

Calibrating a Market Model with Stochastic Volatility to Commodity and Interest Rate Risk

P. Karlsson^{*1}, K.F. Pilz^{†2}, and E. Schlögl^{‡3}

¹Quantitative Analytics, ING Bank, Amsterdam, The Netherlands.

²RIVACON GmbH, Germany.

³University of Technology Sydney, Quantitative Finance Research Centre, Ultimo NSW 2007, Australia.

March 19, 2017

Abstract

Based on the multi-currency LIBOR Market Model (LMM) this paper constructs a hybrid commodity interest rate market model with a stochastic local volatility function allowing the model to simultaneously fit the implied volatility surfaces of commodity and interest rate options. Since liquid market prices are only available for options on commodity futures, rather than forwards, a convexity correction formula for the model is derived to account for the difference between forward and futures prices. A procedure for efficiently calibrating the model to interest rate and commodity volatility smiles is constructed. Finally, the model is fitted to an exogenously given correlation structure between forward interest rates and commodity prices (cross-correlation). When calibrating to options on forwards (rather than futures), the fitting of cross-correlation preserves the (separate) calibration in the two markets (interest rate and commodity options), while in the case of futures a (rapidly converging) iterative fitting procedure is presented. The fitting of cross-correlation is reduced to finding an optimal rotation of volatility vectors, which is shown to be an appropriately modified version of the “orthonormal Procrustes” problem in linear algebra. The calibration approach is demonstrated in an application to market data for oil futures.

Keywords: Calibration; Commodity markets; Derivative pricing; Interest rate modelling; Interest rate derivatives; Oil futures; Energy derivatives

^{*}Email: patrik.karlsson@nek.lu.se. Parts of the work was carried out while Patrik was a PhD Student at Lund University, Sweden and a visiting scholar position at the Quantitative Finance Research Centre (QFRC) at University of Technology Sydney, Australia, Oct 2011 – Jan 2012. He wishes to thank Caroline Dobson and the QFRC for their hospitality. Moreover, thanks to Hans Byström for connecting him with the QFRC.

[†]Email: kay.pilz@rivacon.com. Parts of the work of this paper were carried out while being a visitor at the Quantitative Finance Research Centre (QFRC) at University of Technology Sydney, Australia. Kay Pilz wishes to thank the QFRC for their hospitality.

[‡]Corresponding author. Email: erik.schlogl@uts.edu.au

1 Introduction

The modelling of market risks for the pricing of derivative financial instruments has come a long way since the seminal paper of [1]. In particular, it is widely recognised that such models need to be calibrated to all available liquid market prices, including options of various strikes and maturities, for all relevant sources of risk. For commodity derivatives, the approach presented in this paper represents a step closer to this ideal.

In addition to commodity prices and their stochastic dynamics, the valuation and risk management of positions in commodity derivatives also depend on market interest rates and the stochastic dynamics thereof. The market instruments, to which a model should be calibrated, include the swaption “cube” (swaptions of (1) various maturities, (2) various strikes, on (3) swaps of various lengths) and commodity options of various maturities and strikes. For commodities, futures are more liquid than forwards — consequently (as well as to make the model more realistic) the correlation between commodity prices and interest rates becomes a relevant model input already at the level of calibration.

The model presented in this paper, with its associated calibration method, is fitted to market prices for swaptions in a swaption cube, options on commodity futures for various maturities and strikes, and of course the underlying futures and interest rate term structures. Furthermore, it is fitted to exogenously estimated correlations between interest rates and commodity prices. The construction is based on a LIBOR Market Model (LMM)¹ for the interest rate market and the commodity market, with the two markets linked in a manner analogous to the construction of the multicurrency LMM,² where the convenience yield takes the role of the interest rate in the commodity market (thus convenience yields are assumed to be stochastic). To allow a fit to market-implied volatility smiles (and skews) of commodity and interest rate options, the model is equipped with a stochastic local volatility function (SLV), following [2], [3] and [4].

