



City Research Online

City, University of London Institutional Repository

Citation: Fusai, G., Caldana, R. & Roncoroni, A. (2017). Electricity Forward Curves with Thin Granularity: Theory and Empirical Evidence in the Hourly EPEX Spot Market. *European Journal of Operational Research*, 261(2), pp. 715-734. doi: 10.1016/j.ejor.2017.02.016

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/17018/>

Link to published version: <https://doi.org/10.1016/j.ejor.2017.02.016>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

Accepted Manuscript

Electricity Forward Curves with Thin Granularity: Theory and Empirical Evidence in the Hourly EPEX Spot Market

Ruggero Caldana, Gianluca Fusai, Andrea Roncoroni

PII: S0377-2217(17)30122-4
DOI: [10.1016/j.ejor.2017.02.016](https://doi.org/10.1016/j.ejor.2017.02.016)
Reference: EOR 14253



To appear in: *European Journal of Operational Research*

Received date: 4 May 2016
Revised date: 7 January 2017
Accepted date: 10 February 2017

Please cite this article as: Ruggero Caldana, Gianluca Fusai, Andrea Roncoroni, Electricity Forward Curves with Thin Granularity: Theory and Empirical Evidence in the Hourly EPEX Spot Market, *European Journal of Operational Research* (2017), doi: [10.1016/j.ejor.2017.02.016](https://doi.org/10.1016/j.ejor.2017.02.016)

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Highlights

- Constructive definition of electricity forward price curve with arbitrary granularity
- Forward curve level is jointly consistent to risk-neutral and historical market information
- Curve shape embeds periodical patterns unveiled by past quotes
- Curve level is jointly consistent to baseload and peak-load futures quotations
- Curve path is smooth, monotonicity preserving, cross-sectionally stable, time local

ACCEPTED MANUSCRIPT

Electricity Forward Curves with Thin Granularity

Theory and Empirical Evidence in the Hourly EPEX Spot Market

Ruggero Caldana, Gianluca Fusai, Andrea Roncoroni*

Abstract

We propose a constructive definition of *electricity forward price curve* with cross-sectional timescales featuring hourly frequency on. The curve is jointly consistent with both risk-neutral market information represented by baseload and peakload futures quotes, and historical market information, as mirrored by periodical patterns exhibited by the time series of day-ahead prices. From a methodological standpoint, we combine nonparametric filtering with monotone convex interpolation such that the resulting forward curve is pathwise smooth and monotonic, cross-sectionally stable, and time local. From an empirical standpoint, we exhibit these features in the context of EPEX Spot and EEX Derivative markets. We perform a backtesting analysis to assess the relative quality of our forward curve estimate compared to the benchmark market model of Benth *et al.* (2007).¹

*Ruggero Caldana is with Accenture S.p.A, Via Maurizio Quadrio, 17, 20154 Milan (Italy); Gianluca Fusai (*Corresponding Author*) is with Department of Economics and Business Studies (DiSEI), Università del Piemonte Orientale “Amedeo Avogadro”, via E. Perrone 18, 28100 Novara (Italy), and Faculty of Finance, Cass Business School, 106 Bunhill Row, EC1Y 8TZ London (United Kingdom), Tel.: +44(0)2070408630, E-mail: Gianluca.Fusai.1@city.ac.uk; Andrea Roncoroni is with Finance Department, ESSEC Business School, Avenue Bernard Hirsch, BP 50105, 95021 Cergy-Pontoise (France).

¹*JEL classification*: C31, E43, G13, Q02. *Keywords*: Energy Finance, Forward Pricing, Electricity Markets, Forward Curve Construction.

Acknowledgements: The authors are indebted to Anna Battauz, Francesco Corielli, Rachid Id Brik, Rüdiger Kiesel, Florentina Paraschiv, Francesco Saita, Enrico Telini, Fabio Trojani, Stefan Trueck, Florian Ziel, and the participants to the Energy Finance 2014 Conference (Erice), the 54th Meeting of Euro Working Group for Commodities and Financial Modeling (Milan), the Brown Bag Finance Series 2015 at Bocconi University (Milan), the 2016 International Symposium on Energy and Finance Issues (Paris), and the Energy Finance Christmas Workshop (Essen), for useful conversations and constructive comments. Andrea Roncoroni would like to thank Stefano Caselli and Carlo Ambrogio Favero for a fruitful environment provided during a sabbatical period at Bocconi University and the participants to the Monthly Webinar Series at the Energy and Commodity Finance Research Center, ESSEC Business School (Paris). All remaining errors are ours. We acknowledge financial support by Università degli Studi del Piemonte Orientale, Gruppo Nazionale per l'Analisi Matematica, la Probabilità e le loro Applicazioni (GNAMPA), CERESSEC, and Labex MME-DII.

1 Introduction

The *electricity forward curve* is a key piece of information subsuming future views of power market participants. This quantity is a function mapping each element in a set of future time intervals into the corresponding fair price for forward delivery of one megawatt-hour of electrical power.² A standard setting assumes evenly spaced and continuous time periods spanning the tradable maturity spectrum. Their length defines the *curve granularity*, which we refer to as “thin” for hourly frequency on. We consider the problem of building an electricity forward curve with hourly granularity. This is customarily referred to as the *hourly price forward curve* (HPFC).³ The driving principle underlying our approach is market representativeness. In particular, curve construction ought to use available market prices at best. The proposed analysis considers two major strands of market data.

Risk-neutral information stems from futures price quotations. Forward prices with hourly granularity cannot directly be recovered from these figures: the inverse problem of getting hourly prices out of futures market quotes with longer maturities is by and large undetermined. Market practice tackles the issue to a limited extent by adopting the curve fitting toolkit borrowed from fixed income analysis.⁴ Unfortunately, a direct application of these methods leads to forward curves lacking most of the stylized features required by power market operators. *Historical information* emerges from the past evolution of quoted spot prices. Power quotations stem from the interaction between supply and demand of electricity. This interaction exhibits sharp periodical patterns at frequencies ranging from intra-day up to a season. It turns out that the implied averaging process leading to forward quotes over periods of varying length smoothes these patterns, making them undetectable upon recovering hourly forward prices from futures quotations. In addition, quoted delivery periods usually exhibit varying durations. This fact often entails observing quotations referring to intersecting delivery periods, what leads to an ambiguous forward price assessment over certain time intervals. A further complication is presented by a possible lack of futures quote records for some segments in the maturity domain, a fact which leaves the forward curve undetermined for those portions of time.

In summary, while risk neutral data represented by standing futures prices allows us to partially determine forward curve *level*, it provides little clue about the actual *shape* of the curve. This effect is noticeably exacerbated with increasing time granularity, a phenomenon which reaches its

²Forward price is defined as the fair tariff the market quotes for a zero-cost commitment to future delivery over an entire period of time. For the purpose of this study, we interchangeably use terms “forward” and “futures” based on the assumption of a predictable cost of carry.

³Three major considerations underpin focusing on hourly granularity. First, hourly quotation is a standard for the vast majority of power markets across the world; note, however, that instances of half-hourly granularity exist as well, *e.g.*, UK and New Zealand. Second, published literature on hourly forward prices is particularly scant in contrast to the increasing number of contracts prescribing hourly delivery. Finally, and perhaps most importantly, disposing of hourly forward quotes allows one to price most physical and financially settled power contracts (Burger *et al.* (2014)).

⁴Term structures of interest rates are primarily inferred from bond price data. Existing techniques ultimately combine appropriate bond price stripping methods coping with overlapping payment schedules and interpolation methods allowing for filling in gaps appearing across the maturity dimension (Anderson *et al.* (1996)).

acme in hourly resolution. We argue that historical data offers the missing piece of information. A major contribution of this study is to unveil this finding and devise a way to use historical spot price data to sharpen the assessment of an HPFC compatible to market quotes in a rationally consistent manner.

The paper is organized as follows. Section 2 summarizes existing literature on the subject and points out a direction for this study. Section 3 develops the underlying rationale by inspecting the nature of market information and by introducing the notion of *forward kernel*. Section 4 proposes a metadefinition of *rational forward curve* and offers a corresponding algorithm. Section 5 describes preparatory steps, notably spot data outlier filtering and futures price segmentation. Section 6 implements the model into an algorithm leading to a constructive definition of HPFC. Section 7 empirically investigates the case of the EPEX power market and tests the quality of this proposal against a benchmark market model. Section 8 concludes with a few comments and suggestions for future research work.

2 Review of Literature

Some of the issues described so far have been partially addressed by the existing literature. To the best of our knowledge, Fleten and Lemming (2003) first apply curve fitting methods to electricity markets. Using the proprietary Multiarea Power Scheduling (MPS) model, these authors obtain weekly equilibrium prices and production quantities. They derive electricity forward curves through nonparametric estimates of daily forward prices obtained by minimizing the least-squares distance between target values and MPS equilibrium prices. The problem is solved under price constraints related to bid-ask spreads and curve smoothness properties.

Koekebakker and Os Adland (2004) and Benth *et al.* (2007) propose a model whose terms are made fully available to the user. They assume that forward curves combine seasonal paths with fourth order polynomial splines. Using the maximum smoothness interpolation (MSI) method developed by Adams and van Deventer (1994), these authors derive a daily forward curve fitting a set of market quotes and minimizing a convexity measure of the curve shape. The construction algorithm is effectively applied to Panamax Time Charter freight and Nordpool electricity markets.

Borak and Weron (2008) observe that the two methods described above are sensibly prone to model risk. These authors focus on the tight dependence between seasonal component assignment and the output curve. They propose a dynamic semi-parametric factor model with no explicit representation of periodical terms. As a result, they deliver a parsimonious, smooth, and seasonal forward curve estimate with daily granularity. However, their algorithm may suffer from underfitting market prices and may fail to account for short-term periodical patterns.

The problem of building electricity forward curves with finer granularity than a single day has rarely been addressed by the existing literature. Hildmann *et al.* (2012) propose a framework to

calculate an HPFC by combining parametric estimation and future price prediction on an hourly basis. They estimate a linear model through suitable norm minimization under constraints. These constraints actually ensure model consistency to quoted forward price. They also prevent from arbitrage opportunities in agreements with arbitrage pricing theory. However, this study neither reports any quality assessments, nor does it investigate empirical performance.

Paraschiv *et al.* (2015) further refine the method offered by Benth *et al.* (2007) to account for hourly granularity. They obtain an HPFC by additively superposing an exogenously estimated hourly pattern to a seasonal component, which in turn combines daily and monthly dummies with weather-linked derivatives, in light of their ability to capture temperature related effects. Although peakload price fitting is out of their model scope and temperature forecasting is known to offer significant estimates primarily over a short time horizon, the proposed algorithm is shown to provide reasonable hourly forward prices.

The main contribution of this paper is to put forward a rational definition of HPFC with thin granularity that complies with desirable features unaccounted for by the existing literature. The attribute “rational” refers to the following five properties:

- Raw price series undergo a *filtering* procedure to finely detect and single out data outliers;
- Curve shape incorporates a comprehensive bundle of *periodical patterns* unveiled by past quotes;
- Curve level is jointly consistent to standing *baseload and peakload* futures quotations;
- Curve path satisfies *regularity properties*, including trajectorial smoothness, monotonicity preservation, cross-sectional stability, and time localness (Hagan and West (2006));
- Forward estimate quality is assessed through dedicated empirical tests.

We remark that the proposed definition is fully constructive and self-contained: we make no use of data (*e.g.*, temperature records in Paraschiv *et al.* (2015)⁵) other than electricity spot and futures prices, nor do we adopt any exogenous model (*e.g.*, the MPS in Fleten and Lemming (2003)) to benchmark the proposed construction. As a result, we come up to a flexible algorithm covering virtually all electricity market instances.

