



HAL
open science

Heath-Jarrow-Morton meet lifted Heston in energy markets for joint historical and implied calibration

Eduardo Abi Jaber, Soukaïna Bruneau, Nathan de Carvalho, Dimitri Sotnikov,
Laurent Tur

► **To cite this version:**

Eduardo Abi Jaber, Soukaïna Bruneau, Nathan de Carvalho, Dimitri Sotnikov, Laurent Tur. Heath-Jarrow-Morton meet lifted Heston in energy markets for joint historical and implied calibration. 2025. ⟨hal-04879087v2⟩

HAL Id: hal-04879087

<https://hal.science/hal-04879087v2>

Preprint submitted on 28 Apr 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Heath–Jarrow–Morton meet lifted Heston in energy markets for joint historical and implied calibration

Eduardo Abi Jaber^{*1}, Soukaïna Bruneau¹, Nathan De Carvalho^{†2,3}, Dimitri Sotnikov^{‡1,2}, and Laurent Tur²

¹École Polytechnique, CMAP

²Engie Global Markets

³Université Paris Cité, LPSM

April 28, 2026

Abstract

In energy markets, joint historical and implied calibration is of paramount importance for practitioners, yet notoriously challenging due to the need to align historical correlations of futures contracts with implied volatility smiles from the option market. We address this crucial problem with a multiplicative multi-factor Heath-Jarrow-Morton (HJM) model for forward curves, combined with a stochastic volatility factor coming from the lifted Heston model. We develop a sequential fast calibration procedure leveraging the Kemna-Vorst approximation of futures contracts: (i) historical correlations and the Variance Swap (VS) volatility term structure are captured through Level, Slope, and Curvature factors, (ii) the VS volatility term structure can then be corrected for a perfect match via a fixed-point algorithm, (iii) implied volatility smiles are calibrated using Fourier-based techniques. The main advantage of the proposed calibration framework is the decoupling of the calibration steps: each step tackles a simpler calibration subproblem and guaranties that the previously optimized parameters remain unchanged. Our model displays remarkable joint historical and implied calibration fits on the German power market and enables realistic interpolation within the implied volatility hypercube.

Mathematics Subject Classification (2010): 91G20, 91G60

JEL Classification: C02, Q41, C63

Keywords: Energy markets, Multi-factor HJM, Nelson-Siegel, Stochastic Volatility, Lifted Heston, Kemna-Vorst, Calibration

^{*}eduardo.abi-jaber@polytechnique.edu. I am grateful for the financial support from the Chaires FiME-FDD, Financial Risks, Deep Finance & Statistics and Machine Learning and systematic methods in finance at Ecole Polytechnique.

[†]nathan.decarvalho@engie.com. I am grateful for the financial support provided by Engie Global Markets.

[‡]dimitri.sotnikov@gmail.com. I am grateful for the financial support provided by Engie Global Markets.

Contents

1	The model: HJM with lifted Heston	5
1.1	Correlation structure in the model	8
1.2	The Kemna–Vorst approximation	9
1.3	Variance swap price and volatility	10
1.4	A Nelson–Siegel parametrization for the volatility functions σ	11
1.5	Joint calibration: overview and snapshots	12
2	Market data	14
2.1	Futures contracts’ quotes	14
2.1.1	Stripping optimization to construct rolling futures contracts	15
2.1.2	Estimated covariance of rolling futures’ log returns from past historical data	16
2.2	Options data	17
2.2.1	Typical options quoted in power markets	17
2.2.2	Variance swap volatilities deduced from option prices	17
3	Calibration methodology for historical covariances and variance swap volatilities	19
3.1	Explicit model covariance and VS volatility term structure	20
3.2	Loss function for the historical covariances - VS volatility term structure calibration problem	22
3.3	Solver specification	23
4	Calibration methodology for implied smiles	25
4.1	Fourier option pricing for implied calibration	25
4.2	Exact calibration correction of the VS volatility term structure	26
4.3	Smile shape calibration	31
5	Numerical results	32
5.1	Calibration results	32
5.1.1	Step 1: Joint historical covariances – implied VS variances calibration results	32
5.1.2	Step 2: Term structure exact calibration correction	37
5.1.3	Step 3: Implied smile shapes calibration	38
5.2	A posteriori validation of model approximations	39
5.2.1	Sample path trajectories errors	40
5.2.2	Instantaneous correlations	41
5.2.3	Quality of the Kemna–Vorst approximation on the instantaneous correlations	42
5.2.4	Quality of the Kemna–Vorst approximation on the smiles	42
5.3	Volatility hypercube extrapolation in the calibrated model	43
A	More calibration results	45
A.1	Attainable smiles	45
A.2	Calibration results for TTF market	46
A.3	A word on the additive model	50
B	Additional details for the step 1) calibration	52
B.1	Estimation of historical covariances	52
B.2	Linear cone programming	53
C	Proof of Theorem 4.2	54
D	Monte Carlo simulation scheme	56

Introduction

In power markets, futures contracts deliver electricity continuously over a fixed period, rather than on a fixed delivery date as is typical for commodities like oil. These contracts are settled either physically or financially with respect to the average spot price of electricity over the delivery period. Due to the inability to efficiently store electricity in large quantities, these instruments are often referred to as *swaps*, as the holder effectively exchanges the forward rate against the spot price of the commodity.

The unique characteristics of electricity markets, such as the need to maintain equilibrium between real-time production and consumption, distinguish them from other commodity markets. Real-time delivery is typically managed through intra-day and imbalance markets, which align production and consumption levels and can exhibit frequent price spikes or even negative prices during periods of significant imbalance. Moreover, financial futures with very short-term deliveries often show low correlation with long-term futures, which are highly correlated among themselves. This behavior is accompanied by exponentially increasing realized volatility as the time to delivery reduces, the so-called [Samuelson \(2016\)](#) effect.

Following the deregulation of European electricity markets, an extensive body of literature has emerged on the modeling of electricity markets. We refer to the survey by [Deschatre, Féron, and Gruet \(2021\)](#) and the book by [Benth, Benth, and Koekebakker \(2008\)](#) for an overview of key modeling approaches and market dynamics.

Several approaches to energy market modeling are considered in the literature, differing in the choice of the initial stochastic quantity to be modeled. The first class of models focuses on the spot price process, as seen in works like [Mishura, Ottaviano, and Vargiolu \(2023\)](#); [Schmeck and Schwerin \(2021\)](#); [Cortazar, Lopez, and Naranjo \(2017\)](#). Another approach, inspired by the LIBOR market model, see [Brigo and Mercurio \(2006\)](#), models futures contracts with specific delivery periods (e.g., monthly contracts, as in [Kiesel, Schindlmayr, and Börger \(2009\)](#), [Gardini and Santilli \(2024\)](#)), using these contracts as building blocks to derive the prices of other contracts under no-arbitrage conditions. The third class of models, inspired by the well-known [Heath, Jarrow, and Morton \(1992\)](#) (HJM) interest rate model, takes infinitesimal futures contracts as a starting point and uses them to reconstruct futures contracts with any delivery period. This type of model, adapted to the energy-markets setting by [Clewlow and Strickland \(1999\)](#), ensures consistent dynamics across all futures contracts while also guaranteeing the absence of arbitrage between futures contracts with overlapping delivery periods. The ability to model the dynamics of the entire futures curve, that is, the joint evolution of all futures contracts, is essential for pricing and hedging derivatives that are sensitive to cross-contract correlations, such as swing and spread options ([Carmona and Durrleman, 2003](#)).

Since the HJM approach naturally leads to modeling the entire futures curve, it is common to work with stochastic processes taking values in infinite-dimensional Hilbert spaces. In particular, several classes of infinite-dimensional stochastic volatility models have been developed; see, for instance, [Cox, Karbach, and Khedher \(2022\)](#). [Benth and Simonsen \(2018\)](#) consider an infinite-dimensional Heston-type model, [Benth and Paraschiv \(2018\)](#) employ an NIG Lévy process as the stochastic driver, and [Benth and Harang \(2021\)](#) analyze infinite-dimensional Volterra-type dynamics. For a more systematic treatment of infinite-dimensional modeling, we refer the reader to the monograph by [Benth and Krühner \(2023\)](#).

Early commodity models primarily focus on historical calibration, which involves calibrating the covariance structure of traded futures contracts' returns; see for example [Andersen \(2010\)](#) for the calibration of a HJM model to gas prices, [Edoli, Tasinato, and Vargiolu \(2013\)](#) for a multi-underlyings calibration using the quadratic variations of two-factor models for each market, [Gardini and Santilli \(2024\)](#) where they calibrate a LIBOR market Black-Scholes-type factor model using historical swap prices, and [Féron and Gruet \(2024\)](#) for a historical calibration of a multi-factor HJM model using maximum likelihood and a Kalman filter where they also comment on the number of factors to capture the historical covariance of futures' returns.

Since April 2024, brokers have started quoting smiles for vanilla options on German power. These vanilla options are written on monthly contracts, quarterly contracts, and calendar (yearly) contracts, which can overlap. For example, the first quarter of 2025 and the calendar 2025 can have quoted smiles. Another specificity of the power market is that for one calendar underlying, three or four smiles can be quoted. With the increasing liquidity of the electricity derivatives market in Europe, there has been significant growth in

research in option pricing in power market models, see [Schmeck and Schwerin \(2021\)](#); [Cortazar, Lopez, and Naranjo \(2017\)](#); [Benth, Piccirilli, and Vargiolu \(2019\)](#). However, few studies as in [Piccirilli, Schmeck, and Vargiolu \(2021\)](#); [Musti, Fanelli, and Maddalena \(2016\)](#) address the challenge of implied calibration, which involves calibrating model parameters to fit available option prices. Note that implied calibration in the energy market is exceptionally challenging, as it requires the simultaneous calibration of multiple volatility surfaces associated with different underlyings, which are interconnected through the futures price curve.

In energy markets, the limited liquidity of options and the relatively small set of underlyings for which they are traded make it impossible to calibrate pricing models using option prices alone, a marked contrast to more developed derivatives markets, such as equities or interest rates. Consequently, a crucial requirement for any pricing model in this low-liquidity environment is the ability to produce reasonable interpolations and extrapolations of option prices, not only across strikes and maturities but also across different futures contracts, in order to fill the missing regions of the implied volatility hypercube. Moreover, as discussed above, calibrating the correlation structure is essential, yet cannot be inferred from option prices. Historical calibration therefore becomes indispensable for practical implementation. At the same time, given the gradually increasing liquidity of options, the options market can no longer be neglected, both because of regulatory requirements and due to the arbitrage opportunities that arise under purely historical models. The most significant limitation of the approaches mentioned above is that they allow for either historical calibration or implied calibration, but not both simultaneously. Hence the following questions:

Is joint historical and implied calibration possible in energy markets?

If so, can we do it with a tractable model that provides a consistent interpolation of option prices?

We answer both questions *affirmatively* with a multiplicative multi-factor [Heath, Jarrow, and Morton \(1992\)](#) (HJM) model for forward curves, combined with a stochastic volatility factor coming from the lifted Heston model of [Abi Jaber \(2019\)](#). To the best of our knowledge, this paper is the first to address both calibration problems with a single model, taking into account the entire implied volatility smile.

On the one hand, by capturing historical covariances and implied volatility levels, our model provides insights on the futures correlation structure between liquid and non-liquid futures contracts. On the other hand, implied calibration allows the entire implied volatility smile, not just the At-The-Money (ATM) volatilities, to be represented by the model. Moreover, the proposed implied calibration takes into account potential multiple maturities for a given underlying futures contract. Thus, the joint historical and implied calibration matches both the correlation between futures and quoted smiles — a problem of paramount importance for market practitioners — leading to more accurate pricing of exotic contracts like Asian options or swing options. One potential difficulty is that implied calibration may affect the historical one, but we will see that this influence is negligible. To the best of our knowledge, this paper is the first to address both calibration problems with a single model, taking into account the entire implied volatility smile.

Contributions. More precisely, to solve the joint historical and implied calibration problem, we introduce in Section 1 an HJM model with

- (i) parsimonious parametric Level, Slope and Curvature risk-factors in order to capture both historical covariances of rolling futures’ log returns and implied volatility levels with a few factors,
- (ii) two piece-wise constant functions to perfectly match the implied volatility term structure, including early maturities,
- (iii) a stochastic volatility component coming from the lifted Heston model with three time-scales to match the implied volatility skews.

Such model is by construction arbitrage-free with respect to futures contracts with overlapping delivery periods.

In Section 2, we present our Market data. We recall the typical liquid “absolute” futures quoting on power markets, we detail a stripping algorithm to construct rolling futures contracts satisfying absence of overlapping arbitrage and used to estimate historical covariances. Furthermore, we propose a novel multi-contract SSVI parametrization to extract Variance Swap (VS) volatilities from listed options.

Then, we detail a novel three-step sequential calibration methodology relying on the [Kemna and Vorst \(1990\)](#) (KV) approximation of futures contracts in order to

- 1) jointly capture the historical covariances of rolling futures' daily log returns and the implied VS volatility levels of absolute futures via a non-linear – linear cone program, see [Section 3](#),
- 2) correct and fit perfectly the VS volatility term structure via a fast fixed-point algorithm based on VS prices, see [Section 4.2](#),
- 3) fit the volatility smile's shapes via Fourier inversion techniques as in [Lewis \(2001\)](#) for fast and efficient vanilla option pricing, see [Section 4.1](#).

We emphasize that the calibration procedure is fully decoupled: the three steps are performed sequentially, with each step preserving the calibration results achieved in the previous ones. Moreover, thanks to both the KV approximation and the decoupled structure, each step remains highly tractable. Decomposing the calibration into simpler subproblems yields a procedure that is a priori faster and more efficient than calibrating all parameters simultaneously. Our model displays remarkable joint historical and implied calibration fits to both German power and TTF gas markets. In order to validate our calibration methodology, we show a posteriori how close the KV approximated futures are to the exact arbitrage-free futures in terms of sample trajectories, implied volatility smiles and correlations between futures contracts. Finally, we show that such a fully calibrated model can be used to interpolate the implied volatility hypercube in a consistent manner. The main calibration results are collected in [Section 5](#). Additional model algorithmic insights and calibration results are postponed to the appendices.

Related literature. The paper of [Piccirilli, Schmeck, and Vargiolu \(2021\)](#) is the work most closely related in spirit to our approach, although there are several important differences. The authors propose a two-factor model with Normal Inverse Gaussian Lévy factors, while our model is a stochastic volatility model with a continuous process as the variance. Their calibration procedure considers only one smile per contract, ignoring multiple maturities for a given futures contract. Our model accounts for these “early maturities”, enabling a more refined calibration of the volatility term structure.

Notations. For $N \in \mathbb{N}^*$, we denote by \mathbb{S}_{++}^N (resp. \mathbb{S}_+^N) the set of $N \times N$ definite (resp. semi-definite) positive matrices, and $\|\cdot\|_w$ denotes a weighted Euclidean norm with weights $w := (w_i)_{i \in \{1, \dots, N\}^2} \in \mathbb{R}_+^N$ such that $\|u\|_w := \sqrt{\sum_{i=1}^N w_i u_i^2}$, $u \in \mathbb{R}^N$. We define similarly the weighted Frobenius norm $\|A\|_\Gamma := \sqrt{\sum_{i=1}^N \sum_{j=i}^N \Gamma_{i,j} A_{i,j}^2}$, $A \in \mathbb{R}^{N \times N}$, with matrix weights $\Gamma := (\Gamma_{i,j})_{i,j \in \{1, \dots, N\}^2} \in \mathbb{R}_+^{N \times N}$. We will omit the indices w and Γ for standard Euclidean and Frobenius norms $\|u\|$ and $\|A\|$ respectively. We will also denote by $\|\cdot\|_\infty$ the infinity norm $\|u\|_\infty := \max_{i \in \{1, \dots, N\}} |u_i|$ and $\|A\|_\infty := \max_{i,j \in \{1, \dots, N\}} |A_{ij}|$.

1 The model: HJM with lifted Heston

Fix a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{Q})$ satisfying the usual conditions, where \mathbb{Q} represents the risk-neutral probability. We model the futures price curve under \mathbb{Q} , i.e. the infinitesimal futures contract prices $(f(t, T))_{0 \leq t \leq T < \infty}$, à la [Heath, Jarrow, and Morton \(1992\)](#) (HJM), enhanced with a stochastic volatility component coming from the lifted Heston model of [Abi Jaber \(2019\)](#) in the form

$$\frac{df(t, T)}{f(t, T)} = g(T)h(t)\sqrt{V_t} \sum_{i=1}^N \sigma_i(t, T) dW_t^i, \quad f(0, T) \in \mathbb{R}_+, \quad 0 \leq t \leq T, \quad (1.1)$$

where

- $W := (W_t)_{t \geq 0}$ is an N -dimensional Brownian motion with a correlation matrix $R \in \mathbb{S}_{++}^N$.
- Each $\sigma_i : \mathbb{R}_+^2 \rightarrow \mathbb{R}$, $i \in \{1, \dots, N\}$ is a deterministic continuous and bounded function, capturing the *realized futures contracts' covariances* and the *implied volatility levels*. We set $\sigma = (\sigma_1, \dots, \sigma_N)^\top$.

- V is a nonnegative stochastic variance process responsible for the *implied smile* of the form

$$V_t = 1 + \sum_{i=1}^M c_i U_t^i, \quad t \geq 0 \quad (1.2)$$

where the pseudo-factors $(U^i)_{i \in \{1, \dots, M\}}$, weighted by $c_i \geq 0$, are driven by the same Brownian motion B , but mean-revert at different speeds $0 < x_1 < x_2 < \dots < x_M$ such that

$$dU_t^i = -x_i U_t^i dt + \sqrt{V_t} dB_t, \quad U_0^i = 0, \quad i \in \{1, \dots, M\}. \quad (1.3)$$

Here $B = (B_t)_{t \geq 0}$ is a one-dimensional Brownian motion correlated with W , via $(\hat{\rho}_i)_{i \in \{1, \dots, N\}} \in [-1, 1]^N$, to take into account *the leverage effect* such that

$$B_t = \sum_{i=1}^N \hat{\rho}_i \widehat{W}_t^i + \sqrt{1 - \sum_{i=1}^N \hat{\rho}_i^2} W_t^\perp, \quad t \geq 0,$$

where W^\perp is a scalar Brownian motion independent from \widehat{W} which is a standard N -dimensional Brownian motion constructed from W via the Cholesky decomposition (see, for example, [Horn and Johnson \(2013\)](#))

$$R = L^\top L, \quad \widehat{W} = L^{-1} W, \quad \|\hat{\rho}\| \leq 1, \quad (1.4)$$

so that

$$d\langle B, W^i \rangle_t = \tilde{\rho}_i dt, \quad \tilde{\rho} = L\hat{\rho}. \quad (1.5)$$

- $T \mapsto g(T)$ and $t \mapsto h(t)$ are deterministic bounded positive functions *correcting the implied volatility levels* to match the volatility term structure for the contracts with several maturity dates.

For a fixed $t \in [0, T]$, $f(t, T)$ represents the quote observed at date t of the contract that delivers a unit amount of commodity between dates T and $T + dT$, where $dT > 0$ denotes an infinitesimal amount of time. The special case $t = T$ is seen as the spot price S of such a commodity: $S_t := f(t, T)$, which is well-defined as soon as $\lim_{t \rightarrow T} f(t, T) < \infty$.

We prove in the next theorem that the model (1.1) is indeed well-defined leveraging results from [Abi Jaber \(2019\)](#). In particular, we note that although the different factors U^i can become negative, the variance process V is always nonnegative ([Abi Jaber et al., 2019](#), Theorem 7.1), as illustrated on Figure 1. Furthermore, the variance process (1.2) is Markovian in the state variables $U := (U^i)_{i \in \{1, \dots, M\}}$ with a state space corresponding to the set of $u \in \mathbb{R}^M$ such that

$$1 + \sum_{j=1}^M c_j u_j \geq 0 \quad \text{and} \quad \mu \sum_{j=1}^i \frac{c_j}{x_j} + \sum_{j=1}^i c_j u_j \geq \mu \sum_{j=1}^i \frac{c_j}{x_{i+1}} + \sum_{j=1}^i c_j u_{i+1} \quad \text{for } i \in \{1, \dots, M-1\},$$

where $\mu = \left(\sum_{k=1}^M \frac{c_k}{x_k} \right)^{-1}$ see [Abi Jaber, Bayer, and Breneis \(2024\)](#).

Remark 1.1. *We chose the lifted Heston model for two reasons. First, its affine structure allows for fast vanilla option pricing using Fourier techniques, which is highly beneficial for calibration speed. Second, it offers significantly more flexibility in the shapes of implied volatility surfaces that can be calibrated, compared to the simpler standard Heston model, as we demonstrate in Appendix A.1. Thus, the lifted Heston model strikes a good balance between numerical tractability and universality. However, we emphasize that the modeling framework we propose is not limited to this model and is flexible enough to accommodate any stochastic variance process V_t satisfying $\mathbb{E}[V_t] \equiv 1$.*




Figure 1: Trajectory of the stochastic variance factors $(U^i)_{i=1,2,3}$ (on the left) and the resulting stochastic variance V (on the right) corresponding to the parameters $x = (4.6 \cdot 10^{-6}, 9.712, 20.249)$ and $c = (0.492, 0.68, 2.79)$. The simulation scheme is described in Appendix D.

Theorem 1.2. Fix $T > 0$. Let $g, h : [0, T] \rightarrow \mathbb{R}$ and $\sigma(\cdot, T) : [0, T] \rightarrow \mathbb{R}^N$ be bounded and measurable functions. Then, there exists a unique strong solution $(U^i)_{i \in \{1, \dots, M\}}$ to (1.3) such that V given by (1.2) remains non-negative. Furthermore, the process $f(\cdot, T)$ defined by

$$f(t, T) = f(0, T) \exp \left(-\frac{g^2(T)}{2} \int_0^t h^2(s) V_s \sigma^\top(s, T) R \sigma(s, T) ds + g(T) \int_0^t h(s) \sqrt{V_s} \sigma^\top(s, T) dW_s \right), \quad t \in [0, T], \quad (1.6)$$

is the unique strong solution to (1.1). In particular, $f(\cdot, T)$ is a true martingale.

Proof. The equation (1.1) can be rewritten as a one-dimensional diffusion,

$$\frac{df(t, T)}{f(t, T)} = h(t) \sqrt{V_t} \sqrt{\sigma(t, T)^\top R \sigma(t, T)} d\tilde{W}_t, \quad d\langle B, \tilde{W} \rangle_t = \frac{\sigma(t, T)^\top \tilde{\rho} dt}{\sqrt{\sigma(t, T)^\top R \sigma(t, T)}}, \quad t \in [0, T], \quad (1.7)$$

where the Brownian motion \tilde{W} is given by

$$\tilde{W}_t = \sum_{i=1}^N \frac{\sigma_i(t, T)}{\sqrt{\sigma(t, T)^\top R \sigma(t, T)}} W_t^i, \quad 0 \leq t \leq T.$$

The factor processes $(U^i)_{i \in \{1, \dots, M\}}$ can be written as

$$U_t^i = \int_0^t e^{-x_i(t-s)} \sqrt{V_s} dB_s, \quad i \in \{1, \dots, M\}, \quad 0 \leq t \leq T,$$

so that the variance dynamics (1.2) reads

$$V_t = 1 + \int_0^t \left(\sum_{i=1}^M c_i e^{-x_i(t-s)} \right) \sqrt{V_s} dB_s = 1 + \int_0^t K(t-s) \sqrt{V_s} dB_s, \quad 0 \leq t \leq T, \quad (1.8)$$

where $K(t) := \sum_{i=1}^M c_i e^{-x_i t}$. The equation (1.8) means that V is a Volterra square-root process understood in the sense of (Abi Jaber, Larsson, and Pulido, 2019, Section 6). An application of (Abi Jaber, 2019, Theorem A.1) yields the existence and uniqueness of strong solution $(U^i)_{i \in \{1, \dots, M\}}$, such that the variance process $V \geq 0$, as well as the existence and uniqueness of the strong solution to (1.7) given by (1.6). The proof of the martingality of $f(\cdot, T)$ follows exactly the proof of (Abi Jaber, Larsson, and Pulido, 2019, Lemma 7.3) taking into account that the correlation coefficient in (1.7) is time-dependent. \square

In practice, the infinitesimal futures contracts given by (1.1) are not observed in the market and therefore must be related to market traded futures contracts. Using practitioners' vocabulary, "absolute" futures contracts delivering electricity on a fixed calendar delivery period $[T_s, T_e]$ are typically quoted in organized markets until a few days before their first delivery date i.e. for dates t such that $0 \leq t \leq T_s - \delta$, δ equal to a few days. We recall in Section 2.1 the typical futures quoting in power markets. From the model perspective, the unitary absolute futures contract delivering continuously a unitary power unit of electricity over $[T_s, T_e]$ is given by¹

$$F_t(T_s, T_e) := \frac{1}{T_e - T_s} \int_{T_s}^{T_e} f(t, T) dT, \quad 0 \leq t < T_s. \quad (1.9)$$

ensuring the absence of arbitrage opportunity for futures contracts with overlapping delivery periods.

