



Power Price Simulation using Hybrid Models

By Carlos Blanco, Josh Gray and Marc Hazzard

Introduction

This article is the third in a multi-article series on the FEA Power Sector Model. The first article provided an overview of the issues that must be confronted in the realistic modeling of energy derivatives and the second article focused on the simulation of weather-contingent hourly load, with its inherent historical bent and data calibration issues. On the road leading to the valuation of real options associated with assets in the deregulated energy context, we will now focus in on the simulation of the spot power prices in the context of a hybrid model combining historical and forward looking information on weather, load, power and fuel prices.

The Challenge of Power Price Simulation

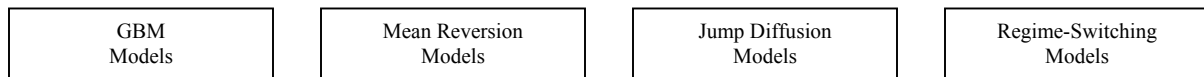
Financial Option vs. Real Options

The influx of Wall Street talent at the dawn of the merchant energy era resulted in a significant increase in the level of sophistication in pricing the embedded optionality in certain contracts that were considered traditional utility structures. Such esoteric contracts as swing, take or pay, peaking, and storage were analyzed as combinations of swaps, put options, call options and priced accordingly. This is the heart of the real options approach, which attempts to map financial option pricing techniques onto physical assets and offer more appropriate valuation and hedging methods. The valuation of these contracts using stochastic models have become industry standard (although approaches may vary).

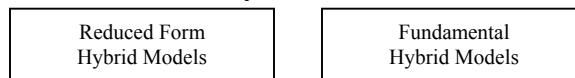
The valuation of generation assets and power-related real options has seen the same jump in sophistication but due to the unique nature and high complexity governing power price behavior, a similar level of industry standardization is yet to occur. A variety of approaches are still used.

Taxonomy of Power Models

Pure Stochastic Models



Hybrid Models



Fundamental Models





Pure Stochastic Models

Due to their ease of use and relatively short development times, pure stochastic models have gained a foothold in the valuation of generation assets and other power-linked deals. The availability of forward curves and the existence of data to estimate parameters assure that a valuation can be performed rapidly and hedge parameters generated.

One problem with the use of this kind of model to price the spread options associated with energy assets is the use of a static linear correlation coefficient to describe the joint dependence of power and fuel. Although energy modelers have long known that correlation may be ill-suited for this purpose, alternatives have not been tractably forthcoming. Since the market for spread options is either illiquid or unobservable, it forces us to use historical data for estimation, a murkier method than that of extracting implied correlation from the market.

Another problem are the models themselves (e.g. mean-reversion, jump diffusion, and regime-switching). Although intellectually satisfying, difficulties arise when it comes to parameter estimation. Since regime probabilities, mean-reversion rates, and jump diffusion parameters cannot be easily implied from the market, it forces us again to naively use historical power price data to estimate parameters or rely on some “fudge” factor to try to calibrate them against observed market parameters. The process of extracting more and more parameters from a limited data set can result in incorrect specification of future power price distributions.

Fundamental Models

Although they are neither easy to use nor do they have a short development time, fundamental models retained popularity from the merchant energy sector’s salad days as regulated utilities. These models attempt to use known natural and operational data to forecast prices, cash flows, and production data. Due to the illiquid nature of most power forward curves, these models actually still have some use. Fundamental models represent the knowledge base of experienced personnel.

The shortcomings of fundamental models are well known in the merchant energy sector. Price “forecasts” based on extending forward curves or equilibrium methods were the basis for many of the purchases of near-bankrupt entities. The lack of applicability to a risk-management framework in both the short and near terms may doom them to the dustbin of modeling tools.

