

Hourly resolution forward curves for power: statistical modeling meets market fundamentals

Michael Coulon, Christian Jacobsson, Jonas Ströjby

May 5, 2014

Notes / Acknowledgments:

This work appears as a book chapter within the publication *Energy Pricing Models: Recent Advances, Methods and Tools*, edited by Marcel Prokopczuk, published in Dec 2014 by Palgrave Macmillan. The first few pages (through Section 1.1) are reproduced here with permission of Palgrave Macmillan. This extract is taken from the authors' original manuscript and has not been edited. The definitive, published, version of record is available here:

<http://www.palgrave.com/page/detail/energy-pricing-models-marcel-prokopczuk/?K=9781137377340>

Contents

1	Introduction	2
1.1	The Challenge: Hourly Forward Curve Construction	4
1.2	Market Fundamentals: The Case of the German Power Market	6
1.3	Statistical Modeling: Existing Techniques and New Ideas	8
2	The Stack Model	10
2.1	Supply Curves, Demand Curves and Price Curves	12
2.2	The Composition of the Stack	16
2.3	The ‘Must-Run’ Stack	17
2.4	The Thermal Stack	19
3	The Fundamental Factors	23
3.1	Demand	24
3.2	Wind	27
3.3	Solar	29
3.4	Nuclear and Combined Heat and Power	30
3.5	Fossil Fuels: Lignite, Coal, Gas and Oil	31
4	Forward Curve Construction	32
4.1	From Spot to Forward	32
4.2	Mixtures of Exponential Stacks	34
4.3	Pairwise Overlap Approximation	36
4.4	The Stochastic Factors	38
4.5	Calibration to the Quotes	40
5	Discussion and Conclusion	42
A	Appendix - Details of Pairwise Overlap Approximation	46

1. Introduction

Electricity markets represent the ultimate challenge for an avid modeler of financial markets and prices. Their complexity and ever-changing structure requires understanding and capturing a huge range of different factors and effects. Thus, weaknesses can always be found in existing models, and consensus on which model is best can rarely be reached. And yet, the fundamental transparency of the electricity price setting mechanism, the large amount of

available data and the rapidly growing literature all tempt the modeler to try a bit harder: perhaps that one elusive model that ticks all the boxes is out there somewhere, just beyond reach!

Of course, we must be realistic, and certainly we do not claim to present the perfect model here. Indeed, as every electricity market is different and every modeling goal different too, the ultimate prize will (and should rightly so) always remain out of reach. However, we do hope to make a contribution in a very promising direction, presenting a flexible stack-based approach, which incorporates as much as possible of the crucial market fundamentals while still retaining user-friendliness and computational tractability. We aim to exploit the growing number of mathematical modeling ideas emerging from academia in this direction, but complement these with a healthy dose of practical industry experience, substantial data analysis, and a realistic focus on what really matters for practitioners in these markets.

Our aim is thus to strike the right balance between the heavy fundamental models traditionally used in the energy industry, and the nimble reduced-form approaches which often migrate across from quantitative finance and econometrics. The former approach can be described for many reasons by the word slow. Data collection is often messy and painfully slow, involving large amounts of information on individual generators and their operational constraints, in order to build up a very detailed supply stack. More importantly, implementation of such giant models tends to rely on large-scale optimization techniques, rendering them too slow to handle stochasticity well, particularly when attempting to calibrate the outputs of spot price scenarios to observed market forward quotes. Analyzing model output or parameter sensitivity for rapid trading, hedging or derivative pricing decisions is not feasible, relegating these slow models instead more towards the realm of scenario testing, price forecasting and long-term investment planning.