On the other hand, a "rolling" futures contract is a contract that depends on the observation date, $t \geq 0$, and maintains a constant time to delivery, $T_s > 0$, with a fixed and contiguous delivery period of duration $T_e - T_s > 0$. As a result, the contract's delivery period adjusts as the observation date changes, ensuring that the contract always quotes. However, its quote is typically not directly observable in the market. Instead, it must be derived using no-arbitrage principles based on market quotes available at each observation date. Its quote is expressed in our model by the formula

$$F_t(t + T_s, t + T_e) = \frac{1}{T_e - T_s} \int_{t+T_s}^{t+T_e} f(t, T) dT, \quad t \geq 0. \quad (1.10)$$

Example 1.3 (Distinction between absolute and rolling futures). *For example, on the 20th of August 2024, the next absolute monthly futures contract quoting on the market corresponds to the contract September 2024 delivering electricity between the 1st to the 30th of September 2024. Fixing T_s to seven days, and $T_e - T_s$ to thirty days, one can define a month-ahead rolling contract delivering electricity from the 27th of August to the 25th of September 2024. On the 21st of August 2024, such rolling contract becomes the one delivering electricity from the 28th of August to the 26th of September 2024, while the forward September 2024 keeps the same delivery period.*

Remark 1.4. *In the discussion above, we assumed that the risk-free rate is $r = 0$. In the general case, however, a similar no-arbitrage argument (Benth and Koekebakker, 2008, Section 4) yields*

$$F_t(T_s, T_e) := \int_{T_s}^{T_e} w(T) f(t, T) dT, \quad 0 \leq t < T_s,$$

where $w(T) := \frac{r e^{-rT}}{e^{-rT_s} - e^{-rT_e}}$. Thus, the extension of our approach to a non-zero constant interest rate is straightforward, although it makes the resulting expressions and computations more cumbersome. For notational convenience, we will therefore assume throughout the paper that $r = 0$.

1.1 Correlation structure in the model

One of the main objectives of the HJM modeling framework is the consistent joint simulation of multiple futures contracts. This capability is essential for pricing derivatives that depend on several underlyings and are sensitive to their correlations, such as spread options. Consequently, the correlation structure between futures contracts plays a crucial role in both pricing and hedging such derivatives and should therefore be carefully calibrated.

We focus on the calibration of *rolling covariances*, that is, the covariances between returns of rolling futures contracts, defined as

$$\text{Cov} \left(r_t^{k_1}(\tau_d), r_t^{k_2}(\tau_d) \right), \quad \text{where} \quad r_t^k(\tau_d) := \log \frac{F_t(t + T_s^k, t + T_e^k)}{F_{t-\tau_d}(t + T_s^k, t + T_e^k)}, \quad (1.11)$$

where the indices k_1 and k_2 correspond to the rolling contracts under consideration, and τ_d denotes the time step used for computing returns, typically set to one day.

¹We assume a zero discount rate ($r = 0$) for simplicity.

The main motivation for calibrating rolling covariances is that they are assumed to be approximately constant over time and can therefore be estimated from historical data. The computation of rolling covariances in the model will be discussed in Section 3, while the historical estimation procedure is outlined in Appendix B.1.

However, although rolling covariances are convenient for calibrating the covariance structure of historical factors, pricing problems require dealing with *instantaneous covariances* of *absolute contracts*, given by

$$\rho_{ij}(t) = \frac{\langle d \log F(T_s^i, T_e^i), d \log F(T_s^j, T_e^j) \rangle_t}{\sqrt{\langle d \log F(T_s^i, T_e^i) \rangle_t} \sqrt{\langle d \log F(T_s^j, T_e^j) \rangle_t}}, \quad (1.12)$$

which have a direct impact on option prices, but cannot be estimated historically. The numerical precision of computing these correlations is discussed in Subsection 5.2.

We now turn to an approximation that renders the model tractable and, in particular, allows us to derive closed-form expressions for (1.11) and (1.12).

1.2 The Kemna–Vorst approximation

The dynamics of the futures contract $F(T_s, T_e)$ cannot be written explicitly in our model, since (1.9) involves an arithmetic mean, and not a geometric one. For this reason, we will use the Kemna and Vorst (1990) approximation such that

$$\begin{aligned} \frac{dF_t(T_s, T_e)}{F_t(T_s, T_e)} &\approx \frac{1}{T_e - T_s} \int_{T_s}^{T_e} \frac{df(t, T)}{f(t, T)} dT \\ &= h(t) \sqrt{V_t} \sum_{i=1}^N \left(\frac{1}{T_s - T_e} \int_{T_s}^{T_e} g(T) \sigma_i(t, T) dT \right) dW_t^i \\ &= h(t) \sqrt{V_t} \Sigma_t(T_s, T_e)^\top dW_t, \end{aligned} \quad (1.13)$$

where

$$\Sigma_\cdot(T_s, T_e) := \frac{1}{T_e - T_s} \int_{T_s}^{T_e} g(T) \sigma(\cdot, T) dT. \quad (1.14)$$

Consequently, the approximated dynamics $\tilde{F}_\cdot(T_s, T_e)$ of the futures contract is given by

$$\frac{d\tilde{F}_t(T_s, T_e)}{\tilde{F}_t(T_s, T_e)} = h(t) \sqrt{V_t} \Sigma_t(T_s, T_e)^\top dW_t, \quad \tilde{F}_0(T_s, T_e) = F_0(T_s, T_e). \quad (1.15)$$

We stress that $\Sigma_\cdot(T_s, T_e) = \Sigma_\cdot(T_s, T_e; g, \sigma)$ depends essentially on g and σ , but for the sake of brevity, we omit these arguments. The equation (1.15) admits

$$\begin{aligned} \tilde{F}_t(T_s, T_e) &= F_0(T_s, T_e) \exp \left(-\frac{1}{2} \int_0^t h^2(r) V_r \Sigma_r^\top(T_s, T_e) R \Sigma_r(T_s, T_e) dr \right. \\ &\quad \left. + \int_0^t h(r) \sqrt{V_r} \Sigma_r^\top(T_s, T_e) dW_r \right), \quad t \in [0, T_s], \end{aligned} \quad (1.16)$$

as the unique strong solution, recall Theorem 1.2 for the existence and uniqueness of V .

Recall that futures contracts' quotes $F := F(T_s, T_e)$ defined by (1.9) are free from arbitrage opportunities in the case of overlapping delivery periods. In contrast, arbitrage may arise in the Kemna-Vorst approximated futures prices $\tilde{F} := \tilde{F}(T_s, T_e)$ from (1.16), as they do not strictly satisfy (1.9). However, such approximated futures' quotes \tilde{F} offer a tractable approach for calibration including

- the explicit computation of their variance swaps' volatilities as shown in the next Section 1.3,
- the explicit computation of their covariances as detailed in Section 3.1,

- the fast pricing of call and put options written on such futures contracts using a Fourier inversion technique as shown in Subsection 4.1.

Once the model (1.1) is calibrated using the Kemna-Vorst approximated futures \tilde{F} , we will illustrate numerically a posteriori that, not only are the trajectories of the arbitrage-free futures' quotes' F similar to the ones of \tilde{F} , but also that their instantaneous futures correlation structure and associated vanilla option prices are nearly identical. Indeed, all these quantities with respect to futures with quotes given by (1.9) can be approximated by Monte-Carlo techniques as detailed in Section 5.2.

Since the numerical analysis of the Kemna–Vorst approximation for our model, presented in Subsection 5.2, demonstrates that the approximated HJM model provides accurate pathwise and weak approximations and preserves the correlation structure well for contracts with the delivery periods of interest (i.e., one month, three months, and one year), we can use the approximated model interchangeably with the postulated model, thereby leveraging its tractability and analytical properties.

From now on, we fix a delivery period $[T_s, T_e]$ and use the approximation \tilde{F} everywhere instead of the exact dynamics. Accordingly, we also use the notation $\Sigma_\cdot := \Sigma_\cdot(T_s, T_e)$, leaving the delivery period implicit when there is no ambiguity.

Remark 1.5. *One may notice that there was no need for the approximation of futures contracts if the initial model (1.1) was additive and not multiplicative. However, the additive specification has several dynamics' disadvantages, which we comment briefly in Appendix A.3.*

1.3 Variance swap price and volatility

By definition, the *variance swap price* with maturity $T > 0$, denoted by VS_T , is the expected integrated quadratic variation of the futures contract following the dynamics (1.15) given by

$$\text{VS}_T := \mathbb{E}\langle \log F \rangle_T \approx \mathbb{E}\langle \log \tilde{F} \rangle_T = \int_0^T h(t)^2 \Sigma_t^\top R \Sigma_t dt \geq 0, \quad (1.17)$$

and the *variance swap volatility* $\sigma_{\text{VS},T}$ satisfies the equation $\sigma_{\text{VS},T}^2 T = \text{VS}_T$, so that

$$\sigma_{\text{VS},T} = \sqrt{\frac{1}{T} \int_0^T h(t)^2 \Sigma_t^\top R \Sigma_t dt}. \quad (1.18)$$

For more profound discussion of variance swaps and their properties, we refer reader to Demeterfi et al. (1999).

Why considering variance swaps for calibration? The variance swap volatility is a zero-order approximation of the ATM volatility level as shown by Bergomi and Guyon (2012), allowing us to formalize the idea of the *implied volatility smile level*. Moreover, formula (1.18) coincides exactly with the implied volatility for maturity T in the HJM model under deterministic volatility, i.e., when $V \equiv 1$.

In fact, in our model (1.1), variance swap prices (1.17) depend the functions h and also σ , g via the deterministic futures' volatilities $\Sigma_\cdot(T_s, T_e)$ (1.14), but not on the parameters of V from (1.2) since we conveniently set $\mathbb{E}V_t = 1$, $t \geq 0$. This allows to first calibrate the variance swap prices, i.e. the smile level with the functions σ , h and g , and then solve independently the smile shape calibration problem with the parameters of V .

Furthermore, the evaluation of the variance swap prices (1.17) turn out to be in practice less time-consuming than the evaluation of ATM volatilities since they are given by the closed formula (1.17), while the ATM volatilities are extracted via the Lewis formula (4.1) which require the computation of the log-price characteristic function, see Section 4.1 for more details.

Although Variance Swaps (VS) are not traded in power markets, their market prices VS_T^{mkt} can be extracted efficiently from the prices of vanilla option quotes as shown in Section 2.2.2. Thus, given a family of futures

contracts whose smiles are quoting, these VS volatilities (1.18) define the implied volatility term structure that we will aim to calibrate by fitting a priori σ , and then perfecting the match with the functions h and g .

We emphasize that our interest is not in variance swap contracts themselves, nor in hedging volatility risk; rather, we use the quantity (1.18), which naturally generalizes the implied volatility in the deterministic volatility case, only to calibrate the implied volatility term structure.

Indeed, under deterministic volatility, the quantity (1.17) reduces to the integrated variance up to maturity T , which was used, for instance, in Kiesel et al. (2009, Section 4) to calibrate the implied volatility term structure. This calibration procedure corresponds to the first step of our joint calibration. However, unlike Kiesel et al. (2009) and other classical HJM calibration approaches, in the second step we also calibrate the functions g and h to achieve an exact fit of the term structure, and in the third step we calibrate the implied volatility smile while keeping the variance swap volatility fixed.

For now, let us introduce a flexible parametrization for the functions σ .

1.4 A Nelson-Siegel parametrization for the volatility functions σ

Back in 1985, Charles B. Nelson and Andrew F. Siegel proposed in Nelson and Siegel (1987) the following “parsimonious” parametrization for the instantaneous forward rate r with $m \geq 0$ days to maturity

$$r(m) := \beta_0 + \beta_1 e^{-\frac{m}{\tau}} + \beta_2 \frac{m}{\tau} e^{-\frac{m}{\tau}}, \quad \beta_0, \beta_1, \beta_2 \in \mathbb{R}, \quad \tau > 0, \quad (1.19)$$

as an alternative to the previously used polynomial fitting techniques to match the yields of US Treasury bills of maturity m , obtained in their formulation by integrating from zero to m the forward rate (1.19) and dividing by m . They showed that such expression of the forward rate has a desirable non-explosive asymptotic behavior, and is capable of reproducing humps, S -shapes, and monotonic curves. We refer to Sepp and Rakhmonov (2023) as an example of a recent application of the HJM approach with Nelson–Siegel parametrization and stochastic volatility to the interest rate market.

In our case, the volatility functions σ in the HJM model (1.1) aim at capturing with parsimony both the realized covariance structure of a family of rolling futures contracts’ daily log returns as well as the VS volatility term structure of the futures contracts whose smiles quote in the market.

Inspired by the Nelson–Siegel parametrization (1.19), we specify three possible forms for the volatility shape functions $(\sigma_i)_{i \in \{1, \dots, N\}}$, which yield respectively three distinct types of factors, named *Level* (L), *Slope* (S) and *Curvature* (C), including

- one L -factor which has a constant volatility function in order to capture *the long-term volatility level of the curve*

$$\sigma_1(t, T) = \sigma_L, \quad 0 \leq t \leq T, \quad (1.20)$$

- S -factors having an exponentially increasing volatility as time to maturity decreases to capture *the Samuelson (2016) effect*

$$\sigma_{i+1}(t, T) = \sigma_{S,i} e^{-\frac{T-t}{\tau_{S,i}}}, \quad 0 \leq t \leq T, \quad i \in \{1, \dots, N_s\}, \quad (1.21)$$

- C -factors aiming at capturing *“humps” in the volatility term structure*, thereby capturing potential *“anti-Samuelson effect”* for long-term deliveries, i.e. a decreasing volatility with time to maturity, such that

$$\sigma_{j+N_s+1}(t, T) = \sigma_{C,j} \frac{T-t}{\tau_{C,j}} e^{-\frac{T-t}{\tau_{C,j}}}, \quad 0 \leq t \leq T, \quad j \in \{1, \dots, N_c\}, \quad (1.22)$$

with $N_s, N_c \in \mathbb{N}$ such that $N := 1 + N_s + N_c$. We display in Figure 2 the three distinct volatility shapes in terms of time to maturity.




Figure 2: Illustrations of the respective Level (plain blue), Slope (plain orange) and Curvature (plain green) volatility shapes with parameters $\sigma_L = 0.1$, $\sigma_S = \sigma_C = 1$, $\tau_S = 0.1$ and $\tau_C = 0.2$. We also plot the approximation of curvature by two opposed slope shapes (dotted green) with $\sigma_{S_1} = \sigma_{S_2} = 8.18$, $\tau_{S_1} = 0.212$ and $\tau_{S_2} = 0.188$.

The use of both L and S factors is not new, see for example [Kiesel, Schindlmayr, and Börger \(2009\)](#), [Gardini and Santilli \(2024\)](#). By contrast, we are not aware of previous works trying to calibrate C factors to power markets. Notice in particular that both S - and C -factors vanish when time to maturity goes to infinity, i.e. $T - t \rightarrow \infty$, hence only the L -factor remains in such regime and therefore captures the long-term volatility level of the energy curve.

Consequently, such L - S - C parametrization writes explicitly as

$$\sigma(t, T)^\top dW_t = \sigma_L dW_t^1 + \sum_{i=1}^{N_s} \sigma_{S,i} e^{-\frac{T-t}{\tau_{S,i}}} dW_t^{i+1} + \sum_{j=1}^{N_c} \sigma_{C,j} \frac{T-t}{\tau_{C,j}} e^{-\frac{T-t}{\tau_{C,j}}} dW_t^{j+N_s+1}. \quad (1.23)$$

In a similar spirit as [Nelson and Siegel \(1987\)](#) who noted that their parametrization (1.19) can be easily fitted to market data by least-squares, given a provisional $\tau > 0$, fixing the parameters $(\tau_{S,i})_{i=1,\dots,N_s}$, $(\tau_{C,i})_{i=1,\dots,N_c}$ in (1.23) leads to an efficient calibration problem for the remaining parameters σ_L , $(\sigma_{S,i})_{i=1,\dots,N_s}$, $(\sigma_{C,i})_{i=1,\dots,N_c}$ of the L - S - C factors, and their correlation matrix R formulated as a linear cone program ensuring that R remains non-negative definite as detailed in Section 3.

Remark 1.6 (Are L - S - C volatility shapes redundant?). *On the one hand, notice that taking $(\sigma_S, \tau_S)|_{\tau_S \rightarrow \infty}$ in a S -factor parametrization (1.21) yields a L -factor parametrization (1.20). Moreover, by considering two perfectly anti-correlated S -factors with identical parameters $\sigma_{S,1}, \sigma_{S,2}$ and distinct mean reversion rates, it is possible to approximate the behavior of a C -factor parametrization (1.22) as illustrated in green dots in Figure 2. Thus, it is indeed reasonable to restrict oneself to S -factors only to identify systematically in the market those distinct level, slope and curvature volatility behaviors but to the price of losing parsimony, as was done for example in [Féron and Gruet \(2024\)](#). On the other hand, a single curvature shape cannot approximate properly either a level or a slope volatility shape, nor a level shape can approximate neither a slope or a curvature shape.*

1.5 Joint calibration: overview and snapshots

The main practical advantage of our model is the possibility to decouple the joint calibration problem into three independent optimization problems to be solved consecutively, without the need to modify the parameters calibrated in previous steps.

- 1) A combined historical and implied calibration of the parameters (σ, R) to capture with parsimony i.e. with a minimum number of deterministic risk factors, the historical correlation of rolling futures

contracts' daily log returns as well as the overall term structure of the implied VS volatility deduced from the market quotes of vanilla options.

- 2) An exact calibration correction of the term structure of the implied VS volatility not captured in step 1), thanks to the functions g and h .
- 3) A calibration of the entire shape of the smile using the parameters of the stochastic variance process V in (1.2), that is $(c, x) = (c_i, x_i)_{i \in \{1, \dots, M\}}$ and $\tilde{\rho} = (\tilde{\rho}_i)_{i \in \{1, \dots, N\}}$.

The flow chart in Figure 3 illustrates the successive calibration procedure with the data used in each step as well as the resulting calibrated model parameters.

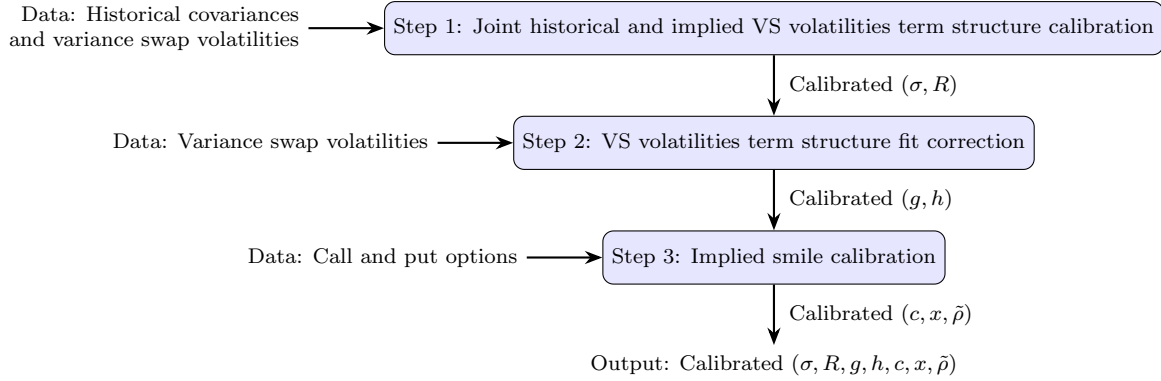


Figure 3: Overview of the calibration methodology in three iterative steps, including input data and calibrated parameters obtained at each step.

In Figure 4, we provide a calibration snapshot that summarizes the results of all three calibration steps for the case of the German power market as of the 1st of July 2024. For this date, $N = 5$ L - S - C risk factors with $N_s = 3$, $N_c = 1$, and $M = 3$ stochastic volatility factors were used, which yields 19² (resp. 11) calibrated parameters for step 1 (resp. step 3). Our model achieves an excellent fit of the historical volatilities and correlations in step 1, a perfect fit of the VS term structures in step 2 and a remarkable fit of the whole implied volatility surface in step 3, see also Figure 15. More detailed specification of model factors and calibrated parameters will be described further in Section 5.1.

²Our choice of 19 parameters (5 factors) remains conservative compared to the literature on historical calibration in power markets: Féron and Gruet (2024) consider 5 S -factors (20 parameters), Gardini and Santilli (2024) use 10 factors with 100 parameters to fit the realized volatility of 144 futures from six different markets, and a PCA study in Koekebakker and Ollmar (2005) supports a similar order of factors.

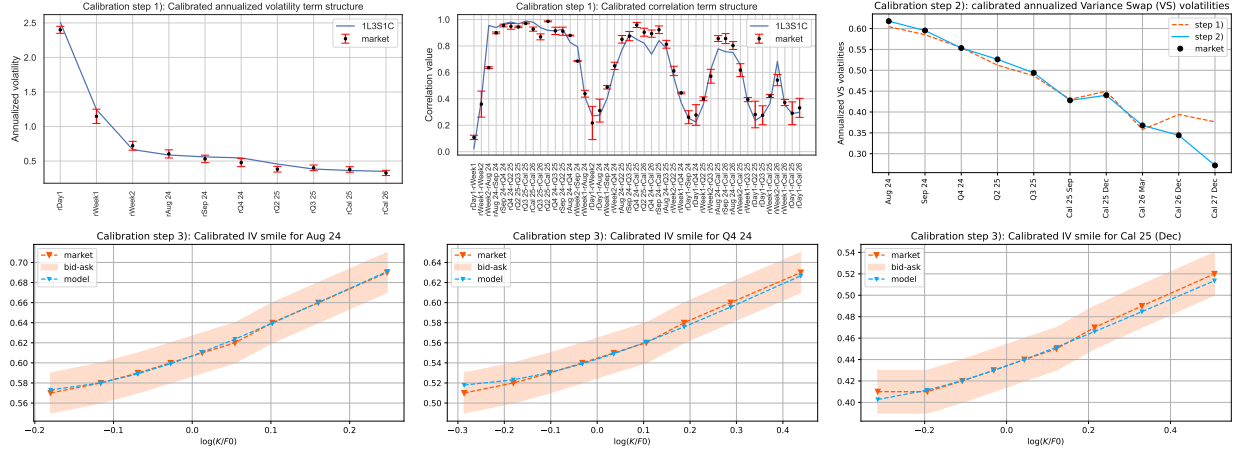


Figure 4: Snapshot of the calibration results on the German power market on the 1st July 2024: fit of the realized volatility term structure (upper left) and correlations (upper middle) of rolling futures in step 1 over one year and a half of historical data; fit of the VS volatility (upper right) term structure in step 1 (orange) and the corrected fit obtained in step 2 (blue); fit of the smile slices in step 3 of the respective futures contracts delivering electricity during August 24 (lower left), the fourth Quarter of 2024 (lower middle) and the whole calendar year 2025 (with maturity December 2024; lower right).

2 Market data

2.1 Futures contracts' quotes

The market data used to calibrate the futures contracts' instantaneous correlation structure are the realized closing market quotes of unitary base-load³ futures contracts i.e. typically the price per MWh for the delivery of electricity every hour of a given delivery period⁴. We present here several types of liquid futures contracts on power markets, these notations that will be used further in Section 4.

- Day-ahead contracts, denoted by DayX with $X \in \mathbb{N}^*$, typically the day-ahead contract (DA1) whose price is set by auction for the delivery of electricity for the following day, with T_s set to the first hour of delivery the next day and T_e equals T_s plus one day.
- Week-ends, denoted by WeekX with $X \in \mathbb{N}^*$, working days or week-ahead contracts, and also balance of month contracts which quotes the price of delivery of electricity such that, in these cases, T_s is set accordingly depending on the observation date, and T_e equals T_s plus, respectively, two days, five days and the number of days until the end of the running month.
- Monthly contracts denoted by Mon YY: delivering during the fixed month “Mon” of the year “20YY”. For these contracts, T_s and T_e are equal to the first and the last day of the delivery month.
- Quarterly contracts denoted by QX YY: delivery period is the “X”-th quarter of the year “20YY” with $X \in \{1, 2, 3, 4\}$. Here, the first quarter corresponds to the first three months of the year, the second one corresponds to the fourth-sixth months, and so on. In this case, T_s is equal to the first date of the first month of the quarter, and T_e is the last day of the last quarter month.
- Calendar contracts denoted by Cal YY: contracts delivering during the whole year “20YY”. T_s and T_e correspond to the first and the last days of the year.

