An arithmetic modeling framework for the term structure of electricity prices*

Kees E. Bouwman[†,1], Eran Raviv[‡,1,2], and Dick van Dijk[§,1]

[1] *Econometric Institute, Erasmus University Rotterdam*
[2] *Tinbergen Institute*

May 13, 2012

Abstract

We propose a tractable class of arbitrage-free models for the term structure of electricity prices, where spot and forward prices are a linear function of latent factors. The modeling approach offers much flexibility in the specification of the factor dynamics by only restricting their risk-neutral drift. We derive a canonical form where the parameters determining the factor loadings for the forward prices can be separated from the parameters describing the factor dynamics. The factor loading parameters can be consistently estimated by directly fitting the cross-section of forward prices. The modeling framework is applied to a panel of daily prices on forward contracts from the Nordpool electricity market, using affine factor dynamics. We find that forward prices *(i)* are mainly driven by changes in the level, slope and curvature of the forward curve; *(ii)* exhibit time-varying volatilities; and *(iii)* incorporate time-varying forward premia.

Keywords: Electricity market; Forwards and futures; Factor models; Affine processes

JEL-code: G12; C50; Q40

* We thank Ronald Huisman, Wing Wah Tham and Gregory Duffee for their comments and suggestions. The paper has benefited from comments received at the 17th International Conference on Computing in Economics and Finance, June 2011, San Francisco and the Energy & Finance Conference, October 2011, Rotterdam.

† *Corresponding author:* K.E. Bouwman, Econometric Institute, Erasmus University Rotterdam, PO Box 1738, 3000 DR Rotterdam, The Netherlands. Tel: +31 10 408 1262. Fax: +31 10 408 9162. Email: kbouwman@ese.eur.nl.

‡ *Address:* PO Box 1738, 3000 DR Rotterdam, The Netherlands. Email: raviv@ese.eur.nl.

§ *Address:* PO Box 1738, 3000 DR Rotterdam, The Netherlands. Email: djvandijk@ese.eur.nl.

1 Introduction

The recent liberalization of energy markets in various countries has led to the institution of organized exchanges where electricity is traded. An important implication of this liberalization is that market participants face a significant increase in price risk as prices now are determined freely by supply and demand. This has spurred the demand for a wide variety of electricity derivatives as well as models and tools for valuation and risk management.

The majority of electricity is not traded in the spot market for immediate delivery, but with forwards and futures contracts for delivery in the future. Together the prices of these contracts define a forward curve or a term structure of electricity prices. The unique features of electricity and the resulting type of traded contracts pose serious challenges to modeling the forward curve. In particular, electricity differs from other commodities in that it is (practically) non-storable. Electricity is therefore often referred to as a *flow* commodity (as opposed to a *stock* commodity) and contracts are specified with delivery taking place over a future time period rather than at a specific point in time.

Existing models for electricity forward prices have a similar structure to models for the term structure of interest rates, see e.g. Hinz and Wilhelm (2006). These models can be roughly classified into two groups. Models in the first group involve a specification for the evolution of the spot price, often featuring a dynamic factor structure. No-arbitrage principles are applied to derive the prices of forwards and futures. Examples are in Clewlow and Strickland (1999), Schwartz and Smith (2000), Lucia and Schwartz (2002), Benth *et al.* (2007), Monfort and Feŕon (2012), among others. The second group of models describes the evolution of the entire forward curve directly, based on the pioneering work of Heath, Jarrow and Morton [HJM] (1992). Examples are found in Koekebakker (2007), Benth and Koekebakker (2008) and Bjerksund *et al.* (2010), among others.

Most of the existing models for electricity pricing, however, yield intractable implications for observed forward prices. Expressions for forward prices and their distributions are often not available in closed-form, which severely complicates estimation and option pricing. A specific complication arises due to the fact that traded forwards and futures contracts do not involve delivery taking place at a specific point in time but instead during longer periods of months, quarters and years. To rule out arbitrage opportunities, forward prices must exhibit an 'additive structure', where e.g the price of a yearly forward contract should equal the average of the corresponding quarterly forward prices. An important reason for the intractability of many pricing models is a non-linear relation between forward prices and the underlying risk factors, which makes them incompatible with this additive structure of forward prices. For example, Benth and Koekebakker (2008) apply the HJM framework to observed contract prices, where particular contract prices with non-overlapping delivery periods are directly modeled by a log-normal process. This approach is however inconsistent for forwards having overlapping delivery periods. For example, when four consecutive quarterly forward contracts are priced by the model, the corresponding price of the yearly forward will not be log-normally distributed (due to the non-additivity of the lognormal distribution).

In this paper we propose a novel class of arbitrage-free electricity pricing models to overcome the above limitation. The key feature of these models is an additive factor structure, such that the arbitrage-free prices of forward and futures contracts are linear in the factors. Our class fits in the linearity-generating processes framework of Gabaix (2009) and offers a great deal of flexibility in specifying the factor dynamics. For instance, factor volatilities can be flexibly specified without affecting the relation between the factors and forward prices. This flexibility allows for factor specifications based on well-known processes with tractable features such as standard affine processes by Duffie and Kan (1996). Moreover, unlike other models, our proposed structure offers a convenient separation between factor loadings (the relation between factors and prices) and factor dynamics. This feature can be exploited to directly estimate the factors and the factor loading parameters by fitting the observed forward prices.

Our modelling approach allows us to fully exploit the information in forward prices by directly estimating models to a full panel of observed forward and futures contracts. We thereby contribute to the extant literature that focusses on estimating models to the dynamics of the observed spot price series. The panel approach offers more precise estimates of the factors and parameters by also exploiting information in the observed shape of the forward curve and the time-series of the forward prices.

In an empirical analysis, we estimate different models using a large panel of spot and forward prices from the Nordic electricity market. We find that forward prices are adequately described by three factors that capture changes in the level, slope and curvature of the forward curve. We explore the class of affine diffusion processes to model factor dynamics and find support for time-varying volatilities. Our findings are however consistent with a trade-off in flexibility between specifying conditional volatilities and correlations as found by Dai and Singleton (2000) as well as a flexibility trade-off between conditional volatility and fit of observed prices. Finally, the models are used to analyse the implications for forward premia. We document a substantial time-varying forward premia, consistent with the findings in Longstaff and Wang (2004).

A few models proposed in the literature fit into our general modeling framework. Benth, Kallsen and Meyer-Brandis (2008) put forward an additive model based on pure-jump Levy processes. Estimation of this model is challenging and is explored in Meyer-Brandis and Tankov (2008). Lucia and Schwartz (2002) and Knittel and Roberts (2005) also consider models with an additive structure, but they only consider Gaussian factor dynamics and focus on the spot price. Koekebakker (2007) advocates the use of an additive structure in the framework of a Gaussian HJM-type model. However, this requires extracting and modeling the entire forward curve first. This can be avoided in our modeling approach by estimating the model directly using observed forwards and futures prices.

2 General modeling framework

This section introduces the arbitrage-free modeling framework that provides tractable prices for electricity forwards and futures. Let S_t denote the *spot price* of electricity for instantaneous delivery at time t, i.e. delivery over the period $[t, t + \mathrm{d}t]$. Similarly, let $f_t(T)$ denote the *forward price* negotiated at $t \leq T$ for instantaneous delivery at time T, i.e. delivery over the period $[T, T + \mathrm{d}t]$. As discussed in the introduction, in practice most electricity forward contracts are specified with delivery periods covering months, quarters or years rather than an instantaneous period. For that purpose, we define $F_t(T_1, T_2)$ as the *average-based forward price* negotiated at $t \leq T_1 < T_2$ for future delivery of electricity over the period $[T_1, T_2]$.

An arbitrage-free model of electricity prices presupposes a frictionless and arbitrage-free market for all financial claims based on the spot price. We achieve this by directly assuming there exists a risk-neutral probability measure $\mathbb{Q}$, equivalent to the objective probability measure $\mathbb{P}$, that prices all such claims. The arbitrage-free price at time t of any contingent claim $\chi(S_T)$ based on the spot electricity price at a future point in time $T > t$ is then given by

$$P_{\chi,t} = e^{-r(T-t)}\,\mathrm{E}_t^{\mathbb{Q}}(\chi(S_T)), \tag{1}$$

where $\mathrm{E}_t^{\mathbb{Q}}$ denotes the expectation under the probability measure $\mathbb{Q}$ conditional on all information available at time t and r denotes the risk-free interest rate, which is assumed to be constant for simplicity.[1] Since a forward contract has no value at initiation, the no-arbitrage pricing relation implies that the instantaneous forward price is given by

$$f_t(T) = \mathrm{E}_t^{\mathbb{Q}}(S_T), \qquad \forall t < T.$$

In words, the forward price is equal to the expected future spot price under the risk-neutral measure, which corrects for a risk premium, see Björk (2004, Ch. 26).

The average-based forward contract that delivers electricity over a period $[T_1, T_2]$ is equivalent to a portfolio of all instantaneous forwards delivering over this period. Hence the average-based forward price is equal to the weighted average of all instantaneous forward prices covering the delivery period $[T_1, T_2]$, i.e.

$$F_t(T_1, T_2) = \int_{T_1}^{T_2} w(s) f_t(s)\,\mathrm{d}s, \tag{2}$$

with $w(s) = \frac{1}{T_2 - T_1}$ if the contract settles at maturity T_2, and $w(s) = \frac{re^{-rs}}{e^{-rT_1} - e^{-rT_2}}$ if settlement

[1] This is a standard assumption. As a consequence there is no difference in price between a forward and a future with equivalent delivery periods. This assumption can be relaxed by assuming interest rates are time-varying but independent of the spot price S_t.

takes place continuously during the delivery period, see Koekebakker (2007).

2.1 An arithmetic linear factor model

The arbitrage pricing framework allows for many different models. By specifying the risk-neutral dynamics of the spot price, we can use no-arbitrage arguments to obtain prices for derivative contracts. A key concern is to specify a model that leads to tractable expressions for commonly traded contracts such as forwards. In our modeling framework, we make two structural assumptions to ensure closed-form expressions for forward prices.

Assumption 1. *The spot price is a linear function of m factors, i.e.*

$$S_t = \delta_0(t) + \boldsymbol{\delta}_1' \boldsymbol{X}_t, \tag{3}$$

where $\boldsymbol{X}_t$ is an m-vector of factors, and $\delta_0(t)$ is a deterministic function capturing potential trends and seasonal components in the spot price.

Assumption 2. *The factors $\boldsymbol{X}_t$ follow a continuous-time vector process with a risk-neutral drift that is linear in $\boldsymbol{X}_t$, i.e.*

$$\mathrm{E}_t^{\mathbb{Q}}(\mathrm{d}\boldsymbol{X}_t) = (\boldsymbol{c}^{\mathbb{Q}} + \boldsymbol{D}^{\mathbb{Q}} \boldsymbol{X}_t)\,\mathrm{d}t. \tag{4}$$

These assumptions ensure that all instantaneous and average-based forward prices are also linear in the factors $\boldsymbol{X}_t$. In particular, the instantaneous forward price is given by[2]

$$f_t(T) = a(t, T-t) + \boldsymbol{b}(T-t)' \boldsymbol{X}_t, \tag{5}$$

$$\boldsymbol{b}(\tau) = e^{\boldsymbol{D}^{\mathbb{Q}\prime} \tau} \boldsymbol{\delta}_1, \tag{6}$$

$$a(t,\tau) = \delta_0(t+\tau) + [\boldsymbol{b}(\tau) - \boldsymbol{\delta}_1]' \, (\boldsymbol{D}^{\mathbb{Q}})^{-1} \boldsymbol{c}^{\mathbb{Q}}, \tag{7}$$

where $\boldsymbol{I}_m$ denotes an $m \times m$ identity matrix and $e^{\boldsymbol{A}}$ denotes the matrix exponential for a

[2]For convenience we assume invertibility of $\boldsymbol{D}^{\mathbb{Q}}$, which is consistent with $\boldsymbol{X}_t$ being mean-reverting under $\mathbb{Q}$.

square matrix $\boldsymbol{A}$.[3] The price of an average-based forward contract that delivers electricity over a period $[T_1, T_2]$ can be found straightforwardly by applying the above expression in (2). In particular, assuming the contract settles at maturity, we have

$$F_t(T_1, T_2) = A(t, T_1 - t, T_2 - t) + \boldsymbol{B}(T_1 - t, T_2 - t)'\boldsymbol{X}_t, \tag{8}$$

$$\boldsymbol{B}(\tau_1, \tau_2) = \tfrac{1}{\tau_2 - \tau_1}\left(e^{\boldsymbol{D}^{\mathbb{Q}\prime}\tau_2} - e^{\boldsymbol{D}^{\mathbb{Q}\prime}\tau_1}\right)(\boldsymbol{D}^{\mathbb{Q}\prime})^{-1}\boldsymbol{\delta}_1, \tag{9}$$

$$A(t, \tau_1, \tau_2) = \tfrac{1}{\tau_2 - \tau_1}\int_{\tau_1}^{\tau_2} \delta_0(t+s)\,\mathrm{d}s + \left[\boldsymbol{B}(\tau_1, \tau_2) - \boldsymbol{\delta}_1\right]'(\boldsymbol{D}^{\mathbb{Q}})^{-1}\boldsymbol{c}^{\mathbb{Q}}. \tag{10}$$

The linear relation between the spot price and the factors contrasts with many existing spot price models that assume a multiplicative relation

$$S_t = e^{\delta_0 + \boldsymbol{\delta}_1{}'\boldsymbol{X}_t}, \tag{11}$$

which typically is motivated by the fact that the resulting spot price is guaranteed to be positive. In our framework, however, we can specify the factor dynamics to ensure positivity of prices. For example, if the factors follow a Cox, Ingersoll and Ross [CIR] (1985) square-root process, we can obtain a positive spot price process under appropriate parameter restrictions.

Assumptions 1 and 2 allow for a great deal of flexibility in specifying the factor dynamics. In particular, we do not impose restrictions on the factor volatilities and the factor drifts under the objective measure $\mathbb{P}$. Equivalence between the risk-neutral measure $\mathbb{Q}$ and the objective measure $\mathbb{P}$ can however impose further regularity conditions on the factor $\mathbb{P}$-drifts.

