GERAD

GROUPE D'ÉTUDES ET DE RECHERCHE
EN ANALYSE DES DÉCISIONS

Les Cahiers du GERAD

G-2013-88
December 2013

CITATION ORIGINALE / ORIGINAL CITATION

D.J. Dupuis, G. Gauthier, F. Godin, Short-term hedging for an electricity retailer, Energy Journal, 37(2) 31-59, 2016.

GERAD HEC Montréal
3000, ch. de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Les Cahiers du GERAD ISSN: 0711–2440

Short-Term Hedging for an Electricity Retailer

D. Dupuis, G. Gauthier, F. Godin

G–2013–88

December 2013

Short-Term Hedging for an Electricity Retailer

Debbie Dupuis
Geneviève Gauthier
Frédéric Godin

GERAD & Department of Management Sciences
HEC Montréal
Montréal (Québec) Canada, H3T 2A7

debbie.dupuis@hec.ca
genevieve.gauthier@hec.ca
frederic.godin@hec.ca

December 2013

Les Cahiers du GERAD
G–2013–88

Abstract: A dynamic global hedging procedure making use of futures contracts is developed for a retailer of the electricity market facing price, load and basis risk. Statistical models reproducing stylized facts are developed for the electricity load, the day-ahead spot price and futures prices in the Nord Pool market. These models serve as input to the hedging algorithm, which also accounts for transaction fees. Backtests with market data from 2007 to 2012 show that the global hedging procedure provides considerable risk reduction when compared to hedging benchmarks found in the literature.

Key Words: Risk management, power markets, energy, load modeling, futures contracts.

1 Introduction

With the recent liberalization of electricity markets and the disentanglement of the vertical integration in the electricity supply chain in Nordic countries, continental Europe and North America, new risks have arisen for some of the participants of the electricity markets. One such risk-facing participant is the retailer buying from wholesalers to sell to end-users. These retailers[1] sign contracts giving them the obligation to supply electricity to consumers. Retailers often need to supply a quantity of electricity at a fixed price while acquiring it at a variable market price (Von der Fehr and Hansen 2010, Johnsen 2011), exposing the retailers to price risk. Furthermore, as the quantity of electricity which must be supplied to consumers is uncertain, retailers also face load (or volumetric) risk (Deng and Oren 2006).

Electricity is not easily storable and retailers cannot build up electricity reserves upon which to draw to cover an unexpectedly high load demand or an electricity price increase. The non-storability of electric power fuels extreme price volatility as highly inelastic demand can cause spot prices to skyrocket when shortages occur. For retailers, the volatility can affect profitability since an unexpected high cost of electricity can lead to major losses. *The profit margin for a retailer is so small in relation to the price risk that the profit margin can quickly disappear if the price risk is not hedged* (NordREG 2010). In some cases, there was eventual bankruptcy as with the Pacific Gas and Electric Company in 2001 and Texas Commercial Energy in 2003. To prevent such events, some government regulatory initiatives were even implemented to force retailers to hedge their obligation to serve electricity loads. For example, the California Public Utility Commission now requires load serving entities (LSE) to use forward contracts and options (with mandatory physical settlement) to reduce their risk exposure (State of California 2004).

It is clear that deficient risk management can lead to financial hardship for retailers and developing effective hedging methodologies in the electricity market has become paramount. Different approaches, using different electricity derivatives, have been proposed in the literature. Deng and Oren (2006) survey available derivatives and list the papers that implement methods pertaining to each. Hedging procedures can be divided into two main categories: (i) static, and (ii) dynamic. For static hedging, hedging instruments are bought at one point in time and the hedging portfolio is never rebalanced. For dynamic hedging, the composition of the hedging portfolio is adjusted through time as additional information becomes available. Dynamic hedging procedures can be divided into two sub-categories, which we refer to as *local* and *global* hedging. Local hedging procedures minimize the risk associated with the portfolio until the next rebalancing whereas global hedging procedures minimize the risk related to the terminal cash flow.

Several papers apply static hedging without considering load uncertainty. Stoft et al. (1998) describe simple hedging strategies with vanilla derivatives. Bessembinder and Lemmon (2002) identify the optimal position in forward contracts for electricity producers and retailers through an equilibrium scheme. Fleten et al. (2010) optimize the static futures contract position of a hydro-power electricity producer in Nord Pool. Other papers studying static hedging incorporate load uncertainty in their model. Wagner et al. (2003) and Woo et al. (2004) investigate static hedges with forward and futures contracts under different risk constraints. Deng and Xu (2009) examine hedging strategies using interruptible contracts in a one-period setting. In a series of papers, Oum et al. (2006), Oum and Oren (2009) and Oum and Oren (2010) propose a static hedging procedure maximizing the expected utility of a LSE using a portfolio of derivatives. Kleindorfer and Li (2005) optimize the expected return of an electricity portfolio corrected by a risk measure (either variance or Value-at-Risk).

To the authors' best knowledge, there exists no paper on dynamic hedging which incorporates load uncertainty. The literature on dynamic hedging strategies includes some local procedures. For example, Ederington (1979) suggests to hedge an underlying asset with its futures in a way to minimize the one-period variance of the total portfolio. Byström (2003), Madaleno and Pinho (2008), Zanotti et al. (2010), Liu et al. (2010), and Torro (2012) adapt this procedure to the electricity market, but with different model specifications for the spot and futures prices. Byström (2003) applies one-week horizon hedges on Nord Pool, comparing conditional and unconditional hedge ratios. The unconditional version of hedge ratios outperforms

[1] Retailers is the term used on Nord Pool for these participants. On the US market, they are referred to as load serving entities.

the conditional models. Madaleno and Pinho (2008) and Zanotti et al. (2010) compare different correlation models for the spot and futures prices to compute optimal hedge ratios on European electricity markets. Liu et al. (2010) use copulas to represent the relationship between the spot and futures prices. Torro (2012) studies the case of early dismantlement of the hedging portfolio in the Nord Pool market.

Alternative dynamic hedging schemes are discussed in Eydeland and Wolyniec (2003). For example, there is delta hedging, a method which consists in building a portfolio with value variations that mimic those of the hedged contingent claim. Eydeland and Wolyniec (2003) apply delta hedging to achieve perfect replication when a LSE hedges the price of a fixed amount of load to be served. When perfect replication cannot be achieved, they propose local mean-variance optimization to tackle hedging problems.

Local procedures are attractive because they are simple to implement. Local risk minimization procedures are myopic however as they do not necessarily minimize the risk through the entire period of exposure (see Rémillard 2013). Global hedging procedures remedy this drawback by taking into account the outcomes of all future time periods at any point in time; they evaluate the adequacy of a hedge by looking at the terminal hedging error, i.e. at the maturity of the hedged contingent claim. The following is a non-exhaustive list of papers which study this methodology in general financial contexts. Schweizer (1995) minimizes the global quadratic hedging error in a discrete-time framework for European-type securities. Rémillard et al. (2010) extend his work for American-type derivatives. Föllmer and Leukert (1999) minimize the probability of incurring a hedging shortfall. Föllmer and Leukert (2000) minimize an expected function of the terminal hedging error. To the authors' best knowledge, developing global dynamic hedging procedures in electricity markets has not yet been attempted.

The current paper therefore seeks to fill the gap in the literature concerning global hedging procedures for electricity markets and offers three main contributions. First, a dynamic global hedging methodology is developed for a retailer trying to hedge itself with futures contracts by considering both its price and load risks. Obtaining global solutions to hedging problems is non-trivial and it often requires advanced numerical schemes. This could explain why this avenue has not yet been explored in electricity markets. We not only show that the approach is feasible, but present the first dynamic hedging strategy to account for load risk, and one of very few to account for transaction costs. Second, as our global hedging algorithm uses *weekly* futures, we develop the required *weekly* load model. We also present a statistical model for the joint dynamics of the spot and futures prices. A statistical approach using multivariate time series analysis is applied. Third, an empirical study which compares the performance of different hedging procedures on the Nord Pool market is presented.

The non-quadratic global hedging procedure developed outperforms the benchmarks in reducing the risk borne by the retailers. Hedging backtests show a significant reduction in several risk metrics applied to the weekly hedging error. Considering the case of a retailer serving 1% of the Nord Pool load, the $\text{TVaR}_{1\%}$ is reduced from 172,900€ to 161,900€ if our load-basis model is used in the delta hedging procedure (see Table 7). When our global hedging procedure is applied, the $\text{TVaR}_{1\%}$ is further shrunk by a considerable amount to 133,100€.

The remainder of the paper is organized as follows. Section 2 presents the price and volumetric risks faced by retailers and describes the hedging procedure. Section 3 describes the data used for modeling purposes and presents the models for the electricity load, the spot price and futures prices. Section 4 describes the numerical experiments which test the efficacy of the hedging methodology. Section 5 presents concluding remarks. Some technical results, estimation details and goodness-of-fit tests are relegated to Appendices.

2 Risk exposure and hedging for retailers

In this section, we describe the risks faced by a retailer and the hedging procedure it can undertake to hedge its exposure with futures contracts.

2.1 Risks faced by retailers

Consider the case of a retailer forced to supply a quantity of electricity at a fixed price to end-users while buying it at a variable price on the market. Market conditions for a retailer might differ across different electricity markets and we look specifically at the Nordic electricity market Nord Pool. This market is chosen since it is one of the first to operate in a liberalized setup, and the mature markets provide some historical data less likely to include structural changes.

Assume all electricity purchases occur on the day-ahead market. In this market, the spot price is set on an hourly basis by balancing supply and demand. Suppose the retailer needs to serve the load $L_{t,d,h}$ during hour h of day d in week t, while $S_{t,d,h}$ is the Nord Pool system spot price for the corresponding period. The total load to be served during week t is thus

$$L_t = \sum_{d=1}^{7} \sum_{h=1}^{24} L_{t,d,h}. \tag{1}$$

The mean price paid by a retailer for the purchase of each unit of load during week t is

$$S_t^* = \frac{\sum_{d=1}^{7} \sum_{h=1}^{24} L_{t,d,h} S_{t,d,h}}{L_t}, \tag{2}$$

the load-weighted average of all hourly prices during the week. Assume the retailer charges a constant price Π for each unit of load. If no hedging is implemented, the retailer cash flow for weekly operations during week t is

$$L_t(\Pi - S_t^*). \tag{3}$$

A retailer thus faces revenue uncertainty due to (i) price risk caused by the variability of the spot price $S_{t,d,h}$, and subsequently of S_t^*, and (ii) volumetric risk caused by randomness in the total volume L_t of electricity to be served.

2.2 Electricity futures contracts

A retailer wishes to hedge its exposure to both price and volumetric risks with derivatives on the electricity markets. Different derivatives are available to hedge those risks: forward and futures contracts, options, weather derivatives and interruptible contracts. With the exception of forward and futures contracts, most derivatives on the Nord Pool market are traded over-the-counter, are illiquid, and are not well-suited for dynamic hedging methodology as they may be unavailable when they are required. More liquid futures and forwards are better suited for dynamic hedging procedures.

For the Nord Pool market, futures and forward contracts are traded on the NASDAQ OMX. Futures contracts provide hedging for shorter horizons (daily and weekly), while forwards cover longer periods (months, quarters and years). The current paper focuses on short-term hedging. Tables 1 and 2 present the percentage of trading days on which non-null trading volumes occur for weekly and daily base load futures.[2] Liquidity is much higher on weekly contracts with 1−, 2− and 3−week maturities. These derivatives will be used in this paper.

Table 1: Liquidity of weekly futures

Weeks-to-maturity	1	2	3	4	5
Percentage	96%	90%	61%	29%	16%

Notes. Percentage of trading days between January 1st, 2007 and December 31th, 2012 with non-null trading volume of Nord Pool weekly futures on NASDAQ OMX.