³The other main traded profile, although less liquid, is the peak-load delivering electricity during working hours of the week. We chose to focus on the base-load data for illustrating our approach.

⁴A precise specification of the contract can be found at <https://www.eex.com/en/markets/power/power-futures> for the German power market and at <https://www.theice.com/products/27996665/Dutch-TTF-Gas-Futures> for futures traded on the TTF gas market.

Similarly, we use the notation rContract to refer to the rolling future contract with same time to delivery and delivery duration as Contract.

2.1.1 Stripping optimization to construct rolling futures contracts

Recall that, except for the first day-ahead contract, a rolling futures contract with time to delivery T_s and delivery duration $T_e - T_s$ is typically not exchanged in standard markets⁵, and therefore its quote cannot be observed directly. Yet, it is possible to construct the quotes of such $P_{\text{hist}} \in \mathbb{N}^*$ rolling futures contracts from the ones of the absolute futures contracts observed on the market via a stripped forward curve defined as follows.

Definition 2.1 (Stripped forward curve). *Fix an observation date $t_0 \geq 0$ and consider a family of $P' \in \mathbb{N}^*$ absolute futures contracts' quotes*

$$(F_{t_0}^{\text{mkt}}(T_s^j, T_e^j))_{j \in \{1, \dots, P'\}},$$

and denote $\bar{T} := \max_{j \in \{1, \dots, P'\}} T_e^j$. Then, a stripped forward curve, at observation date t_0 , is a smooth⁶ price curve $(f^{\text{mkt}}(t_0, T))_{T \in [t_0, \bar{T}]}$ such that the absence of superposition arbitrage opportunities is verified, i.e. such that the equality constraints

$$F_{t_0}^{\text{mkt}}(T_s^j, T_e^j) = \frac{1}{T_e^j - T_s^j} \int_{T_s^j}^{T_e^j} f^{\text{mkt}}(t_0, T) dT, \quad j \in \{1, \dots, P'\} \quad (2.1)$$

are all satisfied. In particular, the quote of any rolling contract at date t_0 can be obtained from the stripped forward curve f^{mkt} by the following formula

$$F_{t_0}^{\text{mkt}}(t_0 + T_s, t_0 + T_e) := \frac{1}{T_e - T_s} \int_{t_0 + T_s}^{t_0 + T_e} f^{\text{mkt}}(t_0, T) dT, \quad (2.2)$$

and we define its realized, or historical, log returns over a period of τ_d days at date t_0 by

$$r_{t_0}^{\text{mkt}}(\tau_d) := \log \frac{F_{t_0}^{\text{mkt}}(t_0 + T_s, t_0 + T_e)}{F_{t_0 - \tau_d}^{\text{mkt}}(t_0 + T_s, t_0 + T_e)}, \quad (2.3)$$

where the delivery period is held fixed, so that their respective sample Fourier spectra (resp. auto-correlations) don't display any significant frequency peak (resp. lag) i.e. absence of weekly effect from working days to week-ends (resp. independent increments assumption) hold valid with such log returns construction, see for example (Cartea and Figueroa, 2005, Figure 4) or (Gardini and Santilli, 2024, Figure 3).

In order to formulate the stripping optimization problem, we use the same notations as in Definition 2.1, and we set the daily time grids

$$\begin{aligned} \mathcal{T}_n(t_0) &:= \{t_0 = T_0 < T_1 < \dots < T_n = \bar{T}\}, \\ \mathcal{T}_{n^j}^j &:= \{T_s^j = T_0^j < T_1^j < \dots < T_{n^j}^j = T_e^j\}, \quad j \in \{1, \dots, P'\}, \end{aligned}$$

where $n \in \mathbb{N}^*$ (resp. $n^j \in \mathbb{N}^*$, $j \in \{1, \dots, P'\}$) denotes the number of days in the time interval $[t_0, \bar{T}]$ (resp. in $[T_s^j, T_e^j]$, $j \in \{1, \dots, P'\}$).

As detailed in Definition 2.1, the stripping optimization problem aims to build a stripped daily forward curve from a family of absolute forwards' quotes $F^{\text{mkt}}(t_0) := (F_{t_0}^{\text{mkt}}(T_s^j, T_e^j))_{j \in \{1, \dots, P'\}}$ observed on the markets such that

⁵Although at some specific observation dates t_0 , a rolling future with time to delivery $\delta > 0$ and delivery duration $\Delta_T > 0$ may coincide with the absolute futures contract delivering on $[T_s, T_e]$ if

$$\begin{cases} \delta = T_s - t_0 \\ \Delta_T = T_e - T_s \end{cases}.$$

⁶ f^{mkt} can be uniquely defined when some additional regularizing criterion is imposed.

- (i) the absence of superposition arbitrage opportunities (2.1) is satisfied;
- (ii) a smoothing criterion is applied ensuring the well-posedness of the algorithm: the sum of daily increments of the curve squared is minimized.

Then, the daily stripped forward curve at observation date t_0 is given by $(f^{\text{mkt}}(t_0, T_i))_{i \in \{0, \dots, n-1\}} := (\hat{s}(t_0, T_i))_{i \in \{0, \dots, n-1\}}$, where \hat{s} is defined as the unique minimizer of

$$\min_{s \in A^n(F^{\text{mkt}}(t_0))} \sum_{i=1}^{n-1} (s_i - s_{i-1})^2, \quad (2.4)$$

with $A^n(F^{\text{mkt}}(t_0))$ denoting the set of admissible curves defined by

$$A^n(F^{\text{mkt}}(t_0)) := \left\{ (s_l)_{l \in \{0, \dots, n-1\}} \in \mathbb{R}^n \mid F_{t_0}^{\text{mkt}}(T_s^j, T_e^j) = \sum_{l=0}^{n^j-1} w_{j, T_l} s_l, j \in \{1, \dots, P'\} \right\},$$

and where, for $j \in \{1, \dots, P'\}$, $(w_{j, T_l})_{l \in \{0, \dots, n^j-1\}} \in [0, 1]^{n^j}$ denotes the family of daily weights associated to the forward $F_{t_0}^{\text{mkt}}(T_s^j, T_e^j)$ depending on *the profiling of the underlying delivery*, e.g. $\frac{1}{365.25}$ for a calendar contract in base profile.

The optimization problem (2.4) can be reformulated as a Quadratic Program with linear equality constraints in the sense of (Boyd and Vandenberghe, 2004, Chapter 4, 4.4) and admits a unique solution by strict convexity. Market data should be without absence of arbitrage to ensure the existence of such solution. Some day-ahead futures contracts may display negative quotes, hence $s_l \in \mathbb{R}$, $l \in \{0, \dots, n-1\}$ in general, yet it is possible to look for a non-negative stripped curve by adding the constraints $s_l \in \mathbb{R}^+$, $l \in \{0, \dots, n-1\}$ if all the quoting market futures contracts are non-negative. The interested reader may look at (Benth and Koekebakker, 2008, Chapter 7) for additional insights on stripping including seasonal effects.

2.1.2 Estimated covariance of rolling futures' log returns from past historical data

Given a family of stripped forward curves $\left((f^{\text{mkt}}(t_h, T_i))_{t_h \leq T_i \leq \bar{T}} \right)_{h \in \{0, \dots, H\}}$ constructed at various observation dates $(t_h)_{h \in \{0, \dots, H\}}$ by solving the stripping optimization problem (2.4), we can reconstitute any family of $P_{\text{hist}} \in \mathbb{N}^*$ rolling forward contracts $(F_{t_h}^{\text{mkt}}(t_h + T_s^k, t_h + T_e^k))_{k \in \{1, \dots, P_{\text{hist}}\}}$ satisfying $t_h + T_e^k \leq \bar{T}$, $k \in \{1, \dots, P_{\text{hist}}\}$ for any $h \in \{0, \dots, H\}$ using (2.2), and then estimate a family of time series of their respective realized daily log returns $\left((r_{t_h}^{\text{mkt}, k}(\tau_d))_{h \in \{0, \dots, H\}} \right)_{k \in \{1, \dots, P_{\text{hist}}\}}$ by setting $\tau_d = 1$ day using (2.3).

Without any surprise, the rolling futures' log returns display heavy tails across all deliveries, and particularly for short-term ones. Since the futures' log returns' historical covariances will be used to calibrate a log normal distribution under the KV approximation (1.13) in the first calibration step, i.e. calibrate (σ, R) , with σ specified in (1.23), we first filter out futures contracts' log returns' large values away from three standard deviations (computed on the whole time-series) and setting them equal to three standard deviations, in the same spirit as in Cartea and Figueroa (2005).

Then, we construct a biased averaged covariance estimator of such daily log returns in order to assign higher weights to newer data points and diminishing weights to older data, and we extract some confidence bounds when weighting differently the past observations. We postpone to Appendix B.1 the detailed estimation procedure.

Finally, we extract the futures contracts' volatilities by taking the square root of the diagonal of the resulting covariance estimator, and their correlations by normalizing their respective covariances by the corresponding volatilities. Figure 5 displays the resulting realized correlation and volatility term structures on the German power market from January 1st 2023 to July 1st 2024.

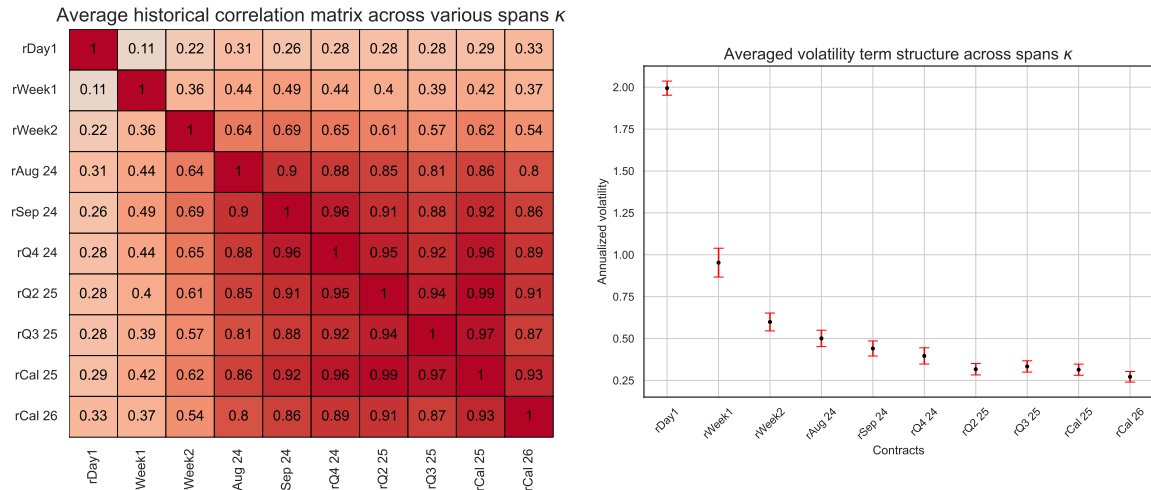


Figure 5: Left: averaged historical correlation matrix of rolling futures contracts’ daily log returns on the German market; right: averaged historical volatility term structure of rolling futures contracts’ daily log returns. The covariance estimator and confidence intervals are derived by averaging estimated covariances over a family of span parameters κ going from 30 to 365 days to put more weights on recent observations, see Appendix B.1.

2.2 Options data

2.2.1 Typical options quoted in power markets

We are interested in calibrating a diverse set of volatility smiles corresponding to different underlying futures contracts. The most liquid options are typically written on three types of underlying futures contracts: the first few months, quarters and calendar contracts. These liquid options are linear combinations of vanilla options like call spreads, put spreads, butterflies or fences. The descriptions of vanilla options can be found on the ice website (ICE).⁷

Vanilla options on monthly or quarterly futures have a single maturity: they expire several business day before the first day of the delivery period. However, vanilla options on calendar futures can have different maturities. The December maturity is generally the most liquid. The other maturities such as September, June and March — referred to as *early expiries* — are also quoted. To avoid confusion, the maturity month (of the year preceding the delivery year) is specified in the option name. For example, “Cal 26 Mar” corresponds to the smile with maturity at the end of March 2025. The precise maturities can be found on the site of the European Energy Exchange (EEX).⁸

For our numerical experiments, we use the following ten implied volatility smiles observed on the on the 1st July 2024 from the German power option market corresponding to eight different underlying futures contracts: Aug 24, Sep 24, Q4 24, Q2 25, Q3 25, Cal 25 Sep, Cal 25 Dec, Cal 26 Mar, Cal 26 Dec, Cal 27 Dec. Note that in energy markets, a large part of the volume is traded over the counter (OTC), unlike in equity markets, where options are typically traded via exchanges in a centralized manner. Consequently, there exist many possible sources of option price data. For our study, we use implied volatility data provided by the broker ICAP.⁹

2.2.2 Variance swap volatilities deduced from option prices

Variance Swap (VS) volatility can be interpreted as an overall smile level and will be used for the first two steps of calibration, recall Figure 3, and can be deduced from VS contracts, recall Section 1.3. As already

⁷<https://www.ice.com/products/65898946/German-Power-Financial-Base-Options>

⁸<https://www.eex.com/en/market-data/power/equity-styled-options>

⁹<https://icap.com/energy-commodities>

mentioned, such contracts are not quoted on power markets. However, we show how variance swap market prices can be extracted from vanilla option prices.

Indeed, these prices can be expressed as prices of European options with logarithmic payoff, further referred to as *log-contracts*, that can be extracted from the implied volatility smile by the Carr and Madan (1998) formula:

$$\text{VS}_T = -2\mathbb{E} \log \frac{F_T}{F_0} = 2 \int_0^{F_0} \frac{P(T, K)}{K^2} dK + 2 \int_{F_0}^{\infty} \frac{C(T, K)}{K^2} dK, \quad (2.5)$$

where $P(T, K)$ and $C(T, K)$ denote respectively the prices of put and call vanilla options.

Notations. The calibration will be performed on call and put options written on the futures contracts $F^i := F(T_s^i, T_e^i)$ for $i \in \{1, \dots, P_{\text{imp}}\}$, where the i -th futures contract is identified with its delivery period $[T_s^i, T_e^i]$. For $i \in \{1, \dots, P_{\text{imp}}\}$, we denote by N_i the number of (sorted) maturities $(T_j^i)_{j \in \{1, \dots, N_i\}}$ i.e. the number of smiles, associated to the underlying F^i . The implied market data contains a family of European call and put option prices

$$\left\{ \text{Call}^{\text{mkt}, i}(T_j^i, K) \right\}_{K \in \mathcal{K}_j^i}, \quad \left\{ \text{Put}^{\text{mkt}, i}(T_j^i, K) \right\}_{K \in \mathcal{K}_j^i} \quad j \in \{1, \dots, N_i\}, \quad i \in \{1, \dots, P_{\text{imp}}\},$$

where $P_{\text{imp}} \in \mathbb{N}^*$ denotes the number of underlyings, \mathcal{K}_j^i denotes the set of strikes corresponding to maturity T_j^i , and $\text{Call}^{\text{mkt}, i}(T_j^i, K)$ (resp. $\text{Put}^{\text{mkt}, i}(T_j^i, K)$) denotes the price of the European call option with maturity T_j^i , strike K , and underlying F^i . We denote by $\sigma_{\text{BS}}^{\text{mkt}, i}(T_j^i, K)$ the corresponding Black implied volatility.

The market price of the variance swap with underlying F^i and maturity T_j^i is deduced from the call and put options by (2.5) as follows

$$\text{VS}_{T_j^i}^{\text{mkt}, i} := 2 \int_0^{F_0^i} \frac{\text{Put}^{\text{mkt}, i}(T_j^i, K)}{K^2} dK + 2 \int_{F_0^i}^{\infty} \frac{\text{Call}^{\text{mkt}, i}(T_j^i, K)}{K^2} dK, \quad (2.6)$$

where F_0^i stands for the initial price of the i -th underlying futures contract. Then the VS volatility is extracted by

$$\sigma_{\text{VS}, T_j^i}^{\text{mkt}, i} = \sqrt{\frac{1}{T} \text{VS}_{T_j^i}^{\text{mkt}, i}}. \quad (2.7)$$

Typically, the number of strikes per smile (no more than 9) present in market data is not enough to apply the discretized formula (2.6) directly. Thus, a consistent smiles arbitrage-free extrapolation model is needed. An example of such a model for surface parametrization is given by the SSVI parametrization of Gatheral and Jacquier (2014). However, in our case, multiple surfaces are present at the same time, and a more general model is needed. Thus, we propose a multi-contract SSVI parametrization of the implied volatility surfaces for all contracts which allows us to interpolate and extrapolate the smiles and compute numerically the integrals in (2.6).

Multi-contract SSVI parametrization. Initially introduced by Gatheral and Jacquier (2014), the Surface SVI (SSVI) parametrization specifies the total implied variance $w(T, k) = \sigma_{\text{BS}}^2(T, k)T$, with corresponding time to maturity T and log-moneyness $k := \log \frac{K}{F_0}$. Namely, the implied total variance of the i -th futures contract is given by

$$w_i(T, k) = \frac{\theta_{T, i}}{2} \left(1 + \rho_i \phi(\theta_{T, i})k + \sqrt{(\phi(\theta_{T, i})k + \rho_i)^2 + (1 - \rho_i^2)} \right), \quad (2.8)$$

where

- $\theta_{T, i} := \sigma_{\text{BS}, i}^2(T, 0)T$ is the at-the-money (ATM) total variance of the underlying F^i ;
- $\rho_i \in [-1, 1]$ is a skew parameter interpreted as a spot-vol correlation for the contract F^i ;

- ϕ is a parametric function satisfying a list of conditions to guarantee the absence of static arbitrage.

Due to the presence of multiple surfaces to be parametrized together, we specify individual correlation parameters $\rho_i \in [-1, 1]$ for each contract F^i , and a common function ϕ since we assume the same nature of volatility across all the underlying futures contracts. This choice appears consistent with observed market data and allows us to simultaneously parametrize all the volatility surfaces using a small number of parameters. We set

$$\phi(\theta) := \frac{\eta}{\theta^\gamma(1+\theta)^{1-\gamma}}, \quad \gamma \in (0, 0.5],$$

and

$$\eta \left(1 + \max_{i \in \{1, \dots, P_{\text{imp}}\}} |\rho_i| \right) \leq 2,$$

which is consistent with the empirically-observed term structure of the volatility skew and such that (2.8) generates static-arbitrage-free surfaces for the respective futures contracts (Gatheral and Jacquier, 2014, Remark 4.4).

A calibrated multi-contract SSVI model allows itself for an efficient smile interpolation and extrapolation for a fixed maturity T as well as for a surface extrapolation for the contracts being in the calibration set. However, in contrast to the HJM stochastic volatility model (1.1), it cannot generate a volatility surface for a futures contract $F^{\bar{i}}$ not present in the calibration set unless its ATM volatility $\sigma_{\text{BS},\bar{i}}(T, 0)$ and correlation parameter $\rho_{\bar{i}}$ are provided. Thus, it can interpolate individual volatility surfaces, but not the volatility hypercube. That is why we limit our use of this parametrization to the evaluation of the implied variance swap prices (2.5). In addition, the calibrated values $(\rho_i)_{i \in \{1, \dots, P_{\text{imp}}\}}$ can be further used to obtain an initial guess for the spot-vol correlations in the smile calibration procedure.

The results of the multi-contract SSVI calibration to the market data are shown in Figure 6. We show a 5% bis-ask spread as well to give an idea of the parametrization quality. The value of 5% spread is consistent with actual market liquidity and will be further used as a benchmark for smile calibration.

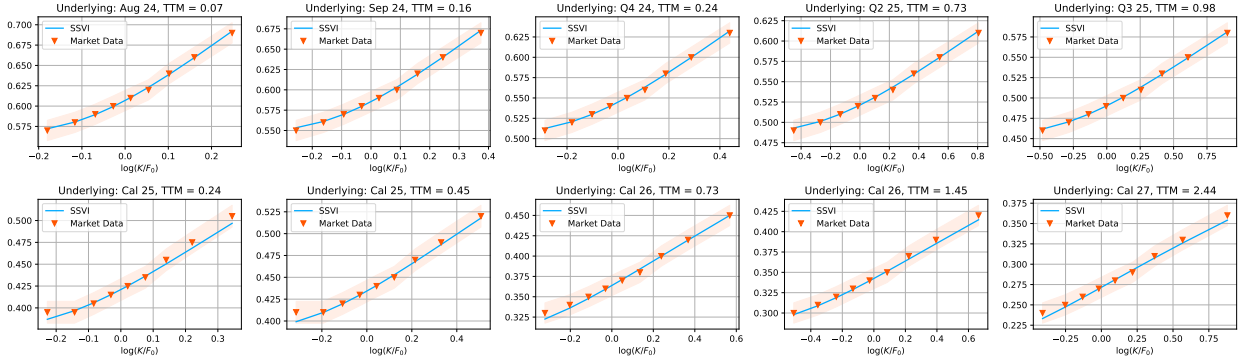


Figure 6: Example of the multi-factor SSVI calibration to the market data as of July 1st, 2024 on the German power market. We mention in the title of each plot the underlying futures contract and the Time To Maturity (TTM) of the smile in years.

3 Calibration methodology for historical covariances and variance swap volatilities

For such calibration step, we set

$$V_t \equiv 1, \quad h(t) \equiv 1, \quad g(T) \equiv 1, \quad 0 \leq t \leq T \leq \bar{T},$$

in (1.1). Our aim is to calibrate the N risk factors' volatility functions σ as well as their correlation matrix R in order to match

- the realized daily log returns' covariances estimated from (B.2) for P_{hist} different *rolling* futures contracts stripped from market quoting *absolute* futures contracts by non-arbitrage arguments using (2.2),
- the VS volatility term structure estimated from market data using the Carr-Madan formula (2.5) for the respective deliveries and maturities associated to P_{smiles} implied volatility smiles corresponding to P_{imp} different underlying futures contracts.

First, we derive explicit formulas of the normalized integrated covariance of rolling futures contracts' log returns and of the VS volatility when the infinitesimal forward rate's dynamics is given by the *L-S-C* specification of the form (1.23).

3.1 Explicit model covariance and VS volatility term structure

Let $H \in \mathbb{N}^*$ denote the number of days considered for the historical past horizon and $\tau_d > 0$ a period in days. In this section, we derive explicit formulas for the covariances of the rolling futures contracts' log-returns as well as their VS volatilities under the KV approximation (1.13).

Stationarity of rolling futures' volatility. To start with, by leveraging the stationary property of the *L-S-C* volatility functions (1.20)–(1.21)–(1.22), that is

$$\sigma(t, T) = \sigma(T - t), \quad 0 \leq t \leq T,$$

see for example Andersen (2010) for additional insights on this property, we readily obtain, by a change of variable, the following stationary property of the volatility (1.14) of the rolling forward contract such that

$$\Sigma_{t+u}(u + T_s, u + T_e) = \Sigma_t(T_s, T_e), \quad t \geq 0, \quad u \in \mathbb{R}. \quad (3.1)$$

Following the definition of the rolling forward contract (1.10) delivering on $[T_s, T_e]$, we define its log returns over a period of τ_d -days under the KV approximation by

$$\begin{aligned} r_{t_h}(\tau_d) := \log \frac{\widetilde{F}_{t_h}(t_h + T_s, t_h + T_e)}{\widetilde{F}_{t_h - \tau_d}(t_h + T_s, t_h + T_e)} &= -\frac{1}{2} \int_{t_h - \tau_d}^{t_h} \Sigma_t(t_h + T_s, t_h + T_e)^\top R \Sigma_t(t_h + T_s, t_h + T_e) dt \\ &\quad + \int_{t_h - \tau_d}^{t_h} \Sigma_t(t_h + T_s, t_h + T_e)^\top dW_t, \quad h \in \{0, \dots, H - \tau_d\}. \end{aligned}$$

Lemma 3.1. Fix two rolling futures contracts with respective time to deliveries T_s^i, T_s^j and delivery durations $T_e^i - T_s^i, T_e^j - T_s^j$. Define their normalized integrated covariance of their respective log returns over a period of τ_d -days $r^i(\tau_d), r^j(\tau_d)$ at observation date $t_h, h \in \{0, \dots, N\}$ by

$$\text{Cov} \left(r_{t_h}^i(\tau_d), r_{t_h}^j(\tau_d) \right) := \frac{1}{\tau_d} \int_{t_h - \tau_d}^{t_h} d \langle r_{t_h}^i(\tau_d), r_{t_h}^j(\tau_d) \rangle.$$

Then, the covariance is explicitly given by

$$\text{Cov} \left(r^i(\tau_d), r^j(\tau_d) \right) = \frac{1}{\tau_d} \int_0^{\tau_d} \Sigma_t(\tau_d + T_s^i, \tau_d + T_e^i)^\top R \Sigma_t(\tau_d + T_s^j, \tau_d + T_e^j) dt, \quad (3.2)$$

and is independent of the observation date t_h .