2.2 Canonical form

The model as introduced above is invariant to affine transformations of the factors $\boldsymbol{X}_t$ and hence unidentified. In particular, defining

$$\tilde{\boldsymbol{X}}_t = \boldsymbol{g} + \boldsymbol{H}\boldsymbol{X}_t \tag{12}$$

[3] The matrix exponential is defined as $e^{\boldsymbol{A}} = \sum_{i=0}^{\infty} \frac{1}{i!}\boldsymbol{A}^i$.

with $\boldsymbol{H}$ an invertible $m \times m$ matrix, the model can be expressed equivalently in terms of $\tilde{\boldsymbol{X}}_t$, satisfying Assumptions 1 and 2, with corresponding parameters

$$\tilde{\delta}_0(t) = \delta_0(t) - \boldsymbol{\delta}_1' \boldsymbol{H}^{-1} \boldsymbol{g}, \quad \tilde{\boldsymbol{\delta}}_1 = (\boldsymbol{H}^{-1})' \boldsymbol{\delta}_1, \quad \tilde{\boldsymbol{c}}^{\mathbb{Q}} = \boldsymbol{H} \boldsymbol{c}^{\mathbb{Q}} - \boldsymbol{H} \boldsymbol{D}^{\mathbb{Q}} \boldsymbol{H}^{-1} \boldsymbol{g}, \quad \tilde{\boldsymbol{D}}^{\mathbb{Q}} = \boldsymbol{H} \boldsymbol{D}^{\mathbb{Q}} \boldsymbol{H}^{-1}.$$

We therefore impose restrictions on the parameters $\delta_0(t), \boldsymbol{\delta}_1, \boldsymbol{c}^{\mathbb{Q}}$ and $\boldsymbol{D}^{\mathbb{Q}}$ to obtain a canonical form that is just-identified and maximally flexible, such that it nests all other models as a special case. In particular, we adopt the canonical form put forward by Joslin *et al.* (2011) that is based on the diagonalization of the matrix $\boldsymbol{D}^{\mathbb{Q}}$ using the Jordan decomposition. The canonical form is given by the restrictions

$$\delta_0(t) \text{ is unrestricted}, \qquad \boldsymbol{\delta}_1 = \boldsymbol{\iota}, \qquad \boldsymbol{c}^{\mathbb{Q}} = \boldsymbol{0}, \qquad \boldsymbol{D}^{\mathbb{Q}} = -\boldsymbol{J}(\boldsymbol{\lambda}^{\mathbb{Q}}), \tag{13}$$

where $\boldsymbol{\iota}$ denotes a vector of ones, and $\boldsymbol{J}(\boldsymbol{\lambda}^{\mathbb{Q}})$ denotes the real Jordan form of a matrix with a vector of ascendingly ordered positive eigenvalues $\boldsymbol{\lambda}^{\mathbb{Q}} \in \mathbb{C}^m$, i.e. $0 < \lambda_1^{\mathbb{Q}} \leq \ldots \leq \lambda_m^{\mathbb{Q}}$. If all elements in $\boldsymbol{\lambda}^{\mathbb{Q}}$ are real and distinct, then $\boldsymbol{J}(\boldsymbol{\lambda}^{\mathbb{Q}}) = \text{diag}(\boldsymbol{\lambda}^{\mathbb{Q}})$.

This canonical form is unique in the sense that it is no longer invariant to linear transformations, hence rendering a well-identified model. It is also maximally flexible as any model satisfying Assumptions 1 and 2 can be linearly transformed to this canonical form (Joslin *et al.*, 2011).

The canonical form offers several advantages for analysis. First, it imposes *all* identifying restrictions on parameters determining the factor loadings in (9). Models with the same canonical parameters can therefore produce the same variation in the shapes of the forward curve. This allows for an easy comparison of different models. More importantly, it allows to identify parameters and factors from a cross-section of forward prices without directly considering the factor dynamics. We return to this issue in more detail in Section 3. Second, the number of free parameters is substantially reduced from $(m+1)^2$ to only $(m+1)$ free parameters, which facilitates estimation. Third, the expressions for the factor loadings in (6), (7), (9) and (10) can be simplified considerably.

2.3 Factor interpretation

Since the model is invariant to affine transformations of the factors, we adopt the canonical form (13) to ensure identification of the factors $\boldsymbol{X}_t$ and parameters. The interpretation of the factors $\boldsymbol{X}_t$ is derived from this canonical form and therefore not necessarily informative. The invariance property of the model, however, allows us to consider more informative transformations (12) of the canonical factors $\boldsymbol{X}_t$. In particular, we propose a transformation of the factors that yields the principal components of the model-implied forward curve.

The instantaneous forward curve of electricity prices as implied by the model is given by (5). For a fixed set of $k > m$ maturities $0 < \tau_1 < \ldots < \tau_k$, the corresponding vector of model-implied instantaneous forward prices $\boldsymbol{f}_t = (f_t(\tau_1), \ldots, f_t(\tau_k))'$ satisfies

$$\boldsymbol{f}_t = \boldsymbol{a}_f + \boldsymbol{B}_f \boldsymbol{X}_t,$$

where $\boldsymbol{a}_f$ and $\boldsymbol{B}_f$ are the correspondingly stacked factor loading coefficients (7) and (6). The unconditional covariance matrix of $\boldsymbol{f}_t$ is given by

$$\boldsymbol{\Sigma}_f = \boldsymbol{B}_f \boldsymbol{\Sigma}_X \boldsymbol{B}_f',$$

where $\boldsymbol{\Sigma}_X$ is the unconditional covariance matrix of factors $\boldsymbol{X}_t$. Note that $\boldsymbol{\Sigma}_f$ is positive semi-definite with rank m and therefore only has m principal components related to its non-zero eigenvalues. The m principal components are linear in $\boldsymbol{X}_t$ and can hence be used as transformed factors $\tilde{\boldsymbol{X}}_t$. This transformation ensures that the transformed factors are unconditionally uncorrelated and normalized to have unit variance. Furthermore, the first transformed factor is the factor that explains most of the variation in the forward curve, whereas subsequent factors explain most variation in the forward curve after accounting for the preceding factors.

2.4 Modeling factor dynamics

Our framework offers a great deal of flexibility in specifying the factor dynamics. The factors can be modeled by diffusions, jump-diffusions or Levy processes and the corresponding volatilities can be specified flexibly. As long as the model produces well-defined factor dynamics satisfying the drift assumption in (4), prices of forwards are linear in the factors as given in (8). In this paper we focus on standard affine diffusion processes to model the factor dynamics. Affine diffusion processes are widely used in finance for their tractability and their properties are well established.

Affine diffusions are characterized by a drift and instantaneous covariance matrix that are linear (affine) functions of the process value. We assume that $\boldsymbol{X}_t$ follows a standard affine diffusion process under both $\mathbb{P}$ and $\mathbb{Q}$ given by

$$\mathrm{d}\boldsymbol{X}_t = (\boldsymbol{c}^{\mathbb{M}} + \boldsymbol{D}^{\mathbb{M}}\boldsymbol{X}_t)\,\mathrm{d}t + \boldsymbol{\Sigma}\sqrt{\mathrm{diag}(\boldsymbol{\alpha} + \boldsymbol{B}\boldsymbol{X}_t)}\,\mathrm{d}\boldsymbol{W}_t^{\mathbb{M}}, \tag{14}$$

where $\mathbb{M} \in \{\mathbb{P}, \mathbb{Q}\}$, $\boldsymbol{W}_t^{\mathbb{M}}$ denotes an m-vector of independent standard Brownian motions under $\mathbb{M}$ and $\mathrm{diag}(\boldsymbol{x})$ denotes a diagonal matrix with the main diagonal given by the vector $\boldsymbol{x}$. The process $\boldsymbol{X}_t$ hence satisfies Assumption 2 and under canonical restrictions in (13) we have $\boldsymbol{c}^{\mathbb{Q}} = \boldsymbol{0}$ and $\boldsymbol{D}^{\mathbb{Q}} = -\boldsymbol{J}(\boldsymbol{\lambda}^{\mathbb{Q}})$.

The class of affine processes is very general. When we take $\boldsymbol{B} = \mathbf{O}$, the volatilities are constant and the process becomes a standard Gaussian process. When $\boldsymbol{B}$ is non-zero, the volatilities vary with the level of the factors.

Further restrictions must be imposed to ensure that the model is admissible and identified (Dai and Singleton, 2000), making it difficult to work with (14) directly. Dai and Singleton (2000) establish a classification of admissible affine processes based on the rank of $\boldsymbol{B}$ into non-nested subclasses. In particular, if the process $\boldsymbol{X}_t$ is admissible with the rank of $\boldsymbol{B}$ equal to k, the process $\boldsymbol{X}_t$ belongs to the subclass $\mathbb{A}_k(m)$. For each subclass, Dai and Singleton (2000) provide a canonical characterization of the maximally flexible model that nests every model in this class as a special case. We examine the performance of the maximally flexible affine diffusion for every subclass $\mathbb{A}_k(m)$.

For a convenient characterization of the maximally flexible affine diffusion, we first linearly transform the factors $\boldsymbol{X}_t$ to $\boldsymbol{Y}_t$, i.e.

$$\boldsymbol{Y}_t = \boldsymbol{g}_{\boldsymbol{Y}} + \boldsymbol{H}_{\boldsymbol{Y}} \boldsymbol{X}_t, \tag{15}$$

and then use a canonical representation of the maximally flexible model for $\boldsymbol{Y}_t$. The canonical process $\boldsymbol{Y}_t$ under both $\mathbb{P}$ and $\mathbb{Q}$ is given by

$$\mathrm{d}\boldsymbol{Y}_t = (\boldsymbol{c}_{\boldsymbol{Y}}^{\mathbb{M}} + \boldsymbol{D}_{\boldsymbol{Y}}^{\mathbb{M}} \boldsymbol{Y}_t)\,\mathrm{d}t + \sqrt{\mathrm{diag}(\boldsymbol{\alpha}_{\boldsymbol{Y}} + \boldsymbol{B}_{\boldsymbol{Y}} \boldsymbol{Y}_t)}\,\mathrm{d}\boldsymbol{W}_t^{\mathbb{M}}, \tag{16}$$

where $\mathbb{M} \in \{\mathbb{P}, \mathbb{Q}\}$ and $\boldsymbol{D}_{\boldsymbol{Y}}^{\mathbb{Q}} = -\boldsymbol{H}_{\boldsymbol{Y}} \boldsymbol{J}(\boldsymbol{\lambda}^{\mathbb{Q}}) \boldsymbol{H}_{\boldsymbol{Y}}^{-1}$ and $\boldsymbol{c}_{\boldsymbol{Y}}^{\mathbb{Q}} = \boldsymbol{H}_{\boldsymbol{Y}} \boldsymbol{J}(\boldsymbol{\lambda}^{\mathbb{Q}}) \boldsymbol{H}_{\boldsymbol{Y}}^{-1} \boldsymbol{g}_{\boldsymbol{Y}}$ by Assumption 2 and the canonical restrictions in (13). We impose the admissibility conditions presented in Appendix A to ensure admissibility and identification of the process $\boldsymbol{Y}_t$. The process for $\boldsymbol{X}_t$ is now completely characterized by the parameters of the transformation $\boldsymbol{g}_{\boldsymbol{Y}}, \boldsymbol{H}_{\boldsymbol{Y}}$ and the parameters $\boldsymbol{c}_{\boldsymbol{Y}}^{\mathbb{P}}, \boldsymbol{D}_{\boldsymbol{Y}}^{\mathbb{P}}, \boldsymbol{\alpha}_{\boldsymbol{Y}}, \boldsymbol{B}_{\boldsymbol{Y}}$.

An important consequence of the admissibility conditions is that they potentially restrict cross-sectional parameters $\boldsymbol{\lambda}^{\mathbb{Q}}$ and hence the specification of the factor dynamics does affect the cross-sectional performance of the model. For example, some affine dynamics rule out complex conjugate pairs for the eigenvalues $\boldsymbol{\lambda}^{\mathbb{Q}}$ (see Appendix A).

2.5 Deterministic trends and seasonal components

The intercept $\delta_0(t)$ in the spot price specification (3) is allowed to be deterministically time-varying to accommodate trends and seasonal components. For example, we may accommodate an annual seasonal pattern using monthly dummies:

$$\delta_0(t) = \sum_{j=1}^{12} M_j(t)\beta_j, \qquad \text{with } M_j(t) = \begin{cases} 1 & \text{if date } t \text{ is in month } j \\ 0 & \text{otherwise.} \end{cases}. \tag{17}$$

Under the canonical form, we have $a(t,\tau) = \delta_0(t+\tau)$ and $A(t,\tau_1,\tau_2) = \frac{1}{\tau_2-\tau_1}\sum_{j=1}^{12}\int_{\tau_1}^{\tau_2} M_j(t+s)\,\mathrm{d}s\beta_j$, so that $A(t,\tau_1,\tau_2) = \beta_i$ for a monthly forward contract delivering in month i,

$A(t, \tau_1, \tau_2) = \frac{1}{3}(\beta_i + \beta_{i+1} + \beta_{i+2})$ for a quarterly contract delivering in months i, $i+1$ and $i+2$ and $A(t, \tau_1, \tau_2) = \frac{1}{12}\sum_{j=1}^{12} \beta_j$ for a yearly contract. Additional components and alternative specifications of the seasonal component, such as in Lucia and Schwartz (2002), can also be incorporated straightforwardly.

3 Estimation

This section describes the procedure for estimating the latent factors and model parameters from a panel of observed forward prices. We describe the estimation procedure for models of the canonical form (13) using the seasonal specification in (17).

Assume that at time $t = 1, \dots, T$ we observe n_t average-based forward prices $F_t(T_{1,t,j}, T_{2,t,j})$ with delivery periods $[T_{1,t,j}, T_{2,t,j}]$ for $j = 1, \dots, n_t$. The number of available contracts and their delivery periods may vary over time such that the panel of observed forward prices is unbalanced. We assume that observed forward prices $F_t(T_{1,t,j}, T_{2,t,j})$ are measured with a Gaussian error $e_{t,j} \sim N(0, \sigma_e^2)$ that is independent across time and across different contracts. By stacking all forward prices observed at time t in an n_t-vector $\boldsymbol{F}_t$, the model can be expressed as

$$\boldsymbol{F}_t = \boldsymbol{A}_t + \boldsymbol{B}_t \boldsymbol{X}_t + \boldsymbol{\varepsilon}_t, \tag{18}$$

where $\boldsymbol{A}_t$ and $\boldsymbol{B}_t$ are the correspondingly stacked factor loading coefficients as specified in (10) and (9) and $\boldsymbol{\varepsilon}_t \sim N(0, \sigma_{\boldsymbol{\varepsilon}}^2 \boldsymbol{I}_{n_t})$ is serially independent.