[2] Base load means that the contracts deliver electricity during all hours of the day, in opposition to peak load contracts that only deliver electricity between 8:00 a.m. and 8:00 p.m.

Table 2: Liquidity of daily futures

Days-to-maturity	1	2	3	4	5
Percentage	64%	16%	8.8%	2.6%	2.0%

Notes. Percentage of trading days between January 1st, 2007 and December 31th, 2012 with non-null trading volume of Nord Pool daily futures on NASDAQ OMX.

Futures on Nord Pool are cash-settled; no exchange of the underlying commodity occurs. The underlying asset S_T of a weekly futures maturing at week T is the arithmetic average of the Nord Pool system spot price observed during week T:

$$S_t = \frac{1}{7 \times 24} \sum_{d=1}^{7} \sum_{h=1}^{24} S_{t,d,h}, \tag{4}$$

where a week starts on Monday and ends on Sunday. Contracts are traded on NASDAQ OMX during weekdays until the Friday of week $T-1$. The futures is thus not traded during its maturity week.

There is a slight mismatch between the average weekly electricity price paid by the retailer, given by (2), and the underlying asset of weekly futures given by (4). The basis ratio

$$\eta_t = \frac{S_t^*}{S_t} \tag{5}$$

links the former and the latter.[3] The basis ratio represents an additional source of risk which must be taken into account by the hedging procedure.

Futures contracts are marked-to-market. This means that (i) their cash flows do not occur strictly at maturity (in opposition to forward contracts) and (ii) the variation of their quote (referred to as the futures price) is reflected by the continuous transfer of funds between the margin accounts of the long and short position holders. Using futures contracts implies having to pay transaction fees and the cost of these will be accounted for in our methods.[4]

2.3 Hedging procedure

Throughout this section, the retailer hedges its week T exposure. A self-financing investment portfolio containing a risk-free asset[5] and futures with maturity week T is set up at t_0 and rebalanced weekly until $T-1$. Since futures are traded during weekdays only, rebalancing occurs on Fridays at closing time. The closing price on Friday of week t (or Sunday if $t = T$) of the risk-free asset is $B_t = \exp(rt)$, and the closing price of the futures is $F_{t,T}$, $t = t_0, \ldots, T-1$. Since futures are cash-settled, the last futures quote on Sunday of week T is automatically set by the clearing house to $F_{T,T} = S_T$. The hedging procedure is summarized by the following algorithm:

At week t_0. An initial amount of capital V_{t_0} is allocated for hedging purposes. The retailer enters into θ_{t_0+1} long positions on the futures contract. A portion $\mathcal{M}_{t_0}$ of the initial capital is placed in the margin account required by the clearing house. Another part is used to pay transaction fees $\mathcal{C}_{t_0}$. The remainder $\mathcal{B}_{t_0}$ is invested in the risk-free asset. If V_{t_0} is insufficient to cover the margin call and fees, the money is borrowed. Since entering positions on futures contracts involves no immediate cash flows besides the amount placed in the margin,

$$V_{t_0} = \mathcal{M}_{t_0} + \mathcal{B}_{t_0} + \mathcal{C}_{t_0}.$$

[3] As shown in Appendix A, the basis ratio η usually evolves between 1 and 1.05. This is explained by a higher spot price during peak hours when electricity consumption is more important.

[4] Transaction fees are described at http://www.nasdaqomx.com/commodities/Marketaccess/feelist/. Fixed annual fees for membership to the Exchange are disregarded in the current study. Variable fees which are proportional to the volume of futures transactions include Exchange fees (for trading positions) and Clearing fees (for clearing positions). Exchange fees are 0.004 EUR/MWh. Clearing fees depend on the volume of futures cleared in the most recent quarter, but they range from 0.0035 EUR/MWh to 0.0085 EUR/MWh. For illustrative purposes, a 0.004 EUR/MWh rate is used. Combining Exchange and Clearing fees, entering or clearing any long or short position is therefore approximated to cost 0.004 EUR/MWh.

[5] Since this paper focuses on short-term hedging, a constant weekly risk-free rate r is assumed.

Capital $\mathcal{M}_{t_0} + \mathcal{B}_{t_0} = V_{t_0} - \mathcal{C}_{t_0}$ (both inside and outside the margin) is invested (or borrowed) at the risk-free rate r.[6]

At week $t+1$, $t \in \{t_0, \ldots, T-2\}$. The total capital available for hedging (the sum of the amount placed in the margin account and in the risk-free asset) at week t before transaction costs are paid is V_t. This capital accrues interest up to week $t+1$ and is now worth

$$(V_t - \mathcal{C}_t)\frac{B_{t+1}}{B_t}.$$

The futures margin account of the retailer is adjusted from marking-to-market,[7] the amount $\theta_{t+1}(F_{t+1,T} - F_{t,T})$ is deposited (withdrawn if negative) in the margin account. The total capital available for hedging at week $t+1$ is therefore

$$V_{t+1} = (V_t - \mathcal{C}_t)\frac{B_{t+1}}{B_t} + \theta_{t+1}(F_{t+1,T} - F_{t,T}).$$

The retailer modifies its portfolio to hold θ_{t+2} long positions on the futures contract. Transactions fees $\mathcal{C}_{t+1}$ are paid. A margin call might be made, but it does not affect the total amount $V_{t+1} - \mathcal{C}_{t+1}$ invested at the risk-free rate.

At week T. The terminal hedging capital is

$$V_T = (V_{T-1} - \mathcal{C}_{T-1})\frac{B_T}{B_{T-1}} + \theta_T(S_T - F_{T-1,T}) - \mathcal{C}_T,$$

where $\mathcal{C}_T$ are clearing fees. Transaction costs are computed following $\mathcal{C}_T = 0.004|\theta_T|$ (final clearing costs) and $\mathcal{C}_t = 0.004|\theta_{t+1} - \theta_t|$ if $t < T$ with $\theta_{t_0} = 0$.

The retailer is at risk of bearing losses when the price it pays to purchase electricity is higher than the price it charges to its customers. To avoid this situation, the hedging algorithm proposed in this paper aims at minimizing risks related to electricity procurement costs incurred by the retailer. Having reliable procurement costs stabilizes the retailer's profitability.[8] Weekly futures, which are used by the retailer to hedge its exposure, allow locking in the payoff of the variable contingent claim S_T to $F_{t_0,T}$ (see Appendix B).[9] However, the retailer has short positions on S_T^* (instead of S_T) because it needs to buy electricity at that price. Since $S_T^* = \eta_T S_T$, the electricity procurement target price for each unit of load bought by the retailer during week T is set to $(S_T^*/S_T)F_{t_0,T} = \eta_T F_{t_0,T}$. The retailer's cash flow at time T, given by Equation (3), can be separated into an unhedged cash flow $L_T(\Pi - \eta_T F_{t_0,T})$, the baseline profit margin, and a more risky component $\mathcal{L}_T(S_T - F_{t_0,T})$, the procurement costs risk:

$$\begin{aligned} L_T(\Pi - S_T^*) &= L_T(\Pi - \eta_T F_{t_0,T}) - L_T(S_T^* - \eta_T F_{t_0,T}) \\ &= L_T(\Pi - \eta_T F_{t_0,T}) - \mathcal{L}_T(S_T - F_{t_0,T}) \end{aligned}$$

where the load-basis $\mathcal{L}_T$ is the product of the load and the basis factor:

$$\mathcal{L}_t = \eta_t L_t. \tag{6}$$

The procurement costs risk can cause large losses when the price S_T peaks way above the futures price $F_{t_0,T}$. The hedging strategy aims at offsetting the variation of the quantity

[6] It is assumed that the retailer can always borrow capital at the risk-free rate. Such an assumption has a limited impact; hedging errors are very insensitive to interest rates because of the short term horizon of the hedge.

[7] To simplify calculations, it is assumed that the futures are marked-to-market weekly. On NASDAQ OMX, marking-to-market is in reality performed daily. However, because maturities are short-term (and therefore accumulation of interest is small), such an approximation has only a minor impact.

[8] Hedging procurement costs does not remove all risks; profits are still proportional to the load. Adequate hedging of procurement costs will however prevent extreme losses.

[9] This is true if transaction fees are disregarded.

$$\Psi_T = \mathcal{L}_T(S_T - F_{t_0,T}) \tag{7}$$

while the retailer determines the fixed price Π to extract an expected but uncertain profit. Considering the load-basis $\mathcal{L}$ (instead of the load L and the basis factor η separately) is convenient since only a single model is required for the former quantity (instead of two models for the latter).

The retailer would like the terminal value of the hedging portfolio V_T to be bigger than the target Ψ_T (or at least as close as possible to it) to offset the procurement costs risk. The global hedging problem that must be solved is thus

$$\min_{(\theta_{t_0+1},\dots,\theta_T)\in\Theta} \mathbb{E}\left[G(\Psi_T - V_T)|\mathcal{G}_{t_0}\right], \tag{8}$$

where $V_T = V_T(\theta_{t_0+1},\dots,\theta_T)$, $\mathcal{G} = \{\mathcal{G}_t | t = t_0, ..., T\}$ is the filtration that defines the information available to the retailer,[10] Θ is the set of all trading strategies available to the retailer[11] and G is a penalty function which weights and sanctions losses. Some integrability and regularity conditions might need to be satisfied to ensure that the solution exists.

There are numerous possibilities for the penalty function G. A standard choice in the literature is the quadratic function, $G(x) = x^2$, since it conveniently leads to semi-analytical formulas (Schweizer 1995) and therefore enhances the tractability and the computational speed of the solution. This approach has two principal caveats: (i) the semi-analytical formulas do not take transaction fees into account; (ii) the quadratic penalty is symmetric, such that gains and losses are equally penalized.[12] To remedy the problem of penalized gains, we also consider a semi-quadratic penalty[13]

$$G(x) = x^2\mathbb{I}_{\{x>0\}}. \tag{9}$$

A retailer using this penalty tries to remove losses as much as possible and disregards gains. A drawback of using this penalty is that it leads to a substantial increase in the numerical burden. The computations are however still feasible for the current framework. The computation of solutions for problem (8) with penalty (9) is discussed in Appendix C. A simulation-based algorithm is proposed to solve the Bellman equation. This algorithm can accommodate a wide class of penalty functions.

3 Models for the state variables

To compute the optimal trading strategy, the dynamics of the state variables $\mathcal{L}_t$ and $F_{t,T}$, the key components in the hedging problem, must be modeled. The proposed models are constructed from historical data.

3.1 Load-basis

We assume that the load the retailer must supply is proportional to the entire system load on the Nord Pool spot market.[14] This proportionality assumption, which is justified by a high correlation between firm load and market load, is also found in Coulon et al. (2012) for the Texas electricity market.

Load forecasting has been studied in the literature. Weron (2006) surveys different load forecasting methods and divides them in two classes: artificial intelligence models (neural networks, fuzzy logic, support vector machines) and statistical models (regression models, exponential smoothing, Box-Jenkins type time series models). Load forecasting methods are split into three different segments: short-term load forecasting (STLF), medium-term load forecasting (MTLF) and long-term load forecasting (LTLF). STLF is interested in hourly forecasts up to one week ahead, MTLF considers forecasts from one week to one year ahead and LTLF considers even longer horizons. The vast majority of the load forecasting literature considers STLF (Hahn et

[10] The retailer is assumed to consider information $\mathcal{G}$ relative to past and contemporaneous load-basis, spot prices and futures prices: $\mathcal{G}_t = \sigma\{\mathcal{L}_u, S_u, F_{u,u+j} | 0 \le u \le t, j = 1, 2, 3\}$.

[11] In the current paper, this consists of all $\mathcal{G}-$predictable trading strategies, meaning that θ_{t+1} is $\mathcal{G}_t-$measurable for all t.

[12] Ni et al. (2012) add a linear term to the quadratic penalty which makes it asymmetric.

[13] This penalty is also considered in a hedging problem by François et al. (2012).

