^P = \begin{bmatrix} 1-\kappa & 0 \\ 0 & 1 \end{bmatrix}, \quad K_1^Q = \begin{bmatrix} 1-\kappa & 0 \\ 0 & 1 \end{bmatrix} \tag{A.2.14}$$

and

$$\Sigma^P = \Sigma^Q = \begin{bmatrix} \sigma_L^2 & 2\rho_{L\epsilon}\sigma_L\sigma_\epsilon \\ 2\rho_{L\epsilon}\sigma_L\sigma_\epsilon & \sigma_\epsilon^2 \end{bmatrix}^{\frac{1}{2}} \tag{A.2.15}$$