Proof. Straightforward calculus yields

$$\begin{aligned} \text{Cov} \left(r_{t_h}^i(\tau_d), r_{t_h}^j(\tau_d) \right) &= \frac{1}{\tau_d} \int_{t_h - \tau_d}^{t_h} \Sigma_t(t_h + T_s^i, t_h + T_e^i)^\top R \Sigma_t(t_h + T_s^j, t_h + T_e^j) dt \\ &= \frac{1}{\tau_d} \int_0^{\tau_d} \Sigma_{t+t_h - \tau_d}(t_h + T_s^i, t_h + T_e^i)^\top R \Sigma_{t+t_h - \tau_d}(t_h + T_s^j, t_h + T_e^j) dt \\ &= \frac{1}{\tau_d} \int_0^{\tau_d} \Sigma_t(\tau_d + T_s^i, \tau_d + T_e^i)^\top R \Sigma_t(\tau_d + T_s^j, \tau_d + T_e^j) dt, \end{aligned}$$

where we used respectively the change of variable $t \mapsto t + t_h - \tau_d$ and the *stationary property of the volatility function* Σ (3.1) to get the second and third equalities. \square

Separability of rolling futures' volatility. Furthermore, notice that the respective L , S and C volatility functions (1.20), (1.21) and (1.22) enjoy the following separable property such that for any $n \in \{1, \dots, N\}$

$$\sigma_n^Y(t, T) = \sigma_{Y,n} \sum_{k=1}^K \sigma_n^{X_k}(t, T) = \sigma_{Y,n} \sum_{k=1}^K a_{X_k,n}^\tau(t) b_{X_k,n}^\tau(T), \quad 0 \leq t \leq T, \quad Y \in \{L, S, C\}, \quad (3.3)$$

with $\sigma_{Y,n} > 0$, $(a_{X_k,n}^\tau)_{k \in \{1, \dots, K\}}$ and $(b_{X_k,n}^\tau)_{k \in \{1, \dots, K\}}$ some explicit functions depending only on parameters τ , and where $K \in \mathbb{N}^*$ denotes the total number of *separable state variables* $(X_k)_{k \in \{1, \dots, K\}}$. Notice that both L and S -factors are readily identified to their unique state variable respectively, while, for $j \in \{1, \dots, N_c\}$ we can decompose the j^{th} C -factor as the sum of two separable state variables C_1 and C_2 such that

$$\sigma_j^C(t, T) := \sigma_{C,j} \frac{T-t}{\tau_{C,j}} e^{-\frac{T-t}{\tau_{C,j}}} = \sigma_{C,j} \left(\sigma_j^{C_1}(t, T) + \sigma_j^{C_2}(t, T) \right), \quad 0 \leq t \leq T,$$

where we define the volatility functions associated respectively to C_1 and C_2 by

$$\begin{aligned} \sigma_j^{C_1}(t, T) &:= e^{\frac{t}{\tau_{C,j}}} \times \frac{T}{\tau_{C,j}} e^{-\frac{T}{\tau_{C,j}}}, \quad 0 \leq t \leq T, \\ \sigma_j^{C_2}(t, T) &:= -\frac{t}{\tau_{C,j}} e^{\frac{t}{\tau_{C,j}}} \times e^{-\frac{T}{\tau_{C,j}}}, \quad 0 \leq t \leq T. \end{aligned}$$

We detail in Table 1 the explicit functions a_X^τ and b_X^τ for the four distinct state variables $X \in \{L, S, C_1, C_2\}$.

Table 1: Explicit functions a_X^τ , b_X^τ from (3.3) and β_X^τ from (3.4) defined in(3.6) for the four state variables $X \in \{L, S, C_1, C_2\}$ from the L - S - C specification.

	L	S	C_1	C_2
$t \mapsto a_X^\tau(t)$	1	$e^{\frac{t}{\tau_S}}$	$e^{\frac{t}{\tau_C}}$	$-\frac{t}{\tau_C} e^{\frac{t}{\tau_C}}$
$T \mapsto b_X^\tau(T)$	1	$e^{-\frac{T}{\tau_S}}$	$\frac{T}{\tau_C} e^{-\frac{T}{\tau_C}}$	$e^{-\frac{T}{\tau_C}}$
$\beta_X^\tau(T_e, T_s)$	1	$\frac{\tau_S}{T_e - T_s} \left(e^{-\frac{T_s}{\tau_S}} - e^{-\frac{T_e}{\tau_S}} \right)$	$\frac{T_s + \tau_C}{T_e - T_s} e^{-\frac{T_s}{\tau_C}} - \frac{T_e + \tau_C}{T_e - T_s} e^{-\frac{T_e}{\tau_C}}$	$\frac{\tau_C}{T_e - T_s} \left(e^{-\frac{T_s}{\tau_C}} - e^{-\frac{T_e}{\tau_C}} \right)$

We are now ready to derive the explicit formulas of the rolling futures contracts log returns and VS volatilities which will play a key role in the formulation of the first step calibration problem.

Proposition 3.2. *Using the same notations as in Lemma 3.1, the normalized stationary covariance (3.2) is explicitly given by*

$$\text{Cov}(r^i(\tau_d), r^j(\tau_d)) = \sum_{p,k=1}^{N+N_c} x_{p,k}^{\sigma,R} \beta_{X_p}^\tau(T_e^i, T_s^i) \beta_{X_k}^\tau(T_e^j, T_s^j) \frac{1}{\tau_d} \int_0^{\tau_d} a_{X_p}^\tau(u) a_{X_k}^\tau(u) du, \quad (3.4)$$

where, for $X_p, X_k \in \{L, S, C_1, C_2\}$, we set the variables

$$x_{p,k}(\sigma, R) := \sigma_{X_p} \sigma_{X_k} R_{p,k}, \quad p, k \in \{1, \dots, N + N_c\}, \quad (3.5)$$

and

$$\beta_{X_p}^\tau(T_e^i, T_s^i) := \frac{1}{T_e^i - T_s^i} \int_{T_s^i}^{T_e^i} b_{X_p}^\tau(T) dT, \quad p \in \{1, \dots, N + N_c\}, \quad (3.6)$$

with $(\tau, \sigma, R) \in (\mathbb{R}_+^*)^N \times (\mathbb{R}_+^*)^N \times \mathbb{S}_{++}^N$.

Similarly, the variance swap variance level (1.18) is explicitly given by

$$\left(\sigma_{\text{VS},T}^{\text{model},i} \right)^2 = \sum_{p,k=1}^{N+N_c} x_{p,k}^{\sigma,R} \beta_{X_p}^\tau(T_e^i, T_s^i) \beta_{X_k}^\tau(T_e^i, T_s^i) \frac{1}{T} \int_0^T a_{X_p}^\tau(u) a_{X_k}^\tau(u) du. \quad (3.7)$$

In particular, all the terms $\left(\beta_{X_p}^\tau(T_e^i, T_s^i)\right)_{p \in \{1, \dots, N\}}$ and $\left(\int_0^{\tau_d} a_{X_p}^\tau(u) a_{X_k}^\tau(u) du\right)_{p, k \in \{1, \dots, N\}}$ are explicit in the L - S - C parametrization, as specified respectively in Tables 1 and 2.

Proof. The explicit formulas (3.4) and (3.7) are readily derived respectively from (3.2) and (1.18), while using the definition (1.14) and the separability property of the L - S - C factors (3.3). \square

L^p	$t_2 - t_1$		
S^p or C_1^p	$\tau_{S/C}^p \left(e^{\frac{t_2}{\tau_{S/C}^p}} - e^{\frac{t_1}{\tau_{S/C}^p}} \right)$	$\tilde{\tau}^{p,k} \left(e^{\frac{t_2}{\tilde{\tau}^{p,k}}} - e^{\frac{t_1}{\tilde{\tau}^{p,k}}} \right),$ $\tilde{\tau}^{p,k} := \frac{\tau_{S/C}^p \tau_{S/C}^k}{\tau_{S/C}^p + \tau_{S/C}^k}$	
C_2^p	$-\left((t_2 - \tau_C^p) e^{\frac{t_2}{\tau_C^p}} - (t_1 - \tau_C^p) e^{\frac{t_1}{\tau_C^p}} \right)$	$-\frac{\tilde{\tau}^{p,k}}{\tau_C^p} \left((t_2 - \tilde{\tau}^{p,k}) e^{\frac{t_2}{\tilde{\tau}^{p,k}}} - (t_1 - \tilde{\tau}^{p,k}) e^{\frac{t_1}{\tilde{\tau}^{p,k}}} \right),$ $\tilde{\tau}^{p,k} := \frac{\tau_{S/C}^p \tau_C^p}{\tau_{S/C}^k + \tau_C^p}$	$\frac{\tilde{\tau}^{p,k}}{\tau_C^a \tau_C^b} \left(\left(2(\tilde{\tau}^{p,k})^2 - 2\tilde{\tau}^{p,k} t_2 + t_2^2 \right) e^{\frac{t_2}{\tilde{\tau}^{p,k}}} - \left(2(\tilde{\tau}^{p,k})^2 - 2\tilde{\tau}^{p,k} t_1 + t_1^2 \right) e^{\frac{t_1}{\tilde{\tau}^{p,k}}} \right),$ $\tilde{\tau}^{p,k} := \frac{\tau_C^k \tau_C^p}{\tau_C^k + \tau_C^p}$
	L^k	S^k or C_1^k	C_2^k

Table 2: Explicit computations for the cross-terms $\int_{t_1}^{t_2} a_{X_p}^\tau(u) a_{X_k}^\tau(u) du$ used in (3.4) for the types of state variables $X_p, X_k \in \{L, S, C_1, C_2\}$, $p, k \in \{1, \dots, N\}$. Notice in particular that S and C_1 state variables share the same function a_{X}^τ as detailed in Table 1.

3.2 Loss function for the historical covariances - VS volatility term structure calibration problem

We consider $P_{\text{smiles}} := \sum_{i=1}^{P_{\text{imp}}} N_i \in \mathbb{N}^*$ Variance Swap (VS) contracts with respective maturities $\left((T_j^i)_{j \in \{1, \dots, N_i\}} \right)_{i \in \{1, \dots, P_{\text{imp}}\}}$ and volatilities $\left(\left(\sigma_{\text{VS}, T_j^i}^{\text{mkt}, i} \right)_{j \in \{1, \dots, N_i\}} \right)_{i \in \{1, \dots, P_{\text{imp}}\}}$ computed using (2.7). Such VS contracts are indeed associated to the implied volatility smiles we aim to calibrate the initial model (1.1) on, whose underlying are respectively the futures contracts with delivery periods $\left([T_s^i, T_e^i] \right)_{i \in \{1, \dots, P_{\text{imp}}\}}$, recall the notations from Section 2.2.2.

Denote by H the number of past observation days until the date of calibration considered for the calibration. Then we introduce a family of $P_{\text{hist}} \in \mathbb{N}^*$ rolling contracts $\left((F(t_h, t_h + T_s^k, t_h + T_e^k))_{h \in \{1, \dots, H\}} \right)_{k \in \{1, \dots, P_{\text{hist}}\}}$ constructed by solving (2.4) and we use their respective daily log returns time-series

$$\left(\left(r_{t_h}^{\text{mkt}, k}(\tau_d) \right)_{h \in \{0, \dots, H\}} \right)_{k \in \{1, \dots, P_{\text{hist}}\}}$$

computed from (2.3), for the estimation of the historical daily log returns' covariances via the estimator (B.2). In our case, we typically chose $P_{\text{hist}} \geq P_{\text{imp}}$, with the set of the rolling forwards' delivery periods including those of the absolute underlying futures contracts, i.e.

$$\left([T_s^i, T_e^i] \right)_{i \in \{1, \dots, P_{\text{imp}}\}} \subset \left([T_s^k, T_e^k] \right)_{k \in \{1, \dots, P_{\text{hist}}\}}$$

to ensure that the historical covariance term structure captures a priori, *somehow*, the *futures correlation structure* of those absolute futures.

We introduce a convex combination of losses between the fit of historical log returns' variance-covariance term structure of rolling forward contracts and the fit of the implied variance swap volatility levels such that

$$J^\lambda(\tau, \sigma, R) := \lambda J_1(\tau, \sigma, R) + (1 - \lambda) J_2(\tau, \sigma, R), \quad \lambda \in [0, 1], \quad (3.8)$$

with

$$J_1(\tau, \sigma, R) := \|C^{\text{mkt}} - C^{\text{model}}(\tau, \sigma, R)\|_{\Gamma}^2, \quad (3.9)$$

$$J_2(\tau, \sigma, R) := \|(\sigma_{\text{VS}}^{\text{mkt}})^2 - (\sigma_{\text{VS}}^{\text{model}})^2(\tau, \sigma, R)\|_w^2, \quad (3.10)$$

where the $(i, j)^{\text{th}}$ entry of C^{model} and l^{th} entry of $(\sigma_{\text{VS}}^{\text{model}})^2$ are respectively given by (3.4) and (3.7), while C^{mkt} and $\sigma_{\text{VS}}^{\text{mkt}}$ are respectively estimated by (B.2) and (2.5). Furthermore, $\Gamma := (\Gamma_{i,j})_{i,j \in \{1, \dots, P_{\text{hist}}\}^2} \in \mathbb{R}_+^{P_{\text{hist}}^2}$ denotes a family of weights associated to the matrix Frobenius norm, while $w := (w_i)_{i \in \{1, \dots, P_{\text{smiles}}\}} \in \mathbb{R}_+^{P_{\text{smiles}}}$ are weights for the vector Frobenius norm, specified respectively as in (B.3) and uniformly in our case. Finally, notice that the hyper-parameter λ conveniently controls the trade-off between the fit of the historical covariances and that of the VS volatility term structure.

Then, for a fixed $\lambda \in [0, 1]$, the joint historical - VS variance term structure calibration is formulated as the following minimization problem under constraints

$$\min_{(\tau, \sigma, R) \in U} J^\lambda(\tau, \sigma, R), \quad (3.11)$$

where the admissible set of parameters is given by

$$U := (\mathbb{R}_+^*)^N \times (\mathbb{R}_+^*)^N \times \mathbb{S}_{++}^N. \quad (3.12)$$

The functional J^λ (3.8) is clearly non-convex in (τ, σ, R) , depending in particular on exponential terms in the parameters τ , so we cannot guarantee a priori the existence of a global minimizer of such optimization problem (3.11).

3.3 Solver specification

The admissible set of parameters (3.12) can be very large, with dimension in $\mathcal{O}(N^4)$, where recall N is the total number of factors in the Nelson-Siegel parametrization (1.23), and include positive constraints as well as a non-trivial positive definite cone constraint for the correlation matrix R . In order to simplify and fasten the numerical implementation of the optimization problem (3.11), we start by noticing the three following facts.

- (i) As soon as the parameters τ are fixed, the functional J^λ (3.8) becomes quadratic in the variables $(x_{p,k}(\sigma, R))_{p,k \in \{1, \dots, N+N_c\}}$ from (3.5) such that

$$J^\lambda(\tau, \sigma, R) = \lambda \sum_{i,j=1}^{P_{\text{hist}}} W_{i,j} \left(C_{i,j}^{\text{mkt}} - \sum_{p,k=1}^{N+N_c} w_{p,k}^\tau(i,j) x_{p,k}^{\sigma,R} \right)^2 \quad (3.13)$$

$$+ (1 - \lambda) \sum_{l=1}^{P_{\text{smiles}}} w_l \left((\sigma_{\text{VS}}^{\text{mkt}})^2 - \sum_{p,k=1}^{N+N_c} w_{p,k}^\tau(l) x_{p,k}^{\sigma,R} \right)^2, \quad (3.14)$$

where, for $p, k \in \{1, \dots, N + N_c\}$, the weights in the squares are respectively identified from (3.4) and (3.7) such that

$$w_{p,k}^\tau(i,j) := \beta_{X_p}^\tau(T_e^i, T_s^i) \beta_{X_k}^\tau(T_e^j, T_s^j) \frac{1}{\tau_d} \int_0^{\tau_d} a_{X_p}^\tau(u) a_{X_k}^\tau(u) du, \quad i, j \in \{1, \dots, P_{\text{hist}}\}, \quad (3.15)$$

$$w_{p,k}^\tau(l) := \beta_{X_p}^\tau(T_e^l, T_s^l) \beta_{X_k}^\tau(T_e^l, T_s^l) \frac{1}{T} \int_0^T a_{X_p}^\tau(u) a_{X_k}^\tau(u) du, \quad l \in \{1, \dots, P_{\text{smiles}}\}. \quad (3.16)$$

- (ii) Note that the loss functional (3.14) is expressed in terms of the state variables, and for each $j \in \{0, \dots, N_c\}$, the j^{th} C -factor has two state variables sharing the same parameter $\sigma_{C,j}$ and the same

Brownian motion so that we need to impose the following $2N_c$ equality constraints on the x variables (3.5) ordered as $\left(\left\{ L, (S_i)_{i \in \{0, \dots, N_s\}}, ((C_{1,j}, C_{2,j}))_{i \in \{0, \dots, N_c\}} \right\} \right)$ such that

$$\begin{cases} x_{l,l}(\sigma, R) = x_{l+1,l+1}(\sigma, R) \\ x_{l,l}(\sigma, R) = x_{l,l+1}(\sigma, R) \end{cases}, \quad l \in \{N_s + 1, \dots, N_s + 1 + N_c\}. \quad (3.17)$$

(iii) S (resp. C -factors) are inter-changeable, which may cause numerical instability.

Successive non-linear – linear cone program formulation. Consequently, instead of solving (3.11) globally, we will solve the following iterative minimization problems

$$\min_{a \in \mathbb{R}^{N_s + N_c}} \min_{x(\sigma, R): (\sigma, R) \in (\mathbb{R}_+^*)^N \times \mathbb{S}_{++}^N} J^\lambda(\tau(a), x(\sigma, R)). \quad (3.18)$$

On the one hand, the outward minimization is performed by an unconstrained non-linear solver (e.g. [Powell minimizer from SciPy](#)), where we optimize on the parameters

$$a := (a_s^1, \dots, a_s^{N_s}, a_c^1, \dots, a_c^{N_c}) \in \mathbb{R}^{N_s + N_c},$$

such that the following change of variables

$$\tau(a) := \left(e^{a_s^1}, e^{a_s^1} + e^{a_s^2}, \dots, \sum_{i=1}^{N_s} e^{a_s^i}, e^{a_c^1}, e^{a_c^1} + e^{a_c^2}, \dots, \sum_{i=1}^{N_c} e^{a_c^i} \right) \quad (3.19)$$

ensures that the $\tau(a)$ parameters are indeed positive and strictly increasing for the N_s S -factors and the N_c C -factors respectively.

On the other hand, for a fixed $a \in \mathbb{R}^{N_s + N_c}$, $\tau(a)$ is fixed and the inward minimization problem in (3.18) can be formulated and solved in terms of the variables $(x_{p,k}(\sigma, R))_{p,k \in \{1, \dots, N + N_c\}}$ from (3.5) as a linear cone program ensuring R is indeed semi-definite positive, as detailed in [Appendix B.2](#).

Solver initialization and optimal parameters extraction. All that remains to do is to initialize properly the outward non-linear solver in the iterative formulation (3.18) and then extract the optimal parameters (σ, R) upon convergence. Indeed, a good initialization $a_0 \in \mathbb{R}^{N_s + N_c}$ is of paramount importance to be able to reach a good local minimum in practice.

Fortunately the solver algorithm runs relatively fast, from a few seconds to a few minutes depending on the number of factors, so we can afford to iterate over various randomized initial guesses. As a rule of thumb, we construct the initial values $a_0 \in \mathbb{R}^{N_s + N_c}$ starting from the mid-points of delivery periods $([T_s^k, T_e^k])_{k \in \{1, \dots, P_{\text{hist}}\}}$ such that the resulting factors with time-scales $\tau(a_0)$ cover reasonably well the futures curve across all deliveries. In our case, since we have $N_s \leq P_{\text{hist}}$ and $N_c \leq P_{\text{hist}}$, we first reduce the number of such mid-points by taking convex combinations of them to the number of S and C factors, and then add up some randomness to obtain the initial positive $\tau(a_0)$.

Once an optimizer $(\hat{a}, \hat{x}(\hat{a}))$ of the iterative optimization problem (3.18) has been reached, then we extract the optimal parameters $\hat{\tau}$ injecting \hat{a} into (3.19), and by positive-definiteness of the calibrated L - S - C factors' covariance $(\hat{x}_{i,j})_{i,j \in I}$ (extracted from an optimal solution $\hat{x}(\hat{a})$ of (B.4), with I denoting the set of indices relative to $\{L, S, C_1\}$ state variables, of cardinal N), we get

$$\hat{\sigma}_i := \sqrt{\hat{x}_{i,i}} > 0, \quad i \in I, \quad (3.20)$$

and

$$\hat{R}_{i,j} := \frac{\hat{x}_{i,j}}{\hat{\sigma}_i \hat{\sigma}_j}, \quad i, j \in \{1, \dots, N\}. \quad (3.21)$$

Remark 3.3. Note that the linear cone program (B.4) only guarantees the matrix $(\hat{x}_{i,j})_{i,j \in I}$ to be semi-definite positive while the inversion formulas (3.20)–(3.21) are well-defined if $(\hat{x}_{i,j})_{i,j \in I}$ is positive definite. If zero happens to belong to the spectrum of $(\hat{x}_{i,j})_{i,j \in I}$, then it means the number of factors could be reduced, either by withdrawing any i^{th} factor associated to $\hat{x}_{i,i} = 0$, or, if \hat{R} is indeed well-defined by (3.21), by summing (resp. subtracting) the volatility shape functions of the perfectly correlated (resp. anti-correlated) factors with respect to all the other factors. See for example (Féron and Gruet, 2024, Table 13-16-19) where such latter phenomenon is repeatedly observed when calibrating historical futures' returns on the Italian, Swiss and UK markets respectively. In practice however, it is always possible to regularize or to tune the hyper-parameters of the non-linear solver such that $(\hat{x}_{i,j})_{i,j \in I} \in \mathbb{S}_{++}^N$ (e.g. the tolerance threshold, the number of iterations, etc).

4 Calibration methodology for implied smiles

Implied calibration focuses on calibrating the model to observed market option prices. We show how the prices of vanilla options can be computed in our model and propose a way to decouple the implied calibration into the correction of the implied volatility term structure given by the VS volatilities to achieve a perfect fit, and the calibration of smile shapes.

4.1 Fourier option pricing for implied calibration

In this section, we describe the vanilla options pricing techniques used in the implied calibration for a futures contract $\tilde{F} = \tilde{F}(T_s, T_e)$ with arbitrary delivery period $[T_s, T_e]$ in our model, where \tilde{F} is given by (1.15) following the KV approximation.

The characteristic function of $\log \tilde{F}$. The European call option prices $C(T, K)$ can be calculated using the Lewis (2001) formula

$$C(T, K) = \mathbb{E}(\tilde{F}_T - K)^+ = F_0 - \frac{K}{\pi} \int_0^\infty \Re \left(e^{i(u - \frac{i}{2}) \log \frac{F_0}{K}} \phi \left(u - \frac{i}{2} \right) \right) \frac{du}{u^2 + \frac{1}{4}}, \quad (4.1)$$

where $\phi(u) := \mathbb{E} \left[e^{iu \log \frac{\tilde{F}_T}{F_0}} \right]$, $u \in \mathbb{R}$ denotes the characteristic function of the normalized log-price. Due to the affine structure of the variance diffusion of the lifted Heston model of [Abi Jaber \(2019\)](#), we show in the following theorem that the characteristic function ϕ of $\log \tilde{F}$ in the model (1.15) is known up to the solution of a Riccati equation.

Theorem 4.1. Fix $v \in \mathbb{C}$ such that $\Re(v) \in [0, 1]$. For $T \leq T_s$, the characteristic function of $\log \frac{\tilde{F}_T}{F_0}$ is given by

$$\mathbb{E} \left[\exp \left(v \log \frac{\tilde{F}_T}{F_0} \right) \right] = \exp \left(\int_0^T G(T-s, v, \psi(s)) ds \right),$$

where $\psi := \sum_{j=1}^M c_j \psi^j$, and ψ^j satisfies the Riccati equation

$$\begin{cases} \dot{\psi}^j(t) &= -x_j \psi^j(t) + G(T-t, v, \psi(t)) \quad \text{a.e.}, \\ \psi^j(0) &= 0 \end{cases}, \quad (4.2)$$

for $j \in \{1, \dots, M\}$, and the function G is defined by

$$G(t, v, \psi) := \frac{h(t)^2}{2} \Sigma_t^\top R \Sigma_t (v^2 - v) + h(t) (\Sigma_t^\top \tilde{\rho}) v \psi + \frac{\psi^2}{2}.$$

The Riccati equation (4.2) admits a unique global solution $\psi \in L_{\text{loc}}^2(\mathbb{R}_+, \mathbb{C})$, which is differentiable a.e. and satisfies $\Re(\psi) \leq 0$.