[14] If the internal load data of the retailer do not support this assumption, the load model should be reworked. For differences between the load consumption patterns across the four countries that are part of the Nord Pool market, see Huovila (2003).

al. 2009), but MTLF has attracted more attention recently. Gonzalez et al. (2008) use a combination of neural networks and Fourier series to represent respectively the trend and the cyclical fluctuation of the monthly load in the Spanish market. In their paper, Fourier series outperform neural networks in their predictive ability for the cyclical load fluctuations. Abdel-Aal (2008) compares the use of neural and abductive networks to forecast the monthly load supplied by a power utility based in Seattle. Abdel-Aal and Al-Garni (1997) compare the use of univariate ARIMA process, abductive networks and multivariate regression models incorporating demographic, economic and weather related covariates to forecast the monthly domestic energy consumption in the Eastern province of Saudi Arabia. ARIMA processes outperformed their competitors in their study. Barakat and Al-Qasem (1998) propose a regression model with time and temperature as covariates to forecast the weekly load on the Riyadh system (Saudi Arabia).

To the authors' best knowledge, no MTLF has been attempted in the literature for the weekly load in Nord Pool. A parametric statistical model for load dynamics on Nord Pool, that supports our hedging methodology, is now presented.

3.1.1 Load-basis data

Time series of hourly load (in MegaWatt-hours, MWh) and hourly day-ahead spot price (in Euros, €) on Nord Pool for the January 1st, 2007 to July 29th, 2012 period are obtained through the Nord Pool FTP server.[15] The hourly load is aggregated as shown in (1) and yields 291 weekly load observations. The resulting load series L and basis ratio series η defined by Equation (5) are then combined to obtain the load-basis series $\mathcal{L}$ in (6).

The most salient feature of the load time series is a seasonal pattern, both in the mean and in the variance. Autocorrelation between consecutive load departures from its trend is also present. The model chosen for the dynamics of the load-basis $\mathcal{L}$ (observed in Figure 1) is thus:

$$\mathcal{L}_t - g(t) = \gamma(\mathcal{L}_{t-1} - g(t-1)) + \sqrt{v(t)}\epsilon_t^{(\mathcal{L})} \tag{10}$$

$$g(t) = \beta_0 + \sum_{j=1}^{P} \beta_j C_t^{(\sin,j)} + \sum_{j=1}^{P} \beta_{j+P} C_t^{(\cos,j)} \tag{11}$$

$$\log v(t) = \alpha_0 + \sum_{j=1}^{Q} \alpha_j C_t^{(\sin,j)} + \sum_{j=1}^{Q} \alpha_{j+Q} C_t^{(\cos,j)} \tag{12}$$

where $\epsilon^{(\mathcal{L})}$ is a strong standardized Gaussian white noise. The g function represents the seasonal trend of the load-basis level and its fitted value is represented by the dashed line in Figure 1. The v function characterizes the trend in the variance of seasonally corrected load-basis observations and the square root of its fitted value is represented by the dashed line in Figure 2. Terms of a Fourier expansion

$$C_t^{(\sin,j)} = \sin\left(\frac{3\pi}{2} + \frac{2\pi jt}{365.25/7}\right), \quad C_t^{(\cos,j)} = \cos\left(\frac{3\pi}{2} + \frac{2\pi jt}{365.25/7}\right)$$

are used to capture yearly cycles (see also Gonzalez et al. (2008)). The γ parameter in Equation (10) represents the autocorrelation in seasonally corrected load-basis observations. To preserve the Markov property, only one lag is considered.[16] Parameters to be estimated are $\gamma, \beta_0, \ldots, \beta_{2P}, \alpha_0, \ldots, \alpha_{2Q}$.

3.1.2 Estimation of model parameters

The model estimation is performed in two steps.[17] The first consists in estimating γ and $\beta_0, \ldots, \beta_{2P}$ by quasi-maximum likelihood under the assumption that $v(t)$ is constant. The optimal number $P = 3$ of Fourier terms in the mean trend is chosen using the cross-validation procedure described in Appendix D.1. Table 3

[15] Nord Pool uses the expression "turnover" to designate the load.
[16] Otherwise each additional lag would have to be included as a state variable.
[17] Results in Appendix D.2 show that the fitted model is good and a more numerically challenging single-step estimation was thus not attempted.

gives estimated parameters and their standard errors for this step. Figure 1 shows the load series L, the load-basis series $\mathcal{L}$ and the estimated load-basis seasonality trend g. Even if the variance v is presumed constant during the estimation of the trend parameters,[18] the overall trend seems reasonably captured. The corrected load is much larger in winter than in summer; this is expected given the winter heating requirements for Scandinavian countries. The autocorrelation parameter γ is estimated at 0.68, this large value indicating a high persistence in load deviations from the trend.

Table 3: Load-basis seasonality trend parameters

Parameter	γ	$\beta_0 \times 10^{-6}$	$\beta_1 \times 10^{-6}$	$\beta_2 \times 10^{-6}$	$\beta_3 \times 10^{-6}$	$\beta_4 \times 10^{-6}$	$\beta_5 \times 10^{-6}$	$\beta_6 \times 10^{-6}$
Estimated Value	0.68	5.70	−1.13	−0.15	0.06	0.24	0.10	0.11
Standard Error	0.04	0.04	0.06	0.05	0.04	0.06	0.05	0.04

Notes. Estimated parameters and their standard error for the load-basis seasonality trend g defined by Equation (11). Observations between January 1st, 2007 and July 29th, 2012. Estimated parameter variance is obtained through the inverse of the observed Fisher information matrix.

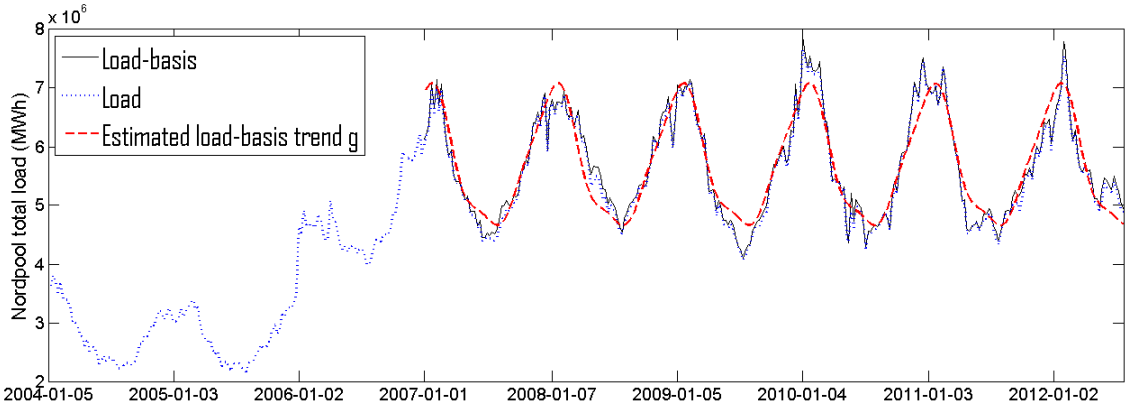

Figure 1: Load-basis seasonality trend curves

Notes. Observed total weekly load on the Nord Pool market as defined by (1), corresponding load-basis $\mathcal{L}$ and fitted seasonality trend $g(t)$. Observations between January 1st, 2007 and July 29th, 2012. Load data before 2007 are also included to show the shift in the overall system load level and justify the use of data starting from January 2007.

Once the trend parameters are estimated, proxy values for $\sqrt{v(t)}\epsilon_t^{(\mathcal{L})}$, denoted $\sqrt{\hat{v}(t)}\hat{\epsilon}_t^{(\mathcal{L})}$, can be computed using Equation (10). Those proxies serve as input in the second step which consists in estimating $\alpha_0, \ldots, \alpha_{2Q}$ by maximum likelihood (ML). The optimal number $Q = 2$ of Fourier terms in the variance trend is selected through the cross-validation procedure described in Appendix D.1. Table 4 presents the estimated parameters for the variance model (12). Figure 2 shows the estimated standard deviation trend $\sqrt{\hat{v}_t}$ (dashed curve) and the absolute value of the $\sqrt{\hat{v}(t)}\hat{\epsilon}_t^{(L)}$ proxies (full curve). The peak in volatility occurs in the beginning of winter, while the lowest volatility is observed during the end of the summer. Goodness-of-fit tests that confirm the adequacy of the load model are found in Appendix D.2.

Table 4: Load-basis variance trend

Parameter	α_0	α_1	α_2	α_3	α_4
Estimated Value	24.40	−0.72	−0.38	0.49	−0.32
Standard Error	0.08	0.11	0.11	0.11	0.12

Notes. Estimated parameters and their standard error for the load-basis variance trend v defined by Equation (12). Observations between January 1st, 2007 and July 29th, 2012. Estimated parameter variance is obtained through the inverse of the observed Fisher information matrix.

[18] At this step, the constant estimated volatility is $\sqrt{\hat{v}} = 2.264 \times 10^5$.

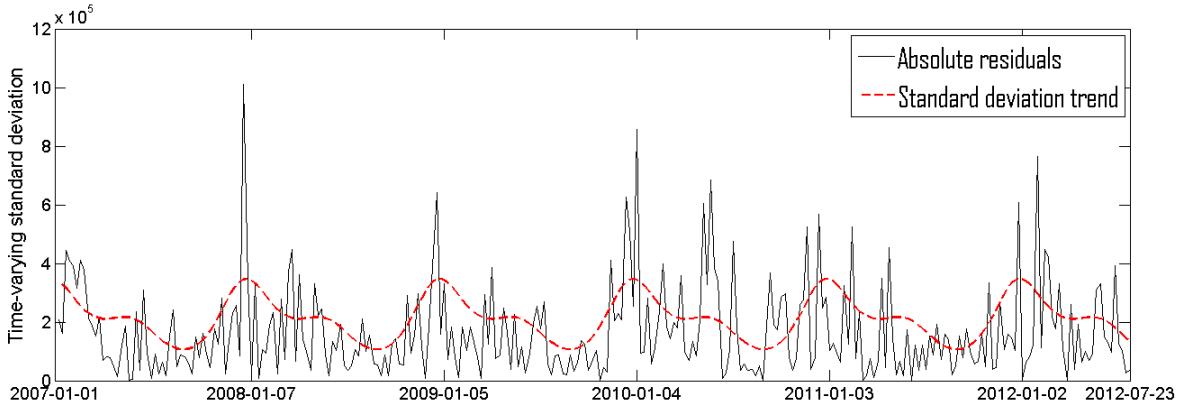

Figure 2: Load-basis standard deviation trend curves

Notes. Realized absolute load-basis volatility $\sqrt{\hat{v}(t)}|\hat{\epsilon}_t^{(\mathcal{L})}|$ and fitted standard deviation trend $\sqrt{\hat{v}}$ as defined by Equation (12). Data between January 1st, 2007 and July 29th, 2012.