The model is completed by a specification of the factor dynamics, such as the affine diffusion processes discussed before. The complete model represents a non-linear state-space system with measurement equation (18) and a transition equation given by the factor dynamics. The parameters and the factors are estimated with Quasi-Maximum Likelihood (QML) using the Kalman filter.

The likelihood of the affine process (14) is only known in closed-form for a few special cases. In particular, processes belonging to $\mathbb{A}_0(m)$ are standard Gaussian processes and hence ML estimators are obtained in closed form. In other cases, exact ML estimation is infeasible and we therefore use QML by approximating the transition density by a normal distribution with the exact first and second moments implied by the affine process. The first and second moments implied by the affine process can be calculated in closed form (Fackler, 2000) and hence the QML estimator is consistent (Fisher and Gilles, 1996). For QML, consistent standard errors are obtained by using the sandwich estimator of White (1982).

3.1 Cross-sectional estimation of factors and loadings

The QML estimation procedure fully utilizes both the cross-sectional information in the fit of the forward prices and the time-series information in the factors. The model however yields a separation of the parameters into *cross-sectional parameters* that only determine the factor loadings and *dynamic parameters* that determine the evolution of the factors. We exploit this structure to estimate the factors and cross-sectional parameters by directly fitting the forward prices as in (18) using Non-linear Least Squares (NLS).

The cross-sectional parameters and the factors are collected in the vector $\boldsymbol{\theta}_c = (\boldsymbol{\lambda}^{\mathbb{Q}}, \boldsymbol{\beta}, \boldsymbol{X}_1, \ldots, \boldsymbol{X}_T)$. In the first step, we estimate $\boldsymbol{\theta}_c$ by minimizing the sum of squared residuals $\boldsymbol{e}_t(\boldsymbol{\theta}_c) = \boldsymbol{F}_t - \boldsymbol{A}_t - \boldsymbol{B}_t\boldsymbol{X}_t$:

$$\hat{\boldsymbol{\theta}}_c = \operatorname{argmin} \sum_{t=1}^{T} \boldsymbol{e}_t(\boldsymbol{\theta}_c)'\boldsymbol{e}_t(\boldsymbol{\theta}_c) \tag{19}$$

using Non-linear Least Squares (NLS). Estimation can be simplified by noting that for a given $\boldsymbol{\lambda}^{\mathbb{Q}}$ the errors $\boldsymbol{e}_t(\boldsymbol{\theta}_c)$ are linear in $\boldsymbol{\beta}$ and $\boldsymbol{X}_t$, $t = 1, \ldots, T$. Hence, the parameters in $\delta_0(t)$ and the factors can be concentrated out by using the corresponding least squares solution. That is, given $\boldsymbol{\lambda}^{\mathbb{Q}}$, the optimal solutions for $\boldsymbol{\beta}$ and the factors $\boldsymbol{X}_t$ are given by the OLS estimates $\hat{\boldsymbol{\beta}}(\boldsymbol{\lambda}^{\mathbb{Q}})$ and $\hat{\boldsymbol{X}}_t(\boldsymbol{\lambda}^{\mathbb{Q}})$ and the optimization problem in (19) now only requires a numerical optimization over $\boldsymbol{\lambda}^{\mathbb{Q}}$. See Appendix B for details. Andreasen and Christensen (2010) show that this NLS estimator consistently estimates the cross-sectional parameters and the factors as the number of observed contracts n_t tends to infinity.

This NLS estimator ignores the time-series information in the factors and is therefore less efficient but also has important benefits. First, it can be calculated fast and straightforwardly with low computational demands. Second, the estimates are independent of the factor dynamics, which facilitates model specification and interpretation. The resulting estimates of factor loadings and factors can be used to guide the subsequent specification of the factor dynamics. Moreover, some specifications of the factor dynamics can restrict the domain of the factor or the cross-sectional parameters, thereby reducing the model fit of observed prices. This can be revealed by comparing estimates of a full model with the unrestricted NLS estimates based exclusively on cross-sectional information.

The NLS estimator does not estimate the dynamic parameters. For forecasting purposes, the dynamic parameters can be estimated in a second step from the estimated factors comparable to Diebold and Li (2006). However, such a two-step procedure fails to account for estimation uncertainty in the factor estimates in the second step and is hence less useful for inference on the complete model.

4 Empirical analysis

The proposed modelling framework is used for an empirical analysis of a large panel of spot and forward prices in the Nordic electricity market. The Nordic electricity market is formed by two related markets, the "Nordpool spot" and "Nordpool ASA" markets. The "Nordpool spot" is a physical market where participants trade electricity for the next day. The market is referred to as a spot market but is in fact a one-day forward market. The market price, also referred to as the "system price", is determined by the intersection of the aggregate demand and supply curves derived from bids by traders. "Nordpool ASA" is the financial market in which derivatives such as futures and forward contracts are traded.[4] Contracts refer to load in megawatt hour (MWh) for a given delivery period. The main difference between the futures and forward contracts is that futures are marked to market, that is, changes in market prices of the contract are settled daily at each participant's account. Forward contracts are settled during the delivery period, not during the preceding trading period.

[4]As of May 2010 "Nordpool ASA" is part of NASDAQ OMX Group, Inc.

4.1 Data description

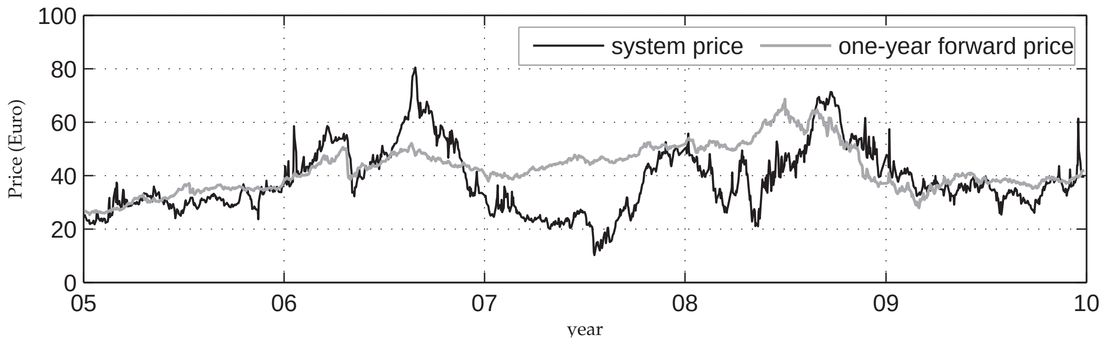

Figure 1: System Price and One-Year Forward Price. The figure shows the daily system price and one-year forward price from January 3, 2005 up to December 28, 2009 in Euros/MWh. The one-year forward price is interpolated by taking the weighted average of the traded one- and two-year forwards.

Our data set covers daily closing prices for monthly, quarterly and yearly forward contracts and the system price over the period from January 3, 2005 up to December 28, 2009. Every contract has a different delivery period. Monthly contracts are available with maturities up till 5 months to start of delivery, while quarterly with and yearly contracts have longer maturities, up to 4 years. We also incorporate the spot price series in the analysis and treat it as a one-day forward with a one-day delivery period, that is $T_2 = T_1 + 1 = t + 1$ in the notation defined in (2). There is no trading during the contracts delivery period. Forward contracts are quoted in Euro while the spot price is quoted in NOK. For consistency, we convert the spot price to Euro using the daily NOK/EUR exchange rate.

Figure 1 shows the system price and the one-year forward price over time. The fluctuations in the system and forward prices are quite volatile, with the system price being more volatile than the forward. The fluctuations in both prices show a tendency to mean-revert towards a level of about 40 Euros/MWh, where the fluctuations in the one-year forward price are more persistent. The system price does exhibit some seasonal fluctuation with prices tending to be higher during winter seasons, which can be attributed to weather conditions.

Table 1 display summary statistics of the system price and monthly, quarterly and yearly forward prices. For comparison, all prices in the table are deseasonalized using monthly dummies by adding $\frac{1}{12}\sum_{j=1}^{12}\beta_j$ to the residuals of a null model with no factors and an intercept given by (17). Forward prices are interpolated by taking the weighted average of the traded forwards with neighbouring maturities to represent prices with a fixed time to delivery. All prices exhibit a positive skew, but there is however no sign of heavy-tails. Price fluctuations are clearly mean-reverting as reflected by the autocorrelation coefficients and the lower price volatility for longer maturities, i.e. longer times to the start of delivery. Price fluctuations are however more persistent for longer maturities, which cannot be explained by a one-factor model. More factors with differences in persistence are needed to capture short-term and long-term variation as in e.g. Schwartz and Smith (2000). The average price tends to be slightly higher for longer maturities, which indicates that forward premia are on average positive.

The next two subsections present the estimation results for different models. First, results for the cross-sectional part of models with different numbers of factors m are discussed, focusing on the model fit of the forward prices and the interpretation of the estimated factors. Next, the estimation results for the affine factor dynamics are presented, where different affine specifications are compared.

4.2 Cross-sectional results

We first estimate the factors $\boldsymbol{X}_t$ and the cross-sectional parameters $\boldsymbol{\beta}$ and $\boldsymbol{\lambda}^{\mathbb{Q}}$ using Non-Linear Least Squares as in (19). Seasonal variation is modeled by monthly dummies as specified in (17). Different eigenvalue structures are considered, in particular both real and complex eigenvalues and both distinct and repeated eigenvalues. The structure with the best fit is reported.

Figure 2 show the fit of models with 1 to 5 factors relative to a model with only seasonal dummies. Models with more factors obviously fit observed forward prices better, with diminishing improvements in fit when including more factors, but can potentially also lead to over-fitting and unstable estimates. Figure 2 indicates that the marginal improvement in fit by including more than three factors is very limited. We therefore restrict further attention to models with one, two and three factors.

The parameters $\boldsymbol{\beta}$ account for the average level of the forward prices and the seasonal variation. Figure 3 plots the estimated values of $\boldsymbol{\beta}$ for the null model, i.e. a model without

Table 1: Descriptive Statistics of Electricity Prices

Descriptive statistics of the daily system price and daily forward prices of monthly (m), quarterly (q) and yearly (y) contracts. Prices are deseasonalized by using monthly dummies covering the delivery period. More specifically, $\frac{1}{12}\sum_{j=1}^{12}\beta_j$ is added to the residuals of a null model with no factors and an intercept given by (17). Forward prices are interpolated by taking the weighted average of the traded forwards with neighbouring maturities to represent prices with a fixed time to delivery. The *start* and *end* of the delivery periods are represented in days (d), months (m), quarters (q) or years (y). All contracts are quoted in Euro/MWh. *no. obs.* denotes the total number of observation days, *std. dev.* the standard deviation and $\widehat{\rho}_{30}$ the 30th order autocorrelation.

delivery		no. obs.	mean	std. dev	skewness	kurtosis	$\widehat{\rho}_{30}$
start	end						
System price							
1d	2d	1245	38.41	12.71	0.74	3.27	0.69
Monthly forward contracts							
1m	2m	1245	40.20	11.92	0.81	3.02	0.75
2m	3m	1245	41.31	11.75	0.79	2.84	0.77
3m	4m	1245	42.23	11.73	0.76	2.70	0.78
4m	5m	1245	42.86	11.70	0.76	2.67	0.79
5m	6m	1245	43.23	11.44	0.75	2.81	0.80
Quarterly forward contracts							
1q	2q	1058	44.50	11.29	0.71	2.59	0.76
2q	3q	1124	44.77	9.41	0.70	2.63	0.80
3q	4q	1185	43.83	8.42	0.67	2.85	0.82
4q	5q	1245	42.73	8.67	0.40	2.81	0.87
5q	6q	1245	42.83	8.61	0.43	2.55	0.89
6q	7q	1245	42.71	8.71	0.37	2.91	0.88
7q	8q	1245	42.45	8.53	0.36	3.19	0.88
Yearly forward contracts							
1y	2y	1245	42.71	8.58	0.43	2.86	0.88
2y	3y	1245	42.52	8.28	0.43	2.92	0.90

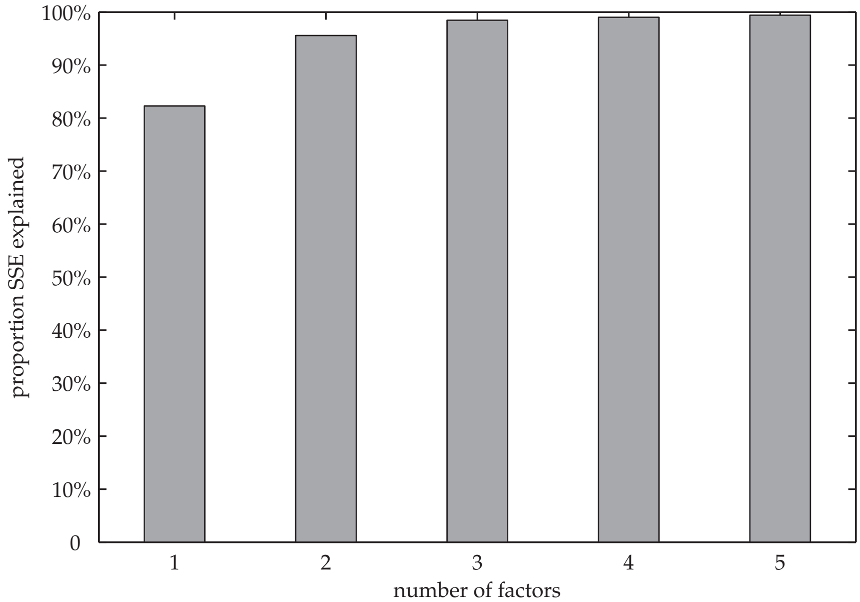

Figure 2: Fit of Forward Prices. The plot shows the percentage of the Sum of Squared Errors (SSE) of a model with seasonal dummies but without factors (i.e. with the spot price given by $\delta_0(t)$ in (17)) explained by a model with m factors for $m = 1, \ldots, 5$.