At the other end of the spectrum, classical reduced-form stochastic models for spot and forward prices are fast in many ways. Writing down a model is fast, estimating parameters usually fast, and if the chosen stochastic processes are sufficiently convenient, calibration to forwards and even options can potentially be very fast. On the other hand, speed and convenience can come at a very high price in power markets. Reduced-form models typically tell you to throw away all your detailed data on weather forecasts, demand

patterns, changes in the generation mix and perhaps most worryingly also your common sense that parameters estimated from price histories are no longer reliable given fundamental market changes underway. Having to wait for years for today's market news to make it into our historical price data and corresponding parameter estimates is an issue on which our so-called 'fast' models are arguably relegated to being the slowest of all! This is not to say that there is no use for either the 'fast' or the 'slow' models, but simply that both have significant disadvantages that can hopefully be limited by looking for a hybrid approach. This idea is certainly not new, and forms part of a growing branch of 'structural' models in the literature, as discussed for example in the recent survey paper of [10].

1.1. The Challenge: Hourly Forward Curve Construction

Before we outline our specific contribution to the literature here, let's briefly address the question of why any of it matters. Why is it fundamentally important for an energy company to have a realistic, reliable, flexible and tractable approach to generating spot price dynamics and forward curves? Why is this particularly crucial nowadays, suggesting the need for new and innovative approaches? Firstly, it is well-known that electricity spot prices can be extremely volatile and difficult to predict, with dramatic but short-lived spikes frequently occurring in both the upwards and downwards directions, as illustrated later in Figure 2a for the German market (EEX). Together with the correlated risk of shocks to load itself, this produces substantial risk management challenges for utilities and power generators, particularly as hedging weather or outage related risk may be impossible. Trading in forward contracts over various maturities is a natural first step to managing electricity spot price risk. However, not all maturity forwards are traded, some suffer from illiquidity, and even the most liquid of contracts have delivery periods of a week, a month or longer, nowhere near the hourly granularity of spot. In sum, the forward curve that the market provides is often far from ideal for many practical applications.

Instead, hourly granularity forward curves can be constructed by averaging over scenarios from an hourly granularity spot price model. This is a very common and yet challenging task for many energy companies. Hourly spot price simulations and hourly forward curves are needed as inputs for a wide

range of applications, from risk evaluation and P&L calculations (mark to market) to virtual contract evaluation or the optimal operation and valuation of a physical asset. Managing hydro plants requires hourly prices to accurately capture the ‘optionality’ embedded in the plant’s operation, whereby water can be pumped up when power is cheap and released again just hours later for a profit. Similarly, flexible gas plants can rapidly switch on and off to exploit hourly price swings and require hourly price curves for both gas and electricity.

Industry experience suggests that there is much room for improvement on this topic, making it a very suitable testing ground for new modeling ideas. The short and long ends of the forward curves require different considerations, which can rarely be managed by a single model. For example, a reduced-form stochastic model can often capture quite effectively the volatility term structure of power forwards (i.e. the Samuelson effect), intra-day and seasonal price patterns, and long-term forward curve shapes. However, to accurately construct a realistic forward curve for the coming days and weeks, it is vital that known weather forecasts are used as inputs, impossible for a pure reduced-form approach. On the other hand, models for demand, wind and solar typically focus on getting the best day-ahead or short-term forecast, and are less amenable to long-term hourly simulations needed for multi-year forward curve. At the long end of the curve, it is often desirable to extrapolate beyond the liquid market quotes to construct a power curve going out five to ten years or more. In this case, the extra information at one’s disposal may be long maturity fuel forward quotes, and growth forecasts for new capacity (both renewable and conventional), which can only be realistically used as inputs in conjunction with a model for the progression of the merit order over time. Moreover, while the heavy full fundamental models can project forward a forecasted merit order over many years, they crucially fail to handle the randomness: various different merit order scenarios are possible over long time horizons, each with different probabilities attached, and each affecting power prices in different ways via the highly non-linear supply stack structure.