Proof. The futures price dynamics (1.15) can be rewritten as a one-dimensional diffusion, i.e. there exists a Brownian motion \tilde{W} such that

$$\frac{d\tilde{F}_t}{\tilde{F}_t} = h(t)\sqrt{V_t}\sqrt{\Sigma_t^\top R\Sigma_t}d\tilde{W}_t, \quad d\langle B, \tilde{W} \rangle_t = \Sigma_t^\top \tilde{\rho} dt, \quad t \in [0, T_s].$$

Note that our model is a particular specification of the Volterra Heston model (Abi Jaber et al., 2019, Section 7) setting the kernel to

$$K(t) := \sum_{j=1}^M c_j e^{-x_j t}, \quad t \geq 0,$$

and $V_0 = 1$, $\theta = 1$, $\lambda = 0$, $\nu = 1$. Repeating the proof of (Abi Jaber et al., 2019, Theorem 7.1 (ii)) adapted to take into account the bounded deterministic volatility component $t \mapsto \theta(t) := h(t)\sqrt{\Sigma_t^\top R\Sigma_t}$, we can prove the existence of $\psi \in L_{\text{loc}}^2(\mathbb{R}_+, \mathbb{C})$ satisfying $\Re(\psi) \leq 0$, which is the solution to the Volterra–Riccati equation

$$\psi(t) = \sum_{j=1}^M c_j \int_0^t e^{-x_j(t-s)} G(T-s, v, \psi(s)) ds, \quad t \geq 0.$$

The solution can be written in the form $\psi = \sum_{j=1}^M c_j \psi^j$, where, for $j \in \{1, \dots, M\}$, ψ^j satisfies

$$\psi^j(t) = \int_0^t e^{-x_j(t-s)} G(T-s, v, \psi(s)) ds, \quad (4.3)$$

which is equivalent to the equation (4.2) by the variation of constants formula. Moreover, (4.3) implies that ψ^j is absolutely continuous, hence ψ is differentiable a.e., which concludes the proof. \square

Numerical scheme for the Riccati equation. For $j \in \{1, \dots, M\}$, the Riccati equation (4.2) can be rewritten in the integral form

$$\psi^j(t) = \int_0^t e^{-x_j(t-s)} G(T-s, iu, \psi(s)) ds, \quad t \geq 0,$$

so that $\psi^j(t)$ satisfies

$$\begin{aligned} \psi^j(t+h) &= e^{-x_j h} \psi^j(t) + \int_t^{t+h} e^{-x_j(t+h-s)} G(T-s, iu, \psi(s)) ds \\ &\approx e^{-x_j h} \psi^j(t) + G(T-t, iu, \psi(t)) \int_t^{t+h} e^{-x_j(t+h-s)} ds \\ &= e^{-x_j h} \psi^j(t) + \frac{1 - e^{-x_j h}}{x_j} G(T-t, iu, \psi(t)), \end{aligned}$$

which is used as a numerical scheme to solve the Riccati equation (4.2). Such numerical scheme is more stable and precise than the explicit Euler discretization scheme, especially if one of the mean-reversion speeds x_j , $j \in \{1, \dots, M\}$ reaches extreme values.

4.2 Exact calibration correction of the VS volatility term structure

The second step of the joint calibration consists of the correction and perfect match of the VS volatility term structure using the functions g and h . Precise calibration of the VS volatilities is indeed crucial for achieving a good overall implied fit in step 3: since $\mathbb{E}[V_t] \equiv 1$, $t \geq 0$, the stochastic volatility parameters responsible for the smile shape do not impact the VS volatilities, and hence cannot correct the VS volatility term structure.

Although the VS volatilities are calibrated by the L - S - C risk factors in step 1, as described in Section 3, their fit is often not exact since the primary goal of this step is to capture the historical futures' correlation structure with as few factors as possible but also match a priori the VS volatilities, where fitting the latter can be considered an *implied regularization* in the loss J^λ (3.8). Yet, at the same time, this regularization is necessary to capture the VS volatility term structure at least approximately in order to avoid over-fitting the functions g and h . Indeed, if the VS term structure were calibrated poorly at step 1, then the functions g and h would vary significantly, which would lead to a strong non-stationarity of the volatility (see Lemma 3.1) and eventually degrade the model's extrapolation ability, mostly induced by step 1. Hence, it is important to keep these functions g and h close to one in practice.

In what follows, we assume that the functions $T \mapsto g(T)$ and $t \mapsto h(t)$ are piece-wise constant. The discontinuity points of the function g are the delivery start and end dates of the underlying futures contracts, while the discontinuities of h are the maturities of the associated vanilla options.

Motivating example. Here, we provide some intuition on the calibration of the functions h and g and clarify why only one of these functions is not enough to fit the VS volatility term structure by considering a simple calibration set containing only five futures contracts: Oct 24, Nov 24, Dec 24, Q4 24, and Cal 25, with observation date the 1st September 2024, see Figure 7. The maturities of options on monthly contracts are five days before the start of delivery (at T_1, T_2 , and T_3 for Oct 24, Nov 24, and Dec24 respectively). The quarterly contract Q4 24 expires in September at T_1 , and the calendar futures Cal 25 has two maturities T_1 and T_4 corresponding to September and December expiries.

The VS term structure calibration consists in matching model and market log-contract prices respectively given by (1.17) and (2.6), i.e. for any IV smile on a futures contract with delivery period $[T_s, T_e]$ and maturity T , one should ensure that

$$\int_0^T h(t)^2 \Sigma_t^\top(T_s, T_e) R \Sigma_t(T_s, T_e) dt = \text{VS}_T^{\text{mkt}}, \quad (4.4)$$

where $\Sigma_t(T_s, T_e) = \frac{1}{T_e - T_s} \int_{T_s}^{T_e} \sigma(\cdot, \tau) g(\tau) d\tau$. This can be achieved in two ways:

- either adjusting the function $T \mapsto g(T)$ on the interval $[T_s, T_e]$;
- or modifying the function $t \mapsto h(t)$ on $[0, T]$.

In our example, each monthly futures can be calibrated with the function $T \mapsto g(T)$ which is constant on each of three months with the values chosen to verify (4.4). However, if all the three monthly futures are calibrated with the function g , there is no more degree of freedom to adjust the variance swap price for the quarterly contract Q4 24. Alternatively, one could calibrate three smiles on Oct 24, Nov 24 and Q4 24 using the value of g on Dec 24 to fit the calibration condition (4.4) for the smile on Q4 24.

Moreover, the function g cannot calibrate the early expiries of the calendar contract Cal 25 (for instance, calibrate simultaneously the log-contract prices for the maturities T_1 and T_4) since g impacts on the overall volatility level of the futures contract: therefore g can fit the smile level for only one maturity per delivery.

For these reasons, the function $t \mapsto h(t)$ needs to be introduced: we could fix the value of h on $[0, T_1]$ to satisfy (4.4) for the smile on Cal 25 with maturity T_1 , then adjust h on $[T_1, T_4]$ to ensure (4.4) holds its smile with maturity T_4 . Alternatively, one could calibrate the value of h on $[0, T_1]$ to fit (4.4) for the smile on Q4 24 while the smiles on monthly futures are calibrated with the function g . However, two problems arise.

First, it is impossible to calibrate simultaneously both smiles written on Q4 24 and on Cal 25 since they share the same maturity T_1 . Hence, one of them should be necessarily calibrated with the function g . In this example, the following solution is possible:

1. the function g is calibrated to the smiles on Oct 24, Nov 24, Dec 24, and Cal 25 with maturity T_1 ,
2. the function h is calibrated to the smiles on Q4 24 and on Cal 25 with maturity T_4 .

Second, once the monthly contracts are calibrated with the function g , any modification of the function h (up to the smile maturity) made during the calibration of smile on Q4 24, will impact the log-contract price

corresponding to the monthly futures, and vice-versa. This is simply due to the fact that the log-contract price in (4.4) depends both on g on $[T_s, T_e]$ and on h on $[0, T]$.

That is why an iterative calibration algorithm is needed. One step of such algorithm should include the calibration of the function g to fit (4.4) for one part of smiles and then the calibration of h taking into account the modifications of g to verify (4.4) for the remaining smiles (possibly breaking this condition for the smiles from the first group). If such iterative algorithm converges to a fixed-point, at the limit, the smiles from both groups will be calibrated. This fixed-point algorithm is the key idea of the proposed calibration method and will be discussed in more detail below.

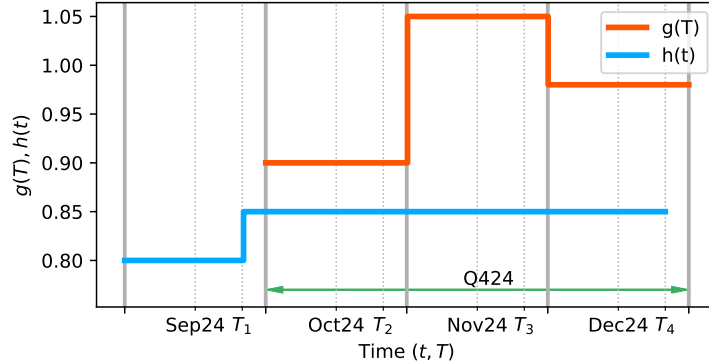


Figure 7: An example of the functions g (in dark orange) and h (in blue) covering the period from September 1, 2024 to January 1, 2025. The function g is constant in October, in November and in December (months are separated by solid gray lines). The function h is constant between the option maturities T_1, T_2, T_3, T_4 corresponding to five days before the start of the next months.

Two groups of smiles. From the discussion above it is clear that the implied volatility smiles should be divided into two groups: for the smiles in the first group the calibration condition is guaranteed by the function g , while the smiles from the second one are calibrated with the function h . It also follows from the example, that these groups should satisfy the following conditions.

Conditions on the smiles calibrated with the function g :

1. **one maturity per futures:** Only one smile for a given futures contract can be calibrated with g ,
2. **linear independence:** The underlying futures corresponding to these smiles should be linearly independent, i.e. no delivery period should be representable as a union of delivery periods of other futures contracts.

Condition on the smiles calibrated with the function h :

- 3 **no coinciding maturities:** for a given maturity, only one smile may be chosen.

Note that such division may be not unique. To fix the division, we decide to calibrate the smiles on monthly and quarterly contracts with the function g and all the smiles on calendar contracts with the function h . Although the latter arbitrary division is not always available, it is almost always possible to find another division satisfying the conditions 1–3.

Calibration of the function $t \mapsto h(t)$. We strip the function h to match the variance swap prices implied from the volatility smiles in (2.5) and the variance swap prices given by the model (1.17) for futures contracts F^i , for each $i \in \{1, \dots, P_{\text{imp}}\}$. Fixing $j \in \{1, \dots, N_i\}$ referring to the j -th smile of the i -th contract, we set

h on $[T_{j,\text{prev}}^i, T_j^i)$ equal to

$$h(t) := \sqrt{\frac{\text{VS}_{T_j^i}^i - \text{VS}_{T_{j,\text{prev}}^i}^i}{\int_{T_{j,\text{prev}}^i}^{T_j^i} \Sigma_s^\top(T_s^i, T_e^i) R \Sigma_s(T_s^i, T_e^i) ds}}, \quad t \in [T_{j,\text{prev}}^i, T_j^i), \quad (4.5)$$

where $T_{j,\text{prev}}^i$ is the previous smile maturity in the sorted list of all maturities and

$$\text{VS}_{T_{j,\text{prev}}^i}^i := \int_0^{T_{j,\text{prev}}^i} h^2(s) \Sigma_s^\top(T_s^i, T_e^i) R \Sigma_s(T_s^i, T_e^i) ds$$

Note that $t \mapsto h(t)$ in (4.5) depends on the function $T \mapsto g(T)$ through the deterministic variance $\Sigma_t(T_s^i, T_e^i)$.

Calibration of the function $T \mapsto g(T)$. The discontinuity points of g coincide with contracts' delivery start and end dates $(T_s^i)_{i \in 1, \dots, P_{\text{imp}}}$, $(T_e^i)_{i \in 1, \dots, P_{\text{imp}}}$. We adopt the following notation for the values of g : if the function g is constant on a delivery period $[T_s^i, T_e^i]$, we will denote its value by \bar{g}_i . Otherwise, if other delivery periods of futures contracts are included in $[T_s^i, T_e^i]$, we will denote by \bar{g}_{i-} the value of g on $[T_s^i, T_e^i] \setminus \cup_{k \in \mathcal{J}_i} [T_s^k, T_e^k]$, where \mathcal{J}_i stands for the indices of futures contracts which smiles are calibrated with the function g and which delivery periods are included in $[T_s^i, T_e^i]$. We also denote by $\hat{\Sigma}$ the volatility of a futures contract delivering over a set I with $g \equiv 1$ such that

$$\hat{\Sigma} \cdot (I) := \frac{1}{\text{Leb}(I)} \int_I \sigma(\cdot, T) dT,$$

where $\text{Leb}(I)$ stands for the Lebesgue measure of I . With a slight abuse of notation, we use interchangeably $\Sigma \cdot ([T_s^i, T_e^i])$ and $\Sigma \cdot (T_s^i, T_e^i)$, recall (1.14).

Hence, the deterministic volatility component of the contract F^i is given by

$$\Sigma \cdot (T_s^i, T_e^i) = \sum_{k \in \mathcal{J}_i} \omega_i^k \bar{g}_k \hat{\Sigma} \cdot (T_s^k, T_e^k) + \left(1 - \sum_{k \in \mathcal{J}_i} \omega_i^k\right) \bar{g}_{i-} \hat{\Sigma} \cdot ([T_s^i, T_e^i] \setminus \cup_{k \in \mathcal{J}_i} [T_s^k, T_e^k]), \quad (4.6)$$

where

$$\omega_i^k = \frac{T_e^k - T_s^k}{T_e^i - T_s^i} \quad (4.7)$$

denotes the relative weight of the contract F^k volatility in the volatility of F^i which is proportional to the length of its delivery period.

For the volatility smiles with underlying futures contract F^i and maturity T_j^i (for some $j \in \{1, \dots, N_i\}$, uniquely determined thanks to condition 1 of smile division), such that $\mathcal{J}_i = \emptyset$, matching the log-contract prices (4.4) leads to

$$g(T) = \bar{g}_i = \sqrt{\frac{\text{VS}_{T_j^i}^i}{\int_0^{T_j^i} h(s)^2 \hat{\Sigma}_s^\top(T_s^i, T_e^i) R \hat{\Sigma}_s(T_s^i, T_e^i) ds}}, \quad T \in [T_s^i, T_e^i], \quad (4.8)$$

since $\Sigma \cdot (T_s^i, T_e^i) = \bar{g}_i \hat{\Sigma} \cdot (T_s^i, T_e^i)$ by (4.6). Note that, since j is uniquely determined for each contract, it can be omitted in the notation \bar{g}_i .

For the contracts F^i not calibrated yet, i.e. such that $\mathcal{J}_i \neq \emptyset$, the function g should be chosen

$$g(T) = \bar{g}_{i-}, \quad T \in [T_s^i, T_e^i] \setminus \bigcup_{k \in \mathcal{J}_i} [T_s^k, T_e^k], \quad (4.9)$$

where \bar{g}_{i-} is a solution of the quadratic equation

$$\text{VS}_{T_j^i} = \int_0^{T_j^i} h(s)^2 \Sigma_s^\top(T_s^i, T_e^i) R \Sigma_s(T_s^i, T_e^i) ds,$$

with the deterministic volatility component given by (4.6) with $\{\bar{g}_k\}_{k \in \mathcal{J}_i}$ already calibrated by (4.8). Note that the set $[T_s^i, T_e^i] \setminus \bigcup_{k \in \mathcal{J}_i} [T_s^k, T_e^k]$ is non-empty thanks to the linear independence condition imposed on the smiles division.

Fixed-point calibration algorithm. Note that the model is calibrated when the equation (4.4) is verified for all smiles simultaneously. Since the functions g and h are interdependent, we propose the following algorithm which is supposed to converge to the desired solution.

Algorithm 1 Fixed-point calibration algorithm

```

set  $h^0(t) \equiv 1$  and  $g^0(T) \equiv 1$ 
 $n \leftarrow 0$ 
while  $\|g^n - g^{n-1}\|_\infty > \epsilon$  do
  recalculate the deterministic volatilities  $(\Sigma_s(T_s^i, T_e^i))_{i \in \{1, \dots, P_{\text{imp}}\}}$ 
  calibrate the function  $h^{n+1}$  using (4.5) with  $g = g^n$ .
  calibrate the function  $g^{n+1}$  using (4.8) and (4.9) with  $h = h^{n+1}$ .
   $n \leftarrow n + 1$ 
end while

```

Denoting one iteration of this algorithm by $\psi : g^n \mapsto g^{n+1}$, we note that model is calibrated if and only if g is a fixed-point of this mapping, i.e. $\psi(g) = g$.

Under additional conditions on the implied volatility data consistency, the following result holds:

Theorem 4.2. *Suppose that all the smiles calibrated with h have the same underlying F^1 and that $\mathcal{J}_i = \emptyset$ for the remaining contracts F^i , $i \in \{2, \dots, P_{\text{imp}}\}$. Suppose also that all the instantaneous correlations between the futures contracts are positive. Then,*

(i) ψ admits a fixed-point if

$$\sum_{i=2}^{P_{\text{imp}}} \frac{N_i \text{VS}_{T_k^i}}{\sum_{k=1}^{N_i} \left(\text{VS}_{T_k^1}^1 - \text{VS}_{T_{k-1}^1}^1 \right) \|Q_k\|} \frac{1}{\int_{[0, T^i] \cap [T_{k-1}, T_k]} \hat{\Sigma}_s^\top(T_s^i, T_e^i) R \hat{\Sigma}_s(T_s^i, T_e^i) ds} < 1, \quad (4.10)$$

where $(Q_k)_{ij} := \omega_1^i \omega_1^j \int_{T_{k-1}^1}^{T_k^1} \hat{\Sigma}_s^\top(T_s^i, T_e^i) R \hat{\Sigma}_s(T_s^j, T_e^j) ds$ and ω_1^i are defined by (4.7), $i, j = \{2, \dots, P_{\text{imp}}\}$.

(ii) All the fixed points are stable.

(iii) If the contract F^1 has only one maturity, the fixed point is unique.

Proof. The proof is given in Appendix C. □

This result is clearly partial as it does not cover the case of multiple underlyings corresponding to smiles calibrated with h and does not admit the nested contracts for smiles calibrated with g . However, the numerical experiments demonstrated the existence of a unique stable fixed point for all the test cases, even the ones not covered by Theorem 4.2. Thus, we believe that the provided result can be proved in a much more general case, though the direct approach used in the proof of Theorem 4.2 is not applicable there.

Remark 4.3. *Note that Theorem 4.2 provides condition (4.10) which seems to be a universal condition ensuring that the model can be calibrated. Indeed, it can be interpreted as a no-arbitrage condition for the*

variance swaps. For example, consider the smile on Cal 25 with maturity T_{Dec} calibrated with h and Q2 25 with maturity T_{Mar} calibrated by g . In this case, (4.10) reads

$$\frac{\text{VS}_{T_{\text{Mar}}}^{\text{Q2 25}} \left(\frac{1}{4}\right)^2 \int_0^{T_{\text{Dec}}} \hat{\Sigma}_s^\top(\text{Q2 25}) R \hat{\Sigma}_s(\text{Q2 25}) ds}{\text{VS}_{T_{\text{Dec}}}^{\text{Cal 25}} \int_0^{T_{\text{Mar}}} \hat{\Sigma}_s^\top(\text{Q2 25}) R \hat{\Sigma}_s(\text{Q2 25}) ds} < 1. \quad (4.11)$$

Since in the calibrated model

$$\text{VS}_{T_{\text{Mar}}}^{\text{Q2 25}} = \int_0^{T_{\text{Mar}}} h(s)^2 \Sigma_s^\top(\text{Q2 25}) R \Sigma_s(\text{Q2 25}) ds = h_1^2 g_{\text{Q2 25}}^2 \int_0^{T_{\text{Mar}}} \hat{\Sigma}_s^\top(\text{Q2 25}) R \hat{\Sigma}_s(\text{Q2 25}) ds,$$

the condition (4.11) is equivalent to

$$\left(\frac{1}{4}\right)^2 \int_0^{T_{\text{Dec}}} h(s)^2 \Sigma_s^\top(\text{Q2 25}) R \Sigma_s(\text{Q2 25}) ds < \text{VS}_{T_{\text{Dec}}}^{\text{Cal 25}}.$$

The integrand in the left-hand side expression is the part of variance of Cal 25 contract corresponding to Q2 25. Since the covariances are supposed to be positive, this part of variance is smaller than the whole variance of Cal 25, and thus, the price of the variance swap on Cal 25 should be greater than the integral on the left. Thereby, (4.10) is a condition describing the consistency of the variance swap data with historically calibrated deterministic volatility functions $\Sigma_s(T_s, T_e)$.

4.3 Smile shape calibration

Parametrization of the correlations. Fix $i \in \{1, \dots, P_{\text{imp}}\}$. For the approximated futures contract \tilde{F}^i , the smile skew is determined by the ‘‘spot-vol’’ correlation

$$\rho_i(t) = \frac{\langle dV, d \log \tilde{F}^i \rangle_t}{\sqrt{\langle dV \rangle_t} \sqrt{\langle d \log \tilde{F}^i \rangle_t}} = \frac{\Sigma_t^\top(T_s^i, T_e^i) \tilde{\rho}}{\sqrt{\Sigma_t^\top(T_s^i, T_e^i) R \Sigma_t(T_s^i, T_e^i)}} = \hat{\rho} \cdot \frac{L^\top \Sigma_t(T_s^i, T_e^i)}{\|L^\top \Sigma_t(T_s^i, T_e^i)\|}, \quad t \geq 0, \quad (4.12)$$

where $\mathbf{1}$ denotes a vector filled with ones, and L is a lower triangular matrix from the Cholesky decomposition (1.4). As $\hat{\rho}$ lies in $\{\rho: \|\rho\| \leq 1\}$, $\tilde{\rho}$ should be in $\{\rho: \|L^{-1}\rho\| \leq 1\}$ for the extended covariance matrix of $(W, B) \in \mathbb{R}^{N+1}$ to be well-defined.

Since the skew of the the implied volatility smile is determined by the correlations (4.12), it is more natural to use them and not $\tilde{\rho}$ in the calibration routine for several reasons. First, if the number of contracts is smaller than the number of historical risk factors N , this will reduce the number of parameters. Second, one spot-vol correlation impacts only one smile related to the corresponding futures contract, whereas the coefficients $\tilde{\rho}$ impact all the smiles in a way difficult to interpret. Thus, we expect to see a better solver behavior when the variables being optimized are the ‘‘spot-vol’’ correlations. Finally, even if one decides to calibrate $\tilde{\rho}$ directly, the ‘‘spot-vol’’ correlations may provide a reasonable initial guess given by the calibrated correlations in the multi-contract SSVI parametrization described in Section 2.2.2.

In order to reconstruct $\tilde{\rho}$ from a set of ‘‘spot-vol’’ correlations $(\rho_1^*, \dots, \rho_P^*)$, we find numerically a solution $\hat{\rho}^*$ to the following optimization problem

$$\min_{\hat{\rho}: \|\hat{\rho}\| \leq 1} \sum_{i=1}^P \left\| \int_0^{T_{C_M}} \hat{\rho}^\top L^\top \Sigma_t(T_s^i, T_e^i) dt - \rho_i^* \int_0^{T_{C_M}} \|L^\top \Sigma_t(T_s^i, T_e^i)\| dt \right\|^2 \quad (4.13)$$

and set $\tilde{\rho} = L\hat{\rho}^*$.

Optimization problem. We denote the parameters being calibrated by $\mathcal{P} = (\rho, c, x)$, where the spot-vol correlations ρ are transformed to $\tilde{\rho}$ by (4.13) and (1.5). The optimization problem then reads

$$\min_{\mathcal{P}} \sum_{i=1}^{P_{\text{imp}}} \sum_{j=1}^{N_i} \sum_{K \in \mathcal{K}_j^i} \left(\frac{\text{Call}^{\text{mkt},i}(T_j^i, K) - \text{Call}^{\text{Model},i}(T_j^i, K)}{\mathcal{V}(K, T, \sigma_{\text{IV}}^{\text{mkt},i}(T_j^i, K))} \right)^2, \quad (4.14)$$

where $\mathcal{V}(K, T, \sigma_{\text{IV}}^{\text{mkt},i}(T_j^i, K))$ denotes the Black-Scholes vega corresponding to the market implied volatility $\sigma_{\text{IV}}^{\text{mkt},i}(T_j^i, K)$.