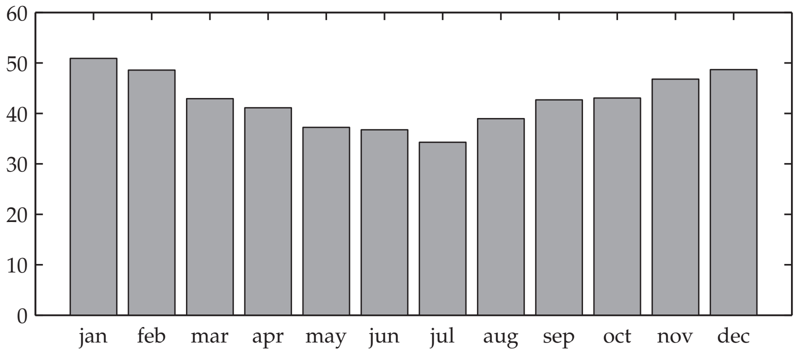

Figure 3: Estimates of Monthly Dummies $\boldsymbol{\beta}$. The plot shows the point estimates of the monthly dummies in the null model, i.e. a model no factors.

factors. The estimates clearly display a seasonal pattern, where prices are on average higher during winter and lower during summer. The estimates for the full set of models are reported in Table 6 in Appendix C. The seasonal pattern is very similar across different models. The average level of the forward curve is however also affected by the average level of the factors and hence the estimates for $\boldsymbol{\beta}$ differ in level across the different models.

Table 2: Cross-Sectional Parameters

The table reports the estimates of the cross-sectional parameters $\boldsymbol{\lambda}^{\mathbb{Q}}$ and σ_ε for models with one, two and three factors. Columns denoted by *cross* report cross-sectional NLS estimates, whereas columns denoted by $\mathbb{A}_k(n)$ report QML estimates based on the corresponding affine factor dynamics. Asymptotic standard errors are shown in parenthesis. Estimates of $\boldsymbol{\lambda}^{\mathbb{Q}}$ are expressed on a yearly basis, i.e. maturities τ are represented in years. Estimates of $\boldsymbol{\beta}$ are reported in Table 6

	$\lambda_1^{\mathbb{Q}}$	$\lambda_2^{\mathbb{Q}}$	$\lambda_3^{\mathbb{Q}}$	σ_ε
Panel A. 1-factor models				
cross	0.16 (0.00)			4.30
$\mathbb{A}_0(1)$	0.16 (0.00)			4.36 (0.01)
$\mathbb{A}_1(1)$	0.16 (0.00)			4.36 (0.07)
Panel B. 2-factor models				
cross	0.09 (0.00)	1.58 (0.02)		2.15
$\mathbb{A}_0(2)$	0.09 (0.00)	1.58 (0.01)		2.22 (0.00)
$\mathbb{A}_1(2)$	0.09 (0.01)	1.58 (0.05)		2.22 (0.03)
$\mathbb{A}_2(2)$	0.10 (0.01)	1.57 (0.05)		2.22 (0.03)
Panel C. 3-factor models				
cross	0.03 (0.00)	$2.28 + 1.84i$ (0.02) (0.02)	$2.28 - 1.84i$ (0.02) (0.02)	1.31
$\mathbb{A}_0(3)$	0.03 (0.00)	$2.21 + 1.84i$ (0.02) (0.01)	$2.21 - 1.84i$ (0.02) (0.01)	1.35 (0.00)
$\mathbb{A}_1(3)$	0.04 (0.01)	3.12 (0.06)	3.15 (0.07)	1.39 (0.02)
$\mathbb{A}_2(3)$	0.04 (0.01)	3.12 (0.08)	3.15 (0.06)	1.39 (0.02)
$\mathbb{A}_3(3)$	0.04 (0.00)	$2.63 + 1.44i$ (0.05) (0.03)	$2.63 - 1.44i$ (0.05) (0.03)	1.36 (0.02)

The estimates of $\boldsymbol{\lambda}^{\mathbb{Q}}$ determine how the factors affect the forward prices and are reported in Table 2. The factor loadings $\boldsymbol{b}(\tau)$ of the instantaneous forward prices in (6) have a simple exponential decay with maturity. Higher values of $\boldsymbol{\lambda}^{\mathbb{Q}}$ produce a stronger decay and hence forwards with longer maturities are less affected by the factors.

In Figure 4, we plot the factor loadings $\boldsymbol{b}(\tau)$ as defined in (6). The factor loadings are plotted for the model-implied principal components $\widetilde{\boldsymbol{X}}_t$ as defined Section 2.3 to facilitate

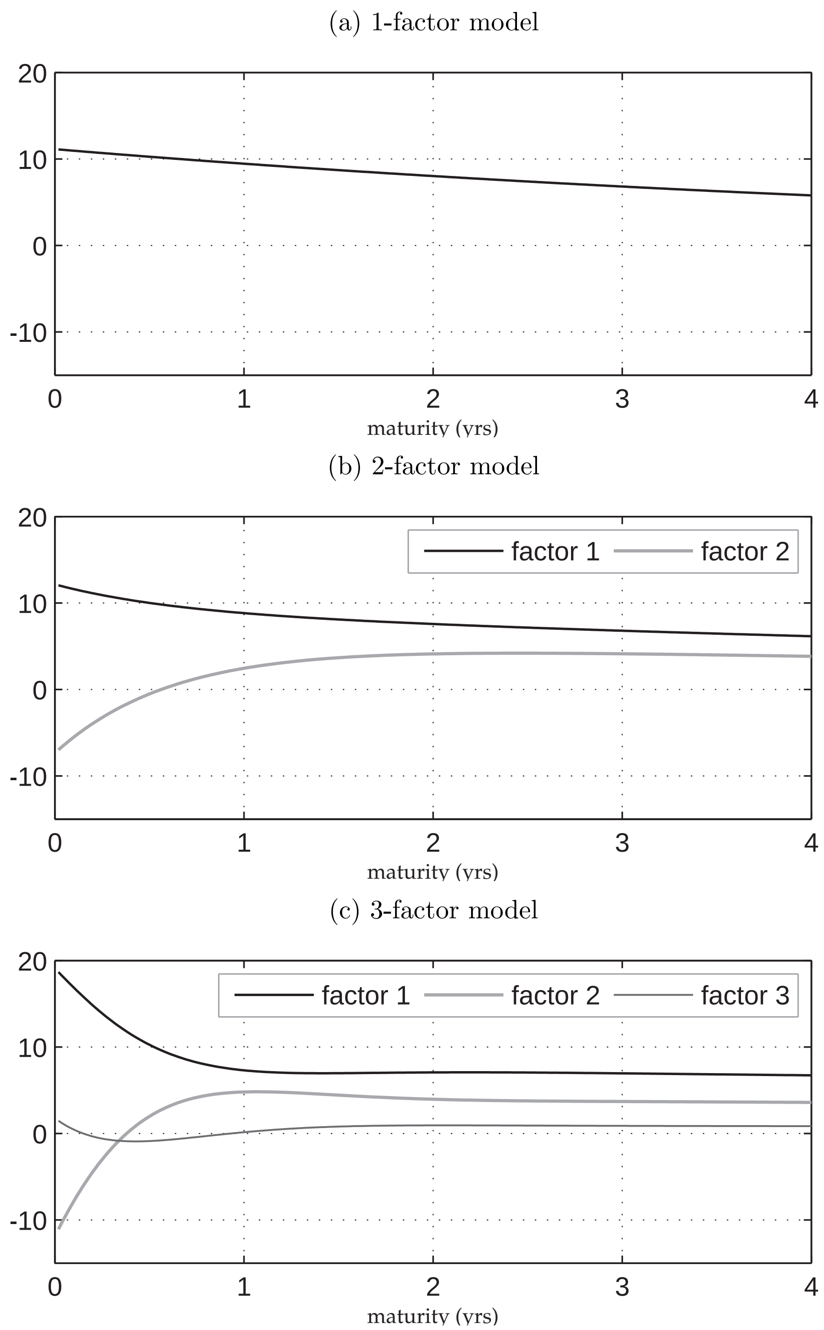

Figure 4: Factor Loadings Panels (a),(b) and (c) plot the estimated factor loadings for the one-, two- and three-factor models. The factors are represented as model-implied principal components $\widetilde{\boldsymbol{X}}_t$ as described in Section 2.3.

interpretation. The factor loadings in Figure 4 show the change in the forward curve due to an increase in the factor by one standard deviation. The estimate of $\lambda_1^{\mathbb{Q}}$ for the one-factor model is low, producing a slowly decaying factor loading for the forward curve, as shown in Panel (a) of Figure 4. An increase in the factor mainly increases the general level of the forward curve, but also tilts the forward curve downward. The estimate $\lambda_1^{\mathbb{Q}}$ decreases when adding extra factors, but the change in the factor loading is small. The estimates of $\lambda_2^{\mathbb{Q}}$ are larger, producing a stronger decay in the loading for the second factor. For the two-factor model, we observe that the second model-implied principal component captures changes in the slope of the forward curve. This pattern in the factor loading is similar in the three-factor model. In the three-factor model, the decay parameter of the second and third factor–$\lambda_2^{\mathbb{Q}}$ and $\lambda_3^{\mathbb{Q}}$–are complex conjugates. The factor loading plot in Panel (c) of Figure 4 shows that the third factor mainly accounts for variations in the short end of the forward curve by producing a hump.

The time-series of the model-implied principal components $\widetilde{\boldsymbol{X}_t}$ are plotted in Figure 5. The graphs indicate that the factor estimates are hardly affected by including extra factors, e.g. the estimated first factor is very similar across the three models. The first and second estimated factor have a similar volatility and persistence, whereas the fluctuations in the third factor are much more transitory.

In Table 2, we also report the QML estimates based on affine factor dynamics. As discussed in Section 2.4, admissible affine diffusions are classified in non-nested subclasses $\mathbb{A}_k(m)$. For each subclass we estimate the maximally flexible model, which for sake of brevity we also indicate by $\mathbb{A}_k(m)$. The QML estimator combines both cross-sectional and time-series information and hence gives up a bit on cross-sectional fit to obtain a better fit of the time-series properties. This small reduction in cross-sectional fit is reflected by the increased estimate for the residual variance σ_ε when compared to the cross-sectional estimates.

The QML estimates are very close to the cross-sectional NLS estimates for the one- and two-factor case. The three-factor models show larger differences. In particular, the estimated eigenvalues for the $\mathbb{A}_1(3)$ and $\mathbb{A}_2(3)$ produce a near-repeated pair of large eigenvalues instead of a complex pair and a relatively poor fit. This can be explained by the admissibility conditions

(a) 1-factor model

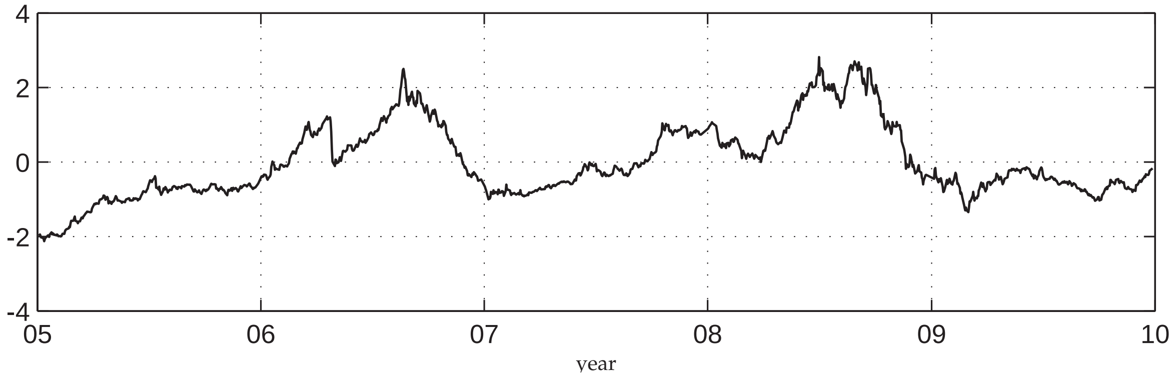

(b) 2-factor model

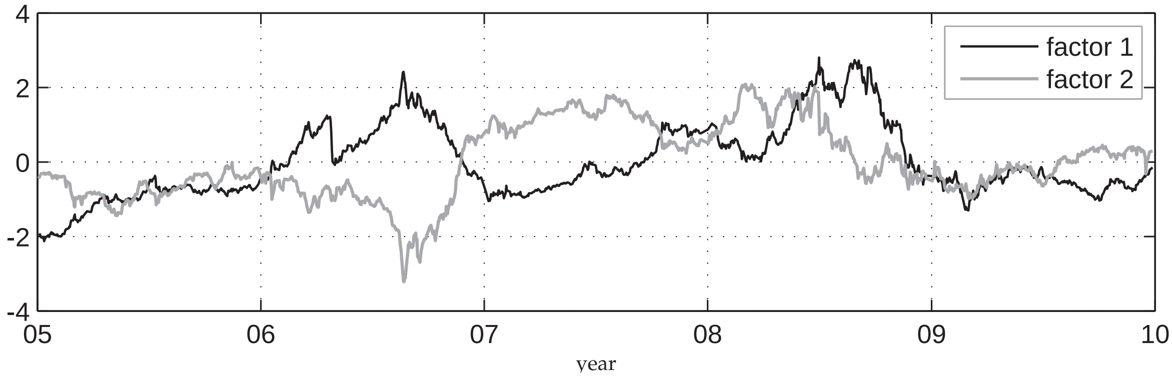

(c) 3-factor model

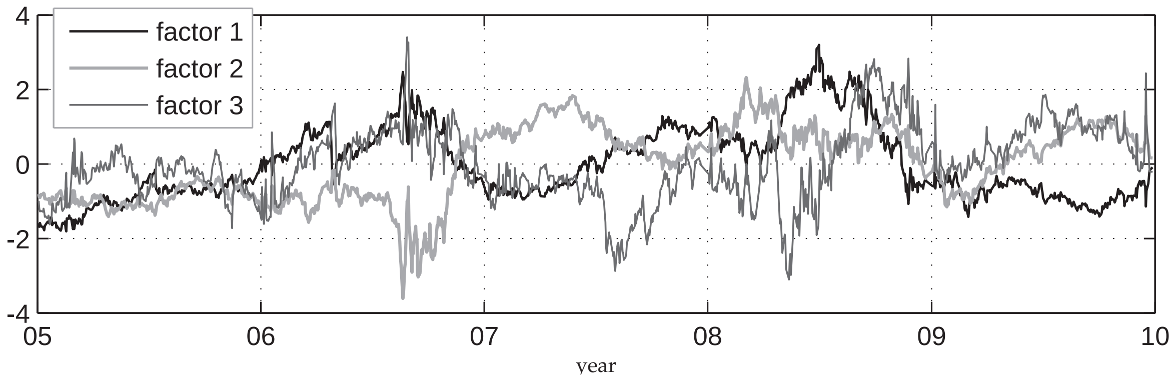

Figure 5: Time-Series of Estimated Factors Panels (a),(b) and (c) show the time series of the estimated factors for the one-, two- and three-factor models. The factors are represented as model-implied principal components $\widetilde{\boldsymbol{X}}_t$ as described in Section 2.3.

in (A.1) that also restrict the risk-neutral drift of the factor dynamics and hence the factor loadings. Still, the factor loadings are very similar when comparing between different factor dynamics and the cross-sectional estimates (see Appendix C for plots of the factor loadings).