Efficient calibration is achieved in two steps, by first calibrating the model to the interest rate market separately, building on a synthesis of the calibration approaches for the LMM in [5] and the SLV-LMM in [6]. In order then to be able to calibrate efficiently to commodity futures, we consider two approximations for calculating the difference between futures and forwards in the proposed model. The separate calibration in the interest rate market in the first step is then followed by an iterative, two-stage calibration to the commodity market, where in the second stage an orthonormal transformation of the commodity volatility vectors is applied, rotating the commodity volatilities relative to the interest rate volatilities in such a manner as to achieve the desired correlations between the two markets. The calibration of this orthonormal transformation to the desired cross-correlations is cast in terms of a modified orthonormal Procrustes problem, permitting an effective solution algorithm to be applied. We illustrate the use of the model on real market data.

For the example, we chose U.S. Dollars (USD) as the “domestic” currency and Brent Crude Oil as the commodity (the “foreign currency”). The “exchange rate” is given by the Brent Crude Oil futures prices denoted in USD. These prices, when converted to forward prices using an appropriate convexity correction, can be interpreted as forward exchange rates between the USD economy and an economy where value is measured in terms of units of Brent Crude Oil (where convenience yields are interpreted as the foreign interest rates). In the example, the model is calibrated to the USD swaptions volatility cube, and the volatility smile of European-style options on Brent Crude Oil futures.

Hybrid modelling combining commodity and interest rate risk was initiated by [7], who modelled interest

¹See the seminal papers by [8], [9] and [10].

²See [11].

³[12] presented an extensions of [13] with stochastic volatility.

rate risk via stochastic dynamics of the continuously compounded short rate, without reference to a full model calibrated to an initial term structure. Subsequently, a number of authors proposed models for stochastic convenience yields, some of whom also incorporated stochastic dynamics of the term structure of interest rates.⁴ In these models, continuously compounded convenience yields (and possibly interest rates) typically are normally distributed, since they are assumed to be driven by a \mathbb{Q} term structure model with generalised (possibly multi-factor) Ornstein/Uhlenbeck dynamics. In such a model, effective calibration to available commodity and interest rate options is difficult when only at-the-money options are considered, and not possible for the full range of available strikes. At-the-money calibration is a strength of lognormal LIBOR Market Models, and \mathbb{Q} construct a hybrid model which exploits this, also making use of an orthonormal rotation of volatility vectors to fit the cross-correlations between the commodity and interest rate markets. By lifting the lognormality assumption, the present paper goes beyond their work to allow calibration to the full swaption cube and commodity volatility surface, and it refines the correlation fitting procedure through rotation by casting it as a modification of the orthonormal Procrustes problem,⁵ which can be solved by a fast numerical algorithm.⁶

From a practical perspective, whether it is worthwhile to move beyond simple models (e.g., in the present context, pricing commodity derivatives using an adaption of the \mathbb{Q} model where the underlying commodity follows geometric Brownian motion and interest rates are deterministic) depends on whether the more sophisticated model produces substantially different derivative prices (and hedges) when compared to a simpler benchmark. Relative to a model along the lines of \mathbb{Q} or \mathbb{Q} , a hybrid model of stochastic commodity prices and interest rates based on a stochastic local volatility LMM represents extensions in two dimensions: fitting the implied volatility surface observed in the market, and integrating commodity and interest rate risk. On the necessity of the former, there is little debate among practitioners: The primary purpose of the calibrated models is to price illiquid derivative products in a manner consistent with observed prices for (typically simpler) products liquidly traded in the market; thus soon after implied volatility smiles (contradicting the Black/Scholes assumption of geometric Brownian motion) appeared in the market, practitioners recognised the need for models consistent with this observation.⁷ For the latter, the impact of adding stochastic interest rates to a (stochastic volatility) commodity model has been studied recently in a series of papers⁸ by Cheng, Nikitopoulos and Schlögl (2016a,b,c,d). \mathbb{Q} find a noticeable impact of interest rate volatility, as well as of correlation between the interest rate process and the (commodity) futures price process, on the prices of long-dated futures options, with this impact becoming less pronounced for shorter maturities. Looking at the problem from a more traditional academic perspective, \mathbb{Q} find that allowing for stochastic interest rates improves the out-of-sample empirical performance of their model. The impact on the hedging effectiveness (as opposed to pricing performance) of incorporating interest rate risk into a model for commodity derivatives is studied in Cheng, Nikitopoulos and Schlögl (\mathbb{Q}) and (\mathbb{Q}), where the first paper considers this question in a simulated model, while the second conducts a back-test of hedging performance on empirical data. The latter study finds that in times of market turbulence, especially during the Global Financial Crisis 2007–2009, augmenting a commodity delta hedge with an interest rate hedge consistently improves hedge performance, more consistently in particular than augmenting the delta hedge by vega or gamma. Thus a hybrid model provides benefits for the pricing and risk management of vanilla products, and this effect would be more pronounced for

⁴See, for example, \mathbb{Q} , \mathbb{Q} , \mathbb{Q} , \mathbb{Q} , and \mathbb{Q} .