5 Numerical results

5.1 Calibration results

In this section, we illustrate and detail all three calibration steps of the HJM model (1.1), cf. Figure 3, onto the German power market, and we postpone to Appendix A.2 the calibration results obtained on the TTF gas market.

To assess the quality of the calibration we consider the differences between model quantities and those used to calibrate the model: historical correlations and volatilities of rolling futures contracts, variance swap variance term structure and the implied volatility smiles. We also validate the quality of the Kemna-Vorst approximation. The calibrated model can be used to represent the market, and especially to deduce market quantities that are not quoted: the smile for all monthly contracts especially those which are not quoted, the at-the-money volatility of daily, monthly and quarterly contracts and finally the instantaneous correlations. Finally, we check the coherence between calibrated and interpolated quantities.

5.1.1 Step 1: Joint historical covariances – implied VS variances calibration results

The first calibration step aims at calibrating the $N := 1 + N_s + N_c$ L - S - C factors' parameters $\sigma_L, (\sigma_{S,i}, \tau_{S,i})_{i \in \{1, \dots, N_s\}}, (\sigma_{C,i}, \tau_{C,i})_{i \in \{1, \dots, N_c\}}$ as well as their correlation matrix R , following the methodology described in Section 3.

We consider the German power i.e. DE PW futures market with historical covariances estimated from daily log returns' time series running from January 1st 2023 to July 1st 2024, and the estimated realized futures correlation structure is displayed on the bottom left corner of Figure 5.

In order to avoid to have more calibratable parameters than market quantities to calibrate on, one needs to ensure that

$$\frac{N(N+3)}{2} - 1 \leq P_{\text{smiles}} + \frac{P_{\text{hist}}(P_{\text{hist}}+1)}{2}. \quad (5.1)$$

Furthermore, although season-averaging effects tend to distort Principal Components Analysis (PCA) results, the number of Principal Components (PCs) required to reach a given threshold of explained variance gives some insights on the number of factors to model the futures curve, see for example [Koekebakker and Ollmar \(2005\)](#), [Andersen \(2010\)](#), [Gardini and Santilli \(2024\)](#). Figure 8 displays the PCA's results for the rolling futures contracts' daily log returns. Note that 5 PCs are sufficient to explain at least 95% of variance, while 7 PCs are required to reach 99%, which is consistent with

1. our quality of fit results for the historical calibration described in Figures 9;
2. the observations made by [Féron and Gruet \(2024\)](#)[Figure 3] when applying BIC and AIC statistical measurements to the German market, where they found that 5 factors appears to be a good trade-off between quality of fit and model complexity.




Figure 8: Principal component analysis performed on the raw data of rolling futures contracts reconstituted from the German market from January 1st 2023 to July 1st 2024.

The results of the joint historical covariance and VS variances calibration is then summarized in the left matrix of Figure 9 where, given there is one L -factor, we display in a single matrix the losses J_1 and J_2 from (3.9)–(3.10) respectively for all possible pairs (N_s, N_c) of S - and C -factors such that the under-parametrization inequality (5.1) is satisfied for the strict lower part of the matrix, and is strictly violated by one additional factor on the main diagonal, corresponding to over-parametrization. For each cell, we used 100 random initializations of a_0 and kept the best calibration fit. Notice that the overall joint loss J^λ from (3.8), for $\lambda = 0.5$, indeed decreases as the number of factor increases when moving from the lower left angle to the main diagonal of the fit matrix. For a fixed number of factors N , we also highlighted by a bold square the model with the best fit along each diagonal.

Furthermore, we show in Figure 10 the respective quality of fits obtained on the selected rolling contracts' historical volatility and correlation term structures, as well as the fit of the VS volatility term structure. Notice that the fit is not perfect, yet it's possible to additionally fine-tune the parameter λ in the joint calibration loss function J^λ in order to better fit either the historical target covariance values (with $\lambda > 0.5$) or the implied VS volatility term structure (with $\lambda < 0.5$). In the next step, we calibrate the functions g and h in order to actually fit perfectly the latter in the case the joint calibration has not succeeded in fitting it well enough to fit the smile shapes in the latter calibration step.

We did the same numerical experiments for the historical covariance fit only by taking $\lambda = 1$ in the loss J^λ , and display the values of the losses in the right hand-side matrix of Figure 9, and the quality of fits in Figure 11. Notice that the quality of fit of both the historical volatility and correlation term structures is much better, but the associated VS volatility term structure is systematically under-estimated by such historical calibration.




Figure 9: Calibration results of step 1 for the German power market for all possible L - S - C model specifications up to height factors. Left (resp. Right): Best fits for the joint historical covariance and VS variances calibration (resp. the historical covariance calibration) when taking $\lambda = 0.5$ (resp. $\lambda = 1$) in the loss function J^λ (3.8), for 100 random initializations of $\tau(a_0)$, varying the number of S and C -factors until violation of the under-parametrization inequality (5.1). In each cell, the above number is the value of J_1 from (3.9) and below is the value of J_2 from (3.10) (which we do not calibrate on in the right-hand side results). For each matrix, we put bold squares on the best fits along each diagonal where the number of risk factors is constant.




Figure 10: Calibration results of step 1 for the German power market for the best $L-S-C$ model specifications up to six factors. The joint calibration fits i.e. $\lambda = 0.5$ in the loss J^λ (3.8), are displayed with respect to (top) the historical realized volatility term structure, (middle) the historical futures correlation structure and (bottom) the VS variance term structure.

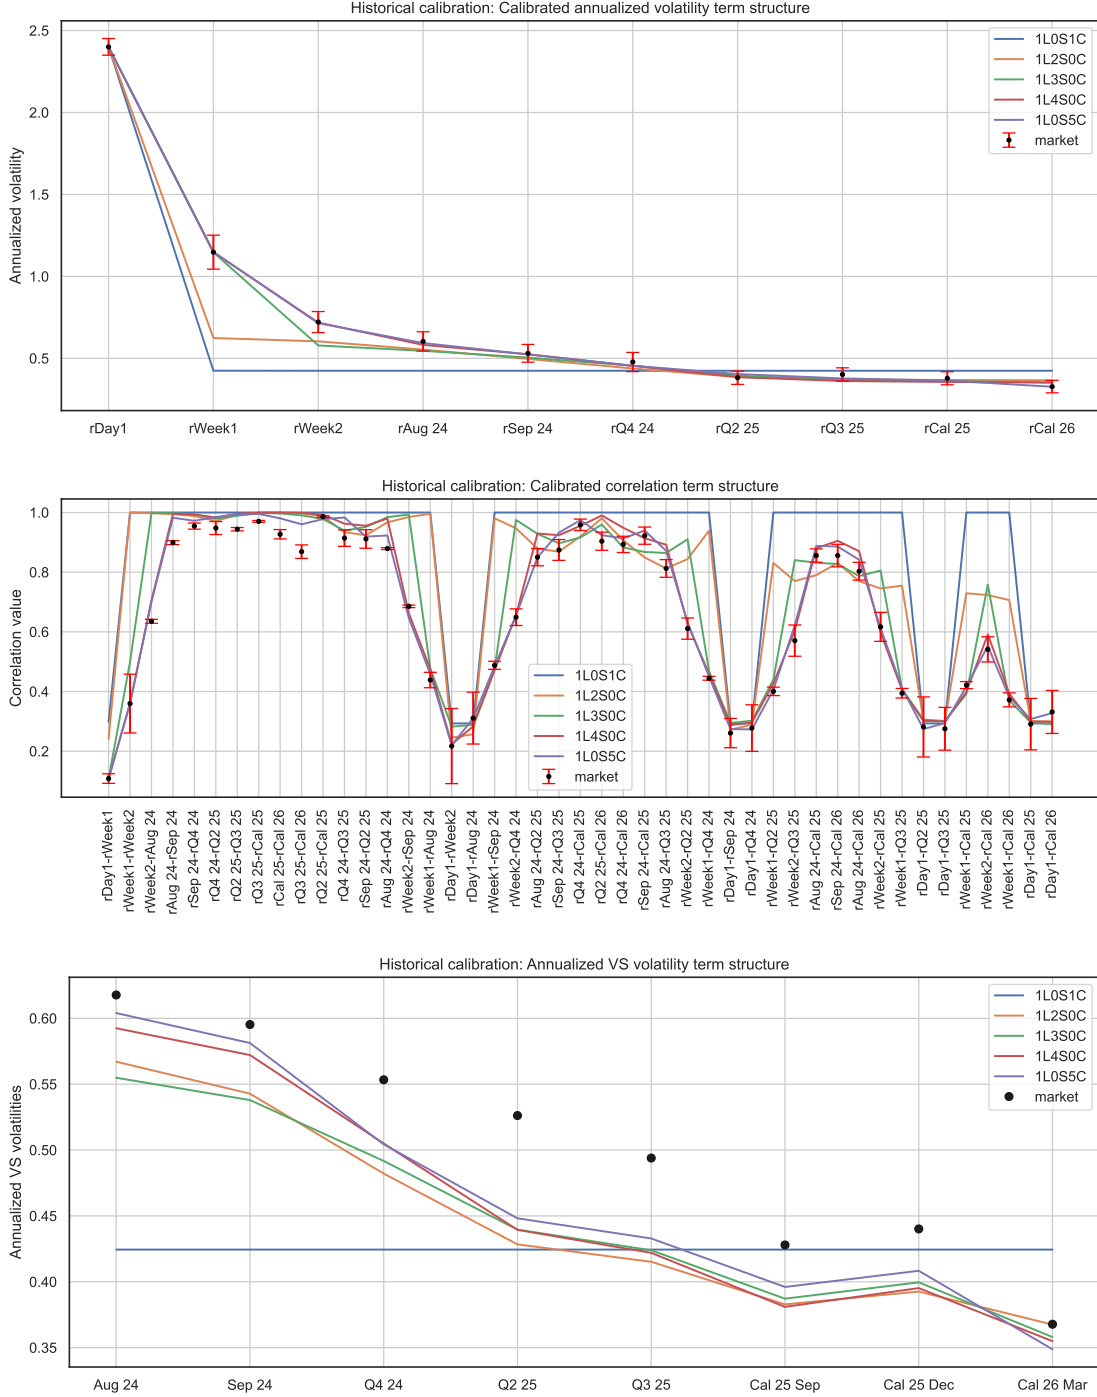


Figure 11: Historical covariance calibration fits i.e. J_2 (3.10) is ignored by taking $\lambda = 1$ in the loss J^λ (3.8), of the best up-to-six-factors models whose loss functions are highlighted in the right-hand side matrix of Figure 9 to (top) the historical realized volatility term structure, (middle) the historical futures correlation structure and (bottom) the VS variance term structure.

For the implied calibration, we choose the model 1L3S1C as it gives a good trade-off between parsimony and the joint calibration quality, and display its calibrated parameters in Figure 12. Notice the L -factor indeed captures the long-term volatility level, with σ_L lying in-between the annualized volatilities of long-term contracts rCal 25 and rCal 26, respectively equal to 0.3784 and 0.3277 while the C factor is placed at the

beginning of the curve to capture the short-term volatility behavior in the case of 1L3S1C. For conciseness, the calibrated parameters of the other models mentioned in Figure 9 are not displayed.

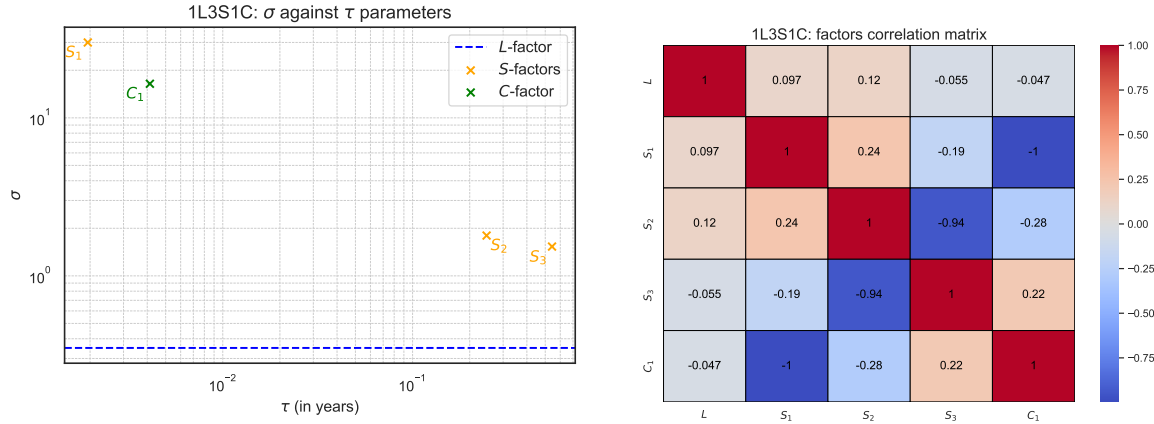


Figure 12: Plots of $(\sigma_i)_i$ against $(\tau_i)_i$ parameters (left) and of the factors' correlation matrix (right) for the 1L3S1C model, when taking $\lambda = 0.5$ in the loss function J^λ (3.8). Model parameters are, in L-S-C order: $(\sigma_i)_i := [0.3499, 29.9849, 1.7992, 1.5325, 16.4317]$, $(\tau_i)_i := [0.0019, 0.2454, 0.5428, 0.0041]$.

It is worth mentioning that the choice of the historical volatility factors does not impact significantly the results of the following calibration steps as long as the fit of the VS term structure is satisfactory. Otherwise, the calibrated functions g and h may deviate significantly from 1 degrading the model's interpolation capabilities.

5.1.2 Step 2: Term structure exact calibration correction

This part addresses the correction of the VS volatility term structure with the functions g and h following the methodology proposed in Section 4.2.

In order to avoid jumps of the day-ahead contracts volatility, both functions g and h are smoothed during calibration. At each step of the fixed point algorithm, we update g and h as piece-wise constant functions and then interpolate it using the monotonicity preserving algorithm PCHIP of Fritsch and Butland (1984). Despite the smoother results, the algorithm is more time-consuming due to the numerical optimization involved in the calibration. On the left-hand side of Figure 13 below, we plot the results of the term structure calibration of piece-wise constant functions g and h , and on the right-hand side, same results with smoothing interpolation. In both cases, the fixed-point algorithm converges in approximately 5–10 iterations, providing an almost perfect fit for the VS volatilities shown in Figure 14.

We also note that thanks to the calibration of the VS volatilities at the first step of the calibration, the functions g and h are needed only to slightly correct the volatility level, so that they remain close to 1 for the months following the observation date. For longer maturities, the function h deviates significantly from 1 since the VS volatilities of smiles Cal 26 Dec and Cal 27 Dec were not calibrated at the first calibration step due to the poor quality of covariance estimations for contracts with long time to delivery.




Figure 13: Calibrated functions h and g with and without smoothing ((a) and (b) correspondingly).




Figure 14: VS term structure fit before the correction (orange) and after the correction (blue).

5.1.3 Step 3: Implied smile shapes calibration

Finally, we calibrate the parameters $(c_i, x_i)_{i \in \{1, \dots, M\}}$ of the stochastic variance process V , as well as the correlations $(\tilde{\rho}_i)_{i \in \{1, \dots, N\}}$, using the procedure described in Section 4.3.

For the stochastic variance component calibration, we choose a model with $M = 3$ pseudo-factors in (1.2) which allow us to cover different volatility timescales and consistently achieve an acceptable fit on various calibration sets. The advantage of using multiple pseudo-factors is demonstrated in Appendix A.1, where we compare the quality of fit between the Heston model and the Lifted Heston model.

The calibrated parameters of the stochastic volatility are provided in Table 3. For the calibrated values, the loss function defined by (4.14) equals 0.001754. It is interesting to notice that the first mean-reversion coefficient is very close to zero and corresponds to very long mean-reversion periods, while the two others correspond approximately to the timescales of one month and 2.5 weeks respectively.

Parameter	Calibrated value
c	(0.492, 0.68, 2.79)
x	($4.6 \cdot 10^{-6}$, 9.712, 20.249)
$\tilde{\rho}$	(0.648, -0.516, 0.16, -0.148, 0.541)

Table 3: Calibrated lifted Heston model parameters

The implied volatility smiles in the calibrated model with smoothed g and h functions, are shown in Figure 15. We also plot the bid-ask spread equal to 5% as a realistic and even conservative proxy to the market spread observed in the power market. Thus, the plot demonstrates sufficiently high quality of fit for the model to be used for practical needs.

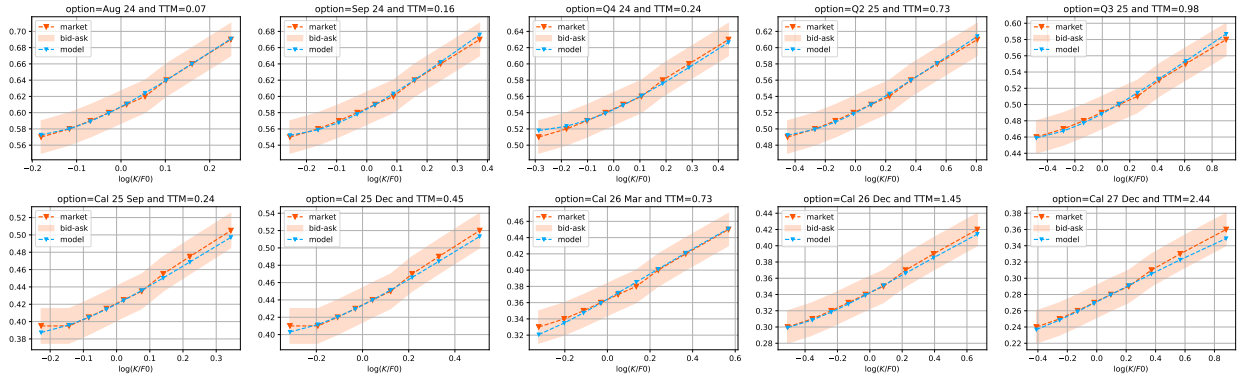


Figure 15: IV smiles in the calibrated model (blue) and market IV smiles (orange).

5.2 A posteriori validation of model approximations

As discussed in Section 1.2, the use of the Kemna-Vorst (KV) approximated model futures \tilde{F} from (1.15) leads to potential arbitrage opportunities for futures contracts with overlapping deliveries, e.g. between monthly futures and the resulting quarter composed by such delivery months. However, we show in the next subsections that the futures prices' trajectories, as well as the futures correlation structures and implied volatility smiles are extremely close in the approximated model \tilde{F} and in the exact model future F from (1.9).

From the forward rate definition (1.1), the exact model is given explicitly by

$$F_t(T_s, T_e) = \frac{1}{T_e - T_s} \int_{T_s}^{T_e} f(0, T) \exp \left(-\frac{g^2(T)}{2} \int_0^t h^2(s) V_s \sigma^\top(s, T) R \sigma(s, T) ds + g(T) \int_0^t h(s) \sqrt{V_s} \sigma^\top(s, T) dW_s \right) dT.$$

Numerically, we consider a one-day discretization step, which corresponds in practice to the futures contracts with the shortest delivery period quoted in the futures market

$$F_t(T_s, T_e) \approx \frac{\delta}{T_e - T_s} \sum_{i=1}^{N-1} f(t, T_i)$$

$$T_s = T_1 < \dots < T_N = T_e \text{ and } T_{i+1} - T_i = \delta = \frac{1}{365}, i = 1, \dots, N - 1.$$

Remark 5.1. In real energy markets, the smallest possible delivery intervals Δ are typically 15 minutes, 30 minutes, or one hour. Consequently, the price of any futures contract $F_t(T_s, T_e)$ can be expressed as the

average

$$F_t(T_s, T_e) = \frac{\Delta}{T_e - T_s} \sum_{i=1}^{(T_e - T_s)/\Delta} f(t, T_s + i\Delta),$$

of the prices $f(t, T_s + i\Delta)$ of all contracts with the smallest granularity Δ included in $[T_s, T_e]$. This implies that the futures curve is inherently discrete, whereas the HJM model provides a continuous approximation. By approximating the contract $F(T_s, T_e)$ using formula (??), we effectively perform an inverse approximation, representing a continuous curve by a discrete sum. For numerical feasibility in the subsequent Monte Carlo simulations, the value $\delta = \frac{1}{365}$, corresponding to one day, was used instead of the finer granularity of one hour or one quarter.

Since the exact daily contracts are computed using a Monte Carlo scheme, the exact model becomes significantly slower in terms of pricing and simulation time than the approximated one.

Besides the KV approximation, it is important to notice that the introduction of the function $g \neq 1$ prevents the futures volatility from being stationary. However, if $g \approx 1$, which is the case in our model as the variance swap volatilities are mostly calibrated at the first calibration step, then the futures contracts dynamics remains almost stationary. Furthermore, the function g impacts the instantaneous correlations between absolute contracts. These correlations are affected by the KV approximation as well. In the following numerical experiments, we demonstrate that neither KV approximation nor the function g have a significant impact on the futures correlation structure.

5.2.1 Sample path trajectories errors

As an illustration of the quality of this approximation, we provide in Figure 16 the trajectories of the exact futures contracts F and their approximations \tilde{F} , simulated with the same random numbers, for three different futures contracts: Sep 24, Q1 25 and Cal 25. Recall that F is defined in (1.9) with $f(t, T)$ explicitly given by (1.6). The mean L^2 distance on $[0, \theta]$ is defined as a sample Root Mean Squared Error (RMSE) over $m \in \mathbb{N}^*$ sample trajectories by

$$\sqrt{\frac{1}{m} \sum_{i=1}^m \|F(\omega_i) - \tilde{F}(\omega_i)\|_{L^2(\theta)}^2} := \sqrt{\frac{1}{m} \sum_{i=1}^m \int_0^\theta (F_t(\omega_i) - \tilde{F}_t(\omega_i))^2 dt} \approx \sqrt{\frac{1}{M} \sum_{i=1}^m \frac{\theta}{N_\theta} \sum_{j=1}^{N_\theta} (F_{t_j}(\omega_i) - \tilde{F}_{t_j}(\omega_i))^2},$$

where N_θ is the number of steps in the daily subdivision of $[0, \theta]$. The estimated distance with $\theta = T_s$ between the approximated and exact trajectories estimated over 10^4 simulations equals 0.0602, 0.0897 and 0.1898 for three futures contracts correspondingly. We observed that this error depends both on the interval length θ as well as on the length of the delivery period. Both dependencies are non-linear, the first one is strictly increasing in θ . As for second one, the error is greater for small (up to one-two months) and for large intervals (two years or more) delivery periods. The worst-case RMSE numerically reaches 0.5% of the initial futures value F_0 .




Figure 16: Sample paths of $\tilde{F}(T_s, T_e)$ and $F(T_s, T_e)$ for the same random numbers for three futures contracts: monthly contract Sep 24 delivering during September 2024, quarterly contract Q1 25 delivering during the first quarter of 2025, and yearly (or calendar) contract Cal 25 delivering during the whole year 2025. In the figures, time to delivery is given in years.