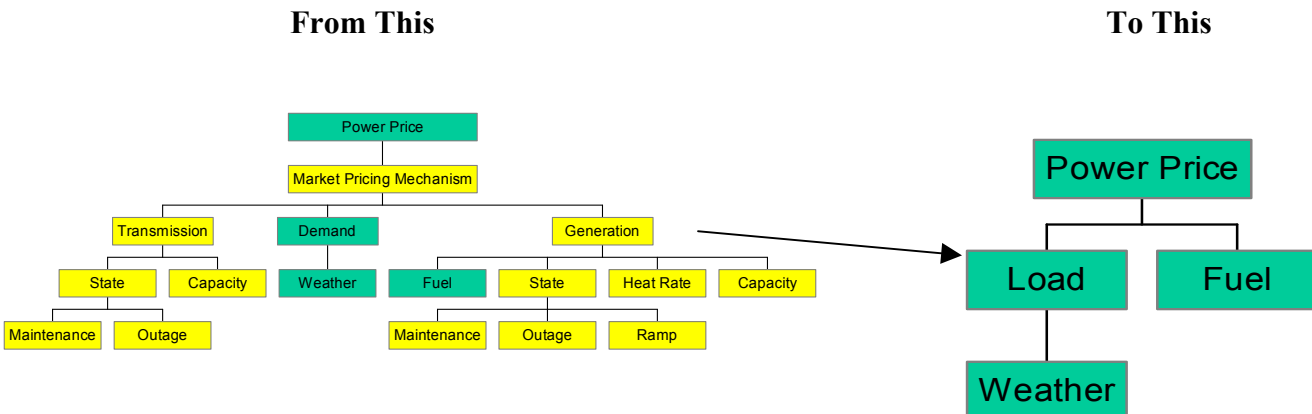
The Next Generation: Modeling the Joint Distribution with Hybrid Models

Hybrid models attempt to blend the best of both fundamental and pure stochastic models. They blend the information richness of the forward curve and the data richness of historical weather and load. They also replace unknown or uncertain parameters in the stochastic formulation and replace them with fundamentally grounded behavior in the hybrid model. Hybrid models share the relatively short calculation times of stochastic models while allowing us to obtain hedging parameters from the stochastic world and use them to protect cash flows.

Hybrid models are rapidly gaining acceptance amongst market practitioners. Eyedland and Wolyniec dedicate an entire chapter of their recently published book to describe hybrid models. Using their words, “Hybrid models can be thought of as not a model of specific volatilities, correlations and the like, but as a model of the relationship among all parameters of the joint distribution”.

The Process and Goals of Building a Hybrid Model

1. Identify the most salient (observable) state variables driving the evolution of power prices. Model the independent state variables as stochastic processes.
2. Conduct a detailed empirical study of the nature and relationships among the model components. Ascertain the functional form of these relationships.
3. Calibrate the proposed models for each of the state variables and power price.
4. Incorporate market conditions.
5. Simulate each stochastic factor via Monte Carlo simulation.





In specifying our hybrid modeling approach to simulating spot power prices, we must throw into the mix the fundamental and stochastic approaches. We will then select the most salient and computationally accessible components of both processes and specify our model.

Weather

In our selection process, we quantitatively and qualitatively identify weather as the fundamental driver of load (demand), which subsequently accounts for a large amount of variance in the price of power. So the joint behavior of load and weather, whose calibration and simulation was covered in a previous article, is chosen to represent the demand component of power prices. The combination of plentiful historical weather data, the stationary nature of weather, and a re-emerging traded market allows us to use weather simulation tools with some degree of confidence.

Load

With the demand component of power prices represented by the load/weather pair, we must choose a simulation method for this parameter. Unlike the future distribution of weather phenomena, the specification of the future distribution of load is an altogether different exercise. In areas where the customer and industrial component of demand are relatively stable, we can assume the relationship between weather and demand will be stable, and that the past provides a reliable indication of the future. Since there is no traded market for load, we will have to use historical load data and shape it in an intelligent manner to form the future distribution.

Fuel

Now that we have chosen the demand component, let's turn to the raw material of power prices, fuel (assuming either natural gas or oil). With its liquid forward and options markets, stochastic simulation can be used to specify the future distribution of fuel prices. There is no need to draw on historical or fundamental "forecasting" approaches here, but the framework that we are presenting also allows for modeling the joint dependence between variables such as weather and fuel.

Power Price

So far we have chosen a hybrid specification for our model, included the load/weather pair to represent demand and fuel price to represent the raw material. From these choices the power spot price follows as a consequence. Spikes do not happen due to a "random" jump laid on top of a stochastic power process but because of response to an extreme temperature, supply, or transmission shock. Power and weather are not related by a correlation coefficient but are bonded together by a demand response and the cost of fuel. Thus, we have taken a step closer to the fundamentals of power pricing.

The Market Marginal Heat Rate

Thus, we have a model that, while not perfect, can explain most phenomena in the specification of power spot prices. Excess demand is driven by temperature extremes. The expected cost of fuel is reflected in the forward price and a seasonally shaped volatility curve. The next step in specifying our hybrid approach is to deal with the components that can drive power prices to extreme levels. This step lies at the intersection of both excess demand and supply scarcity.