⁵See \mathbb{Q} .

⁶This algorithm is given as Algorithm 8.1 in \mathbb{Q} .

⁷This is reported, for example, by \mathbb{Q} .

⁸The stochastic interest rate dynamics used in these papers are simpler than the SLV-LMM dynamics considered here, but nevertheless demonstrate the relevance of incorporating interest rate risk into a commodity derivatives model.

more exotic products which explicitly condition jointly on interest rate and commodity risk.

As noted above, the model considered in the present paper is assembled from individual components based on the stochastic local volatility formulation of the LMM by ???. This represents one major strand of the literature extending the LMM beyond at-the-money calibration. The other major strand is based on the SABR model of ?, laid out in detail in ?. In the present context the choice of the SLV-LMM over the SABR-LMM as the basis of the hybrid interest rate/commodity model is driven by three considerations:

1. Specification of the stochastic volatility dynamics along the lines of ??? avoids the mathematical problems associated with SABR,⁹ i.e. the log-normality of the volatility process and the undesirable behaviour of the stochastic differential equation (SDE) of the underlying for certain values of the “constant elasticity of variance” (CEV) parameter β . The former implies divergence to infinity of volatility almost surely in finite time. The latter involves non-uniqueness of the solution to the SDE (for $0 < \beta < \frac{1}{2}$) and/or the process of the underlying financial variable being absorbed at zero (for $0 < \beta < 1$).
2. The SLV-LMM is directly amenable to calibration by an appropriately modified ? algorithm, in which we directly, exogenously control the correlation structure of the underlying financial variables (as opposed to the correlation structure of the driving Brownian motions). In particular in the absence of liquid market instruments containing useable information on “implied” correlations,¹⁰ correlation is subject to considerable “parameter uncertainty,” and having directly interpretable correlation inputs assists in controlling for this source of “model risk.”
3. In the present paper, the SLV-LMM for each market (interest rates and the commodity) is driven by a vector of independent Brownian motions. Thus correlations are introduced by the way volatility is distributed over these Brownian motions (“factors”) by the vector-valued volatility functions. This permits a two-stage procedure of fitting cross-correlations between markets after fitting the models for the individual markets (though there is some coupling when calibrating to futures), using orthonormal transformations.

The paper is organised as follows. The basic notation, the results of the single- and multi-currency LMM and their interpretation in the context of commodities are presented in Section ???. In Section ??? the calibration of the commodity part of the Commodity LMM to plain vanilla options is discussed. In Section ??? the relationship between futures and forwards in the model is presented, which permits calibration of the model to futures as well as forwards. The calibration of the interest rate part of the hybrid Commodity LMM will not be discussed in detail in this paper, because this problem has already been addressed by many authors, e.g., ???, and most methods should be compatible with our model. However, in Section ??? we discuss how both separately calibrated parts of the model – the interest rate and the commodity part – can be merged in order to have one underlying d -dimensional Brownian motion for the joint model and still match the market prices used for calibration of the particular parts. Section ??? illustrates the application of the model to real market data.

⁹See Section 3.10 of ?.

¹⁰Although swaption prices depend in theory on correlations between forward rates, in practice this dependence is too weak for these correlations to be extracted in a meaningful way; see e.g. ?.