⁵Appendix B.1 discusses the role of temperature records in the construction of forward curves.

Market prices	Delivery period
Day-ahead	Each of twenty-four hours in a day
Baseload futures	All hours in a weekend, week, month, quarter, or calendar year
Peakload futures	On-peak hours (8:00 AM \rightarrow 8:00 PM) in a week, weekend, month, quarter, calendar year
Off-peak futures	Off-peak hours (8:00 PM \rightarrow 8:00 AM) in a week, weekend, month, quarter, calendar year

Table 3.1: **Primitive market data.** Day-ahead prices quote for delivery over each hour on the following day; baseload futures prices quote for delivery over all hours on forthcoming weekends, weeks, quarters, and calendar years; peakload futures prices quote on the same delivery periods as those of baseload futures, while for a restricted number of hours, usually spanning the twelve hours comprised between 8:00 AM and 8:00 PM; off-peak futures prices refer to hours complementary to those of peakload futures, that is from 8:00 PM to 8:00 AM.

3 Market Prices and Patterns

Constructively combining risk-neutral and historical price information into a term structure of forward prices requires that one defines the notion of *forward kernel*, a theoretical quantity generating forward curves with arbitrarily assigned granularity.

3.1 The Forward Kernel

Electrical power is useful for practical purposes provided it is continuously dispatched over a time frame. Deregulated power markets allow for trading contracts entailing physical delivery or financial settlement over a variety of time intervals in the future. As an example, we may consider the general structure of most electricity spot markets. Each day contracts are traded for physical dispatching power on each hour in the following day. By combining agents' bid and offer quotes, the market exchange posts a *day-ahead price* for each contract. Day-ahead prices play the additional role of indices underlying power contracts tariffs, including most financial derivatives written on electricity.

We also consider cash-settled claims written on average day-ahead quotes, of which futures contracts represent the most actively traded instances. These are financial assets whereby the holder receives continuous delivery of electricity over a time period for a fixed price negotiated at inception. Futures exchanges provide traders with quotes for delivery over entire weekends, weeks, months, quarters, and calendar years. *Baseload* futures refer to dispatch periods comprising all the hours on a given day, while *peakload* futures refer to delivery on on-peak hours, namely over a time frame comprising the highest daily consumption figures.⁶ Table 3.1 summarizes all major price data available at these power markets.

The notion of forward kernel arises as a tool to conveniently link an arbitrary number of futures quotes. In broad terms, the time t quoted *continuous forward kernel* $f_t(u)$ is a function recovering

⁶The exact definition of on-peak and off-peak hours depends on the market under consideration. Most market-places define on-peak hours as weekday hours comprised between the ninth hourly block of the day, dispatching power from 8:00 to 9:00 AM, and the twentieth hourly block of the day, dispatching power from 7:00 to 8:00 PM (Burger *et al.* (2014)).

futures prices as arithmetic averages over a period of delivery.⁷ Given an evaluation date t , a day-count convention “-” : $(t, T) \rightarrow T - t$ specifying time duration, and a futures (baseload) quote F^B covering the period comprised between τ^b and τ^e , we stipulate:

$$F^B = \frac{1}{\tau^e - \tau^b} \int_{\tau^b}^{\tau^e} f_t(u) du. \quad (3.1)$$

The existence of forward kernel is traditionally granted by assuming an underlying arbitrage-free pricing model. The argument is as follows. In addition to the quantities introduced above, let $(\Omega, \mathcal{F}, (\mathcal{F}_u)_{u \geq 0}, \mathbb{P})$ be a complete stochastic basis and r be a continuously compounded constant rate of interest representing the time value of money. We consider a stochastic process $(S(u), u \geq 0)$ for spot price dynamics of electricity and assume the resulting market model is free of arbitrage opportunities. In particular, we assume the existence of a risk-neutral probability measure \mathbb{P}^* allowing for pricing financial claims written on S . The fair value of a baseload forward contract delivering power over a period $[\tau^b, \tau^e]$ for a time τ^e cash-settled amount F^B is obtained using the standard risk-neutral expectation formula (Duffie (2001)):

$$V(t) = e^{-r(\tau^e - t)} \mathbb{E}_t^* \left[\int_{\tau^b}^{\tau^e} (S(u) - F^B) du \right]. \quad (3.2)$$

As long as a forward contract is issued for free, *i.e.*, $V(t) = 0$, the corresponding forward price reads as:⁸

$$F^B = \frac{1}{\tau^e - \tau^b} \int_{\tau^b}^{\tau^e} \mathbb{E}_t^* [S(u)] du. \quad (3.3)$$

Hence, (continuous) *forward kernel* may be defined by:

$$f_t(u) := \mathbb{E}_t^* [S(u)], \quad u \geq t. \quad (3.4)$$

Building a forward curve through formula (3.4) amounts to identifying and calibrating an arbitrage-free spot price model for electricity. We instead follow a model-free pathway whereby a forward kernel is inferred from market prices by using a suitable inversion of formula (3.1).⁹ Then, day-ahead and forward prices with thin granularity may easily be obtained by straight-forward integration over the appropriate time period.

Example (Day-Ahead Price). Consider the Actual/365 day-count convention. Let $\Delta := 1 \text{ day} = 1/365$, $\delta := 1 \text{ hour} = \Delta/24$, $[0, \bar{T}]$ be the time horizon under analysis, and \mathcal{D} be the set of all midnight points comprised in that interval, *i.e.*, $\mathcal{D} := \{k\Delta \wedge \bar{T}, k \geq 0\}$, where “ \wedge ” stands for the minimum function.¹⁰ With a slight abuse of notation, we indicate by “ t ” both an element in

⁷A symbol like $f(x)$ is assumed to denote either a function f of variable x or the value assumed by that function at point x , the exact interpretation given by the context. This notational convention allows us to disentangle expressions sharing a common label for two distinct functions, such as $f(x)$ and $f(x, y)$.

⁸A similar relation is obtained for peakload forward prices by integrating over on-peak hours only.

⁹In this framework existence and unicity of a forward kernel is guaranteed by construction. Details are reported in Appendix B.2.

¹⁰A minor amendment allows us to account for either 23 or 25-hour days related to daylight saving standards.

\mathcal{D} and the day beginning that point in time, the exact interpretation of the symbol being clear from the context. For each day $t \in \mathcal{D}$ and hour $h \in \{1, 2, \dots, 23, 24\}$, the day-ahead hourly price $S^h(t, h)$ is defined as the forward price quoted at time 12:00 PM on day $t - 1$ for delivering one megawatt-hour over the h -th hour on day t , *i.e.*,

$$S^h(t, h) = \frac{1}{\delta} \int_{t+(h-1)\delta}^{t+h\delta} f_{t-12\delta}(u) du. \quad (3.5)$$

Example (Hourly Price Forward Curve). In the setting of previous example, the hourly price forward curve quoted at time t is the function of maturity variables “day T ” and “hour h ” defined as:

$$f_t^h : (T, h) \rightarrow f_t^h(T, h) := \frac{1}{\delta} \int_{T+(h-1)\delta}^{T+h\delta} f_t(u) du, \quad (3.6)$$

for $\{T > t\} \cap \mathcal{D}$ and $h \in \{1, 2, \dots, 23, 24\}$.

3.2 Predictable Patterns

Time series of electricity day-ahead prices exhibit a number of stylized patterns. These are primarily driven by the time evolution of balance between supply and demand of power. Correspondingly, they feature periodical recurrence on a variety of frequencies. According to Hildmann *et al.* (2013), any reasonable model for electricity forward curves should include seasonal, weekly, and intraday harmonics. Seasonality is mainly due to the evolution of temperature over the calendar year. This effect is exacerbated at on-peak hours, while off-peak prices exhibit smoother seasonal patterns. Weekly periodicity refers to systematic price discrepancies between weekdays and weekends. Nowotarski *et al.* (2013) detect and filter weekly patterns out of time series of electricity prices. We additionally disentangle working weekdays from non-working days, which include weekends, long weekends, standard holidays, and bank holidays. Daily recurrence refers to asynchronies between supply and demand sides over the course of a day. These paths are affected by a number of factors, which include human activity leading to on-peak hour prices higher than off-peak hour quotes, national supply policy, generating portfolio composition, and feed-in tariffs designed to accelerate investment in renewable energy technologies, among others. Table 3.2 reports a list of periodical patterns exhibited by time series of day-ahead prices and their major determinants.

4 Main Issue

The previous analysis shows that a rational definition of “electricity forward curve” ought to build on two pieces of market information. The first is the time series of day-ahead quotes available at the evaluation point in time. This term seeks to endow the curve with periodical patterns exhibited by these data, a feature that power traders take into account when posting

Period	Major determinants
Season	Temperature and daylight variation during the year
Week	Energy consumption surplus on working weekdays
Day	Human activity pathway (<i>e.g.</i> , meal effect)
	National tariff policy and generating portfolio composition
	Feed-in tariffs related to investment in renewable energy technology

Table 3.2: **Periodical patterns and major determinants.** Seasonal component is driven by temperature and daylight variations over the year. Weekly frequency stem from a systematic excess of working-day energy consumption over weekends. Daily recurrence depends on factors like human activity pathway, national tariff policy, generating portfolio mix, and renewables-linked feed-in tariffs.

futures price bids. The second set of market information is an array of baseload and peakload futures prices corresponding to tradable delivery periods. These elements make the output curve compatible with the standing market view, an approach in keeping with the nature of forward prices intended as estimators of future quotes based on currently available information.

Our task boils down to combining the two aforementioned kinds of market data into a single forward kernel, and then deriving forward curves with arbitrary time granularity using appropriate instances of formula (3.1).¹¹ We propose a *definition* of forward kernel which entails specifying market data, structural components, and exact constraints to comply with. We model the time t continuous-time forward kernel $f(u) := f_t(u)$ as the sum of a backward periodical pattern $\Lambda(u)$, a baseload adjustment term $\varepsilon(u)$, and a peakload adjustment component $\varphi(u)$, *i.e.*,

$$f(u) = \Lambda(u) + \varepsilon(u) + \varphi(u), \quad (4.1)$$

where the three functions are all integrable over a finite time horizon of interest, say $[t, \bar{T}]$. At time t , we dispose of two pieces of price information: one is the 24-dimensional time series of day-ahead hourly prices recorded by time t ; the other is an array of baseload and peakload futures prices, whose delivery periods must span the whole time horizon. A forward kernel estimate is “rational” provided it complies with the five properties put forward in Section 2.

We may constructively specify forward kernel components by using market data according to the following algorithm:

- 1. Market data preparation.** Raw spot price data undergo a nonparametric filtering procedure leading to time series of outlier-free day-ahead figures; in addition, market futures quotes are segmented over non-overlapping delivery periods.
- 2. Historical estimation of day-ahead predictable patterns.** Backward predictable patterns are derived by filtering a macroeconomic trend and periodical components out of day-ahead series using methods borrowed from the theory of signal processing.

¹¹From the perspective of arbitrage pricing theory, this pathway goes in tandem with a similar approach developed in other areas of finance, *e.g.*, Garcia *et al.* (2011), although the explicit adoption is less common in the area of energy markets.

3. **Risk-neutral calibration of quoted futures prices** in the sense of relation (3.1). A first adjustment term combines with the estimated predictable path so as to fit baseload futures quotes.

Given an estimated forward kernel, forward prices with thin granularity may be derived by using formula (3.6) and the corresponding forward curve follows. Figure 4.1 reports a synoptical diagram representing the constructive pathway described so far.