The challenges described above have become increasingly complex in recent years, due to a number of major ongoing changes in the European electricity markets. For example, together with rapid growth of renewables, high gas prices in Europe have kept many gas generators high up the stack and

‘out-of-the-money’, leading to risk of closures. At the same time, ongoing ‘market coupling’ to integrate European electricity markets has led to lower price spreads between countries, as we shall mention again in Section 5. Such developments demand new modeling ideas to handle price fluctuations rarely seen in historical data. Finally, an ever-increasing amount of historical and forward-looking data is now publicly available, making it even more advantageous to build a flexible enough model to allow for more inputs such as detailed demand and capacity forecasts. (e.g. see [18, 12, 5]) Structural models can fairly easily adapt to these market changes and to new information, as each component of the model can be adjusted to reflect current conditions and then bolted back into place.

In this work, we adapt and extend the existing literature in this field by applying a multi-fuel structural stack model to the challenging German power market, illustrating its ability to calibrate to market data and capture key characteristics of the market. We demonstrate a novel approach to obtain a stable fit of the stack model to *both* bid and offer curve data, as required due to the large demand-side elasticity on EEX which we shall explore. Furthermore, we pay specific attention to the rapidly growing wind and solar capacity in Germany, and more generally describe the impact of the various types of generator capacity entering at different points in the stack. Importantly, we also demonstrate how to choose stochastic processes for each of the fuel prices and for residual demand (after removing renewable supply) such that a closed-form approximation to power forward prices can be obtained as a function of fuel forwards. In this manner, we facilitate rapid calibration to market quotes, without needing to resort to a reduced-form price model, and we construct hourly forward curves going out many years. Finally, we discuss future challenges to be faced in the evolving market environment.

1.2. Market Fundamentals: The Case of the German Power Market

The first step to building an effective spot price model is....

Request PDF | On Jan 1, 2014, Michael Coulon and others published Hourly Resolution Forward Curves for Power: Statistical Modeling Meets Market Fundamentals | Find, read and cite all the research you need on ResearchGate. An updating study concerning the analysis and modeling of the European Energy Exchange power market in Germany is urgently called for due to the dramatic transformations of the market in recent years. Based on careful data research, this paper provides a consolidated and detailed view concerning the development of the market and its mechanisms. Hourly Resolution Forward Curves for Power: Statistical Modeling Meets Market Fundamentals more. by Michael Coulon. Publication Date: 2014. We develop a fundamental model for spot electricity prices, based on stochastic processes for underlying factors (fuel prices, power demand and generation capacity availability), as well as a parametric form for the bid stack function more. We develop a fundamental model for spot electricity prices, based on stochastic processes for underlying factors (fuel prices, power demand and generation capacity availability), as well as a parametric form for the bid stack function which maps these price drivers to the power price. Any forward product with delivery length being equal to the market granularity is an unitary product. Definition 2.1.3. We use APX half-hourly prices as a proxy for spot prices when we need to build a model with half-hourly granularity. APX half-hourly prices are in fact volume weighted average traded prices for very short term forwards. Another application will be to build models with daily granularity. In bond pricing applications instead of modelling via spot models the forward rates are specified [5]. Clewlow and Strickland [8] propose a similar approach for energy markets which relies on taking a forward curve and simulating how this forward curve will evolve through time. The Stochastic Differential Equation describing the process followed by the unitary forward curve $F(t, T)$ is Get accurate forward prices built on industry-specific methodologies for crude and refined oil products, coal, power, natural gas and LPG/NGLs. Speak to a specialist. Register below and we will customize a solution that meets your exact needs. When you speak to one of our experts, you may be qualified to sample our industry-leading products on a no-cost basis. Argus produces information and services for global markets, which requires Argus to share your personal information with Argus group companies and service providers that are based both within and outside of the European Economic Area (EEA). Argus Media may use the details submitted to send you information about related Argus products and services which may be of interest to your However, this approach requires extensive statistical analysis of market data. In this paper, we examine if and how this statistical work can be reduced using machine learning. Our paper focuses on two research questions Hourly Resolution Forward Curves for Power: Statistical Modeling Meets Market Fundamentals Hourly Resolution Forward Curves for Power: Statistical Modeling Meets Market Fundamentals. M. Prokopczuk (), Energy Pricing Models. Recent Advances, Methods, and Tools Energy Pricing Models.