5.2.2 Instantaneous correlations

In this section, we assess numerically the impact of the function g on the futures contracts' instantaneous correlations. Indeed, unlike the function h and the stochastic variance component V which only modify the futures' volatility term structure but not their correlations, the function g also impacts the correlations as it changes the approximated futures contract volatilities Σ_t defined by (1.14). Indeed, for $i, j \in \{1, \dots, P\}$, the instantaneous correlation between the approximated futures contracts $\tilde{F}^i := \tilde{F}^i(T_s^i, T_e^i)$ and $\tilde{F}^j := \tilde{F}^j(T_s^j, T_e^j)$ obtained just after the first calibration step, and implied by the historical futures correlation structure of rolling futures, is defined by

$$\tilde{\rho}_{ij}^{\text{hist}}(t) = \frac{\langle d \log \tilde{F}^i, d \log \tilde{F}^j \rangle_t}{\sqrt{\langle d \log \tilde{F}^i \rangle_t} \sqrt{\langle d \log \tilde{F}^j \rangle_t}} = \frac{\hat{\Sigma}_t^\top(T_s^i, T_e^i) R \hat{\Sigma}_t(T_s^j, T_e^j)}{\sqrt{\hat{\Sigma}_t^\top(T_s^i, T_e^i) R \hat{\Sigma}_t(T_s^i, T_e^i)} \sqrt{\hat{\Sigma}_t^\top(T_s^j, T_e^j) R \hat{\Sigma}_t(T_s^j, T_e^j)}}, \quad 0 \leq t \leq \min(T_s^i, T_s^j),$$

where $\hat{\Sigma}_t(T_s^j, T_e^j)$ was defined by (4.2). However, after the second calibration step, the futures contract volatility is given by Σ_t , and the instantaneous correlation equals

$$\tilde{\rho}_{ij}(t) = \frac{\Sigma_t^\top(T_s^i, T_e^i) R \Sigma_t(T_s^j, T_e^j)}{\sqrt{\Sigma_t^\top(T_s^i, T_e^i) R \Sigma_t(T_s^i, T_e^i)} \sqrt{\Sigma_t^\top(T_s^j, T_e^j) R \Sigma_t(T_s^j, T_e^j)}}, \quad 0 \leq t \leq \min(T_s^i, T_s^j). \quad (5.2)$$

Note that these two values coincide if and only if $g(T)$ is constant on $[T_s^i, T_e^i]$ and on $[T_s^j, T_e^j]$, as in this case, $\Sigma_t(T_s^i, T_e^i) = \bar{g}_i \hat{\Sigma}_t(T_s^i, T_e^i)$ and $\Sigma_t(T_s^j, T_e^j) = \bar{g}_j \hat{\Sigma}_t(T_s^j, T_e^j)$, where \bar{g}_i and \bar{g}_j denote the values of the function g on the intervals $[T_s^i, T_e^i]$ and $[T_s^j, T_e^j]$ correspondingly.

However, we can show numerically that the difference between $\tilde{\rho}_{ij}$ and $\tilde{\rho}_{ij}^{\text{hist}}$ is not significant from a practical point of view even if g is not constant. Namely, we examine this difference for the model calibrated in Section 5.1 with the function g given by Figure 13(b). The figure 17 shows the distances $\|\tilde{\rho}_{ij} - \tilde{\rho}_{ij}^{\text{hist}}\|_\infty$ for each pair of contracts for the calibration set. In the worst case, the difference is less than 0.1%. Thus, despite the introduction of the function g , the correlations in the calibrated model are still very close to the ones given by the historical calibration.

Aug 24	0	0.0001	0.0001	0.0003	0	0	0.0001	0.0001	0.0001	0.0001
Sep 24	0.0001	0	0.0008	0.0005	0.0004	0.0004	0.0005	0.0004	0.0006	0.0006
Oct 24	0.0001	0.0008	0	0.0003	0.0008	0.0007	0.0007	0.0007	0.0009	0.0009
Q4 24	0.0003	0.0005	0.0003	0	0.0007	0.0007	0.0007	0.0007	0.0005	0.0004
Q1 25	0	0.0004	0.0008	0.0007	0	0.0003	0.0002	0.0003	0.0001	0.0001
Q2 25	0	0.0004	0.0007	0.0007	0.0003	0	0.0002	0	0.0002	0.0001
Q3 25	0.0001	0.0005	0.0007	0.0007	0.0002	0.0002	0	0.0001	0.0002	0.0003
Cal 25	0.0001	0.0004	0.0007	0.0007	0.0003	0	0.0001	0	0.0002	0.0002
Cal 26	0.0001	0.0006	0.0009	0.0005	0.0001	0.0002	0.0002	0.0002	0	0
Cal 27	0.0001	0.0006	0.0009	0.0004	0.0001	0.0001	0.0003	0.0002	0	0
	Aug 24	Sep 24	Oct 24	Q4 24	Q1 25	Q2 25	Q3 25	Cal 25	Cal 26	Cal 27

Figure 17: Differences $\|\tilde{\rho}_{ij} - \tilde{\rho}_{ij}^{\text{hist}}\|_\infty$ between historically calibrated and model instantaneous correlations.

5.2.3 Quality of the Kemna–Vorst approximation on the instantaneous correlations

As we have just seen, the Kemna–Vorst approximation allows for an explicit computation of instantaneous correlations between futures contracts in (5.2). However, these quantities do not coincide with the correlations between exact futures contracts

$$\rho_{ij}(t) := \frac{\langle d \log F^i, d \log F^j \rangle_t}{\sqrt{\langle d \log F^i \rangle_t} \sqrt{\langle d \log F^j \rangle_t}}, \quad 0 \leq t \leq \min(T_s^i, T_s^j),$$

where F^i and F^j follow (1.9) for $i, j \in \{1, \dots, P\}$. More precisely

$$\rho_{ij}(t) = \frac{\Sigma_t^{\text{exact}^\top}(T_s^i, T_e^i) R \Sigma_t^{\text{exact}}(T_s^j, T_e^j)}{\sqrt{\Sigma_t^{\text{exact}^\top}(T_s^i, T_e^i) R \Sigma_t^{\text{exact}}(T_s^i, T_e^i)} \sqrt{\Sigma_t^{\text{exact}^\top}(T_s^j, T_e^j) R \Sigma_t^{\text{exact}}(T_s^j, T_e^j)}}, \quad 0 \leq t \leq \min(T_s^i, T_s^j),$$

with $\Sigma_t^{\text{exact}}(T_s, T_e) = \frac{\int_{T_s}^{T_e} g(T) \sigma(\cdot, T) f(\cdot, T) dT}{\int_{T_s}^{T_e} f(\cdot, T) dT}$, such that the exact correlation is a stochastic process, since

f is stochastic. However, as shown numerically in Figure 18, that the L^∞ -difference between $\mathbb{E}[\rho_{ij}]$ and $\tilde{\rho}_{ij}$ lies within two standard deviations of ρ_{ij} , which makes the approximated correlation (5.2) a valid and tractable deterministic approximation of the exact instantaneous correlation.

In conclusion, one can control the difference error between the instantaneous correlations $\tilde{\rho}_{ij}^{\text{hist}}$ obtained by the first calibration step from historical futures' correlation structure and the fully calibrated exact model instantaneous correlations ρ_{ij} using the triangle inequality

$$\|\tilde{\rho}_{ij}^{\text{hist}} - \mathbb{E}[\rho_{ij}]\|_\infty = \|(\tilde{\rho}_{ij}^{\text{hist}} - \tilde{\rho}_{ij}) - (\mathbb{E}[\rho_{ij}] - \tilde{\rho}_{ij})\|_\infty \leq \|\tilde{\rho}_{ij}^{\text{hist}} - \tilde{\rho}_{ij}\|_\infty + \|\tilde{\rho}_{ij} - \mathbb{E}[\rho_{ij}]\|_\infty,$$

where the right-hand error quantities are respectively given in Figures 17 and 18 across all liquid contracts. Such numerical measurements illustrate the overall consistency of the instantaneous correlations in the approximated model with respect to the exact one.




Figure 18: Distance $\|\tilde{\rho}_{ij} - \mathbb{E}[\rho_{ij}]\|_\infty$ (on the left) and its standard deviation (on the right) in the calibrated model.

5.2.4 Quality of the Kemna–Vorst approximation on the smiles

Our next goal is to show that the Kemna–Vorst approximation (1.15) provides precise enough pricing results. Since neither closed formula, nor an SDE for the futures price dynamics are available, the option prices can

be computed only with the Monte Carlo scheme presented in Appendix D to generate the trajectories of F and \tilde{F} .

In order to visualize the approximation error, we benchmark the implied volatility smiles for the contracts Sep 24, Q4 24, and Cal 25 Dec with the calibrated model parameters given in Section 4. The same random numbers were used for both Monte Carlo, 10^5 trajectories were simulated and we consider a daily grid for the instantaneous futures used to compute F . In the worst case, the difference in implied volatility between the Monte Carlo estimators is less than 0.2%.

Moreover, we compare the values of the smile calibration loss function introduced further in (4.14) for the calibrated model. The loss function value in the approximated model is given by 0.001754, while the loss function in the exact model is 0.001617.

It allows us finally to conclude that the Kemna-Vorst approximation is a valid tractable approximation for calibration of the exact model which then can be used for pricing and hedging purposes.





Figure 19: IV smiles generated with the same random numbers (10^5 simulations)

5.3 Volatility hypercube extrapolation in the calibrated model


In this section, we address the question of the implied volatility hypercube extrapolation once the model is fully calibrated and provide the volatility term structure induced by the model.

To illustrate the volatility term structure, we compute the ATM volatilities of options on daily contracts expiring in the day preceding the day of the delivery, as well as the ATM volatilities of (generated) monthly, quarterly and calendar contracts expiring three days before the delivery period start. The range of considered days cover the days from July 3, 2024 to July 3, 2025. The resulting ATM volatilities of are shown in Figure 20 (a), (b).


We also generate the implied volatility smiles for monthly futures contracts corresponding to the months between August 2024 and December 2026. The options maturity dates are taken equal to the 25-th day of the month preceding the delivery period. The generated smiles are shown in Figure 20 (c).



(a) ATM daily volatilities.



(b) ATM monthly, quarterly and cal volatilities.



(c) IV smiles from August 2024 to December 2026.

Figure 20: ATM implied volatilities of daily contracts (a) and of monthly, quarterly and calendar contracts (b), and implied volatility smiles for monthly contracts (c) generated by the calibrated model.

We highlight that the initial calibration of the variance swap volatilities at the first calibration step prevents the model from over-fitting the VS volatility term structure when calibrating the functions g and h in step 2). This makes possible a reasonable interpolation and extrapolation of the implied volatilities for the contracts and maturities not included in the calibration set. In Figure 21, we plot the integrated variance for monthly (green) and quarterly (blue) contracts being calibrated (circles) as well as the ones not present in the calibration set and generated within the calibrated model (crosses). We observe that the calibrated model produces consistent interpolation and extrapolation values, and thus capable of completing the volatility hypercube.




Figure 21: Integrated variance for monthly (green) and quarterly (blue) contracts as function of time to maturity T within one year. Circles correspond to the calibrated contracts Aug 24, Sep 24, Q4 24, Q2 25, Q3 25, while crosses correspond to the contracts represented by the calibrated HJM model (monthly contracts from Oct 24 to Jul 25 and Q1 25).

A More calibration results

A.1 Attainable smiles

Since the Lifted Heston model reduces to the standard Heston model when $M = 1$, $c_1 = 1$, and $x_1 = 0$, the Lifted Heston model can reproduce all the smiles that the standard Heston model can generate. The inverse, however, is not true, as the standard Heston model often struggles to capture steep enough smiles for short maturities. We provide a comparison of the calibrations of both models using the previously introduced calibration set. To highlight the differences, only the smiles with the poorest calibration results are shown in Figure 22.

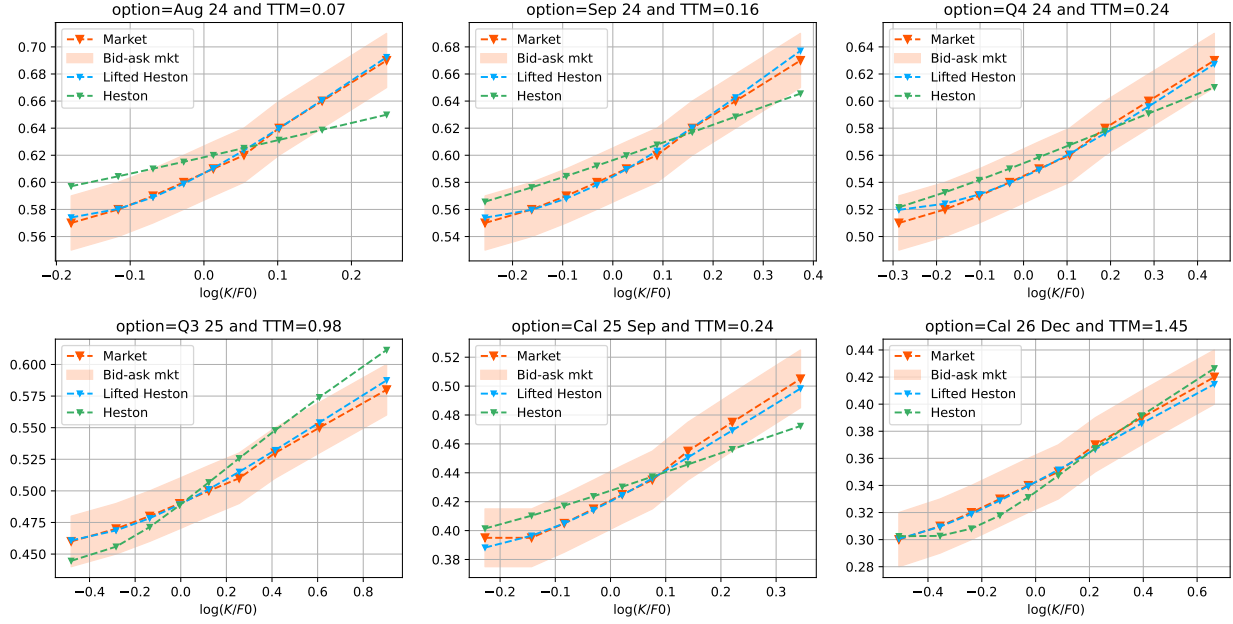


Figure 22: A comparison of the calibration results for the Heston model and the Lifted Heston model with $M = 3$ pseudo-factors on the same calibration set.

A.2 Calibration results for TTF market

In this section, we provide the results of calibration performed on the TTF gas market on the 3rd July 2024.

Note that options on quarterly and calendar gas futures contracts are quoted as strip options (e.g. a price of an option on a quarter futures contract is a mean of option prices on monthly contracts forming this quarter). Thus, it is sufficient to calibrate only the smiles corresponding to the monthly futures contracts. It can be achieved by using only the function h at the term structure calibration step, which is shown in Figure 26. Thus, the historical correlations between futures remain unchanged and coincide perfectly with the correlations calibrated historically, since the function g remains constant. The calibrated VS volatility term structure is presented in Figure 27.

We display in Figure 23 the results of the PC analysis applied to the rolling TTF futures contracts' daily log returns observed from the 3rd of July 2024 to the 3rd of January 2023, and observe only 2 PCs allow to reach 95% of explained variance in this case, and 4 factors to reach 99%, consistent again with the quality of fits we obtain in Figure 24, and the observations by Andersen (2010, Section 6.2.2). In fact, the high futures correlation structure observed in the upper left of Figure 23 stands in sharp contrast to the de-correlation between short-term and long-term observed in the lower left part of Figure 5 for the German power market, and is mainly due to the large storage capacities of gas in chambers, while electricity still cannot be stored at a large scale.

For the implied calibration at steps 2) and 3), we consider the four-factor model $1L2S1C$. The implied calibration results for six smiles on monthly contracts are presented below in Figure 28, and the calibrated parameters of the stochastic volatility are provided in Table 4.

The monthly contracts surface and the ATM volatilities are shown in Figure 29. Note that the term structure of the ATM volatility is not decreasing for short maturities. This is a consequence of the fact that the implied volatility corresponding to the first contract is 5% lower than the volatilities of the other monthly contracts, possibly due to the seasonality. It is also interesting to notice the the difference between daily and monthly volatility is much smaller than in the power market thanks to the very high correlation between all the futures contracts.

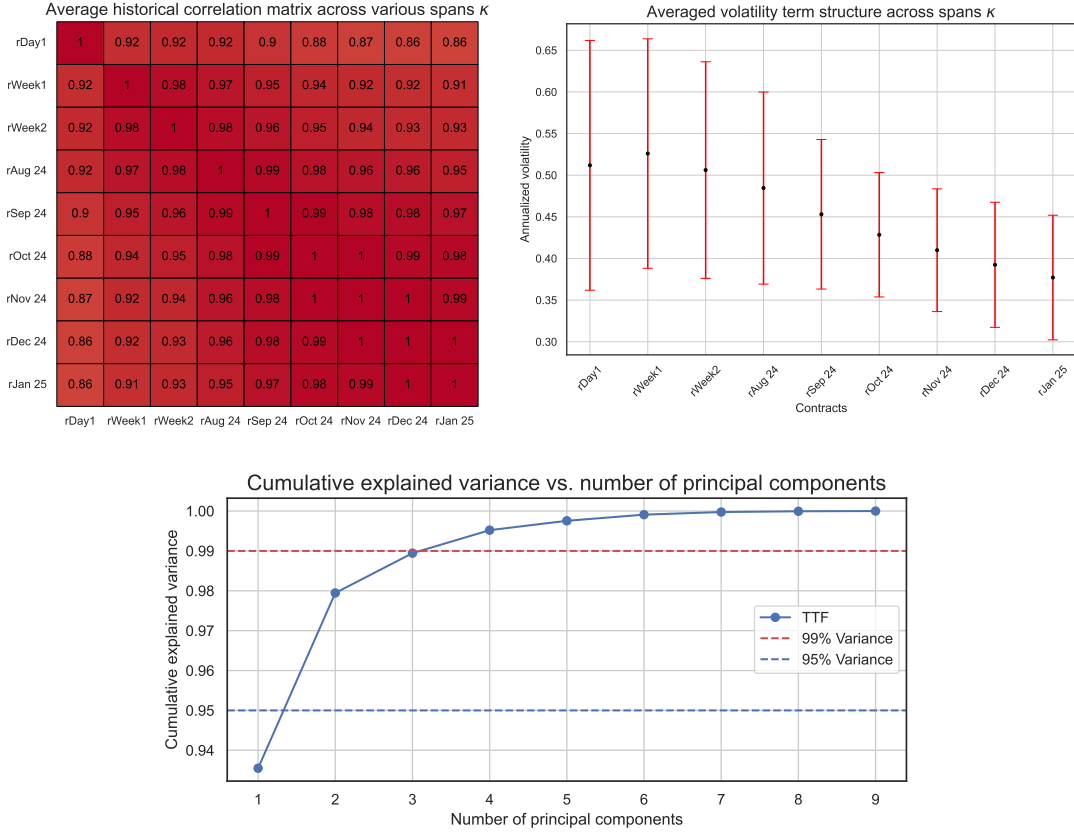


Figure 23: Rolling futures in TTF gas market: historical (upper left) futures correlation structure and (upper right) volatility term structure (both estimated by averaging across span parameters κ going from 30 to 365 days, see Appendix B.1); (bottom) principal component analysis




Figure 24: Historical calibration fits for TTF, with $\lambda = 0.99$, to (top) the historical realized volatility term structure and (bottom) the historical futures correlation structure.




Figure 25: Plots of calibrated $(\sigma_i)_i$ against $(\tau_i)_i$ parameters (left) and of the factors' correlation matrix (right) for the 1L2S1C model on TTF market data, when taking $\lambda = 0.99$ in the loss function J^λ (3.8). Model parameters are, in L - S - C order: $(\sigma_i)_i := [0.4331, 8.8686, 0.3727, 0.5886]$, $(\tau_i)_i := [0.0005, 0.1731, 0.2456]$.




Figure 26: Calibrated function h .




Figure 27: VS term structure fit before the correction (orange) and after the correction (blue).

Parameter	Calibrated value
c	(1.863, 1.155, 3.747)
x	(2.586, 4.919, 27.745)
$\tilde{\rho}$	(0.76, -0.267, -0.222, -0.272)

Table 4: Calibrated lifted Heston model parameters


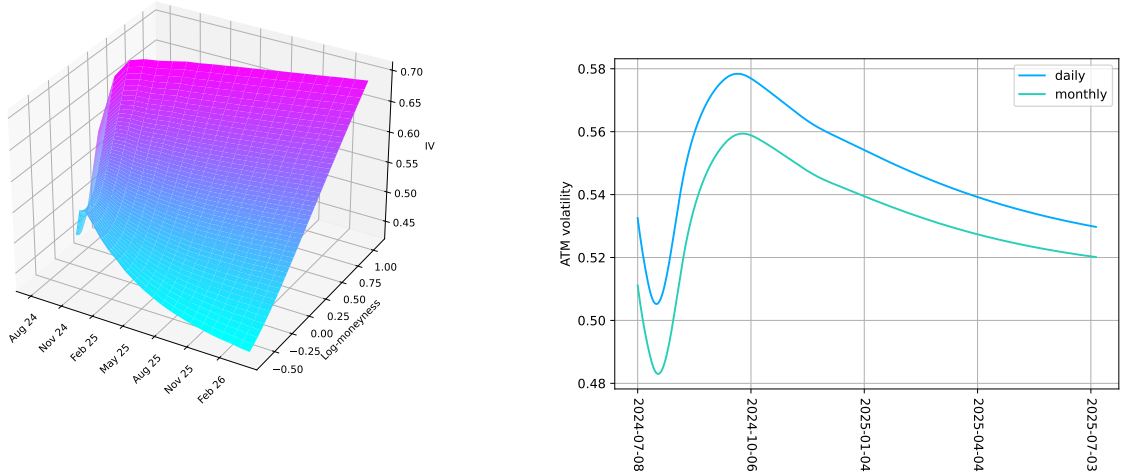


Figure 28: Calibrated IV smiles.



(a) IV smiles from August 2024 to February 2026.

(b) ATM volatilities from July 3, 2024 to July 3, 2027.

Figure 29: Implied volatility smiles for monthly contracts and implied ATM volatilities of daily and monthly contracts generated by the calibrated model. ((a) and (b) correspondingly).

A.3 A word on the additive model

An additive HJM model with stochastic volatility can be considered as well. In this case, the dynamics is given by

$$df(t, T) = g(T)h(t)\sqrt{V_t} \sum_{i=1}^N \sigma_i(t, T) dW_t^i, \quad t \in [0, T].$$

However, we found several disadvantages preventing from its efficient usage. As the market uses the Black-Scholes volatility quotation, instead of the well-known smile flattening effect, the additive model produces negatively skewed smiles for $T \gg 1$, which is hardly acceptable in the commodity market where the skew is typically positive, i.e. the “inverse leverage” as described in Andersen (2010).

We illustrate it with a simple experiment, calibrating the smile shapes of the contracts Aug 24 and Cal 25 Sep and plotting the smile corresponding to Cal 27 Dec, see Figures 30–32. The functions g and h were calibrated to all the three smiles. Since the link between the forward variance and the log-contract is absent in the additive model, g and h were calibrated to match the ATM volatility level in this case, while for the multiplicative model we provide both calibration of the ATM volatilities and of the variance swap volatilities (Figures 31 and 32 correspondingly). We observe that the additive model produces a completely inconsistent shape for the third, extrapolated smile. The fixed-point ATM volatility calibration algorithm stays the same, but volatility term structure and shape calibration cannot be separated and the calibration routine becomes much more time-consuming.




Figure 30: Additive model calibrated to smiles Aug 24 and Cal 25 Sep. Term structure for all the three contracts is calibrated via the ATM volatility.




Figure 31: Multiplicative model calibrated to smiles Aug 24 and Cal 25 Sep. Term structure for all the three contracts is calibrated via the ATM volatility.




Figure 32: Multiplicative model calibrated to smiles Aug 24 and Cal 25 Sep. Term structure for all the three contracts is calibrated via the variance swap volatility.

B Additional details for the step 1) calibration

B.1 Estimation of historical covariances

Given two rolling futures contracts' indices $(i, j) \in \{1, \dots, P_{\text{hist}}\}^2$, we introduce the following normalized covariance estimator with exponentially decaying weights such that

$$\widehat{\text{Cov}}_{\kappa}(r^{\text{mkt},i}(\tau_d), r^{\text{mkt},j}(\tau_d)) := \frac{1}{\tau_d} \frac{1}{1 - \sum_{h=0}^H w_h^2(\kappa)} \sum_{h=0}^H w_h(\kappa) \left(r_{t_h}^{\text{mkt},i}(\tau_d) - \bar{r}_{w(\kappa)}^{\text{mkt},i}(\tau_d) \right) \left(r_{t_h}^{\text{mkt},j}(\tau_d) - \bar{r}_{w(\kappa)}^{\text{mkt},j}(\tau_d) \right),$$

with the exponentially decaying weighted average of log returns given by

$$\bar{r}_{w(\kappa)}^{\text{mkt},n}(\tau_d) := \sum_{h=0}^H w_h(\kappa) r_{t_h}^{\text{mkt},n}(\tau_d), \quad n \in \{1, \dots, P_{\text{hist}}\},$$

and where the weights are computed as follows

$$w_h(\kappa) := \frac{(1 - \alpha_{\kappa})^{H-h}}{\sum_{k=0}^H (1 - \alpha_{\kappa})^{H-k}}, \quad h \in \{1, \dots, H\}, \quad \alpha_{\kappa} := \frac{2}{\kappa + 1}, \quad (\text{B.1})$$

for some time *span* parameter $\kappa \geq 1$ (in number of days) controlling the decay rate of the exponential smoothing. In practice, we consider the covariance estimator of past log returns averaged across a family of $Z \in \mathbb{N}$ time-scale decays $(\kappa_z)_{z \in \{1, \dots, Z\}}$ such that

$$C_{i,j}^{\text{mkt}} := \frac{1}{Z} \sum_{z=1}^Z \widehat{\text{Cov}}_{\kappa_z}(r^{\text{mkt},i}(\tau_d), r^{\text{mkt},j}(\tau_d)) \quad (\text{B.2})$$

as well as the element-wise confidence matrix U whose $(i, j)^{\text{th}}$ entry is the standard deviation of the $(i, j)^{\text{th}}$ entries of the respective covariance matrices $(\widehat{\text{Cov}}_{\kappa_z})_{z \in \{1, \dots, Z\}}$. We then scale $\Gamma \in \mathbb{R}_+^{P_{\text{hist}}^2}$ weighting the Frobenius norm used in the calibration loss (3.8) as

$$\Gamma \propto \frac{1}{\bar{U} + U}, \quad (\text{B.3})$$

where the division is understood element-wise and \bar{U} is the matrix with all entries constant equal to the average entries of U . This choice ensures more weights are given to the entries of the covariance matrix of returns which are less volatile across the different decay rates.