3.1.3 Load-basis forecasting from incomplete information

The selection of θ_{t+1}, the number of futures shares detained in the portfolio from the Friday of week t until the Friday of week $t+1$, is based on $\mathcal{L}_t$, the weekly load-basis on week t. However $\mathcal{L}_t$ is only observed at midnight on Sunday of week t, and not at the closing time of markets on Friday. What is observed at the latter time is the sum of hourly loads from the beginning of week t to 4:00 p.m. on Friday:

$$\tilde{L}_t = \sum_{d=1}^{4}\sum_{h=1}^{24} L_{t,d,h} + \sum_{h=1}^{16} L_{t,5,h}.$$

When θ_{t+1} is selected, the value of $\mathcal{L}_t$ must thus be forecast using $\tilde{L}_t$. The accuracy of several forecasting models were compared through a cross-validation test and the model

$$\hat{\mathcal{L}}_t = \tilde{L}_t \times \left(\ell + c\sum_{j=1}^{q}\frac{\mathcal{L}_{t-j}}{\tilde{L}_{t-j}}\right) \tag{13}$$

produced the lowest out-of-sample forecasting RMSE. The out-of-sample mean absolute percentage error (MAPE) is 1.12%.[19] The parameter ℓ drives the long-term average of the ratio $\mathcal{L}_t/\tilde{L}_t$, while the autoregressive coefficient c characterizes the dependence of the current ratio on previous ratios. The estimated parameters obtained when re-estimating with the full dataset are $\hat{\ell} = 0.716$, $\hat{c} = 0.171$ and $\hat{q} = 3$. The long-term average of the $\mathcal{L}_t/\tilde{L}_t$ ratio is given by $\hat{\ell}/(1-\hat{q}\hat{c}) = 1.47$.[20]

3.2 Futures and spot price

In this section, time series of futures prices are modeled. Modeling the relation between the spot and futures prices in the context of electricity markets is complicated by the fact that electricity is not storable. This prevents the use of the usual cash-and-carry scheme to price futures contracts.

Solving problem (8) requires a model that completely specifies the stochastic dynamics of futures prices. An important strand of the literature studies the risk premium on electricity futures contracts.[21] Although

[19] The out-of-sample MAPE obtained by using the naive benchmark $\hat{\mathcal{L}}_t := \ell\tilde{L}_t$ is 1.21%. Obtaining good load forecasts is crucial to the success of the hedging procedure and the small improvement of model (13) over the naive method justifies its use.

[20] This is consistent with what is expected; since $\mathcal{L}$ and $\tilde{L}$ are respectively approximately the sum of 168 and 112 hourly loads, the long-term average of the ratio should revolve around $168/112 = 1.5$.

[21] For example, Lucia and Torro (2011) study the behavior of the risk premium on Nord Pool weekly futures with an ex-post econometric model taking into account hydropower reservoir levels.

these papers provide relevant information concerning the relation between the spot and futures prices, their models do not directly fit our needs since they do not characterize the futures prices dynamics. Several approaches are however proposed in the literature for this purpose and they will now be discussed.

Despite the fact that electricity is not storable and not openly traded, some authors follow the risk-neutral approach commonly used in finance. The dynamics of the spot price are modeled and a martingale measure is selected to compute futures prices as an expectation of the discounted cash flows. Benth et al. (2008) use a linear combination of non-Gaussian Ornstein-Uhlenbeck processes to represent the stochastic variability of the spot. They use the Esscher transform to compute futures prices. Coulon et al. (2012) propose a structural factor model encompassing natural gas price and electricity load to characterize spot price dynamics on the ERCOT electricity market. They use the Girsanov transform to compute derivatives prices.

Besides the non-storability of electricity, there is another potential pitfall with the risk-neutral approach to price futures. On the Nord Pool market, a principal component analysis applied to weekly futures returns shows that the spot price might be driven by factors different than those driving futures prices (Benth et al. 2008). The martingale measure approach discounting the expected spot price to obtain the futures price might thus be inappropriate. This result is consistent with the study of Koekebakker and Ollmar (2005) who uses principal component analysis to propose a multi-factor model for forward returns. They find that the number of factors necessary to represent the full forward curve is much larger for electricity futures on the Nord Pool market than for other commodities; the correlation between short-term and long-term electricity forward prices is smaller than in other markets.

Benth et al. (2008) also suggest adapting the Heath-Jarrow-Morton framework to electricity markets. Under such a methodology, the dynamics of forward prices that deliver an infinitesimal volume of electricity are directly specified. However, futures prices, which are really swap prices in the context of electricity markets, suffer from severe intractability issues under this model and we did not retain this approach.

The third method proposed in Benth et al. (2008) is to find a statistical model that reproduces the dynamics of the observed futures returns. This approach is followed in the current paper since it better suits our need to fully specify the distribution and the stochastic dynamics of futures and spot prices of electricity. Furthermore, this approach reproduces stylized facts.

Daily prices of futures on NASDAQ OMX are provided by Bloomberg. Since futures prices vary during the day, closing prices are used.

3.2.1 Our model

For a market participant hedging the cost of electricity at maturity week T, the sequence of observed futures prices that must be modeled is $\{F_{T-j,T}|j = 3, 2, 1, 0\}$. We propose a multivariate time series model for the joint dynamics of the spot and futures prices. As with financial assets, futures price returns are modeled (instead of the futures prices) as they are more likely to be stationary. Futures returns defined by

$$\epsilon_{t,T} = \log(F_{t,T}/F_{t-1,T}) \tag{14}$$

are shown in Figure 3 for $t = T, T-1, T-2$.

Futures price returns exhibit autocorrelation, volatility clustering and fat tails. These features suggest a multivariate AR-GARCH process with innovations drawn from a fat-tail distribution. For the latter, we choose a Normal Inverse Gaussian (NIG) distribution. More specifically, for $i = 0, 1, 2$, the trivariate AR(1)-GARCH(1,1) with NIG innovations is

$$\epsilon_{t,t+i} = \mu_i + a_i\,\epsilon_{t-1,t-1+i} + \sigma_{i,t} z_{i,t} \tag{15}$$

$$\sigma^2_{i,t+1} = \min\{\varsigma^2, \kappa_i + \gamma_i \sigma^2_{i,t} + \xi_i \sigma^2_{i,t} z^2_{i,t}\} \tag{16}$$

where $\mathbf{z}_t = (z_{0,t}, z_{1,t}, z_{2,t})$ has the following properties:

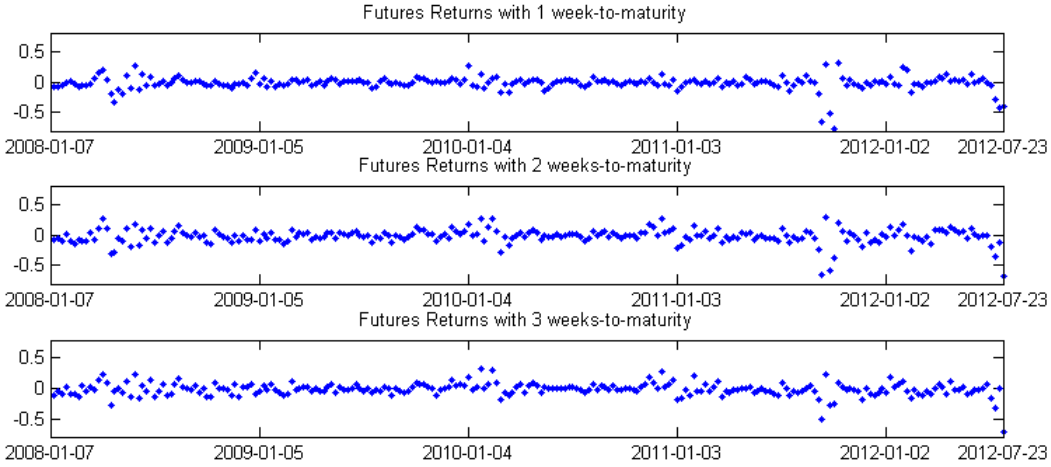

Figure 3: Futures price returns for different times-to-maturity

Notes. Time series for NASDAQ OMX electricity weekly futures returns (on Nord Pool day-ahead spot price) as defined by (14) between January 1st, 2007 and July 29th, 2012. The trivariate time series illustrated contains 290 observations.

if $s \neq t$, $\mathbf{z}_t$ and $\mathbf{z}_s$ are independent;
$z_{i,t}$ are drawn from a standardized[22] NIG(α_i, β_i) distribution;
$z_{0,t}$, $z_{1,t}$ and $z_{2,t}$ are linked by the Gaussian copula.

A bound ζ is used on the volatility to ensure that futures prices are square-integrable.[23] The a_i parameter represents autocorrelation of futures returns while μ_i adjusts their long-term expected level. The κ_i parameter adjusts the long-term level of futures return volatility, the γ_i characterizes the persistence in returns volatility, and ξ_i determines how shocks associated with current returns affect the current volatility. The NIG parameter α_i drives the tail thickness in the distribution of the futures return while the β_i drives its asymmetry.

3.2.2 Model estimation

A two-step procedure is applied. First, the parameters for the three marginal AR(1)-GARCH(1,1) processes ($\epsilon_{t,t+1}, \epsilon_{t,t+2}$ and $\epsilon_{t,t+3}$) are estimated by ML.[24] Plugging the estimated parameters in (15)–(16) yields proxy values $\hat{\mathbf{z}}_t$ for $\mathbf{z}_t$. Then, the proxies are used to estimate the parameters of the Gaussian copula. Letting F_{NIG} denote the cdf associated with the NIG distribution and applying the Rosenblatt (1952) transform to the proxy $\hat{\mathbf{z}}_t$ yields a series of approximatively independent uniformly distributed observations $\mathbf{U}_t = \left(F_{\text{NIG}(\hat{\alpha}_0,\hat{\beta}_0)}(\widehat{z_{0,t}}), F_{\text{NIG}(\hat{\alpha}_1,\hat{\beta}_1)}(\widehat{z_{1,t}}), F_{\text{NIG}(\hat{\alpha}_2,\hat{\beta}_2)}(\widehat{z_{2,t}})\right)$ drawn from the Gaussian copula. ML is used and the closed-form solution is $\hat{\rho}_{i,j} = corr(\Phi^{(-1)}(U_{i,t}), \Phi^{(-1)}(U_{j,t}))$, where $corr$ is the sample correlation and $\Phi^{(-1)}$ is the inverse cdf of a standard Gaussian variable.

Parameter estimates are shown in Tables 5 and 6. The negative mean parameters μ_i indicate the futures market is in contango. The GARCH parameters γ_i and ξ_i are highly significant, confirming the presence of volatility clustering in futures returns. The autocorrelation parameters a_i are also all significant and positive, indicating that futures returns are partially predictable. The α_i parameters are all low (smaller than 2) so

[22] A standardized NIG is a NIG distribution with mean 0 and variance 1. Such a distribution only has two free parameters: α and β. Note that these α and β should not be confused with those used in the load-basis model in Section 3.1.1.

[23] More precisely, the condition

$$\sigma_{i,t} < \frac{\alpha_i - \beta_i}{2} \text{ a.s.} \tag{17}$$

is necessary and sufficient to obtain $\mathbb{E}\left[e^{2\epsilon_{t,t+i}}\right] < \infty$. Thus, the volatility bound ς combined with the additional constraints $\alpha_i > \varsigma$, and $\beta_i \in (-\alpha_i, \alpha_i - 2\varsigma]$ assure (17) is satisfied.

[24] The proxy for the initial value for the volatility $\hat{\sigma}_{i,0}$ is its long-term stationary average. The bound is set at $\varsigma = 0.6$ since such a constraint is not numerically binding with the available data.

Table 5: Futures return parameters

Parameter	$i = 1$	$i = 2$	$i = 3$
$\mu_i \times 10^2$	−0.722 (0.001)	−1.566 (0.003)	−1.265 (0.002)
a_i	0.215 (0.005)	0.143 (0.005)	0.073 (0.004)
$\kappa_i \times 10^2$	0.177 (3×10^{-5})	0.124 (3×10^{-5})	0.106 (2×10^{-5})
γ_i	0.282 (0.018)	0.577 (0.008)	0.603 (0.010)
ξ_i	0.500 (0.027)	0.373 (0.009)	0.340 (0.010)
α_i	1.097 (0.009)	1.270 (0.089)	1.236 (0.022)
β_i	−0.108 (0.001)	−0.056 (0.099)	0.009 (0.006)

Notes. Estimated parameters (standard error) for futures returns model defined in (15)–(16). Observations between January 1st, 2007 and July 29th, 2012 for futures with $i = 1$, 2 and 3 weeks to maturity.