2 The Commodity LIBOR Market Model

2.1 The LIBOR Market Model

For the construction of the LMM for the domestic interest rate market we assume a given probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where the underlying filtration $\{\mathcal{F}_t, t \in [0, T_N]\}$ coincides with the \mathbb{P} -augmentation of the natural filtration of a d -dimensional standard Brownian motion W , and $\mathbb{E}_t^\mathbb{P}[\cdot] := \mathbb{E}^\mathbb{P}[\cdot | \mathcal{F}_t]$ denote the conditional expectation on the information at time t . Let T_N be a fixed time horizon, $P(t, T)$, $T \in [t, T_N]$ the bond price, i.e. the amount that has to be invested at time t to receive one unit of the domestic currency at time T , hence $P(T, T) = 1$ for every $T \in [0, T_N]$. Assuming the discrete-tenor structure, $0 = T_0 < T_1 < \dots < T_N$, with intervals $\tau_n = T_{n+1} - T_n$, the forward LIBOR rate $L(t, T_n)$ with fixing period T_n as seen at time t is given by

$$L(t, T_n) = \tau_n^{-1} \left(\frac{P(t, T_n)}{P(t, T_{n+1})} - 1 \right), \quad q(t) \leq n \leq N-1,$$

where $q(t)$ is the index function of the LIBOR rate with the shortest maturity not fixed at time t , defined as $T_{q(t)-1} \leq t < T_{q(t)}$. The price of the discounted bond maturing at time $T_n > t$ is then given by

$$P(t, T_n) = P(t, T_{q(t)}) \prod_{i=q(t)}^{n-1} \frac{1}{1 + \tau_i L(t, T_i)}.$$

The dynamics of the forward LIBOR rate $L(t, T_n)$ as seen at time $t \in [0, T]$, under the $\mathbb{P}^{T_{n+1}}$ -forward measure¹¹ is given by

$$dL(t, T_n) = \sigma_L(t, T_n)^\top dW^{T_{n+1}}(t), \quad (1)$$

where $\sigma_L(t, T_n)$ is a d -dimensional process, discussed later in this section. From Girsanov's theorem, the dynamics of $L(t, T_n)$ are

$$dL(t, T_n) = \sigma_L(t, T_n)^\top \left(\gamma_L(t, T_n) dt + dW^{T_n}(t) \right),$$

where W^{T_n} is a d -dimensional vector-valued Brownian motion¹² under the \mathbb{P}^{T_n} -forward measure and γ_L is determined by the volatility of the forward bond price process, i.e.

$$\begin{aligned} d \left(\frac{P(t, T_n)}{P(t, T_{n+1})} \right) &= \frac{P(t, T_n)}{P(t, T_{n+1})} \gamma_L(t, T_n) dW^{T_{n+1}}(t) \quad \text{with} \\ \gamma_L(t, T_n) &= \frac{\tau_n \sigma_L(t, T_n)}{1 + \tau_n L(t, T_n)}, \end{aligned} \quad (2)$$

relates $dW^{T_{n+1}}$ to dW^{T_n} by,

$$dW^{T_n}(t) = dW^{T_{n+1}}(t) - \gamma_L(t, T_n) dt. \quad (3)$$

¹¹This forward measure is the equivalent martingale measure associated with taking the zero coupon bond $P(t, T_{n+1})$ as the numeraire, and under this measure (the existence of which is assured under the model assumptions below) forward LIBOR $L(t, T_n)$ is necessarily a martingale, i.e. driftless — see e.g. ?.

¹²Thus W^{T_n} is a d -dimensional vector, each component $W_i^{T_n}$, $1 \leq i \leq d$, is a Brownian motion under the \mathbb{P}^{T_n} -forward measure, and the quadratic covariation between the components is zero: $dW_i^{T_n} dW_j^{T_n} = 0 \forall i \neq j$.

Further results and the connection of this model to the framework of ? can be found in the original LMM literature, commencing with ?, ?, ? and ?.

2.1.1 The Stochastic Local Volatility LMM

For most markets, implied volatilities calculated from traded option prices are strike dependent, i.e. exhibit a volatility smile and skew (slope of the at-the-money volatility). To capture the skew, we assume that the time-dependent volatility functions are of the separable form

$$\sigma_L(t, T_n) = \varphi_L(L(t, T_n)) \lambda_L(t, T_n), \quad (4)$$

where $\lambda_L(t, T_n)$ is a bounded deterministic d -dimensional function and $\varphi_L : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is a time homogenous local volatility function. This is a fairly general setup, and one model allowing for skewed implied volatility is the displaced–diffusion model along the lines of ?, where, following the notation of ? (see their Remark 7.2.13), φ_L is given by

$$\varphi_L(L(t, T_n)) = b_L(t, T_n) L(t, T_n) + (1 - b_L(t, T_n)) L(0, T_n).$$

When $(1 - b_L(t, T_n)) / b_L(t, T_n) < (L(0, T_n) \tau_n)^{-1}$, the existence of path wise unique solutions follow (?, Lemma 14.2.5).