4.3 Factor dynamics

Affine diffusions for the one-factor model are classified into two subclasses $\mathbb{A}_0(1)$ and $\mathbb{A}_1(1)$, corresponding to a Gaussian process and a CIR square-root process respectively. The key difference between these processes is that $\mathbb{A}_0(1)$ has constant volatility, whereas the volatility for $\mathbb{A}_1(1)$ varies with the level of the factor and hence produces conditional heteroscedasticity. For the one-factor model, Table 7 presents the parameter estimates of $\mathbb{A}_0(1)$ and $\mathbb{A}_1(1)$ with the corresponding (quasi) log-likelihood value.

The log-likelihood value for $\mathbb{A}_1(1)$ is substantially higher than for $\mathbb{A}_0(1)$, indicating a much better fit. Since the two models are not nested, we cannot apply standard likelihood ratio tests to infer whether the superior fit of $\mathbb{A}_0(1)$ is significant. Instead we use likelihood ratio tests for non-nested models as developed by Vuong (1989) to test the null hypothesis that two competing models are equally close to the true data generating process. Table 3 reports the tests outcomes. Under the null hypothesis, the test statistic has an asymptotic standard normal distribution. The test statistic for testing $\mathbb{A}_1(1)$ against $\mathbb{A}_0(1)$ equals 3.62, which implies that the superior fit of $\mathbb{A}_0(1)$ is highly significant.

For the two-factor case we have three subclasses $\mathbb{A}_0(2)$, $\mathbb{A}_1(2)$ and $\mathbb{A}_2(2)$. The $\mathbb{A}_0(2)$ is again equivalent to a Gaussian process, where the volatility is constant. At the other end of the spectrum, the model $\mathbb{A}_2(2)$ is a bivariate generalization of the CIR square-root process, also known as the correlated square-root (CSR) process (Dai and Singleton, 2000), where the volatility varies with the levels of the two factors. Admissibility conditions (A.1) on the parameters of $\mathbb{A}_2(2)$ restrict the correlation between the factors. Hence the empirical performance of $\mathbb{A}_2(2)$ versus $\mathbb{A}_0(2)$ will depend on a tradeoff between time-varying volatilities versus flexible correlations, respectively (Dai and Singleton, 2000). The $\mathbb{A}_1(2)$ can be viewed as an intermediate case. The volatilities in $\mathbb{A}_1(2)$ are time-varying, but restricted to vary only with the level of one particular linear combination of the factors. This yields more flexibility for the correlation between the factors.

The parameter estimates and corresponding log-likelihood values for the dynamics in the two-factor model are reported in Table 8. The log-likelihood values indicate a superior fit for models $\mathbb{A}_1(2)$ and $\mathbb{A}_2(2)$ with time-varying volatility over the $\mathbb{A}_0(2)$ with constant volatility. These results are confirmed by the non-nested likelihood ratio tests in Table 3, where the superior fit for $\mathbb{A}_1(2)$ and $\mathbb{A}_2(2)$ over $\mathbb{A}_0(2)$ is indeed significant. Moreover, the model with two correlated square-root diffusions, $\mathbb{A}_2(2)$ also significantly outperforms the $\mathbb{A}_1(2)$. Finally, we see that four parameters are estimated at their boundary values as imposed by the admissibility conditions (A.1).

Four subclasses for the factor dynamics can be distinguished in the three-factor model. Besides the Gaussian model $\mathbb{A}_0(3)$ and the CSR model $\mathbb{A}_3(3)$, we have two intermediate cases $\mathbb{A}_1(3)$ and $\mathbb{A}_2(3)$ with volatility driven by one and two factors, respectively. The parameter estimates and corresponding log-likelihood values for the 3-factor dynamics are reported in Table 9. The log-likelihood values indicate a superior fit of Gaussian model $\mathbb{A}_0(3)$ over all models with time-varying volatility. This difference in fit is indeed significant as indicated by the likelihood ratio in Table 3. Particularly the performance of $\mathbb{A}_1(3)$ and $\mathbb{A}_2(3)$ is weak, which is consistent with the relatively poor fit of the observed forward prices in Table 2. Moreover, Table 9 shows that the admissibility conditions have an important impact on parameter estimates of $\mathbb{A}_3(3)$, where numerous parameters have estimated values equal to their boundary values.

4.4 Forward premia

Forward prices reflect expectations about the average spot price over the delivery period of the contract, but also incorporate a forward premium to compensate for associated risks. Hence the forward price can be decomposed into an expectations part and a forward premium as

$$F_t(T_1, T_2) = \mathrm{E}_t^{\mathbb{P}}\left(\frac{1}{T_2 - T_1}\int_{T_1}^{T_2} S_s \,\mathrm{d}s\right) + fp_t(T_1, T_2).$$

Table 3: Likelihood ratio tests of factor specifications

The table reports the likelihood ratio tests for non-nested models by Vuong (1989) to test the null hypothesis that two competing models are equally close to the true data generating process. The test for two competing models A and B are indicated in the column *hypothesis* by "A vs. B". Under the null hypothesis, the test statistic has an asymptotic standard normal distribution. The test statistic and the corresponding p-value are reported in columns *stat* and *p-value*. A positive test statistic indicates a superior fit of model A over model B.

hypothesis	stat	p-value
$\mathbb{A}_1(1)$ vs. $\mathbb{A}_0(1)$	3.62	0.00
$\mathbb{A}_1(2)$ vs. $\mathbb{A}_0(2)$	1.22	0.11
$\mathbb{A}_2(2)$ vs. $\mathbb{A}_0(2)$	3.64	0.00
$\mathbb{A}_2(2)$ vs. $\mathbb{A}_1(2)$	2.76	0.00
$\mathbb{A}_1(3)$ vs. $\mathbb{A}_0(3)$	-2.87	1.00
$\mathbb{A}_2(3)$ vs. $\mathbb{A}_0(3)$	-2.78	1.00
$\mathbb{A}_3(3)$ vs. $\mathbb{A}_0(3)$	-1.89	0.97
$\mathbb{A}_2(3)$ vs. $\mathbb{A}_1(3)$	0.93	0.18
$\mathbb{A}_3(3)$ vs. $\mathbb{A}_1(3)$	3.45	0.00
$\mathbb{A}_3(3)$ vs. $\mathbb{A}_2(3)$	3.28	0.00

The expectations part represents the expected spot price over the delivery period under the objective measure $\mathbb{P}$ and is given by

$$\mathrm{E}_t^{\mathbb{P}}\left(\frac{1}{T_2-T_1}\int_{T_1}^{T_2} S_s \,\mathrm{d}s\right) = A^{\mathbb{P}}(t, T_1-t, T_2-t) + \boldsymbol{B}^{\mathbb{P}}(T_1-t, T_2-t)'\boldsymbol{X}_t,$$

where $A^{\mathbb{P}}(t,\tau_1,\tau_2)$ and $\boldsymbol{B}^{\mathbb{P}}(\tau_1,\tau_2)$ are obtained from (10) and (9) by replacing the risk-neutral parameters $\boldsymbol{c}^{\mathbb{Q}}$ and $\boldsymbol{D}^{\mathbb{Q}}$ by their objective counterparts $\boldsymbol{c}^{\mathbb{P}}$ and $\boldsymbol{D}^{\mathbb{P}}$, respectively. The forward premium and is given by

$$\begin{aligned} fp_t(T_1,T_2) = &\Big[A(t,T_1-t,T_2-t) - A^{\mathbb{P}}(t,T_1-t,T_2-t)\Big] \\ &+ \Big[\boldsymbol{B}(T_1-t,T_2-t) - \boldsymbol{B}^{\mathbb{P}}(T_1-t,T_2-t)\Big]'\boldsymbol{X}_t \end{aligned}$$

and is a linear function of the factors. Hence the forward premium is generally non-zero and time-varying. However, all forward premia are constant when the risk-neutral and objective mean-reversion matrices are identical, i.e. $\boldsymbol{D}^{\mathbb{P}} = \boldsymbol{D}^{\mathbb{Q}}$. If in addition $\boldsymbol{c}^{\mathbb{P}} = \boldsymbol{c}^{\mathbb{Q}}$, then all forward premia are zero.

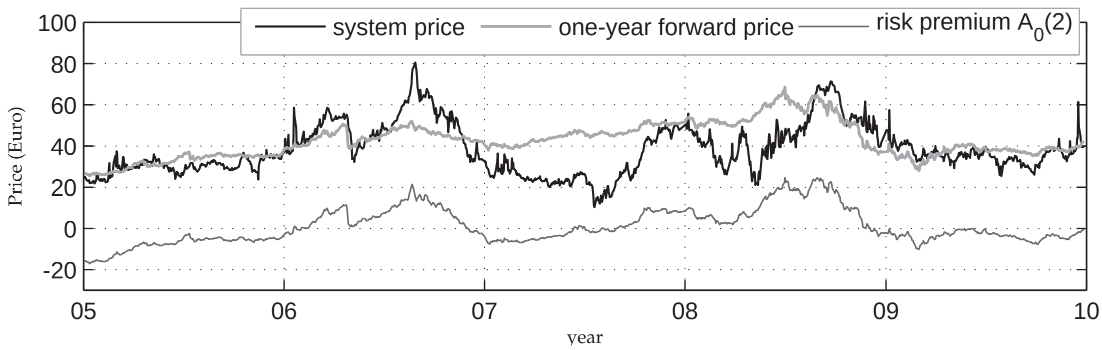

Figure 6: Time-Series of Forward Premium The figure shows the daily system price, one-year forward price together with the forward premium implied by estimated $\mathbb{A}_0(2)$ model from January 3, 2005 up to December 28, 2009 in Euros/MWh. The one-year forward price is interpolated by taking the weighted average of the traded one- and two-year forwards.

We calculate the forward premium on a one-year forward as implied by all estimated models. As an example, Figure 6 plots the forward premium implied by the estimates of the $\mathbb{A}_0(2)$ model together with the system price and the interpolated one-year forward price. The estimated forward premium is clearly time-varying and regularly changes sign, with a negative forward premium in 2005 and 2009 and a predominantly positive forward premium in 2006 and 2008. Moreover, the forward premium strongly comoves with the level of the system and forward prices.

Table 4 presents a comparison of the estimated risk premia between all models. The estimated forward premia differ in level and variability across different models, but show strong comovement as reflected by the high correlation. Correlations between the forward premia and the model-implied principal components reveal a strong comovement between the forward premium and the level of the forward curve; the forward premium is high when the forward curve is high.

Next, we test the hypotheses that *all* forward premia are constant or zero. Constant forward premia implies $\boldsymbol{D}^{\mathbb{P}} = \boldsymbol{D}^{\mathbb{Q}}$, whereas zero forward premia implies $\boldsymbol{D}^{\mathbb{P}} = \boldsymbol{D}^{\mathbb{Q}}$ and $\boldsymbol{c}^{\mathbb{P}} = \boldsymbol{c}^{\mathbb{Q}}$. We construct likelihood ratio tests by reestimating all models under these sets of restrictions. This testing approach has two important advantages. First, the hypotheses about forward premia are tested jointly on all forward prices. Secondly, it fully exploits the information in the panel of observed prices.

Table 4: Properties of the estimated 1-year forward premium

The table reports the properties of the forward premium on a one-year forward contract that starts delivering after one year. $\mathbb{A}_k(m)$ denotes the maximally flexible affine model of the corresponding subclass. The mean (*mean*) and standard deviation (*std.*) of the implied forward premium are expressed in Euro/MWh. The model-implied principal components are based on the cross-sectional NLS estimates of the 3-factor model and denoted by *level*, *slope* and *hump*.

	$\mathbb{A}_0(1)$	$\mathbb{A}_1(1)$	$\mathbb{A}_0(2)$	$\mathbb{A}_1(2)$	$\mathbb{A}_2(2)$	$\mathbb{A}_0(3)$	$\mathbb{A}_1(3)$	$\mathbb{A}_2(3)$	$\mathbb{A}_3(3)$
A. descriptive statistics									
mean	5.60	1.17	1.08	2.00	4.19	-7.96	3.37	9.27	4.56
std	5.11	6.35	8.62	7.07	4.05	6.81	7.81	5.40	6.24
B. correlation matrix									
$\mathbb{A}_0(1)$	1.00								
$\mathbb{A}_1(1)$	1.00	1.00							
$\mathbb{A}_0(2)$	1.00	1.00	1.00						
$\mathbb{A}_1(2)$	0.99	0.99	0.99	1.00					
$\mathbb{A}_2(2)$	0.97	0.97	0.97	0.92	1.00				
$\mathbb{A}_0(3)$	0.94	0.94	0.95	0.88	1.00	1.00			
$\mathbb{A}_1(3)$	0.96	0.96	0.97	0.91	1.00	1.00	1.00		
$\mathbb{A}_2(3)$	0.89	0.89	0.90	0.82	0.98	0.99	0.98	1.00	
$\mathbb{A}_3(3)$	0.97	0.97	0.97	0.92	1.00	0.99	1.00	0.97	1.00
C. correlation with model-implied principal components									
level	0.95	0.95	0.95	0.92	0.95	0.93	0.96	0.89	0.96
slope	0.03	0.03	0.05	-0.10	0.24	0.33	0.25	0.44	0.23
hump	0.31	0.30	0.29	0.37	0.17	0.16	0.15	0.10	0.15

Table 5 reports the results. For all three-factor models, we reject both hypotheses in favour of time-varying risk premia. For the one- and two-factor models, the evidence is more mixed. All in all however, the test results indicate the presence of time-varying risk premia for most models. This finding is consistent with Longstaff and Wang (2004), who find significant, time-varying forward premia in short-maturity electricity forward prices of the Pennsylvania, New Jersey, and Maryland (PJM) electricity market.