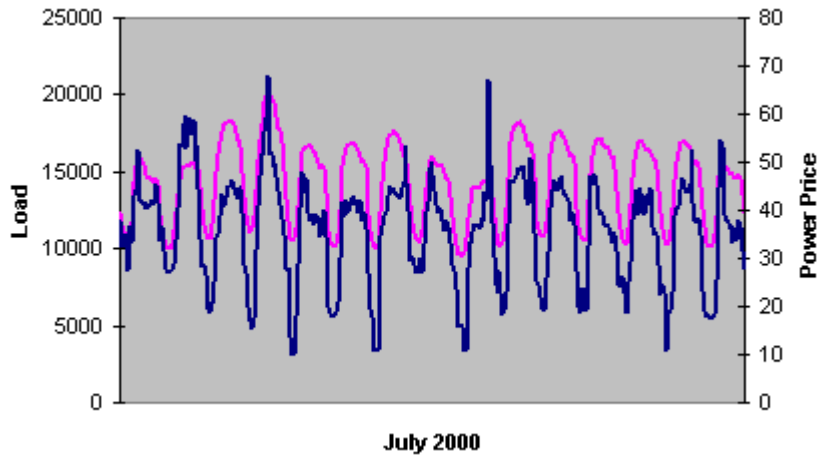


Figure 1: Hourly power price following hourly load level

Given excessive demand, yet plentiful supply, we do not see extreme changes in power prices. Imagine a scenario where a large amount of power plants are not available at a time with mild demand. There you will not find power prices spikes. This is the case in the spring maintenance season, when a variety of plants are inoperative for general upkeep.

On the other hand, a measure of the scarcity in the power supply is given by the market marginal heat rate (MMHR). The MMHR describes a linear relation between the price of power and fuel, which is defined as the ratio of the power price and the fuel price. Plotted against the load, a “hockey stick” type of relationship emerges (See figures 2 and 3, below).

Power Price Formation Process

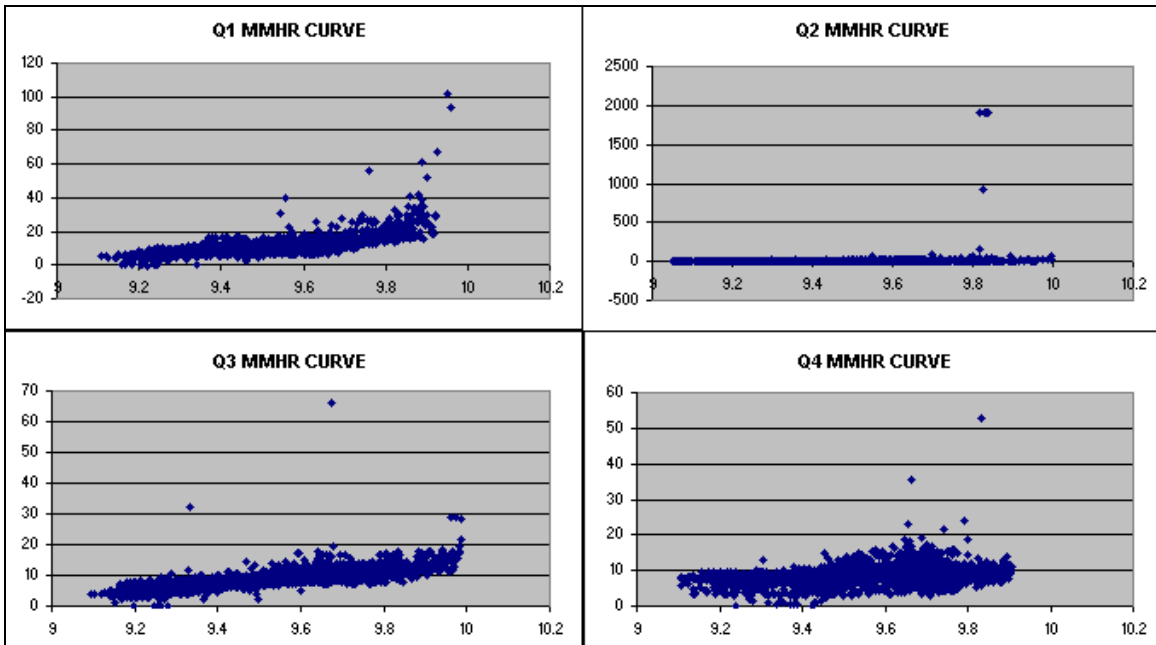
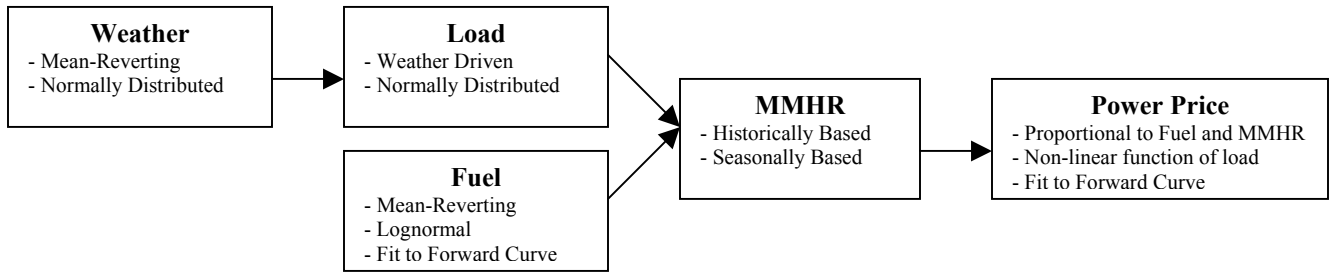