ACCEPTED MANUSCRIPT




Figure 4.1: **Constructive definition of rational forward curve with thin granularity.** Data preparation consists of filtering out day-ahead data outliers and segmenting market baseload and peakload futures quotes over a sequence of non-overlapping delivery periods. Historical estimation of day-ahead predictable patterns entails unveiling a macroeconomic trend and combined seasonal, weekly, and daily periodical components. Risk-neutral calibration of segmented futures prices delivers a continuous-time forward kernel that is jointly consistent with baseload and peakload quotes. Forward curves with thin granularity are obtained by integrating the estimated forward kernel over the interval defining the selected time refinement.

5 Data Preparation

Raw spot and forward data must undergo suitable processing to properly feed our model. First, input time series of day-ahead prices with thin granularity are required to be outlier-free. Second, quoted futures prices must span a sequence of non-overlapping delivery periods covering the whole time horizon under analysis. We provide algorithms to process raw data and obtain market information in keeping with these two requirements.

5.1 Outlier Detection and Filtering

Historical spot price data exhibit a number of undesirable features which may hinder an appropriate estimation of backward pattern Λ and adjustment functions ε and φ . As noted by Truck *et al.* (2007) and Janczura *et al.* (2013), filtering data outliers may sensibly improve the statistical estimation of both weekly and long-term periodical components. To this end, we adopt a nonparametric high-pass filter developed by Hodrick and Prescott (1997) (HP), which we now briefly sketch.¹²

Given a time series $\mathcal{X} = (x_t)_{1 \leq t \leq T}$, HP filtering aims at building a new time series $\mathcal{H} = (h_t)_{1 \leq t \leq T}$ exhibiting smoothing features according to an assigned parameter λ . Specifically, the HP λ -filter is defined as follows:

$$\mathcal{H}^\lambda := \arg \min_{\{h_t\}_{1 \leq t \leq T}} \sum_{t=1}^T (x_t - h_t)^2 + \lambda \sum_{t=1}^T [(h_{t+1} - h_t) - (h_t - h_{t-1})]^2,$$

where the first sum is the variance of the residual time series representing a “business cycle path”,¹³ while the second sum is an acceleration component playing the role of “penalty term”. Clearly, parameter λ tunes the smoothness of filter \mathcal{H}^λ by controlling the extent to which its time-related curvature enters the target functional. In particular, the greater parameter λ , the smoother the output filter \mathcal{H}^λ .

For the purpose of selecting a rational value of input λ , we adopt a method put forward by Pedersen (2001). Starting point is the “ideal business cycle filter” associated to an assigned period τ .¹⁴ Following Baxter and King (1999), this is the high-pass filter singling out components at frequencies lower than $\omega_c = 2\pi/\tau$, namely:

$$H_{\omega_c}^*(\omega) = \mathbb{1}(\omega \geq \omega_c). \quad (5.1)$$

It turns out that no HP filter is high-pass in that it cannot precisely disentangle frequencies above and below an assigned threshold. Hence, one may wish to consider the HP λ -filter coming

¹²Weron and Zator (2015) show that HP filtering is simpler than and equally powerful to wavelets-based methods for the purpose of identifying seasonal components of electricity spot prices.

¹³In economics literature, “business cycle” refers to the fundamental noise term as opposed to predictable patterns in a time series of values.

¹⁴Period is expressed in time units according to the Actual/365 day-count convention.

as close as possible to the ideal business cycle filter. This amounts to selecting coefficient λ in a way to minimize the weighed discrepancy between the power transfer function, or Fourier transform, of HP λ -filter (King and Rebelo (1993)):

$$H_\lambda(\omega) = \left| \frac{4\lambda(1 - \cos(\omega))^2}{4\lambda(1 - \cos(\omega)) + 1} \right|^2$$

and the ideal business cycle filter $H_{\omega_c}^*$ in expression (5.1). Let us introduce $\Delta\omega$ -spaced refinement \mathcal{W} of the interval $[0, \pi]$. The power spectral density $v(\omega)$ of the input time series \mathcal{X} computed over \mathcal{W} is defined as:

$$v(\omega) = 2S(\omega)\Delta\omega \left(\sum_{\omega \in \mathcal{W}} 2S(\omega)\Delta\omega \right),$$

where $S(\omega)$ represents the power spectrum of the root series \mathcal{X} .¹⁵ Then, for an assigned cut-off frequency ω_c , the Pedersen (2001) optimal parameter reads as follows:

$$\lambda^* = \arg \min_{\lambda} \sum_{\omega \in \mathcal{W}} |H_\lambda(\omega) - H_{\omega_c}^*(\omega)| v(\omega), \quad (5.2)$$

where variable ω is expressed in frequency units and $v(\omega)$ plays the role of weighting function.

Our data filtering procedure is as follows. We begin by recording hourly day-ahead prices up to the selected time of evaluation. Next, we build up a time series of daily quotes, each figure being computed as the average hourly price on a single day.¹⁶ We compute the HP filter that selects frequencies lower than a threshold ω_c equal to one month (Truck *et al.* (2007)).¹⁷ Following Cartea and Figueroa (2005), we mark as an “outlier” any data point exceeding three times the standard deviation of the sample distribution of discrepancies between market price and the corresponding point on the HP filter. Last, we collect the entire sequence of twenty-four hourly prices corresponding to each outlier and remove the entire array from the price series. This procedure leads to an outlier-free spot price series.

5.2 Futures Quote Segmentation

Futures prices observed at any point in time may refer to partially overlapping delivery periods. We can transform them into corresponding prices for delivery over a partition of the time horizon. Let us consider a set of baseload and peakload futures prices $F^B(\tau_k^b, \tau_k^e)$ and $F^P(\tau_k^b, \tau_k^e)$ ($k = 1, \dots, n$) quoted on a common date. We assume that delivery intervals span the whole

¹⁵Given time series $\mathcal{X} = (x_t)_{1 \leq t \leq T}$, the corresponding *power spectrum* is defined as: $S(\omega) := g(\exp(-i\omega))$, where g is the *generating function* $g(z) = \sum_{k \in \mathbb{R}} c(k)z^k$ of the *autocovariance function* of \mathcal{X} , i.e., $c(k) = \text{Cov}(x(t), x(t-k))$.

¹⁶Huisman *et al.* (2007) noticed that day-ahead hourly prices exhibit the statistical behavior of a daily cross-section of 24 hourly data, as opposed to a single one-dimensional time series with hourly granularity.

¹⁷Appendix B.3 analyzes the effect of varying levels of ω_c on the corresponding filtered path.

time horizon $[0, \max_k \{\tau_k^e\}]$. We define $\#B$ (resp. $\#P$) as the counting measure of baseload (resp. peakload) delivery days. Next, we sort the $2n$ time points $\tau_1^b, \tau_1^e, \dots, \tau_n^b, \tau_n^e$ in ascending order, delete possible duplicates, and relabel the resulting dates as $t_0 < t_1 < \dots < t_{m-1} < t_m$, where $m \leq 2n$.

We may compute segmented baseload and peakload forward prices $F^B(t_i, t_{i+1})$ and $F^P(t_i, t_{i+1})$ for each segment $[t_i, t_{i+1}]$ ($i = 0, m-1$) by solving a linear system ensuring compliance with observed futures prices $F^B(\tau_k^b, \tau_k^e)$ and $F^P(\tau_k^b, \tau_k^e)$ ($k = 1, \dots, n$). Specifically, if a period $[\tau_k^b, \tau_k^e]$ is covered by intervals $[t_i, t_{i+1}]$ ($i = j(k), j(k) + 1, \dots, J(k)$), then the corresponding market forward price is the arithmetic average of segmented forward prices for delivery on these intervals, *i.e.*,

$$F^X(\tau_k^b, \tau_k^e) = \sum_{i=j(k)}^{J(k)} \frac{\#X(t_i, t_{i+1})}{\#X(\tau_k^b, \tau_k^e)} F^X(t_i, t_{i+1}), \quad (5.3)$$

where $k = 1, \dots, n$ denotes a market delivery period and $X = B, P$ identifies baseload and peakload quotation type, respectively.

Formula (5.3) defines a linear system of equations $A^X \mathbf{y}^X = \mathbf{z}^X$, where $X = B$ or P , $A^X \in \mathbb{R}^{n \times m}$, $\mathbf{y}^X \in \mathbb{R}^m$, and $\mathbf{z}^X \in \mathbb{R}^n$. Under the assumption that sets $[\tau_k^b, \tau_k^e]$ cover the whole time horizon comprised between t_0 and the farthest quoted maturity, existence of a unique vector of non-overlapping futures prices \mathbf{y}^X is granted by assuming that A^X be a squared and non-singular matrix. Should $\text{rank}(A^X) \neq \text{rank}(A^X | \mathbf{z}^X)$, then one may progressively exclude contract covering common periods of time, starting with the least liquid quotations (Jarrow (2014)). The case of infinitely many solutions (*i.e.*, $\text{rank}(A^X) = \text{rank}(A^X | \mathbf{z}^X) \leq m$) may arise when weekly quotations are used. It can be addressed either by excluding some illiquid quotations or by considering additional market products such as “balance of month”, that is a contract delivering energy time until the end of current month, or even by introducing additional constraints on variables. A lack of solutions corresponding to $\text{rank}(A^X) < \text{rank}(A^X | \mathbf{z}^X)$ signals arbitrage opportunities.

6 Model

6.1 Historical Estimation of Day-Ahead Predictable Patterns

Market data preparation leads to outlier-free day-ahead prices $\hat{S}^h(t, h)$ for each hour of the day. We assume the resulting time series additively combines as many as two predictable components:

$$\hat{S}^h(t, h) \simeq \underbrace{\hat{H}(t)}_{\text{macro trend}} + \underbrace{\hat{\Lambda}^h(t, h)}_{\text{periodical trend}}, \quad (6.1)$$

up to a residual term, which we assume to be negligible. Here $t \in \mathcal{T}_{t_0} = \{s : s \in \mathcal{D}, s \leq t_0\}$, $h = 1, \dots, 24$, and $\mathcal{D} := \{k\Delta \wedge \bar{T}, k \geq 0\}$. Component $\hat{H}(t)$ denotes a predictable macroeconomic trend uniformly applying to all hours of day t , while function $\hat{\Lambda}^h(t, h)$ indicates a periodical pattern whose time behavior is idiosyncratic to each single hour h . We first estimate functions $\hat{H}(t)$ and $\hat{\Lambda}^h(t, h)$ using the historical time series of filtered spot prices; then, we use these estimates to infer an optimal continuous-time backward pattern $\hat{\Lambda}^*(u)$ entering forward kernel decomposition (4.1).

Trend $\hat{H}(t)$ can be estimated using a nonparametric filter. Hourly sampled day-ahead prices $\hat{S}^h(t, h)$ are averaged to deliver daily sampled day-ahead quotes $\hat{S}(t) := \frac{1}{24} \sum_{h=1}^{24} \hat{S}^h(t, h)$. The resulting daily series $\hat{\mathbf{S}} = (\hat{S}(t), t \in \mathcal{T}_{t_0})$ enters a Hodrick-Prescott algorithm filtering periods exceeding one year and a half, which corresponds to the period for seasonal pattern (*i.e.*, one year), plus a six-month buffer introduced to avoid distortions near cutoff frequency. Estimation leads to macro trend $\hat{H}(t)$ as well as hourly and daily sampled estimated periodical trend functions $\hat{\Lambda}^h(t, h) := \hat{S}^h(t, h) - \hat{H}(t)$ and $\hat{\Lambda}(t) := \hat{S}(t) - \hat{H}(t)$, respectively.