Table 5: Test of forward premia

The table reports per model the likelihood ratio tests for the *constant forward premia* hypothesis and the *zero forward premia* hypothesis. Each model is reestimated under the corresponding hypothesis and its log-likelihood value is compared against the unrestricted model. The column *stat* reports the likelihood ratio test statistic given by $2(\log L_{unrest} - \log L_{rest})$, where L_{unrest} and L_{rest} denote the likelihood value of the unrestricted and restricted model, respectively. Under the null hypothesis, the test statistic follows an asymptotic χ^2_{df} distribution, with reported degrees of freedom (*df*) and the corresponding p-value (*p-value*). $\mathbb{A}_k(m)$ denotes the maximally flexible affine model of the corresponding subclass.

	constant forward premia $\boldsymbol{D}^{\mathbb{P}} = \boldsymbol{D}^{\mathbb{Q}}$			zero forward premia $\boldsymbol{D}^{\mathbb{P}} = \boldsymbol{D}^{\mathbb{Q}}$ and $\boldsymbol{c}^{\mathbb{P}} = \boldsymbol{c}^{\mathbb{Q}}$		
	stat	df	p-value	stat	df	p-value
$\mathbb{A}_0(1)$	1.56	1	0.21	2.21	2	0.33
$\mathbb{A}_1(1)$	7.14	1	0.01	7.90	2	0.02
$\mathbb{A}_0(2)$	10.57	4	0.03	10.98	6	0.09
$\mathbb{A}_1(2)$	9.14	3	0.03	9.47	5	0.09
$\mathbb{A}_2(2)$	3.79	4	0.43	3.79	5	0.58
$\mathbb{A}_0(3)$	38.87	9	0.00	44.08	12	0.00
$\mathbb{A}_1(3)$	27.07	7	0.00	36.42	10	0.00
$\mathbb{A}_2(3)$	19.24	7	0.01	19.42	10	0.04
$\mathbb{A}_3(3)$	42.00	9	0.00	34.38	12	0.00

5 Option Pricing

Models of electricity prices are widely used for pricing and hedging of electricity derivatives. Beyond forwards and futures a wide variety of other electricity contracts such as options are traded via the exchange or OTC. The Nordpool market offers trading in standardized options on electricity forwards. Our modelling framework generates tractable expressions for option prices by using standard affine factor dynamics. In particular, we exploit the results of Duffie *et al.* (2000) to obtain quasi-analytical pricing formulae for European options on forwards.

Arbitrage-free prices of European options are obtain by applying the pricing equation (1) to the option payoff. Option prices can be expressed as simple functions of the transform

$$H_{\boldsymbol{a},b_0,\boldsymbol{b}_1,\boldsymbol{c},\boldsymbol{X}_t,t,T}(y) = \mathrm{E}_t^{\mathbb{Q}}\left(e^{\boldsymbol{a}'\boldsymbol{X}_T}(b_0 + \boldsymbol{b}_1'\boldsymbol{X}_T)I_{\{\boldsymbol{c}'\boldsymbol{X}_T \leq y\}}\right), \tag{20}$$

where I_A denotes the indicator function. For an affine process $\boldsymbol{X}_t$, Duffie *et al.* (2000) show that the Fourier transform of (20) is know in closed form and given by

$$\mathcal{H}_{\boldsymbol{a},b_0,\boldsymbol{b}_1,\boldsymbol{c},\boldsymbol{X}_t,t,T}(u) \equiv \int_{-\infty}^{\infty} e^{iuy}\,\mathrm{d}H_{\boldsymbol{a},b_0,\boldsymbol{b}_1,\boldsymbol{c},\boldsymbol{X}_t,t,T}(y) = \Phi(\boldsymbol{a} + iu\boldsymbol{c}, b_0, \boldsymbol{b}_1, \boldsymbol{X}_t, t, T),$$

with $i^2 = -1$ and

$$\Phi(\boldsymbol{u}, v_0, \boldsymbol{v}_1, \boldsymbol{X}_t, t, T) \equiv \mathrm{E}_t^{\mathbb{Q}}\left(e^{\boldsymbol{u}'\boldsymbol{X}_T}(v_0 + \boldsymbol{v}_1'\boldsymbol{X}_T)\right) = e^{a_\Phi(t) + \boldsymbol{b}_\Phi(t)'\boldsymbol{x}}\left(c_\Phi(t) + \boldsymbol{d}_\Phi(t)'\boldsymbol{x}\right),$$

where the functions $a_\Phi(t)$, $\boldsymbol{b}_\Phi(t)$, $c_\Phi(t)$ and $\boldsymbol{d}_\Phi(t)$ solve a system of ODEs given in Appendix D. Hence applying the Fourier inversion formula gives

$$\begin{aligned} H_{\boldsymbol{a},b_0,\boldsymbol{b}_1,\boldsymbol{c},\boldsymbol{X}_t,t,T}(y) = {} & \frac{\Phi(\boldsymbol{a}, b_0, \boldsymbol{b}_1, \boldsymbol{X}_t, t, T)}{2} \\ & - \frac{1}{\pi}\int_0^\infty \frac{\mathrm{Im}[\Phi(\boldsymbol{a} + iv\boldsymbol{c}, b_0, \boldsymbol{b}_1, \boldsymbol{X}_t, t, T)e^{-ivy}]}{v}\,\mathrm{d}v, \end{aligned} \tag{21}$$

where $\mathrm{Im}(x)$ denotes the complex part of $x \in \mathbb{C}$. The evaluation of (21) just requires a numerical integration over one dimension.

The price of a European call and put option at time t with expiry date S on a forward that delivers over $[T_1, T_2]$, with $t \leq S \leq T_1 \leq T_2$, are now given by

$$\begin{aligned} \mathcal{C}_t &= e^{-r(S-t)}\,\mathrm{E}_t^{\mathbb{Q}}\left([F_S(T_1,T_2) - K]^+\right) = e^{-r(S-t)} H_{\boldsymbol{0},A-K,\boldsymbol{B}',-\boldsymbol{B}',\boldsymbol{X}_t,t,S}(A - K), \\ \mathcal{P}_t &= e^{-r(S-t)}\,\mathrm{E}_t^{\mathbb{Q}}\left([K - F_S(T_1,T_2)]^+\right) = e^{-r(S-t)} H_{\boldsymbol{0},K-A,-\boldsymbol{B}',\boldsymbol{B}',\boldsymbol{X}_t,t,S}(K - A), \end{aligned}$$

respectively.

6 Conclusion

This paper presents a tractable class of arbitrage-free models for the term structure of electricity prices where forward prices are linear functions of the factors. The class offers much flexibility in the specification of the factor dynamics. By adopting well-known processes such as affine processes for the specification of the factors, we obtain additional tractability for estimation and option pricing.

Empirical results for daily forward prices of the Nordpool market show that forward prices can be adequately modeled with three factors. Changes in the level, slope and curvature are identified as the most important sources of fluctuations in the forward curve of electricity prices. We examine the ability of affine factor processes to describe the dynamics in electricity prices and find that affine factor dynamics that allows for time-varying volatilities fit significantly better than the Gaussian dynamics for the one- and two-factor models. For the three-factor models, we find the opposite. This finding is consistent with the conditional volatility- correlation flexibility trade-off by Dai and Singleton (2000) as well as a flexibility trade-off between volatility and factor loadings.

Our class of models provides a tractable and flexible framework for pricing, hedging and managing risks related to electricity prices. The class imposes only mild conditions on the factor dynamics, thereby offering a great deal of flexibility. Extra tractability is gained by using affine factor dynamics. In particular, the results of Duffie *et al.* (2000) can be used to obtain quasi-analytical prices for common types of electricity derivatives such as e.g. options on forwards. Another important advantage of the framework is allows for a straightforward specification of models featuring unspanned stochastic volatility, i.e. volatility that is not completely spanned by forward prices, since prices of forwards and futures do not depend on the factor volatilities. Trolle and Schwartz (2009) find strong evidence for unspanned volatility in the related commodity market of NYMEX crude oil derivatives. Whether this is the case for electricity markets remains an open question.

References

Andreasen, M. M., & Christensen, B. J. (2010). The SR approach: A new estimation method for non-linear and non-gaussian dynamic term structure models. *CREATES Research Paper 2010-12*.

Benth, F. E., Kallsen, J., & Meyer-Brandis, T. (2007). A non-gaussian Ornstein-Uhlenbeck process for electricity spot price modeling and derivatives pricing. *Applied Mathematical Finance*, *14*(2), 153 - 169.

Benth, F. E., & Koekebakker, S. (2008). Stochastic modeling of financial electricity contracts. *Energy Economics*, *30*(3), 1116 - 1157.

Bjerksund, P., Rasmussen, H., & Stensland, G. (2010). Valuation and risk management in the Norwegian electricity market. In *Energy, natural resources and environmental economics.* Springer.

Björk, T. (2004). *Arbitrage theory in continuous time* (2nd ed.). Oxford, UK: Oxford University Press.

Cheridito, P., Filipovic, D., & Kimmel, R. L. (2007). Market price of risk specifications for affine models: Theory and evidence. *Journal of Financial Economics*, *83*, 123–170.

Clewlow, L., & Strickland, C. (1999). Valuing energy options in a one factor model fitted to forward prices. *Working Paper, Univeristy of Technology Sydney*.

Cox, J. C., Ingersoll, J. E., Jr, & Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, *53*, 385–408.

Dai, Q., & Singleton, K. (2000). Specification analysis of affine term structure models. *Journal of Finance*, *55*, 1943–1978.

Diebold, F. X., & Li, C. (2006). Forecasting the term structure of government bond yields. *Journal of Econometrics*, *130*, 337–364.

Duffie, D., & Kan, R. (1996). A yield-factor model of interest rates. *Mathematical Finance*, *6*, 379–406.

Duffie, D., Pan, J., & Singleton, K. (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, *68*(6), 1343–1376.

Fackler, P. L. (2000). Moments of affine diffusions. *Working Paper, North Carolina State University*.

Fisher, M., & Gilles, C. (1996). Estimating exponential-affine models of the term structure. *Working paper, Federal Reserve Bank of Atlanta*.

Gabaix, X. (2009). Linearity-generating processes: A modelling tool yielding closed forms for asset prices. *Working paper, New York University*.

Greene, W. H. (2003). *Econometric analysis* (5th ed.). New Jersey: Prentice Hall.

Hinz, J., & Wilhelm, M. (2006). Pricing flow commodity derivatives using fixed income market techniques. *International Journal of Theoretical & Applied Finance*, *9*(8), 1299 - 1321.

Joslin, S., Singleton, K. J., & Zhu, H. (2011). A new perspective on Gaussian dynamic term structure models. *Review of Financial Studies*, *24*, 926–970.

Knittel, C. R., & Roberts, M. R. (2005). An empirical examination of restructured electricity prices. *Energy Economics*, *27*(5), 791 - 817.

Koekebakker, S. (2007). *An arithmetic forward curve for the electricity market.*

Longstaff, F. A., & Wang, A. W. (2004). Electricity forward prices: A high-frequency empirical analysis. *The Journal of Finance*, *59*(4), 1877–1900.

Lucia, J. J., & Schwartz, E. S. (2002). Electricity prices and power derivatives: Evidence from the nordic power exchange. *Review of Derivatives Research*, *5*(1), 5–50.

Meyer-Brandis, T., & Tankov, P. (2008). Multi-factor jump-diffusion models of electricity prices. *International Journal of Theoretical & Applied Finance*, *11*(5), 503 - 528.

Monfort, A., & Feŕon, O. (2012). Joint econometric modeling of spot electricity prices, forwards and options. *Review of Derivatives Research*, *forthcoming*.

Schwartz, E., & Smith, J. E. (2000). Short-term variations and long-term dynamics in commodity prices. *Management Science*, *46*, 893–911.

Trolle, A. B., & Schwartz, E. S. (2009). Unspanned stochastic volatility and the pricing of commodity derivatives. *Review of Financial Studies*, *22*(11), 4423-4461.

Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, *57*(2), 307–333.

White, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica*, *50*, 1–25.

A Admissibility conditions

The process $\boldsymbol{Y}_t$ follows a canonical representation of the maximally flexible by assuming that $\boldsymbol{H}_{\boldsymbol{Y}}$ lower triangular if $k = 0$ and

$$\begin{aligned}
c^{\mathbb{P}}_{\boldsymbol{Y},i}, c^{\mathbb{Q}}_{\boldsymbol{Y},i} \geq \frac{1}{2}, \quad & \alpha_{\boldsymbol{Y},i} = 0, \quad B_{\boldsymbol{Y},ii} = 1, && i = 1,\dots,k, \\
D^{\mathbb{P}}_{\boldsymbol{Y},ij}, D^{\mathbb{Q}}_{\boldsymbol{Y},ij} \geq 0, \quad & B_{\boldsymbol{Y},ij} = 0, && j = 1,\dots,k, \quad j \neq i, \\
D^{\mathbb{P}}_{\boldsymbol{Y},ij}, D^{\mathbb{Q}}_{\boldsymbol{Y},ij} = 0, \quad & B_{\boldsymbol{Y},ji} \geq 0, \quad \alpha_{\boldsymbol{Y},j} = 1, && j = k+1,\dots,m, \\
B_{\boldsymbol{Y},ij} = 0 \quad & && i = 1,\dots,m, \quad j = k+1,\dots,m.
\end{aligned} \tag{A.1}$$

These restrictions ensure admissibility and identification of the process $\boldsymbol{Y}_t$ with well-behaved risk premia (Cheredito *et al.*, 2007).[5]

A.1 Effect on factor loadings

We now demonstrate how affine factor dynamics that allow for stochastic volatility can restrict the factor loadings for forwards. As discussed in Section 2.2, the factor loadings for forwards in (6) and (9) are determined by the eigenvalue structure of the risk-neutral mean-reversion matrix $\boldsymbol{D}^{\mathbb{Q}}$ and hence that of $\boldsymbol{D}^{\mathbb{Q}}_{\boldsymbol{Y}}$. The admissibility conditions for affine diffusions in (A.1) however impose conditions on $\boldsymbol{D}^{\mathbb{Q}}$ that can restrict its eigenvalue structure.