Table 6: Futures return copula parameters

Parameter	$\rho_{0,1}$	$\rho_{0,2}$	$\rho_{1,2}$
Estimate (Standard Error)	0.76 (0.03)	0.67 (0.04)	0.88 (0.01)

Notes. Estimated parameters (standard errors) for the Gaussian copula linking futures returns. $\rho_{i,j}$ links returns on futures with respectively $i + 1$ and $j + 1$ weeks to maturity. Observations between January 1st, 2007 and July 29th, 2012.

the kurtosis of futures returns is more pronounced than in a Gaussian distribution (which corresponds to an infinite α). The correlation parameters of the Gaussian copula are all higher than 0.65, indicating a somewhat high correlation between futures returns across the time-to-maturity dimension.

Goodness-of-fit tests that confirm the adequacy of the futures return model are found in Appendix D.3. Futures returns and load-basis innovations are assumed to be independent. Statistical tests in Appendix D.4 validate this assumption.

4 Performance assessment

We carry out numerical experiments to assess the performance of the hedging strategy given by solutions of problem (8). We propose two different hedging procedures: (i) the hedging methodology which solves problem (8) with $G(x) = x^2$ is referred to as quadratic dynamic global hedging (QDGH); (ii) the methodology solving that same problem but without penalizing the gains, i.e. using (9), is called semi-quadratic dynamic global hedging (SQDGH). The benchmarks are described in Section 4.1 while the backtests are explained in Section 4.2.

4.1 Benchmarks

4.1.1 Delta Hedging

If the load to be served by the retailer is known with certainty and no transaction fees exist, the delta hedging strategy proposed by Eydeland and Wolyniec (2003) completely eliminates the price risk borne by the retailer by locking in the spot price to $F_{t_0,T}$ (see Appendix B). This strategy is adapted to the case of a stochastic load by hedging the expected load-basis, i.e. the retailer enters into

$$\theta_{t+1} = \frac{B_{t+1}}{B_T}\mathbb{E}\left[\mathcal{L}_T|\mathcal{G}_t\right] \tag{18}$$

long positions in the futures contract at time t to cover its exposure at time T. Improved delta hedging (IDH) uses the load-basis model (10)–(12) to compute $\mathbb{E}\left[\mathcal{L}_T|\mathcal{G}_t\right]$ in (18).

To quantify the impact of using the (10)–(12) load-basis model in the hedging algorithm, alternative load-basis models are also proposed to compute $\mathbb{E}\left[\mathcal{L}_T|\mathcal{G}_t\right]$. For example, one may state that a good prediction of

the expected load-basis in a near future is the last observed load-basis. This points to the first alternative, the naive delta hedging (NDH), which uses the naive prediction model

$$\mathbb{E}[\mathcal{L}_{t+1}^{(NDH)}|\mathcal{G}_t] = \mathcal{L}_t^{(NDH)}.$$

The second alternative, referred to as delta hedging (DH), uses a load-basis model inspired from Wagner et al. (2003) where the latent variable found in their model is removed for simplicity. Their model specifies the load dynamics, but is applied here to the load-basis. More specifically, the load-basis model for DH is

$$\mathcal{L}_{t+1}^{(DH)} = \mathcal{L}_t^{(DH)} + \gamma^{(DH)}(\bar{\mathcal{L}}_{m_{t+1}} - \mathcal{L}_t^{(DH)}) + \mathcal{E}_{t+1}$$

where $\mathcal{E}$ is a Gaussian white noise, $\bar{\mathcal{L}}_m$ is the mean value of the load-basis during the m^{th} month of the year ($m = 1, \ldots, 12$) in the estimation set, m_{t+1} is the month associated with the week $t+1$ and $\gamma^{(DH)}$ is estimated by ML. We find $\hat{\gamma}^{(DH)} = 0.3477$.

4.1.2 Local Minimal Variance Hedging (LMVH)

The objective of this strategy, which is based on the Ederington (1979) scheme, is to construct a portfolio of futures whose variation mimics the variation of the spot price as closely as possible for the current period. More precisely, for each unit of load to serve, the retailer would detain ϑ_{t+1} units of futures at time t, where ϑ_{t+1} minimizes $\mathbb{V}ar\left[(S_{t+1} - S_t) - \vartheta_{t+1}(F_{t+1,T} - F_{t,T})|\mathcal{G}_t\right]$. This yields the solution $\vartheta_{t+1} = \mathbb{C}ov\left[S_{t+1}, F_{t+1,T}|\mathcal{G}_t\right]/\mathbb{V}ar\left[F_{t+1,T}|\mathcal{G}_t\right]$. To adapt this scheme to the case of stochastic load, the retailer hedges its expected load-basis by detaining at time t,

$$\theta_{t+1}^{(LMVH)} = \mathbb{E}\left[\mathcal{L}_T|\mathcal{G}_t\right] \frac{\mathbb{C}ov\left[S_{t+1}, F_{t+1,T}|\mathcal{G}_t\right]}{\mathbb{V}ar\left[F_{t+1,T}|\mathcal{G}_t\right]} \tag{19}$$

long positions in the futures contract to cover its exposure at time T. Many different models are used in the literature to compute the conditional variance and covariance in (19). We compute these quantities with the futures model (15) for consistency and refer to the approach as local minimal variance hedging (LMVH).

4.1.3 Static Hedging

Since many papers are devoted to static hedging procedures, we include them in our study. To apply static hedging (SH), the retailer identifies the solution to problem (8) under the constraint $\theta_{t_0+1} = \ldots = \theta_T$. We use the semi-quadratic penalty (9) and identify the optimal trading strategy through simulation.

4.2 Backtests

In all tests, the initial value of the portfolio V_{t_0} is set to 0 and the annualized continuously compounded risk free rate is $r = 0.0193$.[25] The case of a retailer serving 1% of the Nord Pool load is considered.

4.2.1 In-sample backtest

In this experiment, our global hedging and the benchmarks are applied to historical data during the 287 weeks over the January 29th, 2007 to July 23th, 2012 period. Hedging errors $\Psi_T - V_T$ are recorded at the end of week T and the performance of the various approaches are compared through the following metrics:

$$\text{RMSE} = \sqrt{\frac{1}{287}\sum_{T=1}^{287}(\Psi_T - V_T)^2}, \tag{20}$$

$$\text{Semi-RMSE} = \sqrt{\frac{1}{287}\sum_{T=1}^{287}\left((\Psi_T - V_T)\mathbb{I}_{\{\Psi_T > V_T\}}\right)^2}, \tag{21}$$

[25] The average overnight EURO LIBOR rate between January 1st, 2007 and July 29th, 2012.

$$\text{TVaR}_\alpha \quad = \quad \frac{\sum_{T=1}^{287}(\Psi_T - V_T)\mathbb{I}_{\{\Psi_T - V_T \geq q_{(1-\alpha)}\}}}{\sum_{T=1}^{287}\mathbb{I}_{\{\Psi_T - V_T \geq q_{(1-\alpha)}\}}}.$$

where $\text{VaR}_\alpha = q_{(1-\alpha)}$ is the quantile of level $1-\alpha$ of hedging errors $\Psi_T - V_T$. Results are reported in Table 7.

Table 7: In-sample backtest results

Model	SQDGH	QDGH	IDH	DH	NDH	STAH	LMVH	NOH
Mean	6.107	7.641	7.696	8.176	7.069	8.240	−18.95	−70.73
RMSE	26.20	27.65	27.36	31.37	35.03	28.05	274.8	474.1
Semi-RMSE	23.68	26.30	26.10	29.75	32.15	26.99	208.4	319.5
$\text{VaR}_{5\%}$	48.97	52.41	54.53	60.56	73.85	52.93	414.0	688.7
$\text{VaR}_{1\%}$	115.4	130.3	129.2	142.7	153.5	146.9	1206	1471
$\text{TVaR}_{5\%}$	89.97	100.3	99.31	115.3	121.0	103.5	784.5	1201
$\text{TVaR}_{1\%}$	133.1	163.1	161.9	172.9	180.9	159.2	1273	1831

Notes. Hedging error risk metrics for the in-sample backtest (in 1000€). Semi-quadratic dynamic global hedging (SQDGH), quadratic dynamic global hedging (QDGH), Delta Hedging (DH), Improved Delta Hedging (IDH), Naive Delta Hedging (NDH), Static Hedging (STAH), Local Minimal Variance Hedging (LMVH) and No hedging (NOH), i.e. $\theta_t^{(NOH)} = 0$ for all t.

The main result is that the semi-quadratic SQDGH outperforms all other methods in terms of risk reduction; it reduces the semi-RMSE, the $\text{TVaR}_{5\%}$ and the $\text{TVaR}_{1\%}$ by 2,420€, 9,340€ and 28,800€, respectively (i.e. by 9.3%, 9.4% and 17.8% in relative measurement), with respect to IDH, the best benchmark. Those improvements can be attributed to using global hedging procedures instead of delta hedging since both approaches share the same load model. To put these numbers in context, the mean weekly procurement costs of electricity (the average of $\mathcal{L}_T S_T$ for the January 2007 to August 2012 period) for the considered retailer is $2.35M$ €. Von der Fehr and Hansen (2010) identify a retail price mark-up ranging between 7.2% and 13% over the wholesale price for fixed-price contracts in Norway. Using a 10% mark-up for ballpark calculations, this leaves the retailer with an average weekly margin of 235,000€ to cover expenses and profit; average profits will be a fraction of that amount. SQDGH reduces the 1% worst-scenarios average loss with respect to IDH by 28,800€, a substantial fraction of average profits.

Note that IDH benefits from our load-basis model (10)–(12). The added value of the latter model is isolated by comparing IDH with DH and NDH. The $\text{TVaR}_{1\%}$ is reduced from 180,900€ for the NDH to 172,900€ for the DH, and further reduced to 161,900€ for the IDH. This illustrates the importance of having an accurate load-basis model and the benefits provided by the model (10)–(12) in terms of risk reduction.

The combined reduction in $\text{TVaR}_{1\%}$ due to methodology presented in this paper obtained by comparing SQDGH and NDH is 47,800€, with a combined reduction in semi-RMSE of 8,470€.

It is also interesting that the mean hedging error is lower for SQDGH than for all other models except LVMH and NOH. This indicates the risk reduction yielded by the SQDGH method is not obtained at the expense of a lesser profitability. The LMVH and NOH methods are the two most profitable on average, but they yield extremely poor results in terms of risk. The poor performance of the LMVH method is explained by positions in the futures that are significantly too low. Indeed, since the cash-and-carry relationship of futures price and the spot price does not hold in this market, the correlation between spot price and futures price variations are much lower than in other markets. This reduces the $\theta^{(LMVH)}$ position and produces under-hedging. Because the Nord Pool electricity futures market is in contango,[26] under-hedging produces higher average profits than full hedging.

In terms of semi-RMSE, STAH underperforms IDH, QDGH and SQDGH, showing the benefits of dynamic hedging over a static procedure.

[26] The average 3-weeks futures price is 7.7% higher than the arithmetic average spot price for the January 2007 to July 2012 period.

4.2.2 Out-of-sample backtest

In this experiment, parameters for the state variables models are estimated using only data for the 2007 to 2011 period (the in-sample). Hedging is then performed for exposure weeks in 2012 (the out-of-sample). This replicates the more realistic application conditions where future observations cannot be used to estimate state variable models. As our out-of-sample set only comprises 29 observations however, the experiment is mainly illustrative. Risk metrics (20)–(21) applied to hedging errors for the out-of-sample backtest are given Table 8. TVaRs are not given because of the low number of observations.