Figure 33: Left: exponentially decaying weights (B.1) for various span parameters $\kappa \geq 1$; right: resulting weight matrix from (B.3) used in the loss J_1 (3.9).

B.2 Linear cone programming

The inward minimization problem in (3.18) is formulated as a linear cone program in terms of the variables

$$\bar{x} := \left(\bar{x}_0, (\bar{x}_p)_{p \in \{1, \dots, \frac{(N+N_c+1)(N+N_c)}{2}\}} \right),$$

where $(\bar{x}_p)_{p \in \{1, \dots, \frac{(N+N_c+1)(N+N_c)}{2}\}}$ is identified as a column-major ordered vector corresponding to the lower triangular part of the symmetric matrix variables $(x_{p,k}(\sigma, R))_{p,k \in \{1, \dots, N+N_c\}}$ from (3.5) such that

$$\min_{\bar{x} \in \bar{U}} c^T \bar{x} \quad (\text{B.4})$$

$$\bar{U} := \left\{ \bar{x} \in \mathbb{R}^{1 + \frac{(N+N_c+1)(N+N_c)}{2}} : G^\lambda(\tau(a)) \bar{x} + s = h, A\bar{x} = b, s \in C \right\}, \quad (\text{B.5})$$

with $C := C_0 \times C_1 \times C_2$ such that

$$C_0 := \{u \in \mathbb{R}^N : u_i \geq 0, i \in \{1, \dots, N\}\}, \quad (\text{B.6})$$

$$C_1 := \left\{ (u_0, u_1) \in \mathbb{R} \times \mathbb{R}^{\frac{(N+N_c+1)(N+N_c)}{2}} : u_0 \geq \|u_1\|_2 \right\}, \quad (\text{B.7})$$

$$C_2 := \{\text{vec}(u) \mid u \in \mathbb{S}_+^N\},$$

and where $\|\cdot\|_2$ the standard Euclidean norm, and $\text{vec}(u)$ denotes a symmetric matrix u stored as a vector in column major order.

Indeed, the linear cone program (B.4)–(B.5) handles the various constraints on (σ, R) stated in the inner optimization problem from (3.18) such that:

- the non-negative orthant C_0 (B.6) captures constraints of the form $\sigma_i^{\text{upper}} \geq \sigma_i \geq 0$, $i \in \{1, \dots, N\}$, where typically σ_1^{upper} is chosen to be a small multiple of the rolling futures contract’s annualized volatility associated to the latest delivery to constraint σ_L ,
- we specify $c := \left(1, 0_{\mathbb{R}^{\frac{(N+N_c+1)(N+N_c)}{2}}}\right)$ so that “ \bar{x}_0 plays the role of u_0 ” in the second-order cone C_1 from (B.7), and we define accordingly the $\tau(a)$ -dependent quantity $G^\lambda(\tau(a))$ using the weights (3.15)–(3.16), and h with the respective target market values from the loss definition (3.13) so as to ensure both the quadratic loss function J^λ is minimized and $(x_{i,j})_{i,j \in I} \in \mathbb{S}_+^N$, where I denotes the set of indices relative to $\{L, S, C_1\}$ state variables, i.e. $(x_{i,j})_{i,j \in I}$ is obtained by withdrawing the columns and rows involving the state variable C_2 from the matrix $(x_{p,k})_{p,k \in \{1, \dots, N+N_c\}}$,

- finally, we specify (A, b) accordingly to ensure the equality constraints (3.17) with respect to the state variables of the C -factors.

In practice, (B.4) can be solved numerically by any linear cone program solver such as *conelp* from CVXOPT, see <https://cvxopt.org/userguide/coneprog.html> for additional details, and from which we borrowed the notations.

C Proof of Theorem 4.2

Proof. Without loss of generality, suppose that $[T_s^j, T_e^j] \subset [T_s^1, T_e^1]$ for $j \in \{2, \dots, P_{\text{imp}}\}$, since the delivery periods non included in $[T_s^1, T_e^1]$ will not impact the function h . Hence, they can be eliminated from the fixed-point iteration algorithm to be calibrated once at the end.

Given the vector $g = (\bar{g}_1, \dots, \bar{g}_{P_{\text{imp}}-1}) \in \mathbb{R}^{P_{\text{imp}}-1}$, we construct the vector $h = (h_1, \dots, h_{N_1}) \in \mathbb{R}^{N_1}$ by (4.5):

$$h_k^2(g) = \frac{\text{VS}_{T_k^1}^1 - \text{VS}_{T_{k-1}^1}^1}{\int_{T_{k-1}^1}^{T_k^1} \Sigma_s^\top(T_s^1, T_e^1) R \Sigma_s(T_s^1, T_e^1) ds}, \quad k \in \{1, \dots, N_1\}. \quad (\text{C.1})$$

Since $\Sigma_s(T_s^1, T_e^1)$ is affine in g

$$\Sigma_s(T_s^1, T_e^1) = \sum_{i=2}^{P_{\text{imp}}} \omega_1^i \bar{g}_i \hat{\Sigma}_s(T_s^i, T_e^i) + \left(1 - \sum_{i=2}^{P_{\text{imp}}} \omega_1^i\right) \hat{\Sigma}_s\left([T_s^1, T_e^1] \setminus \cup_{i=2}^{P_{\text{imp}}} [T_s^i, T_e^i]\right),$$

the equation (C.1) can be rewritten as

$$h_k^2(g) = \frac{\text{VS}_{T_k^1}^1 - \text{VS}_{T_{k-1}^1}^1}{g^\top Q_k g + 2g^\top p_k + r_k}, \quad k \in \{1, \dots, N_1\},$$

where $Q_k \in \mathbb{R}^{(P_{\text{imp}}-1) \times (P_{\text{imp}}-1)}$, $p_k \in \mathbb{R}^{P_{\text{imp}}-1}$, and $r_k \in \mathbb{R}$ are defined by

$$\begin{aligned} (Q_k)_{ij} &= \omega_1^i \omega_1^j \int_{T_{k-1}^1}^{T_k^1} \hat{\Sigma}_s^\top(T_s^i, T_e^i) R \hat{\Sigma}_s(T_s^j, T_e^j) ds, \quad i, j \in \{2, \dots, P_{\text{imp}}\}, \\ (p_k)_i &= \omega_1^i \left(1 - \sum_{l=2}^{P_{\text{imp}}} \omega_1^l\right) \int_{T_{k-1}^1}^{T_k^1} \hat{\Sigma}_s^\top(T_s^i, T_e^i) R \hat{\Sigma}_s\left([T_s^1, T_e^1] \setminus \cup_{l=2}^{P_{\text{imp}}} [T_s^l, T_e^l]\right) ds, \quad i \in \{2, \dots, P_{\text{imp}}\}, \\ r_k &= \left(1 - \sum_{l=2}^{P_{\text{imp}}} \omega_1^l\right)^2 \int_{T_{k-1}^1}^{T_k^1} \hat{\Sigma}_s^\top\left([T_s^1, T_e^1] \setminus \cup_{l=2}^{P_{\text{imp}}} [T_s^l, T_e^l]\right) R \hat{\Sigma}_s\left([T_s^1, T_e^1] \setminus \cup_{l=2}^{P_{\text{imp}}} [T_s^l, T_e^l]\right) ds \end{aligned}$$

Since we have assumed that the contracts have positive instantaneous correlations, we obtain $(Q_k)_{ij} > 0$ and $(p_k)_i > 0$ for all $i, j \in \{2, \dots, P_{\text{imp}}\}$. Moreover, by the linear independence hypothesis on the division, $[T_s^1, T_e^1] \setminus \cup_{i=2}^{P_{\text{imp}}} [T_s^i, T_e^i] \neq \emptyset$, so that $r_k > 0$.

The new iteration of g calculated by (4.8) is then given by

$$\psi_i(g) = \sqrt{\frac{\text{VS}_{T_i}^i}{\sum_{k=1}^{N_1} h_k^2(g) \int_{[0, T^i] \cap [T_{k-1}, T_k]} \hat{\Sigma}_s^\top(T_s^i, T_e^i) R \hat{\Sigma}_s(T_s^i, T_e^i) ds}} = \sqrt{\frac{\text{VS}_{T_i}^i}{S_i^\top h^2(g)}}, \quad i \in \{2, \dots, P_{\text{imp}}\}. \quad (\text{C.2})$$

where we have denoted $S_i = \left(\int_{[0, T^i] \cap [T_{k-1}, T_k]} \hat{\Sigma}_s^\top(T_s^i, T_e^i) R \hat{\Sigma}_s(T_s^i, T_e^i) ds \right)_{k \in \{1, \dots, N_1\}}$. Note all that the elements of S_i are positive.

(i) Existence. We will apply the Brouwer's theorem to establish the existence of the fixed-point. Since all the coefficients are positive, the mapping $\psi : \mathbb{R}_+^{P_{\text{imp}}-1} \rightarrow \mathbb{R}_+^{P_{\text{imp}}-1}$ is continuous. Thus, it is sufficient to find a convex compact $K \subset \mathbb{R}_+^{P_{\text{imp}}-1}$, such that $\psi(K) \subset K$. We will look for a compact of the form $K_R = \{x \in \mathbb{R}_+^{P_{\text{imp}}-1} : \|x\| \leq R\}$ and we will show that there exists an $R > 0$ big enough, such that $\psi(K_R) \subset K_R$.

Indeed, if $g \in K_R$, then

$$h_k^2(g) \geq \frac{\text{VS}_{T_k^1}^1 - \text{VS}_{T_{k-1}^1}^1}{\|Q_k\| \|g\|^2 + 2\|g\| \|p_k\| + r_k} \geq \frac{\text{VS}_{T_k^1}^1 - \text{VS}_{T_{k-1}^1}^1}{\|Q_k\| R^2 + 2\|p_k\| R + r_k}.$$

The norm of ψ can be bounded:

$$\|\psi(g)\|^2 \leq \sum_{i=2}^{P_{\text{imp}}} \frac{\text{VS}_{T^i}^i}{\sum_{k=1}^{N_1} S_i^k \frac{\text{VS}_{T_k^1}^1 - \text{VS}_{T_{k-1}^1}^1}{\|Q_k\| R^2 + 2\|p_k\| R + r_k}}$$

Hence,

$$\lim_{R \rightarrow \infty} \frac{\|\psi(g)\|^2}{R^2} \leq \sum_{i=2}^{P_{\text{imp}}} \frac{\text{VS}_{T^i}^i}{\sum_{k=1}^{N_1} S_i^k \frac{\text{VS}_{T_k^1}^1 - \text{VS}_{T_{k-1}^1}^1}{\|Q_k\|}} < 1,$$

where we have used (4.10). Thus, there exists R , such that for $g \in K_R$, we have $\|\psi(g)\| < R$, and $\psi(K_R) \subset K_R$.

(ii) Stability. To prove stability of the fixed-point, we linearize ψ and study its Jacobian matrix $\frac{\partial \psi}{\partial g}$.

Taking the derivative of (C.2), we obtain

$$\frac{\partial \psi_i(g)}{\partial g_j} = -\frac{1}{2} \sqrt{\frac{\text{VS}_{T^i}^i}{S_i^\top h^2(g)}} \frac{1}{S_i^\top h^2(g)} \sum_{k=1}^{N_1} S_i^k \frac{\partial}{\partial g_j} h_k^2(g) = \frac{\psi_i(g)}{S_i^\top h^2(g)} \sum_{k=1}^{N_1} S_i^k h_k^2(g) \frac{2g^\top Q_k^j + 2(p_k)_j}{2g^\top Q_k g + 4g^\top p_k + 2r_k},$$

where Q_k^j denotes the j -th row of Q_k . If g^* is a fixed-point of ψ , then $\psi_i(g^*) = g_i^*$, and we have

$$\sum_{j=2}^{P_{\text{imp}}} \frac{1}{g_i^*} \frac{\partial \psi_i(g^*)}{\partial g_j} g_j^* = \frac{1}{S_i^\top h^2(g^*)} \sum_{k=1}^{N_1} S_i^k h_k^2(g^*) \underbrace{\frac{2g^{*\top} Q_k g^* + 2g^\top p_k}{2g^{*\top} Q_k g^* + 4g^{*\top} p_k + 2r_k}}_{< 1} < 1, \quad i \in \{2, \dots, P_{\text{imp}}\}, \quad (\text{C.3})$$

as all the coefficients in this expression are positive. (C.3) implies that the matrix $A := \text{diag}(g^*)^{-1} \frac{\partial \psi}{\partial g}(g^*) \text{diag}(g^*)$

is a contraction in $(\mathbb{R}^{P_{\text{imp}}-1}, \|\cdot\|_\infty)$. Thus, there exists a neighborhood U of g^* such that the iterations g^n converge for any $g^0 \in U$. Indeed, defining $e_n = g_n - g^*$, we obtain in U

$$e_{n+1} = \text{diag}(g^*) A \text{diag}(g^*)^{-1} e_n,$$

so that $\|\text{diag}(g^*)^{-1} e_n\|_\infty$ tends to 0 as $n \rightarrow \infty$, and $\|e_n\|_\infty \rightarrow 0$.

(iii) Uniqueness. We establish uniqueness only for the case $N_1 = 1$. By (C.2), the fixed-point equation $\psi(g) = g$ reads

$$\psi_i(g) = \sqrt{\frac{1}{S_i^1} \frac{\text{VS}_{T^i}^i}{S_i^1 h_1^2(g)}} = \sqrt{\frac{1}{S_i^1} \frac{\text{VS}_{T^i}^i}{\text{VS}_{T_1^1}^1 - \text{VS}_{T_0^1}^1}} \sqrt{g^\top Q_1 g + 2g^\top p_1 + r_1} = \alpha(g) e_i,$$

where

$$e := \left(\sqrt{\frac{1}{S_i^1} \frac{\text{VS}_{T^i}^i}{\text{VS}_{T_1^1}^1 - \text{VS}_{T_0^1}^1}} \right)_{i \in \{2, \dots, P_{\text{imp}}\}}, \quad \alpha(g) := \sqrt{g^\top Q_1 g + 2g^\top p_1 + r_1}.$$

Is is clear that the fixed point g , if exists, should be of the form

$$g = \beta \mathbf{e}, \quad \beta > 0,$$

so that the fixed-point problem in $\mathbb{R}^{P_{\text{imp}}-1}$ is reduced to the one-dimensional fixed-point problem for the mapping

$$\psi_\beta : \mathbb{R}_+ \rightarrow \mathbb{R}_+, \quad \beta \mapsto \alpha(\beta \mathbf{e}) = \sqrt{\mathbf{e}^\top Q_1 \mathbf{e} \beta^2 + 2\mathbf{e}^\top p_1 \beta + r_1},$$

which is equivalent to the quadratic equation

$$(1 - \mathbf{e}^\top Q_1 \mathbf{e})\beta^2 - 2\mathbf{e}^\top p_1 \beta - r_1 = 0. \quad (\text{C.4})$$

The condition (4.10) guarantees that $\mathbf{e}^\top Q_1 \mathbf{e} < 1$, so that the equation (C.4) admits a unique positive root β^* and the unique fixed-point of ψ is given by $g^* = \beta^* \mathbf{e}$. \square

D Monte Carlo simulation scheme

The European call option prices $C(T, K)$ can also be computed using the Monte Carlo method. At each step t of our discrete-time grid, given $(U_t^i)_{i \in \{1, \dots, M\}}$, V_t , and F_t and a time-step Δt , we simulate $(U_{t+\Delta t}^i)_{i \in \{1, \dots, M\}}$, then $V_{t+\Delta t}$, and finally $F_{t+\Delta t}$. For the processes U_t^i , we start with the following semi-implicit Euler discretization scheme of (1.3):

$$U_{t+\Delta t}^i - U_t^i = -x_i U_{t+h}^i \Delta t - \lambda V_t \Delta t + \nu \sqrt{V_t} (B_{t+\Delta t} - B_t),$$

which leads to

$$U_{t+\Delta t}^i = \frac{1}{1 + x_i \Delta t} (U_t^i - \lambda V_t \Delta t + \nu \sqrt{V_t} (B_{t+\Delta t} - B_t))$$

We choose this semi-implicit Euler scheme following [Abi Jaber \(2019\)](#), since it gives more stable results than standard Euler scheme which explodes for large mean-reversion coefficients x_i .

We obtain the variance $V_{t+\Delta t}$ directly by (1.2) ensuring that the variance process is floored at zero to prevent negative values. For $\log F_{t+\Delta t}$, standard Euler scheme is used. We sum up the simulation scheme with the following formulae:

$$\begin{aligned} U_{t+\Delta t}^i &= \frac{1}{1 + x_i \Delta t} (U_t^i - \lambda V_t \Delta t + \nu \sqrt{V_t} (B_{t+\Delta t} - B_t)), \quad i \in \{1, \dots, M\}, \\ V_{t+\Delta t} &= \left(m_0(t + \Delta t) + \sum_{i=1}^M c_i U_{t+\Delta t}^i \right)^+, \\ \log F_{t+\Delta t} &= \log F_t - \frac{1}{2} h(t)^2 V_t \Sigma_t^\top R \Sigma_t \Delta t + h(t) \sqrt{V_t} \Sigma_t^\top (W_{t+\Delta t} - W_t). \end{aligned}$$

References

- Eduardo Abi Jaber. Lifting the Heston model. *Quantitative Finance*, 19(12):1995–2013, 2019.
- Eduardo Abi Jaber, Martin Larsson, and Sergio Pulido. Affine Volterra processes. *The Annals of Applied Probability*, 29(5):3155–3200, 2019.
- Eduardo Abi Jaber, Christian Bayer, and Simon Breneis. State spaces of multifactor approximations of nonnegative Volterra processes. *arXiv:2412.17526*, 2024.
- Leif B. G. Andersen. Markov models for commodity futures: Theory and practice. *Quantitative Finance*, 10(8), 2010.
- Fred Espen Benth and Fabian A Harang. Infinite dimensional pathwise volterra processes driven by Gaussian noise—probabilistic properties and applications. *Electronic Journal of Probability*, 26:1–42, 2021.

- Fred Espen Benth and Steen Koekebakker. Stochastic modeling of financial electricity contracts. *Energy Economics*, 30(3):1116–1157, 2008.
- Fred Espen Benth and Paul Krühner. *Stochastic Models for Prices Dynamics in Energy and Commodity Markets*. Springer, 2023.
- Fred Espen Benth and Florentina Paraschiv. A space-time random field model for electricity forward prices. *Journal of Banking & Finance*, 95:203–216, 2018.
- Fred Espen Benth and Iben Cathrine Simonsen. The Heston stochastic volatility model in Hilbert space. *Stochastic Analysis and Applications*, 36(4):733–750, 2018.
- Fred Espen Benth, Jurate Saltyte Benth, and Steen Koekebakker. *Stochastic modelling of electricity and related markets*, volume 11. World Scientific, 2008.
- Fred Espen Benth, Marco Piccirilli, and Tiziano Vargiolu. Mean-reverting additive energy forward curves in a Heath–Jarrow–Morton framework. *Mathematics and Financial Economics*, 13(4):543–577, 2019.
- Lorenzo Bergomi and Julien Guyon. Stochastic volatility’s orderly smiles. *Risk*, 25(5):60, 2012.
- Stephen Boyd and Lieven Vandenbergh. *Convex optimization*. Cambridge university press, 2004.
- Damiano Brigo and Fabio Mercurio. *Interest Rate Models — Theory and Practice: With Smile, Inflation and Credit*. Springer Verlag, 2006.
- René Carmona and Valdo Durrleman. Pricing and hedging spread options. *Siam Review*, 45(4):627–685, 2003.
- Peter Carr and Dilip Madan. Towards a theory of volatility trading. *Volatility: New estimation techniques for pricing derivatives*, 29:417–427, 1998.
- Alvaro Cartea and Marcelo G Figueroa. Pricing in electricity markets: a mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance*, 12(4):313–335, 2005.
- Les Clewlow and Chris Strickland. Valuing energy options in a one factor model fitted to forward prices. research paper series 10. *Quantitative Finance Research Centre, University of Technology, Sydney*, 10, 1999.
- Gonzalo Cortazar, Matias Lopez, and Lorenzo Naranjo. A multifactor stochastic volatility model of commodity prices. *Energy Economics*, 67:182–201, 2017.
- Sonja Cox, Sven Karbach, and Asma Khedher. An infinite-dimensional affine stochastic volatility model. *Mathematical Finance*, 32(3):878–906, 2022.
- Kresimir Demeterfi, Emanuel Derman, Michael Kamal, and Joseph Zou. More than you ever wanted to know about volatility swaps. *Goldman Sachs quantitative strategies research notes*, 41:1–56, 1999.
- Thomas Deschatre, Olivier Féron, and Pierre Gruet. A survey of electricity spot and futures price models for risk management applications. *Energy Economics*, 102:105504, 2021.
- Enrico Edoli, Davide Tasinato, and Tiziano Vargiolu. Calibration of a multifactor model for the forward markets of several commodities. *Optimization*, 62(11):1553–1574, 2013.
- Olivier Féron and Pierre Gruet. Estimation of the number of factors in a multi-factorial Heath–Jarrow–Morton model in power markets. In *Quantitative Energy Finance: Recent Trends and Developments*, pages 3–39. Springer, 2024.
- Frederick N Fritsch and Judy Butland. A method for constructing local monotone piecewise cubic interpolants. *SIAM journal on scientific and statistical computing*, 5(2):300–304, 1984.
- Matteo Gardini and Edoardo Santilli. A Heath–Jarrow–Morton framework for energy markets: review and applications for practitioners. *Decisions in Economics and Finance*, 2024. ISSN 1129-6569.

- Jim Gatheral and Antoine Jacquier. Arbitrage-free SVI volatility surfaces. *Quantitative Finance*, 14(1): 59–71, 2014.
- David Heath, Robert Jarrow, and Andrew Morton. Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica*, 60(1):77–105, 1992. ISSN 00129682, 14680262.
- Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Matrix Analysis. Cambridge University Press, 2013. ISBN 9780521839402.
- Angeliem Kemna and A.C.F. Vorst. A pricing method for options based on average asset values. *Journal of Banking & Finance*, 14(1):113–129, 1990.
- Rüdiger Kiesel, Gero Schindlmayr, and Reik H Börger. A two-factor model for the electricity forward market. *Quantitative Finance*, 9(3):279–287, 2009.
- Steen Koekebakker and Fridthjof Ollmar. Forward curve dynamics in the nordic electricity market. *Managerial Finance*, 31(6):73–94, 2005.
- Alan Lewis. A simple option formula for general jump-diffusion and other exponential Levy processes. *SSRN Electronic Journal*, 2001.
- Yuliya Mishura, Stefania Ottaviano, and Tiziano Vargiolu. Gaussian Volterra processes as models of electricity markets. *SSRN Electronic Journal*, 2023.
- Silvana Musti, Viviana Fanelli, and Lucia Maddalena. Modelling electricity futures prices using seasonal path-dependent volatility. *Applied Energy*, 173:92–102, 2016.
- Charles Nelson and Andrew Siegel. Parsimonious modeling of yield curves. *The Journal of Business*, 60: 473–89, 1987.
- Marco Piccirilli, Maren Diane Schmeck, and Tiziano Vargiolu. Capturing the power options smile by an additive two-factor model for overlapping futures prices. *Energy Economics*, 95:105006, 2021.
- Paul A Samuelson. Proof that properly anticipated prices fluctuate randomly. In *The world scientific handbook of futures markets*, pages 25–38. World Scientific, 2016.
- Maren Diane Schmeck and Stefan Schwerin. The effect of mean-reverting processes in the pricing of options in the energy market: An arithmetic approach. *Risks*, 2021.
- Artur Sepp and Parviz Rakhmonov. Stochastic volatility for factor HJM framework. *SSRN Electronic Journal*, 01 2023. doi: 10.2139/ssrn.4646925.