To capture the volatility smile, we follow ? and scale the Brownian motions with a mean-reverting stochastic volatility process given by,

$$dz_L(t) = \theta(z_{L,0} - z_L(t)) dt + \eta \sqrt{z_L(t)} dZ_L(t), \quad (5)$$

where θ and η are positive constants, $z_L(0) = z_{L,0} = 1$ and Z_L is a Brownian motion under the spot measure \mathbb{Q}^B .¹³ The quadratic covariation of Z_L and each component of W is assumed to be zero. Assuming the LIBOR dynamics in (??) with the separable volatility function in (??) and stochastic volatility (??), the stochastic local volatility LIBOR market model (SLV–LMM) specifies the dynamics of the forward LIBOR rates for $n = 1, \dots, N$ by

$$dL(t, T_n) = \sqrt{z_L(t)} \varphi_L(L(t, T_n)) \lambda_L^\top(t, T_n) dW^{T_{n+1}}(t). \quad (6)$$

When calibrated to interest rate option market data, the model matches at-the-money volatilities through λ_L , the skews (slope of the Black/Scholes implied volatilities) through b_L and curvatures of the volatility smiles through the volatility of variance η . The speed of mean reversion κ determines how fast the spot volatility converges to the forward volatility, or more specifically, how fast $z_L(t)$ is pulled back to its long–term mean level $z_{L,0}$.

The relationship between spot measure \mathbb{Q}^B and forward measures is given by standard results for the LMM dating back to ? and ?:

$$dW^{T_{n+1}}(t) = \sqrt{z_L(t)} \mu_n(t) dt + dW^B(t), \quad (7)$$

$$\mu_n(t) = \sum_{j=q(t)}^n \frac{\tau_j \varphi_L(L(t, T_j))}{1 + \tau_j L(t, T_j)} \lambda_L(t, T_j). \quad (8)$$

Moreover, we assume that the Brownian motion $Z_L(t)$ of the variance process $z_L(t)$ is independent of

¹³See for instance Section 4.2.3 of ?.

the d -dimensional Brownian motion W and that all forward LIBORs and factors are driven by the same scaling $\sqrt{z_L(t)}$.

2.2 The Commodity Market

The approach incorporating a commodity market corresponds largely to the approach as described for the log-normal case in Section 2.2 of ?, which in turn is based on the multi-currency extension of LIBOR market models introduced in ?. Therefore, we focus on the aspects related to the stochastic local volatility extensions of the model. The setup for the commodity parallels the one for interest rates in the previous section, and the corresponding volatility functions are denoted by σ_F and γ_F . As explained in ?, the commodity market can be seen as a “foreign interest market” with the commodity (e.g. crude oil) as currency. “Foreign bond prices” $C(t, T)$ can be interpreted as “convenience yield discount factors” for the commodity, defined as the amount of the commodity today which is equivalent to the discounted (using domestic interest rates) value of receiving one unit of the commodity (e.g. one barrel of crude oil) at time T , taking into account any storage costs and convenience yields.¹⁴ The same logic as for the domestic interest rate market can be used to derive “forward rates” for the commodity market, but since such “convenience yield instruments” are not traded for commodities, we construct the model by specifying domestic interest rate dynamics on the one hand, and the dynamics of forward commodity prices on the other hand. Then, as noted in ?, this implicitly determines the “foreign interest rate dynamics,” i.e. the convenience yield dynamics in the present interpretation.

As in ?, the existence of a spot price process $S(t)$ for the commodity is assumed, denoted in the local currency (e.g. USD per barrel crude oil). Then, its forward value is given by

$$F(t, T_n) = \frac{C(t, T_n)S(t)}{P(t, T_n)}, \quad (9)$$

for all $n = 0, \dots, N$. We assume the same tenor structure $\tau_n = T_{n+1} - T_n$ for interest and commodity markets. If this assumption needs to be lifted in order to reflect market reality, an interpolation on either of the forwards can be applied. Since LIBORs have typically 3-month or 6-month tenors, and exchange traded futures¹⁵ often have expiries with 1-month or 3-month time difference, the interpolation has to be made for forward interest rates in most cases. See for instance ? on forward interest rate interpolation.