According to the analysis reported in Section 3.2, periodical trend $\hat{\Lambda}$ features seasonal, weekly, and daily patterns. We thus assume a parametric form combining the following terms:

- A *seasonal pattern* with annual recurrence:

$$a \cos\left(\frac{2\pi}{365}t + b\right), \quad (6.2)$$

where $0 \leq b \leq 2\pi$ defines the corresponding point of maximum and minimum.

- A *weekly pattern* represented by seven dummy variables, one for each day of the week:

$$\sum_{j=1}^7 c_j \mathbb{1}(t \in \text{day}_j), \quad (6.3)$$

where “ day_j ” is the j -th daily time interval and “1” stand for “Sunday”.¹⁸

- A *daily pattern* modeled through a (4×24) -array, each term representing the deviation of

¹⁸Any day in a public holiday period or a long weekend is treated as a Sunday.

hourly price from the corresponding daily average:

$$\sum_{h=1}^{24} \sum_{l=1}^4 d_{h,l} \mathbb{1}(t \in \text{hour}_h \cap \mathcal{C}_l), \quad (6.4)$$

where “ hour_h ” is the h -th hourly time interval and \mathcal{C}_l is defined as follows: for the purpose of capturing daily effects related to temperature and sunlight exposure upon warm and cold seasons, we segment all hours in a calendar year into four disjoint time sets, which we define as:

$$\begin{aligned} \mathcal{C}_{1/2} &:= \text{working/nonworking days in the cold season (i.e., from October to March),} \\ \mathcal{C}_{3/4} &:= \text{working/nonworking days in the warm season (i.e., from April to September).} \end{aligned}$$

The resulting continuous-time periodical component reads as:

$$\Lambda(t) = \underbrace{a \cos\left(\frac{2\pi}{365}t + b\right)}_{\text{Seasonal}} + \underbrace{\sum_{j=1}^7 c_j \mathbb{1}(t \in \text{day}_j)}_{\text{Weekly}} + \underbrace{\sum_{h=1}^{24} \sum_{l=1}^4 d_{h,l} \mathbb{1}(t \in \text{hour}_h \cap \mathcal{C}_l)}_{\text{Daily}}, \quad (6.5)$$

where coefficient a allows us to tune the seasonal macroeconomic trend, term c_j assigns a weight to a dummy variable related to the j -th day in the week, and parameter $d_{h,l}$ measures the discrepancy between the price of hour h and the corresponding daily average, provided that t belongs to time cluster \mathcal{C}_l . As an additional condition we require a balance constraint whereby the twenty-four hourly deviations from daily mean price sum up to zero on each time cluster, *i.e.*,

$$\sum_{h=1}^{24} d_{h,l} = 0, \quad l = 1, \dots, 4. \quad (6.6)$$

Parametric function Λ defined in formula (6.5) can be estimated using hourly and daily sampled trend patterns $\hat{\Lambda}^h(t, h)$ and $\hat{\Lambda}(t)$. For the sake of clarity, we write $\Lambda(t) = \Lambda_{\theta, D}(t)$, with:

$$\begin{aligned} \theta &= [a, b, c_1, c_2, c_3, c_4, c_5, c_6, c_7], \text{ (seasonal + weekly)} \\ D &= \{d_{h,l}, h = 1, \dots, 24; l = 1, \dots, 4\}, \text{ (daily)} \end{aligned}$$

and pursue a two-step estimation exercise.

First, we jointly estimate the combined seasonal-weekly parametric set θ . Let $\Lambda_{\theta}^d(t)$ be the daily-sampled parametric periodical trend corresponding to the continuous-time periodical trend $\Lambda_{\theta, D}(t)$:

$$\Lambda_{\theta}^d(t) := \frac{1}{\Delta} \int_t^{t+\Delta} \Lambda_{\theta, D}(u) du, \quad (6.7)$$

where $\Delta = 1$ day under the assumed time unit and $t \in \mathcal{T}_{t_0}$. Note that in light of assumption (6.6), integrating over a full day amounts to eliminating dependence on the parametric set D , hence

the left-hand side in expression (6.7) depends on θ only. We fit the daily sampled parametric periodical trend $\Lambda_\theta^d(t)$ to the daily sampled estimated periodical trend $\hat{\Lambda}(t)$ by solving a least-squares minimization problem:

$$\theta^* = \arg \min_{\theta} \sum_{s \in \mathcal{T}_{t_0}} e^{-\alpha(t_0-s)} \left(\Lambda_\theta^d(s) - \hat{\Lambda}(s) \right)^2, \quad (6.8)$$

where parameter α allows us to tune the fitting quality over the time-to-maturity dimension.

We are now in a position to estimate the daily parametric set D . Let $\Lambda_{\theta,D}^h(t, h)$ denote the hourly sampled parametric periodical trend corresponding to the continuous-time periodical trend $\Lambda_{\theta,D}(t)$, *i.e.*,

$$\Lambda_{\theta,D}^h(t, h) := \frac{1}{\Delta} \int_{t+(h-1)\delta}^{t+h\delta} \Lambda_{\theta,D}(u) du,$$

where time lag $\delta := 1$ hour under the assumed time unit, variable $t \in \mathcal{T}_{t_0}$, and hour $h = 1, \dots, 24$. We fit hourly sampled parametric periodical trend $\Lambda_{\theta,D}^h(t, h)$ to the hourly sampled estimated periodical trend $\hat{\Lambda}^h(t, h)$ by solving the following constrained least-squares minimization problem:

$$\begin{cases} D^* = \arg \min_D \sum_{(s,h) \in \mathcal{T}_{t_0} \times \{1, \dots, 24\}} e^{-\alpha(t_0-s)} \left(\Lambda_{\theta^*, D}^h(s, h) - \hat{\Lambda}^h(s, h) \right)^2, \\ \sum_{h=1}^{24} d_{h,l} = 0, \quad l=1, \dots, 4. \end{cases} \quad (6.9)$$

The optimal parametric setting (θ^*, D^*) leads to a historical estimation of day-ahead predictable pattern $\hat{\Lambda}^*(t) := \Lambda_{\theta^*, D^*}(t)$ entering the forward kernel definition (4.1).

6.2 Risk-Neutral Calibration of Quoted Futures Prices

Model calibration aims at devising adjustment terms $\varepsilon(t)$ and $\varphi(t)$ which additively combine with the historically estimated day-ahead predictable pattern $\hat{\Lambda}^*(t)$ and deliver a market forward kernel through formula (4.1). In what follows, all statements involving i -indexed quantities apply for all $i = 0, \dots, n-1$. Data preparation offers an array of time segmented futures prices F_i^B and F_i^P quoted at the evaluation point in time. Curve form (4.1) combined with baseload consistency condition (3.3) leads to:

$$F_{i+1}^B = \frac{1}{t_{i+1} - t_i} \int_{t_i}^{t_{i+1}} (\Lambda(u) + \varepsilon(u) + \varphi(u)) du.$$

Our plan consists of fitting term $\varepsilon(t)$ to baseload prices F_i^B and then calibrating term $\varphi(t)$ to peakload quotes F_i^P .

A key point in our development is the ability to select term φ such that it matches peakload prices F^P , while leaving fitting properties of term ε unaffected. We split the period ranging from Monday to Friday into a set P comprising twelve daily on-peak hours and a set O gathering as

many daily off-peak hours. Next, we define a piecewise constant function as follows:

$$\varphi(t) := c_{i+1} [\mathbf{1}_P(t) - \mathbf{1}_O(t)], \quad (6.10)$$

over each time interval $[t_i, t_{i+1}]$.¹⁹ As long as weekend hours are all off-peak, we may agree on identifying them as on-peak and off-peak at the same time. Hence, function $\varphi(t)$ vanishes on weekends, while the integral does so on each weekday. Consequently, it shows the required property:²⁰

$$F_{i+1}^B = \frac{1}{t_{i+1} - t_i} \int_{t_i}^{t_{i+1}} (\Lambda(u) + \varepsilon(u)) du.$$

We now calibrate term $\varepsilon(t)$ to segmented baseload futures prices. Given estimated backward pattern $\hat{\Lambda}^*(t)$, let us define adjusted baseload futures quotes as:

$$\varepsilon_{i+1}^A := F_i^B - \frac{1}{t_{i+1} - t_i} \int_{t_i}^{t_{i+1}} \hat{\Lambda}^*(u) du = \frac{1}{t_{i+1} - t_i} \int_{t_i}^{t_{i+1}} \varepsilon(u) du. \quad (6.11)$$

These quantities can easily be computed by numerically integrating the optimal backward pattern $\hat{\Lambda}^*(t)$. Searching for a function $\varepsilon(t)$ that matches the risk-neutral baseload fitting condition (6.11) amounts to solving an inverse problem, which possibly admits infinitely many solutions. We adopt one based on the monotone convex interpolator put forward by Hagan and West (2006) in light of desirable regularity properties it brings with. Specifically, we set:

$$\hat{\varepsilon}^*(t) = \begin{cases} g_i \left(\frac{t-t_{i-1}}{t_i-t_{i-1}} \right) + \varepsilon_i^A, & t \in [t_{i-1}, t_i], i = 1, \dots, n, \\ g_n \left(\frac{t-t_{n-1}}{t_n-t_{n-1}} \right) + \varepsilon_n^A, & t \geq t_n, \end{cases} \quad (6.12)$$

where g_i is the monotone convex interpolator corresponding to points ε_{i+1}^A over the interval $[0, 1]$ as defined in Appendix A. Peakload calibration reads as:

$$F_{i+1}^P = \frac{1}{\int_{t_i}^{t_{i+1}} \mathbf{1}_P(u) du} \int_{t_i}^{t_{i+1}} [\Lambda(u) + \varepsilon(u) + \varphi(u)] \mathbf{1}_P(u) du,$$

where indicator function $\mathbf{1}_P$ within the integral operator restricts integration over on-peak

¹⁹Term $\mathbf{1}_A$ denotes the indicator function of a set A , i.e., $\mathbf{1}_A = 1$ on A and zero otherwise.

²⁰Another function admissible to our framework is:

$$\varphi(t) := a_{i+1} \mathbf{1}_P(t) + b_{i+1} \mathbf{1}_{\tilde{O}}(t),$$

where set \tilde{O} comprises all off-peak hours over the week, including all weekend hours as well. This selection would lead to a system of algebraic equations:

$$\begin{cases} a_{i+1} \int_{t_i}^{t_{i+1}} \mathbf{1}_P(u) du + b_{i+1} \int_{t_i}^{t_{i+1}} \mathbf{1}_{\tilde{O}}(u) du = 0 \\ a_{i+1} = F_{i+1}^P - \frac{\int_{t_i}^{t_{i+1}} [\Lambda(u) + \varepsilon(u)] \mathbf{1}_P(u) du}{\int_{t_i}^{t_{i+1}} \mathbf{1}_P(u) du} \end{cases}$$

for each time step $i = 0, \dots, n-1$. In general, selecting a kind of “optimal” peakload adjustment function meeting the requirement: $\int_{t_i}^{t_{i+1}} \varphi(u) du = 0$, for $i = 0, \dots, n-1$, is a subject beyond the scope of the present study. We thank one of the referees for pointing out this alternative solution.

hours only. By plugging definition (6.10) into this expression and rearranging terms, we obtain coefficients:

$$c_{i+1} = \frac{12\delta \#P(t_i, t_{i+1}) F_i^P - \int_{t_i}^{t_{i+1}} (\Lambda(u) + \varepsilon(u)) \mathbf{1}_P(u) du}{12\delta \#P(t_i, t_{i+1})},$$

which in turn define peakload adjustment component $\hat{\varphi}^*(t)$ via formula (6.10). This concludes the risk-neutral calibration step.