Figure 2: Quarterly MMHR Curves

Figure 1 depicts the “hockey stick” relation between the logarithm of hourly load and the hourly MMHR (hourly power price divided by daily fuel price) for each of the four quarters. The figures clearly demonstrate the seasonality of the MMHR curve through the range and clustering of the data points.

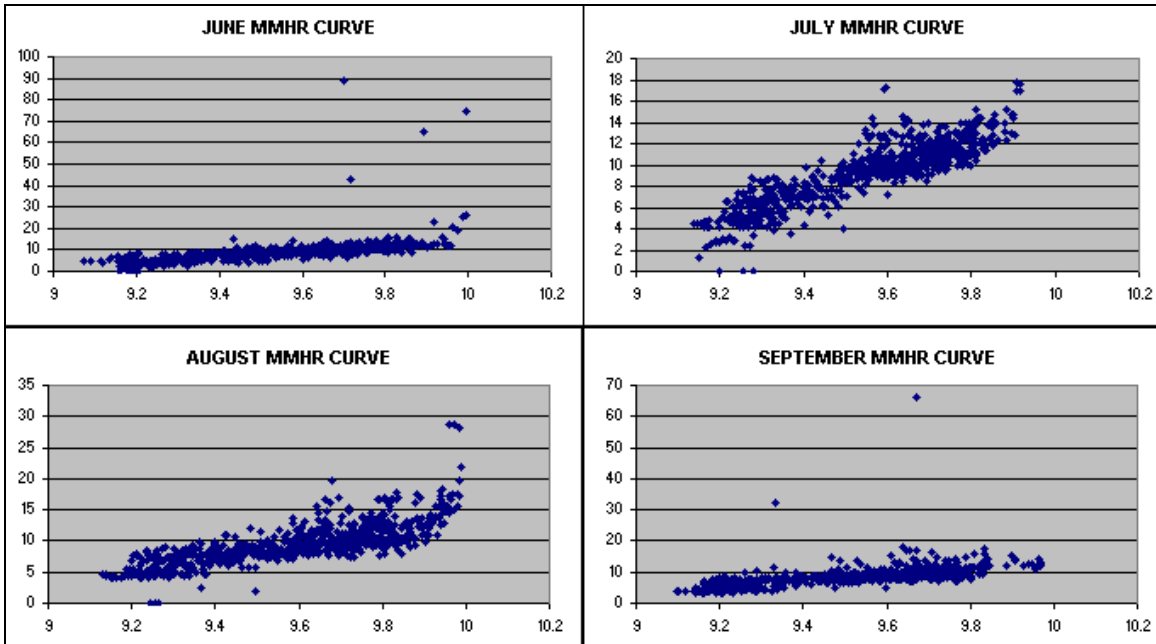


Figure 3: Summer MMHR Curves

Figure 2 represents a finer granularity snapshot of the MMHR curve, shown on a monthly basis instead. Notice the similarity in shape between July and August, and the distinction with the two shoulder months.

As the load increases and the demand starts to outstrip the supply generated from power plants, a sharp rise in the MMHR can be observed. Since this rise can be severe and unconnected to the usual equilibrium dynamics of power price formulation, a jump diffusion process is then used to model these extreme movements.

While naively using jump diffusion coefficients from the limited amount of historical power prices may be a cause for concern, we have expanded our methodology to account for the amount and quality of available historical weather-load data. The jump parameters - the jump probability, the expected jump size, and the jump standard deviation - are conditioned on the historical load level. When the weather-contingent, load simulation is made, a determination of whether a jump occurs, given that load level, is addressed via this conditioning. Figure 4 shows a historical justification of this jump calibration along with a Power Sector simulation.