Table 8: Out-of-sample backtest results

Model	SQDGH	QDGH	IDH	DH	NDH	STAH	LMVH	NOH
Mean	13.64	15.99	15.99	18.30	25.66	16.77	31.75	17.70
RMSE	36.30	38.47	38.65	44.58	50.28	44.49	194.6	517.8
Semi-RMSE	35.18	38.27	38.41	44.09	50.09	43.56	166.9	431.4

Notes. Hedging error risk metrics for the out-of-sample backtest (in 1000€). Semi-quadratic dynamic global hedging (SQDGH), quadratic dynamic global hedging (QDGH), Delta Hedging (DH), Improved Delta Hedging (IDH), Naive Delta Hedging (NDH), Static Hedging (STAH), Local Minimal Variance Hedging (LMVH) and No hedging (NOH), i.e. $\theta_t^{(NOH)} = 0$ for all t.

Once again, SQDGH outperforms all the benchmarks, reducing the half-RMSE by 8.1% with respect to QDGH, its closest competitor. Furthermore, SQDGH proved to be the most profitable method with an average hedging error lower than every other method, even LMVH and NOH. The latter should produce lower hedging errors than the other benchmarks on average since they give under-hedging in a market in a contango situation. This did not materialize however given the higher volatility of their hedging errors and the small number of out-of-sample observations. The SQDGH is therefore the best hedging method among all the proposed methods in both the in-sample and out-of-sample backtests.

5 Conclusion

A dynamic global hedging methodology involving futures contracts is developed to allow retailers to cover their exposure to price and load risk. Global hedging procedures have received little or no attention in the electricity markets literature because they often yield solutions which are computationally more complex than their local counterparts. We show that the approach is not only feasible but easily allows us to account for load uncertainty, basis risk and transaction costs when seeking the optimal trading strategy.

Statistical models were proposed for the load to be served by the retailer, the electricity spot price and futures contract prices on the Nord Pool market. Those models were built from weekly historical data and reproduce their stylized facts. The load basis model accounts for seasonality in the mean and the variance, as well as autocorrelation in seasonally corrected shocks. The proposed model for futures price returns, a multivariate AR(1)-GARCH with NIG innovations, exhibits stochastic volatility, partially predictable returns and fat tails. Multiple goodness-of-fit tests validate the adequacy of all models developed.

Backtests using historical market data show the superiority of the semi-quadratic global hedging procedure compared to various benchmarks of the literature in terms of risk reduction.

A Basis ratio

The weekly average price S_t^* paid for electricity differs from the weekly arithmetic average price S_t, which is the underlying asset of weekly futures. The extent to which S_t^* and S_t differ is represented by basis ratio η_t in (5). Figure 4 shows the observed ratio over the January 1st, 2007 and July 23th, 2012 period.

As η_t is larger than one in all but one instance, S_t^* overestimates S_t. Such a departure has not yet been considered in the literature. This departure is due to the fact that more electricity is consumed during peak hours when its price is higher.

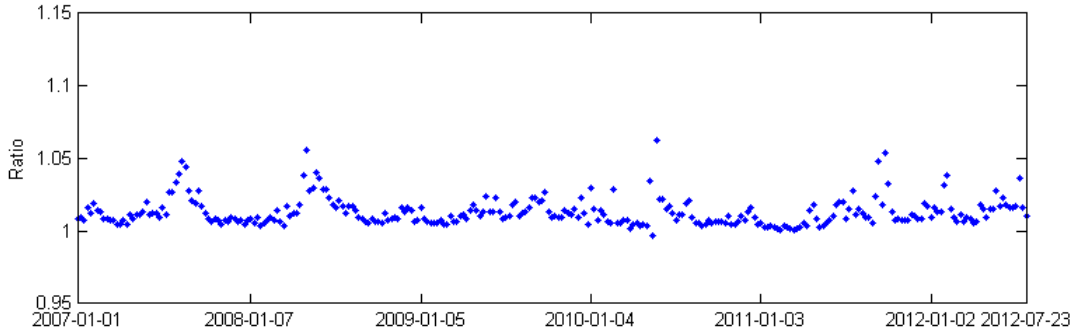

Figure 4: Basis ratio time series

Notes. Observed ratio of the load weighted mean spot price to the arithmetic mean spot price as defined by (5). Observations between January 1st, 2007 and July 23th, 2012.

B Delta-hedging with futures

If transaction costs are disregarded, the terminal value of the self-financing hedging portfolio with an initial value of 0 is given by $V_T = \sum_{j=t_0+1}^{T} \theta_j (B_T/B_j)(F_{j,T} - F_{j-1,T})$. Setting $\theta_j = B_j/B_T$, the terminal value of the portfolio becomes

$$V_T = \sum_{j=t_0+1}^{T} (F_{j,T} - F_{j-1,T}) = F_{T,T} - F_{T,t_0} = S_T - F_{T,t_0}.$$

Therefore, holding one unit of this portfolio for each unit of load sold (in the case where the load to serve is known with certainty) permits to lock in the price of electricity to F_{T,t_0}.

C Solving problem (8)

The optimal trading strategy $(\theta^*_{T-2}, \theta^*_{T-1}, \theta^*_T)$ solving problem (8) with the semi-quadratic penalty (9) is obtained through dynamic programming (Bertsekas 1995):

$$\psi_{t,T} = \min_{\theta_{t+1}} \mathbb{E}\left[\psi_{t+1,T} | \mathcal{G}_t\right] \quad \text{with the terminal condition } \psi_{T,T} = G(\Psi_T - V_T), \tag{22}$$

$$\theta^*_{t+1} = \arg\min_{\theta_{t+1}} \mathbb{E}\left[\psi_{t+1,T} | \mathcal{G}_t\right]. \tag{23}$$

This optimization problem is tackled using backward induction over time. The traditional approach used for solving (22) is based on a lattice which includes all state variables of the problem; these include the current value of the load-basis and futures prices, current futures return volatilities, the current hedging portfolio value, lagged futures returns and the past portfolio composition. Such an approach is not viable due to its large dimension. Our approach is a stochastic tree which is feasible because the hedging portfolio is only rebalanced three times. The optimization of the trading position θ_t is performed numerically by discretizing its possible values.

C.1 Simulation of the stochastic tree

Since the terminal condition $\psi_{T,T} = G(\Psi_T - V_T) = G(\mathcal{L}_T (F_{T,T} - F_{T-3,T}) - V_T)$ depends on the state variables (the load-basis $\mathcal{L}$ and the futures contracts related variables) and some endogenous variables (the portfolio value V_T and consequently the corresponding portfolio positions θ_{T-1}, θ_{T-2}, and θ_{T-3}), the random tree must account for all these dimensions.

At time $T-3$, M_{T-3} scenarios for the state variables are simulated from Equations (10)–(12) and (15)–(16).[27] These scenarios are combined with all the possible portfolio positions[28] $\theta_{T-2} \in \Theta_{T-2}$ to generate $N_{T-3} = M_{T-3}\text{Card}\{\Theta_{T-2}\}$ simulated values for endogenous variables (V_{T-2}, θ_{T-2}).

At time $T-2$, these N_{T-3} scenarios for the state and endogenous variables are subdivided into $N_{T-2} = M_{T-2}\text{Card}\{\Theta_{T-1}\}$ branches corresponding to all combinations of simulated state variables and possible portfolio positions. A similar iteration occurs at time $T-1$, leading to $N_{T-3} \times N_{T-2} \times N_{T-1}$ terminal nodes.

C.2 Backward induction

The algorithm solving (22) starts by computing the final hedging penalty at each terminal node[29] of the tree:

$$\begin{aligned}\hat{\psi}_T & \begin{pmatrix} m_{T-3}, m_{T-2}, m_{T-1} \\ \theta_{T-2}, \theta_{T-1}, \theta_T \end{pmatrix} \\ = \; & G\left(\mathcal{L}_T \begin{pmatrix} m_{T-3}, m_{T-2}, m_{T-1} \\ \theta_{T-2}, \theta_{T-1}, \theta_T \end{pmatrix} (F_{T,T}(m_{T-3}, m_{T-2}, m_{T-1}) - F_{T-3,T}) - V_T \begin{pmatrix} m_{T-3}, m_{T-2}, m_{T-1} \\ \theta_{T-2}, \theta_{T-1}, \theta_T \end{pmatrix}\right).\end{aligned}$$

Equations (22)–(23) are then approximated using the following backward recursion for each node of the tree:

$$\begin{aligned}\hat{\theta}^*_T \begin{pmatrix} m_{T-3}, m_{T-2} \\ \theta_{T-2}, \theta_{T-1} \end{pmatrix} &= \underset{\theta \in \Theta_T}{\arg\min} \frac{1}{M_{T-1}} \sum_{m=1}^{M_{T-1}} \hat{\psi}_T \begin{pmatrix} m_{T-3}, m_{T-2}, m \\ \theta_{T-2}, \theta_{T-1}, \theta \end{pmatrix}, \\ \hat{\psi}_{T-1} \begin{pmatrix} m_{T-3}, m_{T-2} \\ \theta_{T-2}, \theta_{T-1} \end{pmatrix} &= \min_{\theta \in \Theta_T} \frac{1}{M_{T-1}} \sum_{m=1}^{M_{T-1}} \hat{\psi}_T \begin{pmatrix} m_{T-3}, m_{T-2}, m \\ \theta_{T-2}, \theta_{T-1}, \theta \end{pmatrix}, \\ \hat{\theta}^*_{T-1} \begin{pmatrix} m_{T-3} \\ \theta_{T-2} \end{pmatrix} &= \underset{\theta \in \Theta_{T-1}}{\arg\min} \frac{1}{M_{T-2}} \sum_{m=1}^{M_{T-2}} \hat{\psi}_{T-1} \begin{pmatrix} m_{T-3}, m \\ \theta_{T-2}, \theta \end{pmatrix}, \\ \hat{\psi}_{T-2} \begin{pmatrix} m_{T-3} \\ \theta_{T-2} \end{pmatrix} &= \min_{\theta \in \Theta_{T-1}} \frac{1}{M_{T-2}} \sum_{m=1}^{M_{T-2}} \hat{\psi}_{T-1} \begin{pmatrix} m_{T-3}, m \\ \theta_{T-2}, \theta \end{pmatrix}, \\ \hat{\theta}^*_{T-2} &= \underset{\theta \in \Theta_{T-2}}{\arg\min} \frac{1}{M_{T-3}} \sum_{m=1}^{M_{T-3}} \hat{\psi}_{T-2} \binom{m}{\theta}, \\ \hat{\psi}_{T-3} &= \min_{\theta \in \Theta_{T-2}} \frac{1}{M_{T-3}} \sum_{m=1}^{M_{T-3}} \hat{\psi}_{T-2} \binom{m}{\theta}.\end{aligned}$$

In the experiments of Section 4, the number of scenarios are $M_{T-3} = M_{T-2} = 1000$ and $M_{T-1} = 100$. Fewer scenarios are required at the final step since the conditional expectations can partially be solved analytically. More precisely, Equations (22)–(23) involve double integrals (one over the load innovation and the other over the futures return innovation with a one-week maturity). Fortunately, the load innovation is Gaussian, so the first integral can be computed analytically. Therefore, instead of using a regular Monte-Carlo simulation for the futures innovation, a quadrature in a single dimension is applied.

The discrete sets of portfolio positions are $\Theta_{T-2} = \{0.96, 0.965, \ldots, 1.04\}$ and $\Theta_{T-1} = \Theta_T = \{0.93, 0.94, \ldots, 1.07\}$, implying that $\text{Card}\{\Theta_{T-2}\} = 17$ and $\text{Card}\{\Theta_{T-1}\} = \text{Card}\{\Theta_T\} = 15$.

Variance reduction techniques improve the precision of the Monte Carlo estimates and compensate for small sample sizes. Antithetic variables are used in the simulation for load-basis innovations $\epsilon^{(\mathcal{L})}$. The first half of scenarios are simulated by regular Monte-Carlo methods. In the last half of scenarios, the futures return innovations are identical to the ones in the first half. Load innovations are however set equal to their antithetic counterparts.