Once equipped with historical backward pattern $\hat{\Lambda}^*(t)$ and risk-neutral baseload and peakload adjustment terms $\hat{\varepsilon}^*$ and $\hat{\varphi}^*$, respectively, we finally arrive at an optimally estimated forward price kernel as:

$$\hat{f}_t^*(u) = \Lambda^*(u) + \varepsilon^*(u) + \varphi^*(u). \quad (6.13)$$

The corresponding HPFC estimate $\hat{f}^{h*}(T, h)$ is obtained by inserting kernel (6.13) into HPFC formula (3.6).

The curve we obtain exhibits all of the desirable features laid out in Section 2. First, it stems from outlier-filtered market data, so that forward estimates are free from spurious effects unrelated to actual market views. Second, the curve shape embeds a full bundle of periodical patterns shown by the time series of spot prices with thin granularity. Third, the curve values jointly fit an arbitrary number of baseload and peakload quotes. Finally, the estimated curve exhibits a number of regularity properties, including trajectorial smoothness, continuity over the maturity spectrum, monotonicity preservation,²¹ and cross-sectional localness. This latter implies that any change affecting a baseload futures price F_i^B entails the sole modification of output forward kernel on the corresponding interval $[t_{i-1}, t_i]$, as well as on the two adjacent segments $[t_{i-2}, t_{i-1}]$ and $[t_{i+1}, t_{i+2}]$.²² We now turn to the empirical assessment of these, as well as other properties of the forward curves resulting from our constructive definition.

7 Empirical Evidence

The most liquid electricity market in continental Europe is operated by EPEX Spot SE and EEX Power Derivatives. EPEX Spot SE is the short-term power exchange resulting from the 2008 merger of German EEX and French Powernext. It is a daily spot market covering the geographical areas of France, Germany, Austria, and Switzerland. EEX Power Derivatives is the related power derivative marketplace. It runs trading and quotation services for electricity futures and options, including those written on EPEX Spot prices. These markets offer time series of day-ahead prices, as well as arrays of baseload and peakload futures quotes. Consequently,

²¹As many as three consecutively increasing (resp. decreasing) seasonal adjusted futures prices always come with a monotonically increasing (resp. decreasing) forward kernel over the corresponding portion of the time-maturity horizon (Hyman (1983)).

²²In the case of peakload futures prices, which turn out to exhibit a relatively more pronounced variation over observational time, cross-sectional localness is even tighter than for baseload quotes: a change in peakload futures price F_i^P entails the sole modification of output forward kernel on the corresponding interval $[t_{i-1}, t_i]$.

Contract type	Delivery begin	Delivery end	Baseload price (EUR/MWh)	Peakload price (EUR/MWh)
Day	November 27, 2015	November 27, 2015	36.77	41.83
Day	November 28, 2015	November 28, 2015	27.87	32.75
Day	November 29, 2015	November 29, 2015	11.25	16.99
Day	November 30, 2015	November 30, 2015	24.50	34.75
Week	November 30, 2015	December 6, 2015	28.76	37.32
Week	December 7, 2015	December 13, 2015	34.34	45.00
Week	December 14, 2015	December 20, 2015	32.13	41.88
Week	December 21, 2015	December 27, 2015	22.22	29.38
Month	December 1, 2015	December 31, 2015	28.93	36.75
Month	January 1, 2016	January 31, 2016	29.87	38.72
Month	February 1, 2016	February 29, 2016	32.62	40.70
Quarter	January 1, 2016	March 31, 2016	30.37	37.61
Quarter	April 1, 2016	June 30, 2016	27.75	33.06
Quarter	July 1, 2016	September 30, 2016	28.74	34.73
Calendar	January 1, 2016	December 31, 2016	29.38	36.25

Table 7.1: **EEX forward market prices.** On November 27, 2015, as many as fifteen forward-looking prices are quoted for delivery periods including four days, four weeks, three months, three quarters, and one calendar year. For each contract, first and last day of delivery, baseload price, and peakload price are reported. Prices are expressed in Euros per megawatt-hour (EUR/MWh). Bold character highlights (partially) overlapping periods whose futures prices require passing through quote segmentation.

they represent an appropriate market context to apply and test our constructive definition of electricity forward price curve we propose.

7.1 Curve Construction

We consider 2,522 daily spot price arrays, each containing 24 day-ahead hourly quotes. They span the time period from January 1, 2009 to November 27, 2015, which we take as the evaluation date t_0 . We also record an array of baseload and peakload futures quotes collected at that point in time. Table 7.1 reports market data including as many as fifteen quotes: they encompass four day-ahead, four week, three month, three quarter contracts, and one calendar. For each item, we indicate delivery period, baseload, and peakload prices. Time unit is assumed equal to one day, which corresponds to setting $\Delta := 1$ and thus $\delta := 1/24$. We highlight partially overlapping periods whose futures prices require passing through quote segmentation. In contrast to most commodities, no clear pathwise feature (*e.g.*, backwardation, contango, periodical or mixed patterns) seems to emerge from visual inspection of these data, which seems to merit a deeper analysis based on time series.

The filtering method detailed in Section 5.1 allows us to identify and single out as many as twenty-seven data outliers. They result from points in time featuring sharp discrepancies between price path and the Hodrick-Prescott filter corresponding to the optimal parametric as-




Figure 7.1: Spot price outlier detection and filtering. For the period comprised between January 1, 2009, and November 27, 2015, processed data are plotted for the time series of Hodrick-Prescott filter (red plain line) corresponding to the optimal parametric assessment $\lambda^* = 1,128.7$; price outlier flags (magenta colored stars), defined as any datum exceeding three times the standard deviation of filtered price sample distribution; and, the resulting time series of outlier-free day-ahead prices (blue plain line).

assessment $\lambda^* = 1,128.7$ computed using formula (5.2).²³ Figure 7.1 depicts the time series of Hodrick-Prescott filter, price outliers, and the resulting outlier-free day-ahead prices. Outliers turn out to account for about 1.07% of the overall sample. More interestingly, an exhibit of flags for removed outliers qualitatively shows that their frequency of occurrence fades away along with price level. This last point represents a signal of EPEX Spot market maturity going in pair with an increasing level of competitiveness. Our analysis leads to an outlier-free price series $\hat{S} = (\hat{S}(t, h), t \in \mathcal{T}_{t_0}, h = 1, \dots, 24)$.

We disentangle futures market quotes in keeping with the method laid out in Section 5.2. This process leads to segmented baseload and peakload quotes spanning a series of non-overlapping time intervals, which represent prices F_i^B and F_i^P feeding the model. As an example, Table 7.1 shows that the frontline quarter covers the period ranging from January 1 to March 31, which overlaps with both January and February monthly contracts. Price segmentation provides implied segmented baseload and peakload quotes for March. They amount to 28.76 EUR/MWh

²³Optimal parameter λ^* stems from assuming an ideal business cycle filter at a benchmark frequency $\omega_c = 30$ days. An experiment not reported here derives sixteen forward curves over as many filtering frequencies ranging from one week up to a year. Mean absolute deviations of each of these curves from the one corresponding to the benchmark frequency show negligible discrepancies. They range from 0.0238 to 0.0471 for frequencies greater than 30 days and from 0.0429 to 0.0817 for frequencies lower than 30 days. In monetary terms, the most relevant difference results from setting $\omega_c = 7$ days and accounts for as many as 8 Euro cents per megawatt-hour.

Contract type	Begin delivery date	End delivery date	Baseload price (EUR/MWh)	Peakload price (EUR/MWh)
Week	November 30, 2015	December 6, 2015	28.76	37.32
Segmented	December 1, 2015	December 6, 2015	29.47	37.96
Month	December 1, 2015	December 31, 2015	28.93	36.75
Segmented	December 28, 2015	December 31, 2015	24.79	28.67
Quarter	January 1, 2016	March 31, 2016	30.37	37.61
Segmented	March 1, 2016	March 31, 2016	28.76	33.77
Calendar	January 1, 2016	December 31, 2016	29.38	36.25
Segmented	October 1, 2016	December 31, 2016	30.65	38.27

Table 7.2: **Segmented forward market prices.** On November 27, 2015, four futures contracts show overlapping intervals of delivery. These correspond to contract series switching from one type to the following. That is from last week to first month, from last month to first quarter, and from last quarter to calendar. Their quotations undertake a maturity segmentation procedure leading to baseload and peakload prices spanning a series of non-overlapping time intervals filling in the whole time-to-maturity spectrum.

and 33.77 EUR/MWh, respectively. A full array of segmented forward quotes coupled with the market prices they replace is reported in Table 7.2. We use outlier-free spot quotes to estimate a macroeconomic trend component $\hat{H}(t)$. According to the proposed definition, this quantity is obtained as the Hodrick-Prescott filter with an optimal parameter $\lambda^* = 8.322 * 10^7$, a value allowing for filtering frequencies lower than one and a half years. Figure 7.2 shows the outlier-free price path together with the corresponding macroeconomic trend with daily granularity. We clearly see that filtered trend exhibits a time behavior combining a variety of periodical cycles with long-term frequencies, each exceeding the selected 1.5-year threshold. The two-step algorithm for estimating periodical patterns leads to parametric values for seasonal, weekly, and daily periods gathered into expression (6.5). Throughout the analysis, we assume a constrained optimization weighting parameter $\alpha = 0.4$.²⁴

Table 7.3 reports numerical values for seasonal and weekly components, together with corresponding standard deviations. Coefficients c_2 and c_6 are smaller than coefficients c_3 , c_4 , and c_5 because they refer to working days where hourly prices tend to be lower than elsewhere. As noted in Ziel et al. (2015), day-ahead prices quoted in the EPEX market exhibit lower quotes near the weekend than during other periods in the working segment of the week. That is on Friday afternoon and evening as well as on Monday morning. Specifically, Friday afternoon prices show a declining path towards the level attained early on Saturday. This feature reflects a relatively low level of power demand near the early closing of offices and factories at the end of the week. A similar behavior occurs on Monday morning, mirroring the inertia of power demand transition from the lowest level reached during weekend to the highest level on working days.

²⁴Appendix B.4 discusses the choice of smoothing parameter α in great detail.

Parameter	a	b	c_1	c_2	c_3	c_4	c_5	c_6	c_7
Value	-2.1856	3.8235	-9.8507	1.8837	3.6655	3.6856	3.9474	2.8900	-4.7847
St Dev.	(0.0347)	(0.0078)	(0.1074)	(0.1286)	(0.1237)	(0.1242)	(0.1281)	(0.1329)	(0.1262)

Table 7.3: **Seasonal and weekly parameters.** Coefficients a and b assign seasonal macroeconomic trend component, while each term c_j defines a weighting parameter for the dummy variable corresponding to the j -th day in the week. These latter jointly define the weekly pattern. Numbers within round brackets represent standard deviation estimates for the selected terms.




Figure 7.2: **Estimated macroeconomic trend $\hat{H}(t)$.** For a time series of outlier-free prices $\hat{S}(t)$ (blue line), a macroeconomic trend component $\hat{H}(t)$ is estimated as the Hodrick-Prescott filter with coefficient $\lambda^* = 8.322 * 10^7$. According to Pedersen (2001), this figure entails **to(-cut)** filtering frequencies lower than one year and a half, a number corresponding to the period of seasonal patterns plus a six-month distortion smoothing window.