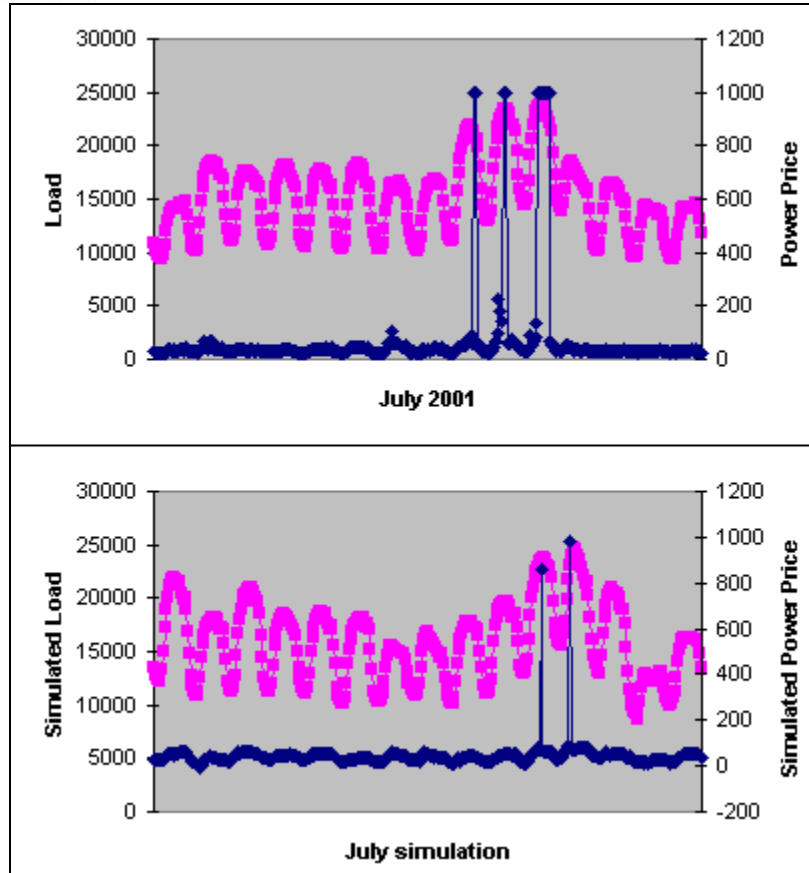


Figure 4: Historical and Power Sector-based simulation of hourly load and power price jumps

With plentiful historical weather data and structurally similar demand response, we have moved the estimation of these parameters from the power price time series to the more stationary weather-load pair. The base power price can then be recovered by multiplication with the fuel price. Given a supply stack, the historical MMHR can be used as a proxy for the stack and a measure of its inherent scarcity.

Conclusion

In FEA’s Power Sector model, the challenging problem of power price simulation has been addressed through a “hybrid” modeling process. This “best practices” approach allows us to focus on the strengths of both the stochastic and fundamental models, without relying on the weaknesses or limitations of any one model.

The Market Marginal Heat Rate provides an effective, and simple tool to scaling the load-following hourly price of power. Additionally, calibrating jump parameters conditioned on load allows for an explicit context to determine the timing of extreme events. We can then use FEA’s Power Sector implementation to value contracts whose value depends on the price of power, such as generation assets and load contracts. In the next two papers, we will examine how this is done.



Bibliography

- Soronow, D., Pierce, M., and Wang, K. 2002. The Power Sector Model. New Frontiers. January.
- Blanco, C., Gray, J., and Hazzard, M. 2003. Finally....A realistic model of power prices: The FEA Power Sector Model. The Risk Desk. Vol. III, No. 3.
- Blanco, C., Gray, J., and Hazzard, M. 2003. Weather Contingent Load Simulation. The Risk Desk. Vol. III, No. 4.
- Eyedland A., and Wolyniec, K. 2003. Energy and Power Risk Management: New Developments in Modeling, Pricing, and Hedging.
- Embrechts, P., McNeil, A., Straumann, D. 2002. Correlation and dependence in risk management: properties and pitfalls. In: Risk Management: Value at Risk and Beyond

Dr. Carlos Blanco directs the Global Support and Professional Services Group at FEA, a Barra Company. Dr. Josh Gray is a Senior Derivatives Specialist and Marc Hazzard is a Derivatives Specialist at the same group. The authors thank Craig Jimenez, David Soronow, Michael Pierce, Dr. King Wang, and Dr. Angelo Barbieri for advice and assistance.