[27] Simulating a scenario at time t involves simulating the values of the load-basis and futures price innovations, respectively $\epsilon^{(\mathcal{L})}_{t+1}$ and $\epsilon_{t+1,t+j}$, $j = 1, 2, 3$.

[28] A discretize subset Θ_{T-2} of the possible positions is considered. $\text{Card}\{\Theta_{T-2}\}$ represents the number of elements it contains.

[29] The terminal nodes are identified with the set of indices corresponding to the branches constituting the path:

$$\begin{pmatrix} m_{T-3}, m_{T-2}, m_{T-1} \\ \theta_{T-2}, \theta_{T-1}, \theta_T \end{pmatrix}.$$

C.3 Re-simulation

The previous algorithm determines the optimal hedging strategy $\left(\theta^*_{T-2}, \theta^*_{T-1}, \theta^*_T\right)$ as seen from time $T-3$. At time $T-2$, the retailer holds θ^*_{T-2} long futures positions and has to select θ^*_{T-1} to perform the rebalancing. The realization of the state variables at time $T-2$ will not exactly fall on one particular node of the random tree. The standard approach used to solve this issue is to interpolate between the nodes of the tree to determine the optimal hedging position θ^*_{T-1}. We opted for a re-simulation to obtain simulated data which incorporates the newly observed realization of state variables. More precisely, a two-period random tree is simulated from time $T-2$ up to time T to update the optimal hedging strategy $\left(\theta^*_{T-1|T-2}, \theta^*_{T|T-2}\right)$. Since this tree is smaller than the previous one, we opted for a thinner discretization of the portfolio positions: $\Theta_{T-1} = \{0.93, 0.9325 \ldots, 1.07\}$, and $\Theta_T = \{0.93, 0.94 \ldots, 1.07\}$ while keeping $M_{T-2} = 1000$ and $M_{T-1} = 100$.

Finally, at time $T-1$, a one-period random tree with $\Theta_T = \{0.900, 0.901 \ldots, 1.100\}$ is simulated to update the final hedging position $\theta^*_{T|T-1}$.

D Load-basis model estimation

D.1 Cross-validation procedure for load model selection

To determine the number P of Fourier terms in step 1 of the load-basis model estimation (or Q in step 2), a cross-validation procedure is implemented. The load-basis data are from 2007 to 2012. Data from year y are removed and retained as out-of-sample, while remaining data are in-sample. For each value of P (or Q), the model is estimated in-sample. Denote $\mathcal{J}^y_{1,P} = (\gamma, \beta_0, \ldots \beta_{2P+1})$ and $\mathcal{J}^y_{2,Q} = (\alpha_0, \ldots \alpha_{2Q+1})$. f denotes the pdf function.

$$\begin{aligned}
\hat{\mathcal{J}}^y_{1,P} &= \underset{\mathcal{J}^y_{1,P}}{\operatorname{argmax}} \sum_{t, year(t) \neq y} \log f_{\mathcal{L}_t|\mathcal{L}_{t-1}}(\mathcal{L}_t|\mathcal{L}_{t-1}) \quad \text{(under assumption that } v(t) \text{ is constant)} \\
\hat{\mathcal{J}}^y_{2,Q} &= \underset{\mathcal{J}^y_{2,Q}}{\operatorname{argmax}} \sum_{t, year(t) \neq y} \log f_{v_t(\mathcal{J}^y_{2,Q})\epsilon_t}(\sqrt{\tilde{v}(t)}\tilde{\epsilon}^{(\mathcal{L})}_t)
\end{aligned}$$

where $\tilde{g}(t)$ and $\tilde{v}(t)$ are obtained by respectively plugging $\hat{\mathcal{J}}^y_{1,P}$ in (11) and $\mathcal{J}^y_{2,Q}$ in (12). The $\tilde{\epsilon}^{(\mathcal{L})}_t$ are calculated by replacing $g(t)$ and $v(t)$ by $\tilde{g}(t)$ and $\tilde{v}(t)$ in (10).

Then, a test statistic assessing the goodness-of-fit (MSE for P, log-likelihood for Q) is calculated out-of-sample:

$$\begin{aligned}
\text{MSE}^P_y &= \frac{1}{n_y} \sum_{t, year(t)=y} (\mathcal{L}_t - \text{Pred}(\mathcal{L}_t, \hat{\mathcal{J}}^y_{1,P}))^2 \\
\text{log-l}^Q_y &= \sum_{t, year(t)=y} \log f_{\sqrt{v_t(\hat{\mathcal{J}}^y_{2,Q})}\epsilon^{(\mathcal{L})}_t}(\sqrt{\hat{v}(t)}\hat{\epsilon}^{(\mathcal{L})}_t)
\end{aligned}$$

where n_y is the number of observations in year y. $\hat{g}(t)$ and $\hat{v}(t)$ are obtained by respectively plugging $\hat{\mathcal{J}}^y_{1,P}$ in (11) and $\hat{\mathcal{J}}^y_{2,Q}$ in (12). The $\hat{\epsilon}^{(\mathcal{L})}_t$ are calculated by replacing $g(t)$ and $v(t)$ by $\hat{g}(t)$ and $\hat{v}(t)$ in (10). The predicted load-basis is $\text{Pred}(\mathcal{L}_t, \hat{\mathcal{J}}^y_{1,P}) = \hat{g}(t) + \hat{\gamma}(\mathcal{L}_{t-1} - \hat{g}(t-1))$ where $\hat{g}$ is calculated by plugging $\hat{\mathcal{J}}^y_{1,P}$ in (11) and $\hat{\gamma}$ is the first component of $\hat{\mathcal{J}}^y_{1,P}$. The prediction is obtained by applying a conditional expectation on (10). This operation is repeated for all years y and the test statistic is aggregated across all years:

$$\text{RMSE}^P_{total} = \sqrt{\frac{\sum_{y=2007}^{2012} n_y \text{MSE}^P_y}{\sum_{2007}^{2012} n_y}} \quad \text{or} \quad \text{log-l}^Q_{total} = \sum_{y=2007}^{2012} \text{log-l}^Q_y.$$

Parameters $\hat{P}$ and $\hat{Q}$ are selected to optimize the corresponding test statistic

$$\hat{P} = \underset{P}{\operatorname{argmin}}\ \text{RMSE}^P_{total} \quad \text{and} \quad \hat{Q} = \underset{Q}{\operatorname{argmax}}\ \text{log-l}^Q_{total}.$$

Results are shown in Tables 9 and 10 and suggest $\hat{P} = 3$ and $\hat{Q} = 2$.

Table 9: Cross-validation test results for the load-basis seasonality trend

Value for P	1	2	3	4	5
Cross-validation RMSE ($\times 10^5$)	2.386	2.376	**2.360**	2.364	2.367

Notes. Out-of-sample cross-validation prediction root-mean-square-error for the load-basis model with different numbers of Fourier terms P in the load-basis seasonality trend g defined by (11).

Table 10: Cross-validation test results for the load-basis variance trend

Value for Q	1	2	3	4	5
Cross-validation log-likelihood ($\times 10^{-3}$)	−3.973	**−3.967**	−3.973	−3.974	−3.976

Notes. Out-of-sample cross-validation log-likelihood for the load-basis model with different numbers of Fourier terms Q in the load-basis variance trend v defined by (12).

D.2 Goodness-of-fit for the load model

In this section, the properties of the standardized residuals $\hat{\epsilon}_t^{(\mathcal{L})}$ are analyzed to determine the adequacy of the load-basis model (10)–(12). Figure 5 shows a boxplot of residuals by quarter of the year, a QQ-plot and a kernel density plot. Residuals look reasonably uniform across quarters; there is thus no obvious evidence that the seasonal trend is not properly being captured. The Gaussian distribution seems to be a suitable candidate for residuals, even if the empirical left tail of the load residuals is slightly heavier. A bootstrap Cramer-Von-Mises goodness-of-fit test for the adequacy of the Gaussian distribution is applied to the residuals and the p-value is 27%, not rejecting the Gaussian distribution. A Ljung-Box test for autocorrelation of residuals has a p-value of 92% and does not reject $\hat{\epsilon}_t^{(L)}$ as white noise. The presence of a GARCH effect in the residuals is tested through the McLeod-Li test (p-value of 18%) and Lagrange Multiplier test (p-value of 16%); there is no significant presence of a GARCH effect. Therefore, the $\epsilon^{(\mathcal{L})}$ load-basis innovations are modeled by a strong Gaussian white noise.

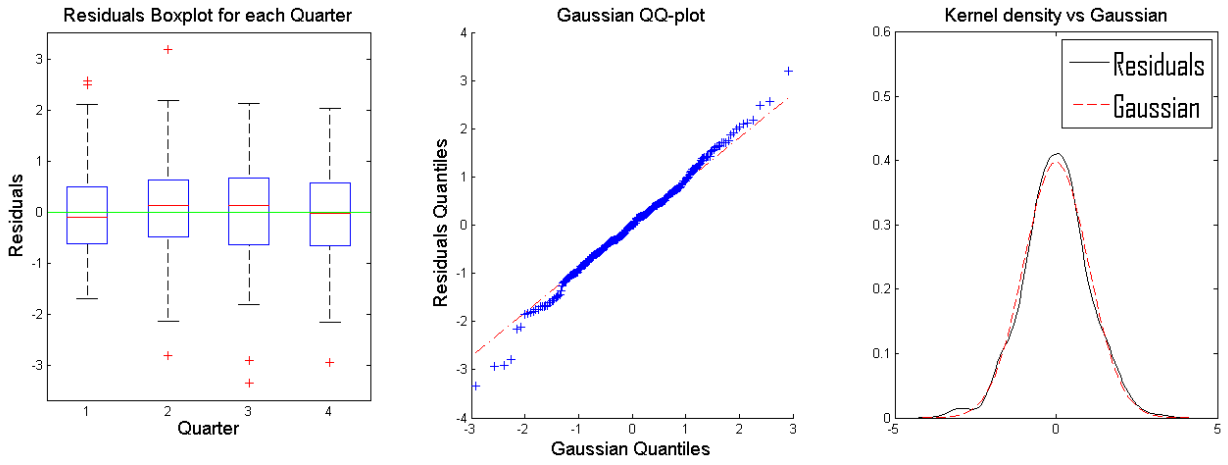

Figure 5: Load-basis model residuals

Notes. Boxplot, Gaussian QQ-plot and kernel plot for load-basis residuals $\hat{\epsilon}^{(\mathcal{L})}$ between January 1st, 2007 and July 29th, 2012.

D.3 Goodness-of-fit of futures return model

Ljung-Box and McLeod-Li tests for strong white noise are carried out on the scaled residuals $\widehat{z}_{j,t}$, $j = 0, 1, 2$. P-values are obtained through simulation (usual p-value formulas incorrectly assume Gaussianity). P-values are given in Table 11 and none of the tests reject the white noise hypothesis.

Table 11: Autocorrelation tests for futures return innovations

Series	$z_{0,t}$	$z_{1,t}$	$z_{2,t}$
Ljung-Box p-value	0.36	0.35	0.41
McLeod-Li p-value	0.97	0.33	0.72

Notes. Bootstrapped $p-$values for the Ljung-Box and McLeod-Li tests applied on futures return innovations. Observations between January 1st, 2007 and July 29th, 2012 for futures with 1, 2 and 3 weeks to maturity.

The choice of the NIG distribution for the innovations must be validated. Figure 6 compares the kernel density of the $\hat{z}_{i,t}$, its fitted NIG distribution and a corresponding Gaussian distribution. The NIG distribution represents more adequately the shape of the empirical residuals distribution than the Gaussian distribution, the latter is unable to capture the peakedness of the empirical futures returns distribution. Cramer-Von-Mises tests (with simulated p-value) are applied to assess the adequacy of the fit of the NIG distribution for the $z_{i,t}$ innovations. P-values are found in Table 12 for each univariate $z_{i,t}$, $i = 0, 1, 2$ series. The p-value for the joint trivariate series is 0.82. The NIG distribution thus provides an acceptable fit.