$d_{h,l}$	$l = 1$	$l = 2$	$l = 3$	$l = 4$	$d_{h,l}$	$l = 1$	$l = 2$	$l = 3$	$l = 4$
$h = 1$	-12.3310	-3.2827	-7.7044	2.1429	$h = 13$	2.1028	1.2196	1.4131	-1.4627
$h = 2$	-13.7391	-5.8830	-9.8285	-1.2566	$h = 14$	1.7368	-1.2723	-0.3860	-4.4729
$h = 3$	-15.3356	-7.0224	-11.6334	-2.9295	$h = 15$	1.6981	-1.9584	-1.7240	-6.4422
$h = 4$	-15.8044	-8.0886	-13.0099	-4.2086	$h = 16$	3.2985	-0.9233	-1.7391	-6.4117
$h = 5$	-14.8031	-8.0257	-12.4374	-4.7724	$h = 17$	5.0629	1.9364	-1.2453	-5.0588
$h = 6$	-11.7646	-7.2103	-9.1018	-4.8059	$h = 18$	12.5352	9.3095	2.4956	-0.6392
$h = 7$	-1.4978	-7.7414	0.4200	-4.1806	$h = 19$	16.5081	14.7238	6.0516	5.1848
$h = 8$	9.7413	-4.9989	8.5658	-2.1025	$h = 20$	14.2991	13.6426	9.0492	8.9070
$h = 9$	11.0394	-1.1507	10.7903	0.3571	$h = 21$	5.1588	7.1647	8.1074	9.6660
$h = 10$	8.2239	1.9200	7.9841	1.3708	$h = 22$	-1.6372	2.5741	5.1669	8.5080
$h = 11$	5.7082	2.3385	5.5551	0.2843	$h = 23$	-4.5236	2.7480	2.5351	8.5621
$h = 12$	4.7328	3.0086	4.7980	0.1970	$h = 24$	-10.4094	-3.0281	-4.1224	3.5637

Table 7.4: **Daily parameters.** Coefficients $d_{h,l}$ assess the discrepancy between hour h price and the corresponding daily average, assuming that day belongs to time cluster \mathcal{C}_l . These numbers jointly define the daily pattern.

Figure 7.3 depicts the estimated seasonality function for sample year 2014 on a daily basis. As expected, average prices follow seasonal paths, showing their highest values in wintertime and lowest during mild seasons. In addition, we detect a significant discrepancy between values quoted on nonworking days and expected figures on weekdays. Table 7.4 indicates estimated parameters for daily terms, while Figure 7.4 exhibits the hourly profile over a full day for each of the four combinations of working *vs.* nonworking days and cold *vs.* warm seasons. A number of stylized facts featuring hourly profiles emerge: on-peak hours show higher prices than off-peak hours during weekdays throughout the year, whereas a similar property only holds for weekends during cold season. Indeed, the warm season exhibits smoother intraday price dispersion than cold seasons, with a marked effect at weekends. A change in season affects on-peak hours more sharply than off-peak hours. The same effect is more pronounced during weekdays than weekends. Wintertime shows a spike in the evening across the whole week.




Figure 7.3: **Estimated periodical trend $\hat{\Lambda}^*(t)$** resulting from additively combining seasonal, weekly, and daily patterns under the corresponding estimated parametric set. Values are shown on a daily granularity in the period comprised between January 1, 2014 and December 31, 2014.




Figure 7.4: **Hourly dummy profiles.** Dummy variable coefficients $d_{h,l}$ are shown against daily hours $h = 1, \dots, 24$. Upper panels depict series for time clusters corresponding to working and nonworking days during the cold season (October to March). Lower panels report series for analogous clusters during the warm season (April to September).




Figure 7.5: **Day-ahead predictable pattern, peakload and baseload adjustment terms.** Left panel: estimated day-ahead predictable pattern $\hat{\Lambda}^*(t)$ as of November 27, 2015 (blue line) over a maturity spectrum ranging up to December 31, 2016; calibrated baseload adjustment term $\varepsilon(t)$ (red line); and calibrated peakload adjustment term $\varphi(t)$ (black line). Right panel: short-term components up to January 31, 2016.

Figure 7.5 superposes historically estimated day-ahead predictable pattern $\hat{\Lambda}^*(t)$, baseload and peakload adjustment terms $\varepsilon(t)$ and $\varphi(t)$ computed as of November 27, 2015. The former exhibits a mixture of daily, weekly, and seasonal recurrence features, while adjustment component $\varepsilon(t)$ shows a continuous, smooth, and positive-valued path, and adjustment term $\varphi(t)$ unveils a piecewise constant trajectory featuring a symmetrical behavior around the zero-level abscissa.

By combining estimated pattern $\hat{\Lambda}^*(t)$ with risk-neutral components $\hat{\varepsilon}^*(t)$ and $\hat{\varphi}^*(t)$, we come up to an estimate for the forward kernel implied by the spot price series and forward quotes on the evaluation date t_0 . By using formula (3.6), we finally arrive at time t_0 HPFC spanning the time horizon up to December 31, 2016. Figure 7.6 reports two exhibits of the resulting forward curve with hourly granularity. The left panel draws the full curve path, while the right panel zooms in over the first two months of delivery. Hourly forward prices range from 2.49 EUR/MWh to 55.11 EUR/MWh. As expected, price variability is larger during the cold season than the warm season. The full path shows that the long-term curve is primarily determined through the estimated annual periodical components, while forward quotes indicate long-period average values exhibiting no specific shape. The short-term detail reported in the right panel of the exhibit highlights the relevance of weekly and daily patterns for shaping the curve. We may qualitatively assess curve consistency to baseload and peakload quotes by visually inspecting their level depicted as red and magenta colored segments.

7.2 Property Assessment

The curve we derived for the German power market shows a number of interesting features. By construction, it perfectly fits input baseload and peakload quotes reported in the futures segment




Figure 7.6: **Hourly price forward curve vs. quoted forward quotes.** Left panel: estimated hourly price forward curve as of November 27, 2015 (blue line) over a maturity spectrum ranging up to December 31, 2016. Right panel: short-term forward curve up to January 31, 2016. Red segments represent segmented baseload forward prices, while pink segments refer to peakload quotes.

of the market and embeds periodical patterns shown by spot quotes with hourly granularity. From a functional perspective, it preserves monotonicity properties displayed by market prices while showing trajectorial smoothness in agreement with traders' expectations. In addition, the curve is cross-sectionally local: changes occurring on any segmented baseload quote exclusively affect the curve estimate on the corresponding delivery interval as well as on the two adjacent ones. This property becomes particularly relevant when evaluating physical and financial assets and analyzing their sensitivity. For instance, a power plant's revenue over a calendar year is expected to be largely insensitive to any reasonable variation of a single baseload segment, a property ensured by cross-sectional localness.

We may further appreciate this property by considering quotations as of February 28, 2013, which we report on the last column of Table 7.5, and assess the effect on the forward curve generated by a sudden drop in the spot price by 20 EUR/MWh. We analyze relative fitting

Type	Delivery begin	Delivery end	February 26	February 27	February 28
Day	February 26, 2013	February 26, 2013	55.49	-	-
Day	February 27, 2013	February 27, 2013	59.67	59.67	-
Day	February 28, 2013	February 28, 2013	-	53.96	53.96
Month	March 1, 2013	March 31, 2013	38.88	39.38	40.83
Month	April 1, 2013	April 30, 2013	37.94	37.98	38.60
Quarter	April 1, 2013	June 30, 2013	37.17	37.21	37.68
Quarter	July 1, 2013	September 30, 2013	38.80	38.94	39.35
Quarter	October 1, 2013	December 31, 2013	44.21	44.15	44.45
Quarter	January 1, 2014	March 31, 2014	44.96	44.89	45.40
Year	January 1, 2014	December 31, 2014	41.87	41.75	42.12

Table 7.5: **Time stability assessment.** EEX baseload futures quotes as of February 26, 27, and 28, 2013. Data comprise three spot prices, two month contracts, four frontline quarters, and the next calendar quotation. Settlement prices are expressed in EUR/MWh.




Figure 7.7: Cross-sectional localness assessment. Left panel: Electricity forward curves computed using MSI method adopted in Benth *et al.* (2007): the blue curve refers to market data as of February 28, 2013; the red curve stems from an artificial downward drop by 20 EUR/MWh at the shortest end in the time-to-maturity spectrum. We clearly see that a local drop propagates over the entire cross-sectional dimension. Right panel: Electricity forward curves computed using the proposed method: the two curves shows a slight discrepancy right on and next to the spot date affected by the price drop; all the remaining portions of the time-to-maturity spectrum exhibit a common path of forward prices.

quality compared to the Maximum Smoothness Interpolation (MSI) method adopted by Benth *et al.* (2007), which we take as a benchmark. Figure 7.7 shows the daily forward curve produced with the perturbed data setting using either construction methodologies. The left panel exhibits forward curves derived using MSI. The blue curve stems from market data quoted on February 28, 2013. A drop in the spot price generates forward estimate variations over the entire maturity spectrum as the red curve shows. However, Poletti Laurini and Ohashi (2015) show that non-local interpolation behavior in forward curve construction may hinder the significance of statistical assessments, such as a Principal Component Analysis, based on the resulting data arrays. In addition, Andersen and Piterbarg (2010) underline the importance that the curve construction exclusively produces a local perturbation upon shifting benchmark values. Finally, time localness complies with the fact that electricity long-term contracts cannot be hedged using short term forwards. These considerations reinforce the appropriateness of this feature. Indeed, the right panel of Figure 7.7 shows forward curves computed on the same date and scenarios using the model we propose. Cross-sectional localness is reflected into a negligible price discrepancy throughout the following delivery period; more importantly, the remaining portions of the curve are left unaffected by the local drop. As an example, forward price estimates derived using the benchmark MSI applied to market perturbed input data for delivery on December 31, 2014, are 44.68 EUR/MWh and 46.90 EUR/MWh, respectively. In contrast, our model offers a single quote at 45.83 EUR/MWh, therefore showing no drop propagation effect beyond a limited and well-defined neighborhood of the spot date.

Time stability refers to the property of slight curve variation following small perturbations in the input data. A lack of stability may result from the adopted interpolation method. Time unstable fitting entails unreliable forward prices and biased forward volatility estimates based on historical




Figure 7.8: Time stability assessment. Left panel: Electricity forward curves are computed on February 26, 27, and 28, 2013, using MSI method adopted in Benth *et al.* (2007): curves across the three-day period show pronounced variation along with maturity, a fact featuring a particularly marked lack of time stability on the shortest term segment. (Right panel) Electricity forward curves are obtained on the same time frame using the proposed method: curves across the three days are virtually indistinguishable, a fact mirroring a similar property owned by futures quotes.

series of extrapolated forward curves, a fact that may introduce spurious risk terms upon asset pricing and risk assessment. We may appreciate time stability from an empirical standpoint. Consider futures prices quoted on February 26, 27, and 28, 2013 (values are reported in Table 7.5). We note that data vary slightly from one evaluation date to the next. The largest variation refers to futures contracts delivering on March 2013, which shift from 38.88 EUR/MWh to 39.38 EUR/MWh, up to 40.83 EUR/MWh on the last quotation day. Overall, this move amounts to an almost 5% variation over the period. Each day, we compute a market HPFC using the proposed method and MSI, then we exhibit the resulting curves on a combined graph. Left panel of Figure 7.8 shows that MSI leads to a time unstable triplet of forward curves. This is reflected by important fluctuations of estimated forward prices corresponding to certain maturities. For instance, the forward price for delivery on March 31, 2013, varies from 7.42 EUR/MWh, up to 34.63 EUR/MWh, and then down to 25.41 EUR/MWh. The right panel reports the three curves obtained using the methodology developed in this study. It shows a negligible perturbation which corresponds with the way input market forwards vary over the three days under exam. The forward price for delivery on March 31, 2013, now varies from 26.71 EUR/MWh to 26.97 EUR/MWh, then up to 28.00 EUR/MWh on the last day of quotation.