For example the $\mathbb{A}_2(2)$ model requires that $D^{\mathbb{Q}}_{\boldsymbol{Y},12}, D^{\mathbb{Q}}_{\boldsymbol{Y},21} \geq 0$ in

$$\boldsymbol{D}^{\mathbb{Q}}_{\boldsymbol{Y}} = \begin{pmatrix} D^{\mathbb{Q}}_{\boldsymbol{Y},11} & D^{\mathbb{Q}}_{\boldsymbol{Y},12} \\ D^{\mathbb{Q}}_{\boldsymbol{Y},21} & D^{\mathbb{Q}}_{\boldsymbol{Y},22} \end{pmatrix}.$$

Using the quadratic formula, the eigenvalues of $\boldsymbol{D}^{\mathbb{Q}}_{\boldsymbol{Y}}$ are given by

$$\lambda_i = \tfrac{1}{2}(D^{\mathbb{Q}}_{\boldsymbol{Y},11} + D^{\mathbb{Q}}_{\boldsymbol{Y},22}) \pm \tfrac{1}{2}\sqrt{(D^{\mathbb{Q}}_{\boldsymbol{Y},11} - D^{\mathbb{Q}}_{\boldsymbol{Y},22})^2 + 4D^{\mathbb{Q}}_{\boldsymbol{Y},12}D^{\mathbb{Q}}_{\boldsymbol{Y},21}},$$

[5] Our canonical characterization is almost identical to that in Cheridito *et al.* (2007), but for $k = 0$ we assume $\boldsymbol{H}_{\boldsymbol{Y}}$ lower triangular rather than $\boldsymbol{D}^{\mathbb{P}}_{\boldsymbol{Y}}$. This still yields a well-identified model, but no longer requires the assumption that $\boldsymbol{D}^{\mathbb{P}}$ is diagonalizable.

which are restricted to be real since $D^{\mathbb{Q}}_{\boldsymbol{Y},12}, D^{\mathbb{Q}}_{\boldsymbol{Y},21} \geq 0$. By similar reasoning it follows that e.g the model $\mathbb{A}_2(3)$ does not allow for complex eigenvalues.

B Estimation of $\boldsymbol{\beta}$ and $\boldsymbol{X}_t$

Using (17), we can represent (18) as the following regression equation

$$\boldsymbol{F}_t = \boldsymbol{Q}_t\boldsymbol{\beta} + \boldsymbol{B}_t(\boldsymbol{\lambda}^{\mathbb{Q}})\boldsymbol{X}_t + \boldsymbol{\varepsilon}_t, \qquad t = 1, \ldots, T \tag{B.1}$$

where $\boldsymbol{Q}_t$ is the matrix of monthly dummies corresponding to the contracts in $\boldsymbol{F}_t$. For a given $\boldsymbol{\lambda}^{\mathbb{Q}}$, the estimates of $\boldsymbol{\beta}$ and $\boldsymbol{X}_t$ defined in (19) are equivalent to OLS estimates $\hat{\boldsymbol{\beta}}(\boldsymbol{\lambda}^{\mathbb{Q}})$ and $\hat{\boldsymbol{X}}_t(\boldsymbol{\lambda}^{\mathbb{Q}})$ of the regression given by (B.1). Hence $\boldsymbol{\beta}$ and $\boldsymbol{X}_t$ can be concentrated out of the criterion function in (19) by plugging in the OLS solutions. In particular, the estimate of $\boldsymbol{\lambda}^{\mathbb{Q}}$ defined in (19) is given by

$$\hat{\boldsymbol{\lambda}}_{\mathbb{Q}} = \text{argmin} \sum_{t=1}^{T} \boldsymbol{e}_t^*(\boldsymbol{\lambda}^{\mathbb{Q}})'\boldsymbol{e}_t(\boldsymbol{\lambda}^{\mathbb{Q}}), \tag{B.2}$$

$\boldsymbol{e}_t^*(\boldsymbol{\lambda}^{\mathbb{Q}}) = \boldsymbol{F}_t - \boldsymbol{Q}_t\hat{\boldsymbol{\beta}}(\boldsymbol{\lambda}^{\mathbb{Q}}) + \boldsymbol{B}_t(\boldsymbol{\lambda}^{\mathbb{Q}})\hat{\boldsymbol{X}}_t(\boldsymbol{\lambda}^{\mathbb{Q}})$. The corresponding estimates of $\boldsymbol{\beta}$ and $\boldsymbol{X}_t$ are obtained by $\hat{\boldsymbol{\beta}}(\hat{\boldsymbol{\lambda}}^{\mathbb{Q}})$ and $\hat{\boldsymbol{X}}_t(\hat{\boldsymbol{\lambda}}_{\mathbb{Q}})$.

The regression in (B.1) will be large for typical sample sizes as it includes many regressors. The calculation of the OLS estimates for $\boldsymbol{\beta}$ and $\boldsymbol{X}_1, \ldots, \boldsymbol{X}_T$ can be simplified considerable by using the Frisch-Waugh theorem (see e.g. Greene (2003)) and reduces to running a series of sequential small-scale regressions. The procedure is given by the following steps:

1. Regress $\boldsymbol{F}_t$ and $\boldsymbol{Q}_t$ on $\boldsymbol{B}_t$ for all $t = 1, \ldots, T$ to obtain the residuals

$$\boldsymbol{v}_t = \boldsymbol{F}_t - \boldsymbol{B}_t \left(\boldsymbol{B}_t{}'\boldsymbol{B}_t\right)^{-1} \boldsymbol{B}_t{}'\boldsymbol{F}_t \qquad\qquad \boldsymbol{W}_t = \boldsymbol{Q}_t - \boldsymbol{B}_t \left(\boldsymbol{B}_t{}'\boldsymbol{B}_t\right)^{-1} \boldsymbol{B}_t{}'\boldsymbol{Q}_t$$

   and stack them as $\bar{\boldsymbol{v}} = (\boldsymbol{v}_1', \ldots, \boldsymbol{v}_T')'$ and $\bar{\boldsymbol{W}} = (\boldsymbol{W}_1', \ldots, \boldsymbol{W}_T')'$.

2. Regress the stacked residuals $\bar{\boldsymbol{v}}$ on the stacked residuals $\bar{\boldsymbol{W}}$ to obtain the OLS estimate

$\hat{\boldsymbol{\beta}}(\boldsymbol{\lambda}^{\mathbb{Q}})$ and the residuals $\boldsymbol{e}_t^*(\boldsymbol{\lambda}^{\mathbb{Q}})$:

$$\hat{\boldsymbol{\beta}}(\boldsymbol{\lambda}^{\mathbb{Q}}) = \left(\bar{\boldsymbol{W}}'\bar{\boldsymbol{W}}\right)^{-1}\bar{\boldsymbol{W}}'\bar{\boldsymbol{v}} \qquad\qquad \boldsymbol{e}_t^*(\boldsymbol{\lambda}^{\mathbb{Q}}) = \boldsymbol{v}_t - \boldsymbol{W}_t\hat{\boldsymbol{\beta}}(\boldsymbol{\lambda}^{\mathbb{Q}}).$$

3. (optional) Regress $\boldsymbol{F}_t - \boldsymbol{Q}_t\hat{\boldsymbol{\beta}}(\boldsymbol{\lambda}^{\mathbb{Q}})$ on $\boldsymbol{B}_t$ to estimate the factors $\boldsymbol{X}_t$:

$$\hat{\boldsymbol{X}}_t(\boldsymbol{\lambda}^{\mathbb{Q}}) = \left(\boldsymbol{B}_t{}'\boldsymbol{B}_t\right)^{-1}\boldsymbol{B}_t{}'\left(\boldsymbol{F}_t - \boldsymbol{Q}_t\hat{\boldsymbol{\beta}}(\boldsymbol{\lambda}^{\mathbb{Q}})\right).$$

C Estimation results

The estimates of the cross-section parameters $\boldsymbol{\lambda}^{\mathbb{Q}}$ and σ_ε for all models are reported in Table 2. Table 6 reports the complementary estimates of the coefficients for the seasonal dummies $\boldsymbol{\beta}$. Figures 7, 8 and 9 plot the corresponding factor loadings for all models. The dynamic parameters and the quasi log-likelihood values for the one-, two- and three-factor models are reported in Tables 7, 8 and 9.

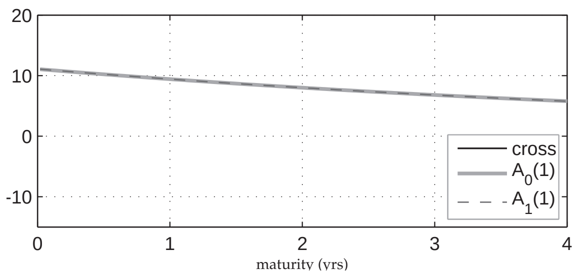

Figure 7: Factor loadings 1-factor models Panels (a),(b) and (c) plot the estimated factor loadings for the one-, two- and three-factor models. The factors are represented as model-implied principal components $\widetilde{\boldsymbol{X}}_t$ as described in Section 2.3.

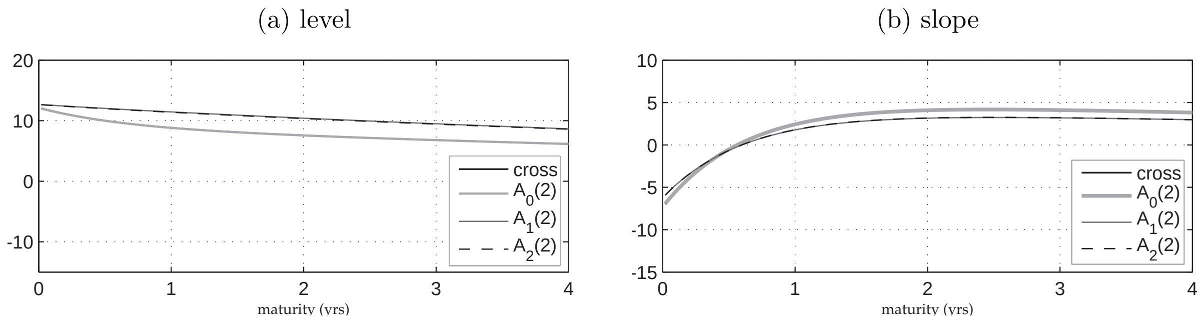

Figure 8: Factor loadings 2-factor models Panels (a),(b) and (c) plot the estimated factor loadings for the one-, two- and three-factor models. The factors are represented as model-implied principal components $\widetilde{\boldsymbol{X}}_t$ as described in Section 2.3.

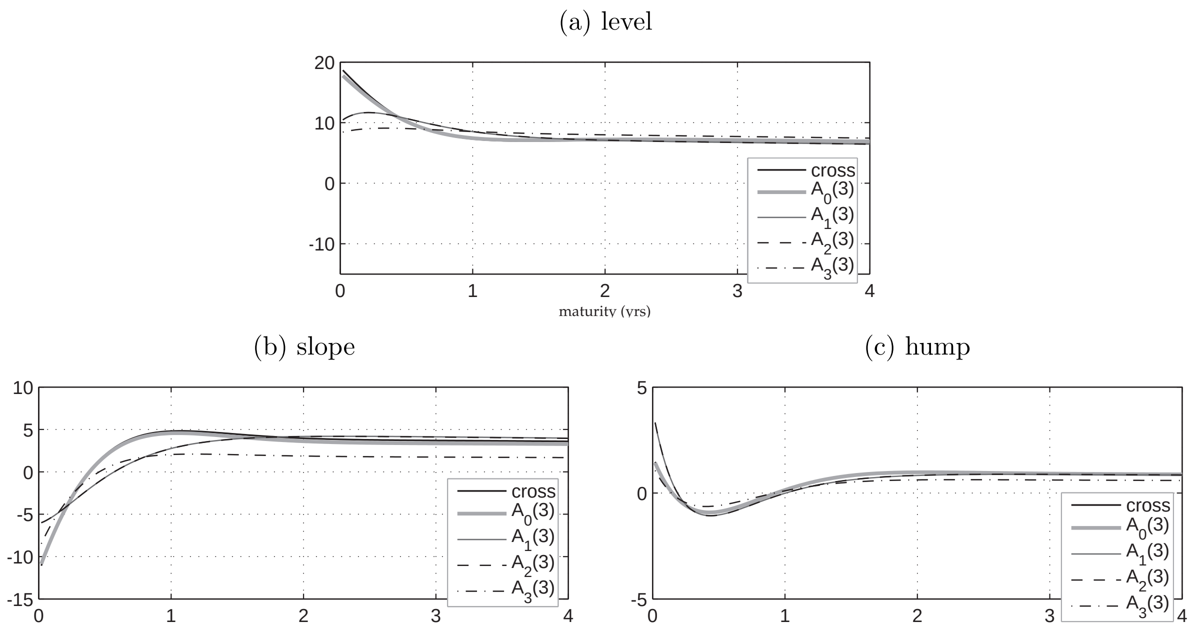

Figure 9: Factor loadings 3-factor models Panels (a),(b) and (c) plot the estimated factor loadings for the one-, two- and three-factor models. The factors are represented as model-implied principal components $\widetilde{\boldsymbol{X}}_t$ as described in Section 2.3.

Table 6: Estimates of Seasonal Dummies $\boldsymbol{\beta}$

The table reports the estimates of the cross-sectional parameters $\boldsymbol{\beta}$ for models with one, two and three factors. Columns denoted by *cross* report cross-sectional NLS estimates, whereas columns denoted by $\mathbb{A}_k(n)$ report QML estimates based on the corresponding affine factor dynamics. Asymptotic standard errors are shown in parenthesis. Estimates of $\boldsymbol{\lambda}^{\mathbb{Q}}$ are expressed on a yearly basis, i.e. maturities τ are represented in years.