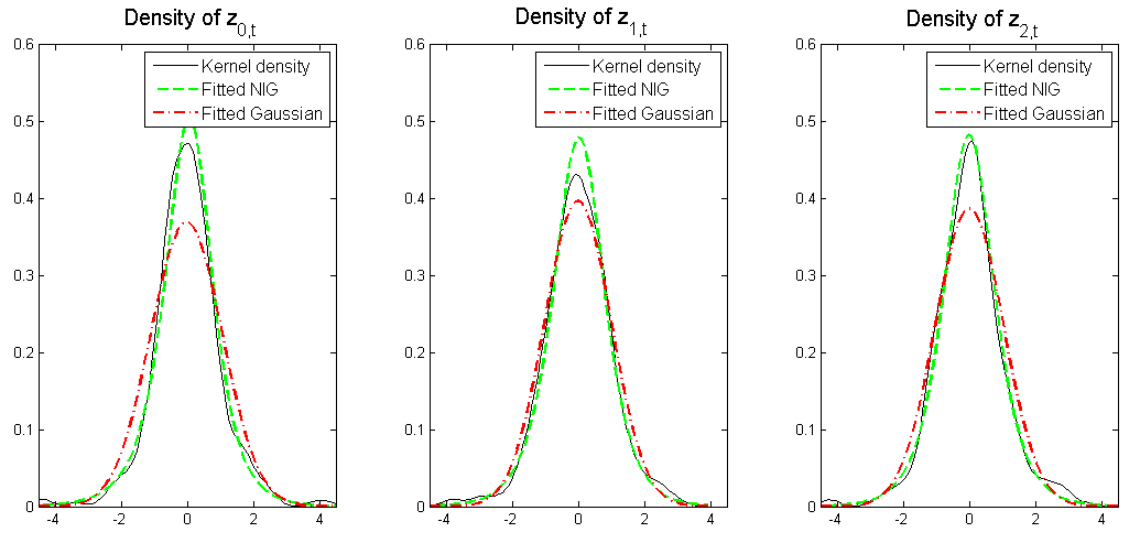

Figure 6: Futures return distribution

Notes. Kernel density plots of futures return innovations and fitted NIG and Gaussian distributions. Observations between January 1st, 2007 and July 29th, 2012 for futures with 1, 2 and 3 weeks to maturity.

Table 12: Goodness-of-fit of the futures return distribution

Series	$z_{0,t}$	$z_{1,t}$	$z_{2,t}$
p-value	0.09	0.88	0.65

Notes. Bootstrapped $p-$values for the Cramer-Von-Mises goodness-of-fit test on the NIG distribution for futures return. Observations between January 1st, 2007 and July 29th, 2012 for futures with 1, 2 and 3 weeks to maturity.

To validate the choice of the copula, Cramer-Von-Mises goodness-of-fit tests are applied for the Gaussian copula on the three following pairs of processes: $(z_{0,t}, z_{1,t})$, $(z_{0,t}, z_{2,t})$ and $(z_{1,t}, z_{2,t})$.[30] The p-values for the three tests are given in Table 13. Since p-values are all high, the Gaussian copula provides an acceptable fit.

Table 13: Goodness-of-fit of the futures return copula

Innovation Pair	$(z_{0,t}, z_{1,t})$	$(z_{0,t}, z_{2,t})$	$(z_{1,t}, z_{2,t})$
p-value	0.90	0.67	0.97

Notes. Bootstrapped $p-$values for the Cramer-Von-Mises goodness-of-fit test applied to the Gaussian copula linking futures returns. Tests are applied on pairs of returns instead of the triplet $(z_{0,t}, z_{1,t}, z_{2,t})$. Observations between January 1st, 2007 and July 29th, 2012 for futures with 1, 2 and 3 weeks to maturity.

D.4 Independence of futures return and load-basis innovations

Independence tests for load-basis residuals $\hat{\epsilon}_t^{(\mathcal{L})}$ and futures return innovation proxies $\hat{z}_{t,i}$ are applied for each of the three futures return maturities: $i = 0, 1, 2$. We use a Cramer-Von-Mises goodness-of-fit test on the independence copula. The p-values are obtained through simulation and are given in Table 14. Large p-values allow us to assume that the load-basis residuals and the futures return innovations are independent.

Table 14: Independence test for load-basis and futures return innovations

Futures Returns Series	$i = 0$	$i = 1$	$i = 2$
p-value	0.28	0.66	0.50

Notes. Bootstrapped $p-$values for the Cramer-Von-Mises goodness-of-fit test applied to the independence copula linking the load-basis observations and futures returns. Three tests are applied separately for the three maturities of futures returns. Observations between January 1st, 2007 and July 29th, 2012 for futures with 1, 2 and 3 weeks to maturity.

References

Abdel-Aal, R.E. 2008. Univariate modeling and forecasting of monthly energy demand time series using abductive and neural networks. Comput. & Industrial Engineering 54(4) 903–917.

Abdel-Aal, R.E., A.Z. Al-Garni. 1997. Forecasting monthly electric energy consumption in eastern Saudi Arabia using univariate time-series analysis. Energy 22(11) 1059–1069.

Barakat, E.H., J.M. Al-Qasem. 1998. Methodology for weekly load forecasting. IEEE Trans. Power Systems 13(4) 1548–1555.

Bessembinder, H., M.L. Lemmon. 2002. Equilibrium pricing and optimal hedging in electricity forward markets. J. Finance 57(3) 1347–1382.

Benth, F.E., J.S. Benth , S. Koekebakker, 2008. Stochastic Modelling of Electricity and Related Markets. World Scientific, Singapore.

Bertsekas, D.P. 1995. Dynamic Programming and Optimal Control. Athena Scientific, Belmont, MA.

Byström, H.N.E. 2003. The hedging performance of electricity futures on the Nordic power exchange. Applied Econom. 35(1) 1–11.

Coulon, M., W.B. Powell, R. Sircar. 2012. A model for hedging load and price risk in the Texas electricity market. Working paper, Princeton University, NJ.

Deng, S.J., S.S. Oren. 2006. Electricity derivatives and risk management. Energy 31(6-7) 940–953.

Deng, S.J., L. Xu. 2009. Mean-risk efficient portfolio analysis of demand response and supply resources. Energy 34(10) 1523–1529.

Ederington, L.H. 1979. The hedging performance of the new futures markets. J. Finance 34(1) 157–170.

Eydeland, A., K. Wolyniec. 2003. Energy and Power Risk Management: New Developments in Modeling, Pricing, and Hedging. Wiley, Hoboken, NJ

[30] The test was not carried on the triplet $(z_{0,t}, z_{1,t}, z_{2,t})$. The numerical burden associated with such a test is very high.

Fleten, S.E., E. Brathen, S.E. Nissen-Meyer. 2010. Evaluation of static hedging strategies for hydropower producers in the Nordic market. J. Energy Markets 3(4) 1–28.

Föllmer, H., P. Leukert. 1999. Quantile hedging. Finance and Stochastics 3(3) 251–273.

Föllmer, H., P. Leukert. 2000. Efficient hedging: cost versus shortfall risk. Finance and Stochastics 4(2) 117–146.

François, P., G. Gauthier, F. Godin. 2012. Optimal hedging when the underlying asset follows a regime-switching Markov process. Working paper, HEC, Montréal, Canada

González-Romera, E., M.A. Jaramillo-Morán, D. Carmona-Fernández. 2008. Monthly electric energy demand forecasting with neural networks and Fourier series. Energy Conversion and Management 49(11) 3135–3142.

Hahn, H., S. Meyer-Nieberg, S. Pickl. 2009. Electric load forecasting methods: Tools for decision making. Eur. J. Oper. Res. 199(3) 902–907.

Huovila, S. 2003. Short-term forecasting of power demand in the Nord Pool market. Master thesis, Lappeenranta University of Technology, Lappeenranta, Finland.

Johnsen, T.A., O.J. Olsen. 2011. Regulated and unregulated Nordic retail prices. Energy Policy 39(6) 3337–3345.

Koekebakker, S., F. Ollmar. 2005. Forward curve dynamics in the Nordic electricity market. Managerial Finance 31(6) 73–94.

Kleindorfer, P.R., L. Li. 2005. Multi-period VaR-constrained portfolio optimization with applications to the electric power sector. Energy 26(1) 1–26.

Madaleno, M., C. Pinho. 2008. The hedging effectiveness of electricity futures. Working paper, University of Aveiro, Portugal.

Liu, S.D., J.B. Jian, Y.Y. Wang. 2010. Optimal dynamic hedging of electricity futures based on copula-GARCH models. IEEM Conf. Proc., Macau. China. 2498–2502.

Lucia, J.J., H. Torro. 2011. On the risk premium in Nordic electricity futures prices. Internat. Rev. Econom. & Finance 20(4) 750–763.

Ni, J., L.K. Chu, F. Wu, D. Sculli, Y. Shi. 2012. A multi-stage financial hedging approach for the procurement of manufacturing materials. Eur. J. Oper. Res. 221(2) 424–431.

Nordic Energy Regulators (NordREG). 2010. The Nordic financial electricity market. Report 8/2010, November 2010.

Oum, Y., S.S. Oren, S.J. Deng. 2006. Hedging quantity risks with standard power options in a competitive wholesale electricity market. Naval Res. Logist. 53(7) 697–712.

Oum, Y., S.S. Oren. 2009. VaR constrained hedging of fixed price load-following obligations in competitive electricity markets. Risk and Decision Analysis 1(1) 43–56.

Oum, Y., S.S. Oren. 2010. Optimal static hedging of volumetric risk in a competitive wholesale electricity market. Decision Analysis 7(1) 107–122.

Rémillard, B. 2013. Statistical Methods for Financial Engineering. CRC Press, Boca Raton, FL.

Rémillard, B., H. Langlois, A. Hocquard, N. Papageorgiou. 2010. Optimal hedging of American options in discrete time. R. A. Carmona, P. Del Moral, P. Hu, N. Oudjane, eds. Numerical Methods in Finance. Springer, Berlin, Germany, 145–170.

Rosenblatt, M. 1952. Remarks on a multivariate transformation. Annals Math. Statist. 23(3) 470–472.

Schweizer, M. 1995. Variance-optimal hedging in discrete time. Math. Oper. Res. 20(1) 1–32.

State of California (2004). Order Instituting Rulemaking to Establish Policies and Cost Recovery Mechanisms for Generation Procurement and Renewable Resource Development Interim Opinion. Public Utility Commission Decision No. 04-01-050. http://docs.cpuc.ca.gov/PUBLISHED/FINAL_DECISION/33625.htm

Stoft, S., T. Belden, C. Goldman, S. Pickle. 1998. A primer on electricity futures and other derivatives. Technical Report, Laurence Berkeley National Laboratory, University of California, Berkeley.

Torro, H. 2011. Assessing the influence of spot price predictability on electricity futures hedging. J. Risk 13(4) 31–61.

Von der Fehr, N.H.M., P.V. Hansen. 2010. Electricity retailing in Norway. Energy 31(1) 25–45.

Wagner, M., P. Skantze, M. Ilic. 2003. Hedging optimization algorithms for deregulated electricity markets. Proc. 12th Conf. Intelligent Systems Appl. Power Systems, Lemnos, Greece.

Weron, R. 2006. Modeling and Forecasting Electricity Loads and Prices. Wiley, Chichester, England.

Woo, C.K., R. Karimov, I. Horowitz. 2004. Managing electricity procurement cost and risk by a local distribution company. Energy Policy 32(5) 635–645.

Zanotti, G., G. Gabbi, M. Geranio. 2010. Hedging with futures: Efficacy of GARCH correlation models to European electricity markets. J. Internat. Financial Markets, Institutions and Money 20(2) 135–148.