A last test aims at measuring the quality of forward price estimates derived using the proposed model. Electricity forward prices are known to be biased estimates of expected spot prices (Longstaff and Wang (2004)). An absolute quality check would require a reliable valuation model of the underlying market price of risk. Although this approach is undoubtedly rigorous from a financial theory perspective, it would necessarily introduce an additional source of model risk. Our goal here is to evaluate the relative quality of the proposed model compared to the benchmark MSI. Hence, we propose a test quantifying the goodness of fit in a fictitious framework featuring no market price of risk. This method allows us to assess the estimated

curve as well as each of the three components $\hat{\Lambda}^*(t)$, $\hat{\varepsilon}^*(t)$, and $\hat{\varphi}^*(t)$, each taken separately from the others. We consider as many as sixteen evaluation dates, each corresponding to the first day of January, April, July, and October of years ranging from 2010 to 2013, respectively. Historical spot prices preceding an evaluation date are used to estimate periodical patterns at that point in time. In addition to the standard form put forward in formula (6.5), we consider two alternative specifications. One is obtained from the former, by replacing the seasonal term with a monthly component, namely:

$$\Lambda(u) = \sum_{i=1}^{12} g_i \mathbb{1}(u \in month_i) + \sum_{j=1}^7 c_j \mathbb{1}(u \in day_j) + \sum_{h=1}^{24} \sum_{l=1}^4 d_{h,l} \mathbb{1}(u \in hour_h \cap u \in \mathcal{C}_l). \quad (7.1)$$

The other functional forms combine seasonal and daily terms only, that is:

$$\Lambda(u) = a \cos\left(\frac{2\pi}{365}u + b\right) + \sum_{h=1}^{24} \sum_{l=1}^4 d_{h,l} \mathbb{1}(u \in hour_h \cap u \in \mathcal{C}_l). \quad (7.2)$$

We denote these three curves as SWD (seasonal, weekly, daily), MWD (monthly, weekly, daily), SD (seasonal, daily), respectively. Historical spot prices following each evaluation date are arithmetically averaged over a number of settlement periods to generate fictitious baseload and peakload futures quotations. They represent fictitious forward prices under the historical probability, hence based on the assumption of vanishing market price of risk. Specifically, we build up risk-neutral forward prices for day-ahead, the first three consecutive months, and the following three consecutive quarters. Seasonality functions (7.1) and (7.2), combined with fictitious risk-free forward prices enter the proposed curve model. Corresponding to each combination of the aforementioned sixteen evaluation dates, three periodical patterns, and two computational methods, we derive a forward curve with hourly granularity spanning the following one-year baseload delivery period. This curve is compared to the time series $\hat{\mathbf{S}}$ of hourly spot prices actually realized in the period following the corresponding evaluation date. Comparison is performed using a median absolute deviation defined as:

$$\text{MAD} = \frac{1}{24\#\mathcal{S}} \sum_{T \in \mathcal{S}} \sum_{h=1, \dots, 24} \left| \hat{f}(T, h) - \hat{S}(T, h) \right|,$$

over the time frame $\mathcal{S} = \{t_0, t_0 + 1\text{day}, \dots, 1\text{year}\}$. Here $\#\mathcal{S}$ denotes the number of days comprised in the set \mathcal{S} . As long as backtesting is run using either benchmark MSI or our model, we obtain as many as ninety-six MAD figures, which we report in Table 7.6. The new model outperforms benchmark MSI on 34 out of 48 scenarios, or 70.83% of the sample. Last, periodical patterns are best captured by the functional form (6.5).

Periodical trend	SWD	SWD	MWD	MWD	SD	SD
Model	New	MSI	New	MSI	New	MSI
1-Jan-2011	4.6042	5.0555	5.0313	5.5224	5.3240	5.9655
1-Apr-2011	4.8767	4.8303	5.4364	5.2706	5.5582	5.4705
1-Jul-2011	5.0922	4.9534	5.5507	5.4632	6.0746	5.9421
1-Oct-2011	4.8753	5.0682	4.9590	5.5031	6.2292	6.3474
1-Jan-2012	5.5590	6.5753	5.6216	6.4068	7.1990	7.6437
1-Apr-2012	5.3410	5.5142	5.4939	5.6720	7.3011	7.7397
1-Jul-2012	5.4710	5.5494	5.5975	5.6378	7.1937	7.6346
1-Oct-2012	5.8395	5.7809	5.9352	5.9276	7.6463	7.5514
1-Jan-2013	5.8953	5.9180	5.8241	6.0907	7.8100	7.8299
1-Apr-2013	6.1887	6.3537	6.2484	6.3643	8.0009	8.3227
1-Jul-2013	5.5793	5.6755	5.7439	5.7600	7.5643	7.5830
1-Oct-2013	5.5776	5.5472	5.7322	5.6595	6.8063	6.8080
1-Jan-2014	5.3186	5.3206	5.3515	5.4954	6.5844	6.4773
1-Apr-2014	5.4140	5.4028	5.3760	5.4036	6.5418	6.5450
1-Jul-2014	4.9721	5.0259	5.0257	5.0965	6.2374	6.2063
1-Oct-2014	5.4118	5.4859	5.4465	5.5648	6.5062	6.6435
Average	5.3760	5.5035	5.5234	5.6774	6.7861	6.9194

Table 7.6: **Fitting quality assessment.** On each evaluation date corresponding to the first day in a month, ranging from January 1, 2011, to October 1, 2014 at a quarterly pace, we compute the median absolute deviation (MAD) between hourly spot price series which occurred during the front year (verify 2nd instance) and hourly forward quotes computed by(-cut) using fictitious figures obtained through future spot price averaging. The forward curve is built using either our model or the MSI adopted in Benth *et al.* (2007), each one implemented using a periodical pattern of functional form SWD (seasonal, weekly, daily), MWD (monthly, weekly, daily), or SD (seasonal, daily).

8 Conclusion

A reliable electricity forward curve is vital for a number of precious uses. These include marking energy portfolios to market quotes and accounting for their fair values (*e.g.*, Teixeira Lopes (2007)); calibrating arbitrage models for pricing electricity positions and assessing their exposure (*e.g.*, Islyayev and Date (2015)); analyzing risk premia and volatility of swap price returns (*e.g.*, Frestad *et al.* (2010)); conceiving and testing proprietary trading rules involving futures contracts (*e.g.*, Furio and Lucia (2009)); planning optimized long-term consumption strategies for energy-intensive processes (*e.g.*, Lima *et al.* (2015)); and performing physical asset valuation using the theory of real options (*e.g.*, Näsäkkälä and Fleten (2005)), among others.

We propose a constructive definition for an electricity price forward curve with thin granularity and develop the corresponding algorithm for the case of hourly frequency. By using filtered spot price data, we identify and assess periodical price patterns, which in turn combine with futures quotes, leading to our estimate of the underlying forward kernel. The resulting curve is jointly consistent to baseload and peakload prices. It also features regularity properties which include trajectorial smoothness, monotonicity preservation, cross-sectional localness, and time stability. From an empirical standpoint, we build forward curve estimates for the EPEX Spot - EEX

marketplace and evaluate their relative performance against the benchmark model of Benth *et al.* (2007).

Future research may envisage assessing the interplay between cross-sectional granularity and the underlying market price of risk; building and analyzing forward curves for other commodity markets exhibiting periodical patterns, *e.g.*, natural gas and most agriculturals; and, addressing the traditional issue of expectation hypothesis under price quotation with thin granularity.

ACCEPTED MANUSCRIPT

References

- Adams, K.J., van Deventer, D.R (1994). Fitting yield curves and forward rate curves with maximum smoothness. *Journal of Fixed Income*, 4, 53-62.
- Anderson, N., Breedon, F., Deacon, M., Derry, A., Murphy, G. (1996). *Estimating and interpreting the yield curve*. London: John Wiley & Sons.
- Baxter, M., King, R.G. (1999). Measuring business cycles: Approximate band-pass filters for economic time series. *The Review of Economics and Statistics* 81, 575-593.
- Benth, F.E., Koekebakker, S., Ollmar, F. (2007). Extracting and applying smooth forward curves from average-based commodity contracts with seasonal variation. *Journal of Derivatives*, 15(1), 52-66.
- Björk, T. (2009). *Arbitrage theory in continuous time* (3rd ed.). Oxford University Press.
- Borak, S., Weron, R. (2008). A semiparametric factor model for electricity forward curve dynamics. *The Journal of Energy Markets*, 1(3), 3-16.
- Burger, M, Graeber, B., Schindlmayr, G. (2014). *Managing energy risk*. (2nd ed.). Chichester: John Wiley & Sons.
- Cartea, A., Figueroa, M. (2005). Pricing in electricity markets: a mean reverting jump-diffusion model with seasonality. *Applied Mathematical Finance*, 12(4), 313-335.
- Duffie, D. (2001). *Dynamic asset pricing theory*. (3rd ed.). Princeton: Princeton University Press.
- Fleten, S., Lemming, J. (2003). Constructing forward price curves in electricity markets. *Energy Economics*, 25, 409-424.
- Frestad, D., Benth, F.E., Koekebakker, S. (2010). Modeling term structure dynamics in the nordic electricity swap market. *Energy Journal*, 31(2), 53-86.
- Furió, D., Lucia, J.J. (2009). Congestion management rules and trading strategies in the Spanish electricity market. *Energy Economics*, 31(1), 48-60.
- Garcia, R., Lewis, M.A., Pastorello, S., Renault, E. (2011). Estimation of objective and risk-neutral distributions based on moments of integrated volatility. *Journal of Econometrics*, 160(1), 22-32.
- Hagan, P.S., West, G. (2006). Interpolation methods for curve construction. *Applied Mathematical Finance*, 13(2), 89-129.
- Hildmann, M., Kaffe, E., He, Y., Andersson, G. (2012). Combined estimation and prediction of the hourly price forward curve. *IEEE Power and Energy Society General Meeting*.

- Hildmann, M., Andersson, G., Caro, G., Daly, D., Rossi, S. (2013). What makes a good hourly price forward curve? *European Energy Markets 10th Intl. Conference*.
- Hodrick, R.J., Prescott, E.C. (1997). Postwar US business cycles: an empirical investigation. *Journal of Money, Credit and Banking*, 29(1), 1-16.
- Huisman, R. Huurman, C., Mahieu, R. (2007). Hourly electricity prices in day-ahead markets. *Energy Economics*, 29. 240-248.
- Hyman, J.M. (1983). Accurate monotonicity preserving cubic interpolation. *SIAM Journal of Scientific and Statistical Computing*, 4(4), 645-654.
- Islyayev, S., Date, P. (2015). Electricity futures price models: calibration and forecasting. *European Journal of Operational Research*, 247(1), 144-154.
- Janczura, J., Truck, S., Weron, R., Wolff, R. (2013). Identifying spikes and seasonal components in electricity spot price data: a guide to robust modeling. *Energy Economics*, 38, 96-110.
- Jarrow, R. (2014). Forward rate curve smoothing. *The Annual Review of Financial Economics*, 6, 443-458.
- King, R.G., Rebelo, S.T. (1993). Low frequency filtering and real business cycles. *Journal of Economic Dynamics and Control* 17, 207-231.
- Koekebakker, S., Os Adland, R. (2004). Modelling forward freight rate dynamics: empirical evidence from time charter rates. *Maritime Policy and Management*, 31(4), 319-335.
- Lima, R.M., Novais, A.Q., Conejo, A.J. (2015). Weekly self-scheduling, forward contracting, and pool involvement for an electricity producer. An adaptive robust optimization approach. *European Journal of Operational Research*, 240(2), 457-475.
- Longstaff, F., Wang, A.W. (2004). Electricity forward prices: a high-frequency empirical analysis. *The Journal of Finance*, 59(4), 1877-1900.
- Näsäkkälä, E., Fleten, S.E. (2005). Flexibility and technology choice in gas fired power plant investments. *Review of Financial Economics*, 14(3-4), 371-393.
- Nowotarski, J., Tomczyk, J., Weron, R. (2013). Robust estimation and forecasting of the long-term seasonal component of electricity spot prices. *Energy Economics* 39, 13-27.
- Paraschiv, F., Fleten, S.E., Schürle (2015). A spot-forward model for electricity prices with regime shifts. *Energy Economics*, 47, 142-153.
- Pedersen, T.M. (2001). The Hodrick-Prescott filter, the Slutsky effect, and the distortionary effect of filters. *Journal of Economic Dynamics and Control*, 25, 1081-1101.
- Poletti Laurini, M., Ohashi, A. (2015). A noisy principal component analysis for forward rate curves. *European Journal of Operational Research*, 246(1), 140-153.