	1-factor			2-factor				3- factor				
	cross	$\mathbb{A}_0(1)$	$\mathbb{A}_1(1)$	cross	$\mathbb{A}_0(2)$	$\mathbb{A}_1(2)$	$\mathbb{A}_2(2)$	cross	$\mathbb{A}_0(3)$	$\mathbb{A}_1(3)$	$\mathbb{A}_2(3)$	$\mathbb{A}_3(3)$
β_{jan}	57.25 (0.16)	57.32 (0.26)	57.26 (0.53)	52.71 (0.09)	52.70 (0.17)	52.72 (0.52)	52.71 (0.51)	70.97 (0.04)	81.60 (0.86)	72.09 (2.86)	71.09 (2.78)	76.45 (2.97)
β_{feb}	56.43 (0.15)	56.49 (0.27)	56.42 (0.51)	52.18 (0.07)	52.16 (0.16)	52.18 (0.52)	52.17 (0.51)	70.67 (0.04)	81.29 (0.86)	71.81 (2.78)	70.79 (2.79)	76.14 (2.97)
β_{mar}	52.22 (0.15)	52.27 (0.32)	52.21 (0.49)	48.05 (0.07)	48.04 (0.16)	48.06 (0.53)	48.05 (0.52)	66.55 (0.04)	77.27 (0.87)	67.67 (2.87)	66.68 (2.79)	72.09 (2.97)
β_{apr}	49.84 (0.15)	49.90 (0.26)	49.84 (0.51)	45.61 (0.07)	45.60 (0.17)	45.62 (0.53)	45.61 (0.52)	64.03 (0.04)	74.78 (0.87)	65.16 (2.88)	64.16 (2.78)	69.59 (2.96)
β_{may}	46.57 (0.14)	46.62 (0.23)	46.57 (0.49)	42.55 (0.06)	42.53 (0.16)	42.55 (0.54)	42.54 (0.52)	61.10 (0.04)	71.82 (0.85)	62.22 (2.86)	61.22 (2.80)	66.63 (2.98)
β_{jun}	45.82 (0.15)	45.87 (0.28)	45.81 (0.47)	41.72 (0.07)	41.71 (0.16)	41.72 (0.54)	41.72 (0.53)	60.30 (0.04)	71.03 (0.86)	61.42 (2.86)	60.42 (2.80)	65.83 (2.98)
β_{jul}	43.56 (0.16)	43.62 (0.22)	43.56 (0.48)	39.64 (0.07)	39.63 (0.15)	39.65 (0.54)	39.64 (0.52)	58.23 (0.04)	68.96 (0.87)	59.35 (2.87)	58.36 (2.79)	63.77 (2.97)
β_{aug}	46.57 (0.20)	46.63 (0.24)	46.57 (0.47)	42.29 (0.09)	42.28 (0.15)	42.30 (0.55)	42.30 (0.54)	61.02 (0.05)	71.77 (0.86)	62.14 (2.87)	61.14 (2.79)	66.57 (2.97)
β_{sep}	48.81 (0.17)	48.87 (0.33)	48.81 (0.50)	44.10 (0.08)	44.09 (0.15)	44.11 (0.56)	44.11 (0.55)	62.67 (0.05)	73.45 (0.86)	63.79 (2.87)	62.79 (2.80)	68.22 (2.98)
β_{oct}	49.23 (0.16)	49.29 (0.32)	49.23 (0.49)	44.57 (0.08)	44.56 (0.13)	44.58 (0.58)	44.58 (0.57)	63.16 (0.05)	73.96 (0.86)	64.29 (2.88)	63.30 (2.79)	68.73 (2.98)
β_{nov}	52.69 (0.17)	52.75 (0.30)	52.69 (0.52)	48.10 (0.08)	48.08 (0.15)	48.11 (0.55)	48.10 (0.54)	66.41 (0.05)	77.15 (0.86)	67.53 (2.87)	66.54 (2.79)	71.95 (2.98)
β_{dec}	54.81 (0.16)	54.87 (0.30)	54.81 (0.51)	50.52 (0.08)	50.50 (0.15)	50.52 (0.53)	50.52 (0.51)	68.76 (0.04)	79.45 (0.87)	69.88 (2.87)	68.89 (2.79)	74.27 (2.97)

Table 7: Parameter Estimates of the 1-Factor Dynamics

The table reports the estimates of the parameters of affine factor dynamics for the 1-factor model. The models, denoted by $\mathbb{A}_k(1)$ for $k = 0, 1$ denote the maximally flexible model of the corresponding subclass in the the classification of Dai and Singleton (2000). Parameter estimates and corresponding asymptotic standard errors are reported in the column denoted *estim.* and *s.e.*, respectively. The (quasi) log-likelihood value is reported in the row denoted by *log L*.

	$\mathbb{A}_0(1)$		$\mathbb{A}_1(1)$	
	estim.	s.e.	estim.	s.e.
$c_{\boldsymbol{Y}}^{\mathbb{Q}}$	0.00	-	0.50	0.00
$D_{\boldsymbol{Y}}^{\mathbb{Q}}$	-0.16	0.00	-0.16	0.00
$\delta_{\boldsymbol{Y}}$	15.15	0.50	10.33	0.18
$c_{\boldsymbol{Y}}^{\mathbb{P}}$	-0.75	0.66	2.41	0.82
$D_{\boldsymbol{Y}}^{\mathbb{P}}$	-0.77	0.59	-1.07	0.33
Log L	-73882.43		-73835.68	

Table 8: Parameter Estimates of the 2-Factor Dynamics

The table reports the estimates of the parameters of affine factor dynamics for the 2-factor model. The models, denoted by $\mathbb{A}_k(2)$ for $k = 0, 1, 2$ denote the maximally flexible model of the corresponding subclass in the the classification of Dai and Singleton (2000). Parameter estimates and corresponding asymptotic standard errors are reported in the column denoted *estim.* and *s.e.*, respectively. The (quasi) log-likelihood value is reported in the row denoted by *log L*.

	$\mathbb{A}_0(2)$		$\mathbb{A}_1(2)$		$\mathbb{A}_2(2)$	
	estim.	s.e.	estim.	s.e.	estim.	s.e.
$c_{\boldsymbol{Y},1}^{\mathbb{Q}}$	0.00	-	7.84	2.88	0.50	0.00
$c_{\boldsymbol{Y},2}^{\mathbb{Q}}$	0.00	-	0.00	-	0.50	0.00
$D_{\boldsymbol{Y},11}^{\mathbb{Q}}$	-0.09	0.00	-1.58	0.05	-0.30	0.04
$D_{\boldsymbol{Y},21}^{\mathbb{Q}}$	0.15	0.09	-0.67	0.27	0.36	0.04
$D_{\boldsymbol{Y},12}^{\mathbb{Q}}$	0.00	0.00	0.00	-	0.72	0.12
$D_{\boldsymbol{Y},22}^{\mathbb{Q}}$	-1.58	0.01	-0.09	0.01	-1.37	0.14
$\delta_{\boldsymbol{Y},1}$	10.47	1.14	9.14	2.00	0.63	0.16
$\delta_{\boldsymbol{Y},2}$	21.24	0.84	1.62	0.25	16.42	1.10
$c_{\boldsymbol{Y},1}^{\mathbb{P}}$	-0.52	0.62	6.41	3.77	4.07	6.42
$c_{\boldsymbol{Y},2}^{\mathbb{P}}$	0.34	0.68	7.22	5.10	0.50	0.00
$D_{\boldsymbol{Y},11}^{\mathbb{P}}$	-1.11	0.64	-1.26	0.51	-0.68	0.75
$D_{\boldsymbol{Y},21}^{\mathbb{P}}$	0.78	0.79	-9.81	2.66	0.43	0.13
$D_{\boldsymbol{Y},12}^{\mathbb{P}}$	-2.75	1.23	0.00	-	0.00	0.00
$D_{\boldsymbol{Y},22}^{\mathbb{P}}$	-1.00	1.08	-1.04	0.46	-1.79	0.50
β_{21}	0.00	-	12.26	4.06	0.00	-
Log L	-57489.01		-57461.18		-57407.47	

Table 9: Parameter Estimates of the 3-Factor Dynamics

The table reports the estimates of the parameters of affine factor dynamics for the 3-factor model. The models, denoted by $\mathbb{A}_k(3)$ for $k = 0, 1, 2, 3$ denote the maximally flexible model of the corresponding subclass in the the classification of Dai and Singleton (2000). Parameter estimates and corresponding asymptotic standard errors are reported in the column denoted *estim.* and *s.e.*, respectively. The (quasi) log-likelihood value is reported in the row denoted by *log L*.

	$\mathbb{A}_0(3)$		$\mathbb{A}_1(3)$		$\mathbb{A}_2(3)$		$\mathbb{A}_3(3)$	
	estim.	s.e.	estim.	s.e.	estim.	s.e.	estim.	s.e.
$c^{\mathbb{Q}}_{\boldsymbol{Y},1}$	0.00	-	1.08	0.42	0.50	0.00	0.50	0.00
$c^{\mathbb{Q}}_{\boldsymbol{Y},2}$	0.00	-	0.00	-	0.50	0.00	0.50	0.00
$c^{\mathbb{Q}}_{\boldsymbol{Y},3}$	0.00	-	0.00	-	0.00	-	0.50	0.00
$D^{\mathbb{Q}}_{\boldsymbol{Y},11}$	-0.03	0.00	-3.12	0.06	-3.05	0.09	-1.98	0.10
$D^{\mathbb{Q}}_{\boldsymbol{Y},21}$	0.29	0.11	2.67	1.03	3.14	6.90	0.00	0.00
$D^{\mathbb{Q}}_{\boldsymbol{Y},31}$	-1.73	0.25	21.46	9.53	71.11	17.84	8.98	1.11
$D^{\mathbb{Q}}_{\boldsymbol{Y},12}$	0.00	0.00	0.00	-	0.06	0.02	0.38	0.04
$D^{\mathbb{Q}}_{\boldsymbol{Y},22}$	-0.39	0.12	-0.04	0.00	-0.11	0.13	-1.29	0.07
$D^{\mathbb{Q}}_{\boldsymbol{Y},32}$	-3.06	0.23	0.10	0.05	-1.29	0.41	0.00	0.00
$D^{\mathbb{Q}}_{\boldsymbol{Y},13}$	0.00	0.00	0.00	-	0.00	-	0.00	0.00
$D^{\mathbb{Q}}_{\boldsymbol{Y},23}$	2.19	0.09	0.07	0.00	0.00	-	1.42	0.13
$D^{\mathbb{Q}}_{\boldsymbol{Y},33}$	-4.03	0.12	-3.15	0.08	-3.15	0.06	-2.04	0.14
$\delta_{\boldsymbol{Y},1}$	28.97	2.25	-9.64	3.59	-7.95	6.53	0.30	0.04
$\delta_{\boldsymbol{Y},2}$	42.64	1.93	3.48	0.30	2.59	1.81	-2.58	0.29
$\delta_{\boldsymbol{Y},3}$	25.52	0.85	10.07	5.36	2.54	0.52	5.87	0.56
$c^{\mathbb{P}}_{\boldsymbol{Y},1}$	-7.83	2.48	1.86	0.94	0.54	0.30	4.31	14.41
$c^{\mathbb{P}}_{\boldsymbol{Y},2}$	-0.41	3.53	50.15	50.17	7.43	29.49	10.64	57.13
$c^{\mathbb{P}}_{\boldsymbol{Y},3}$	5.98	3.01	14.84	25.57	11.75	8.01	22.24	32.19
$D^{\mathbb{P}}_{\boldsymbol{Y},11}$	0.37	1.07	-3.82	1.90	-4.16	3.80	-2.79	9.99
$D^{\mathbb{P}}_{\boldsymbol{Y},21}$	0.61	1.02	24.07	22.47	9.87	35.16	0.00	0.00
$D^{\mathbb{P}}_{\boldsymbol{Y},31}$	-4.11	1.26	40.12	20.22	125.78	53.31	22.74	8.88
$D^{\mathbb{P}}_{\boldsymbol{Y},12}$	1.92	1.84	0.00	-	0.11	0.15	0.41	1.86
$D^{\mathbb{P}}_{\boldsymbol{Y},22}$	0.20	1.76	-1.94	1.34	-0.93	0.45	-0.59	3.05
$D^{\mathbb{P}}_{\boldsymbol{Y},32}$	-7.81	1.97	-0.68	1.42	-4.22	1.72	2.52	3.34
$D^{\mathbb{P}}_{\boldsymbol{Y},13}$	8.66	2.35	0.00	-	0.00	-	0.00	0.00
$D^{\mathbb{P}}_{\boldsymbol{Y},23}$	3.14	3.16	-6.80	4.21	0.00	-	0.00	0.00
$D^{\mathbb{P}}_{\boldsymbol{Y},33}$	-13.28	2.51	-6.21	1.77	-6.40	3.11	-9.86	4.30
β_{21}	0.00	-	13.70	4.69	0.00	-	0.00	-
β_{31}	0.00	-	9.98	11.25	115.36	61.55	0.00	-
β_{32}	0.00	-	0.00	-	0.84	1.09	0.00	-
Log L	-45995.29		-46715.73		-46696.32		-46232.82	

D Option pricing details

The functions $a_\Phi(t)$ and $\boldsymbol{b}_\Phi(t)$ solve the following recursive system of ODEs

$$\frac{\mathrm{d}\boldsymbol{b}_\Phi}{\mathrm{d}t}(t) = -\boldsymbol{D}^{\mathbb{Q}\prime}\boldsymbol{b}_\Phi(t) - \frac{1}{2}\sum_{i=1}^{n}\left[\boldsymbol{b}_\Phi(t)'\boldsymbol{\Sigma}\right]_i^2 \boldsymbol{\beta}_i,$$
$$\frac{\mathrm{d}a_\Phi}{\mathrm{d}t}(t) = -\boldsymbol{c}^{\mathbb{Q}\prime}\boldsymbol{b}_\psi(t) - \frac{1}{2}\sum_{i=1}^{n}\left[\boldsymbol{b}_\Phi(t)'\boldsymbol{\Sigma}\right]_i^2 \alpha_i,$$

with α_i such that $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_m)'$, $\boldsymbol{\beta}_i$ such that $\boldsymbol{B} = (\boldsymbol{\beta}_1, \ldots, \boldsymbol{\beta}_m)'$ and $[\boldsymbol{x}]_i$ denotes the i-th element of vector $\boldsymbol{x}$. The functions $a_\Phi(t)$ and $\boldsymbol{b}_\Phi(t)$ satisfy the boundary conditions $\boldsymbol{b}_\Phi(T) = \boldsymbol{u}$, $a_\Phi(T) = 0$, $\boldsymbol{d}_\Phi(T) = \boldsymbol{v}_1$ and $c_\Phi(T) = v_0$.

The functions $c_\Phi(t)$ and $\boldsymbol{d}_\Phi(t)$ solve the following recursive system of ODEs

$$\frac{\mathrm{d}\boldsymbol{d}_\Phi}{\mathrm{d}t}(t) = -\boldsymbol{D}^{\mathbb{Q}\prime}\boldsymbol{d}_\Phi(t) - \frac{1}{2}\sum_{i=1}^{n}\left[\boldsymbol{d}_\Phi(t)'\boldsymbol{\Sigma}\right]_i^2 \boldsymbol{\beta}_i,$$
$$\frac{\mathrm{d}c_\Phi}{\mathrm{d}t}(t) = -\boldsymbol{c}^{\mathbb{Q}\prime}\boldsymbol{d}_\psi(t) - \frac{1}{2}\sum_{i=1}^{n}\left[\boldsymbol{d}_\Phi(t)'\boldsymbol{\Sigma}\right]_i^2 \alpha_i,$$

$\boldsymbol{d}_\Phi(T) = \boldsymbol{v}_1$ and $c_\Phi(T) = v_0$.