Taylor, J.W., Buizza, R. (2004). A Comparison of temperature density forecasts from GARCH and atmospheric models. *Journal of Forecasting* 23, 337-355.

Teixeira Lopes, P. (2007). Accounting for electricity derivatives under IAS 39. *Journal of Derivatives & Hedge Funds*, 13(3), 233-246.

Truck, S., Weron, R., Wolff, R. (2007). Outlier treatment and robust approaches for modeling electricity spot prices. *Technical Report*, Hugo Steinhaus Center, Wroclaw Univ. Tech.

Weron R., Zator, M. (2015). A note on using the Hodrick-Prescott filter in electricity markets. *Energy Economics*, 48, 1-6.

Ziel, F., Steinert, R., Husmann, S. (2015). Efficient modeling and forecasting of electricity spot prices. *Energy Economics* 47, 98-111.

ACCEPTED MANUSCRIPT

A Monotone Convex Interpolator

Given scalars ε_i^A ($i = 1, \dots, n$), let us define constants:

$$\varepsilon_i = \begin{cases} \text{bound}(0, \varepsilon_1^A - (\varepsilon_1 - \varepsilon_1^A)/2, 2\varepsilon_1^A), & i = 0 \\ \text{bound}\left(0, \frac{t_i - t_{i-1}}{t_{i+1} - t_{i-1}} \varepsilon_{i+1}^A + \frac{t_{i+1} - t_i}{t_{i+1} - t_{i-1}} \varepsilon_i^A, 2 \min(\varepsilon_i^A, \varepsilon_{i+1}^A)\right), & i = 1, \dots, n-1, \\ \text{bound}(0, \varepsilon_n^A - (\varepsilon_{n-1} - \varepsilon_n^A)/2, 2\varepsilon_n^A), & i = n, \end{cases} \quad (\text{A.1})$$

where operator $\text{bound}(a, x, b) := \min(\max(a, x), b)$. Set $g_i(0) := \varepsilon_{i-1} - \varepsilon_i^A$ and $g_i(1) := \varepsilon_i - \varepsilon_i^A$. Hagan and West (2006) define convex linear interpolators as functions $g_i : [0, 1] \rightarrow \mathbb{R}$ ($i = 1, \dots, n$) with analytical expressions given by:

case 1: [$g_i(0) > 0$ and $-0.5g_i(0) \geq g_i(1) \geq -2g_i(0)$] or [$g_i(0) < 0$ and $-0.5g_i(0) \leq g_i(1) \leq -2g_i(0)$]:

$$g_i(x) := g_i(0)(1 - 4x + 3x^2) + g_i(1)(2x + 3x^2),$$

case 2: [$g_i(0) < 0$ and $g_i(1) > -2g_i(0)$] or [$g_i(0) > 0$ and $g_i(1) < -2g_i(0)$]:

$$g_i(x) := \begin{cases} g_i(0) & \text{for } 0 < x \leq \eta \\ g_i(0) + (g_i(1) - g_i(0)) \left(\frac{x-\eta}{1-\eta}\right)^2 & \text{for } \eta < x < 1 \end{cases}, \quad \eta = \frac{g_i(1) + 2g_i(0)}{g_i(1) - g_i(0)},$$

case 3: [$g_i(0) > 0$ and $0 > g_i(1) > -0.5g_i(0)$] or [$g_i(0) < 0$ and $0 < g_i(1) < -0.5g_i(0)$]:

$$g_i(x) := \begin{cases} g_i(1) + (g_i(0) - g_i(1)) \left(\frac{\eta-x}{\eta}\right)^2 & \text{for } 0 < x \leq \eta \\ g_i(1) & \text{for } \eta < x < 1 \end{cases}, \quad \eta = \frac{3g_i(1)}{g_i(1) - g_i(0)},$$

case 4: [$g_i(0) \geq 0$ and $g_i(1) \geq 0$] or [$g_i(0) \leq 0$ and $g_i(1) \leq 0$]

$$g_i(x) := \begin{cases} A + (g_i(0) - A) \left(\frac{\eta-x}{\eta}\right)^2 & \text{for } 0 < x \leq \eta \\ A + (g_i(1) - A) \left(\frac{x-\eta}{1-\eta}\right)^2 & \text{for } \eta < x < 1 \end{cases}, \quad \begin{aligned} \eta &= \frac{g_i(1)}{g_i(1) - g_i(0)} \\ A &= -\frac{g_i(0)g_i(1)}{g_i(0) + g_i(1)} \end{aligned}.$$

B Model and Implementation Details

B.1 On the role of temperature time series.

In their role of predictors within time series models, temperature records are undoubtedly useful at forecasting energy prices and their periodical patterns. In particular, accurate predictions of forward temperatures may represent significant factors in forecasting models. However, temperature predictions are accurate for time horizons up to a few days, say from one to ten at best. If we focus on longer time frames, temperature estimates experience a sharp drop in their forecasting quality (Taylor and Buizza (2004)). Our definition of forward curve spans a time period typically covering one or two years from the date of assessment. Correspondingly, it seems that temperature series do not offer a relevant piece of information for estimating price patterns. Fortunately, seasonal effects are by definition long-term phenomena grounded on price averages. Which in turn are fully captured and shown by spot price series. In addition, using time series of temperature records would require having estimates with hourly granularity, which may prove difficult or excessively costly to obtain. Finally, one would also have to consider a reasonable way to combine temperatures related to a large number of regions in a way that the resulting average mimics the same periodical patterns as those of the gas market covering those areas. These considerations underpin our endogenous approach, whereby long-term pathways are estimated by using prices only. We would hasten to note that our method undergoes a number of tests, all showing significant performance at pricing.

B.2 On the unicity of the forward kernel.

The forward kernel is a mathematically convenient object to build forward curves of arbitrary granularity. It basically conveys market information into a channel leading to the desired term structure of forward prices. Theoretically, an infinite number of forward kernels, and thus curves with thin granularity, can generate any finite set of observed prices. This is a well-known issue in the finance literature dealing with curve extrapolation and model calibration. See, e.g., Anderson et al. (1996). There are essentially two ways to overcome this challenge. One consists of putting forward an arbitrage-free model for the commodity price in question; then, a model-based forward curve is analytically derived and fit to market observables (Björk (2009)). The other strategy is model-free. It uses optimization techniques to fit suitable functional forms of the curve to actual market quotes. We adopt the second approach in light of a number of advantages. First, the output curve kernel is optimized with respect to rational criteria drawn from considerations about properties required by the typical final user. This comes as opposed to abstract fitting of arbitrage-free models to market prices. Second, the existence and uniqueness of the forward kernel are ensured by the corresponding properties of the underlying optimization problems. As long as they entail optimizing continuous and concave target functions on compact domains, both properties are guaranteed. Third, any resulting forward kernel is compatible to

whatever arbitrage-free pricing model is sufficiently flexible to host the output curve. Finally, forward prices with arbitrary granularity may easily be obtained by straightforward integration over the appropriate period of time. These considerations show that forward kernel unicity is not an issue to the extent that we do not need to rely on any previously selected arbitrage model of commodity prices.

B.3 On the selection of the Hodrick-Prescott parameter λ (or ω_c).

We tested the stability of the procedure within the numerical example developed in Section 7.1. Specifically, we computed sixteen forward curves over as many threshold periods in the outlier identification step (See pictures reported below). Cutoff periods vary from a single week to a full year along with increasing steps. The optimal Hodrick-Prescott parameter λ is computed on a daily basis according to Pedersen's procedure, as described in the paper. We calculated deviations from the benchmark curve with period ω_c set equal to thirty days as a mean absolute deviation between the resulting hourly price forward curves. Resulting figures are reported in the following Table.


Days	Lambda Opt	Nr Outlier	Deviation
7	2	26	0,0817
15	116	27	0,0429
30	1129	27	0,0000
45	6501	25	0,0471
60	16941	26	0,0084
75	47629	28	0,0181
90	99813	29	0,0228
120	230058	31	0,0274
150	721855	31	0,0274
180	1297804	31	0,0254
210	2744315	31	0,0254
240	3665992	32	0,0256
270	6472759	31	0,0240
300	12364353	31	0,0240
330	12425597	31	0,0240
365	42023717	30	0,0238

We observe that the discrepancies between the benchmark curve and all other curves computed using alternative values for λ^* 's are rather small. The greatest difference stems from ω_c equal to one week and accounts for as many as 8 Euro cents per megawatt-hour.

B.4 On the selection of the smoothing parameter α .

Smoothing parameter α appearing in equations 6.8 and 6.9 represents an exponentially decaying factor in the target functional. This number allows the modeler to tune the relative importance attributed to data recorded on varying periods of time in the past. We provide insight on the relevance of setting $\alpha = 0.4$ related to the typical time scale involved in the underlying price

dynamics. Following is a graph showing the cumulated percentage of information weight seen as a function of the recording period expressed in days. Vertical lines mark yearly periods. For instance, by considering the 2,521 days under analysis and by using $\alpha = 0.4$, data in the most recent year receive a weight equal to 35% of the total information.







We originally performed a number of empirical experiments across several European markets, assuming a variety of seasonal shapes. Retained selection $\alpha = 0.4$ worked quite well in all cases. However, we made no assessment for these parameters based on criteria of optimality. We performed a test on the stability of the procedure under the setting of Section 7.1. Specifically, we computed twenty forward curves under as many parameter “alpha” adopted upon model calibration. We considered values between 0.1 to 2.0, lagged by a step equal to 0.1. Deviations from our curve (i.e., $\alpha = 0.4$ days) taken as a benchmark are computed in terms of mean absolute deviation and reported in the following table.



Alpha	Deviation
0.1	0.918
0.2	0.526
0.3	0.230
0.4	0.000
0.5	0.181
0.6	0.326
0.7	0.443
0.8	0.540
0.9	0.623
1	0.694
1.1	0.757
1.2	0.812
1.3	0.862
1.4	0.907
1.5	0.947
1.6	0.985
1.7	1.020
1.8	1.052
1.9	1.082
2	1.111

Differences between each curve and the benchmark are rather modest across varying values for parameter α . The greatest discrepancy occurs at $\alpha = 2.0$, showing an approximate deviation of 1 Euro and 11 cents. By combining the aforementioned qualitative and quantitative considerations, we may reasonably retain the selected value for alpha as a rational assessment for this parameter.



ACCEPTED MANUSCRIPT

*Additional Figures.***Days: 7 and 15****Days: 30 and 45**



Days: 60 and 75





Days: 90 and 120





Days: 150 and 180



Days: 210 and 240



Days: 270 and 300



Days: 330 and 365