The forwards in (??) are necessarily martingales under the \mathbb{P}^{T_n} -forward measure, i.e.

$$dF(t, T_n) = \sigma_F^\top(t, T_n) dW^{T_n}(t), \quad (10)$$

for all $n = 0, \dots, N$. To account for a stochastic local volatility dynamics for the commodity prices, the dynamics for the commodity forward prices are set to

$$dF(t, T_n) = \sqrt{z_F(t)} \varphi_F(F(t, T_n)) \lambda_F(t, T_n) dW^{T_n}(t), \quad (11)$$

where

$$dz_F(t) = \theta_F(z_{F_0} - z_F(t)) dt + \eta_F \sqrt{z_F(t)} dZ_F(t), \quad (12)$$

$$\varphi_F(F(t, T_n)) = b_F(t, T_n) F(t, T_n) + (1 - b_F(t, T_n)) F(0, T_n), \quad (13)$$

¹⁴Thus, the $C(t, T)$ represent the effect of the convenience yield net of storage cost.

¹⁵The futures versus forward relation will be discussed in Section ??.

and θ_F, η_F positive constants, $z_F(0) = z_{F_0} = 1$ and $b_F(t, T)$ a deterministic function mapping from $\mathbb{R}^+ \rightarrow \mathbb{R}^+$. The Brownian motions $W^{T_{n+1}}$, for $n = 0, \dots, N - 1$, are the same as in (??), $Z_F(t)$ is a Brownian motions under the spot measure \mathbb{Q}^B , and there is no correlation between the underlying drivers and volatility drivers in the sense that for all $n = 0, \dots, N - 1$

$$dW_i^{T_{n+1}} dZ_L(t) = dW_i^{T_{n+1}} dZ_F(t) = dZ_F(t) dZ_L(t) = 0 \quad \forall 1 \leq i \leq d.$$

Note that the structure of the dynamics for the commodity forwards $F(t, T_n)$ is the same as for the interest forward rates $L(t, T_n)$, except that they are martingales under different forward measures.

As demonstrated in ?, this fully specifies the hybrid model: Denote by $\gamma_F(t, T_n)$ the volatility of the quotient $C(t, T_n)/C(t, T_{n+1})$ of convenience yield discount factors, then $\gamma_F(t, T_n)$ is determined by the no-arbitrage relation of the multi-currency LMM,¹⁶

$$\sigma_F(t, T_n) = \gamma_F(t, T_n) - \gamma_L(t, T_n) + \sigma_F(t, T_{n+1}). \quad (14)$$

where $\gamma_L(t, T_n)$ is defined by (??) and $\sigma_F(t, T_n)$ (and analogously $\sigma_F(t, T_{n+1})$) by (??) above.

In addition to the no-arbitrage condition (??), cross-correlations specify linkages between the interest rates and commodity forwards markets. Their form and calibration will be discussed in detail in Section ??.

3 Calibration with Time Dependent Parameters

This section discusses aspects of the calibration of the hybrid model, which consists in our approach of two parts. The first part calibrates the (LIBOR) interest rate forward market and the commodity market separately to their market instruments. The second part merges the two separate calibrations with due regard to cross-correlations and the no-arbitrage condition. Readers who have their own preferred individual calibration routines for the stochastic local volatility LMM of Section ??, as well as for the stochastic local volatility commodity model of Section ??, may skip this and the following section and continue directly with Section ?. As mentioned in the introduction, we focus on the calibration of the commodity leg, since a calibration of the LIBOR market model in the context of stochastic local volatility has already been addressed by many other authors, for instance ? and ?.

Since the Commodity LMM is based on commodity forwards, we have to calibrate to forward implied volatilities or plain vanilla option prices written on forwards. However, commodities futures rather than forwards are most liquidly traded (consider, for example, the Brent Crude Oil futures in the market data example in Section ??) and thus forward prices have to be deduced from futures prices. As we are working in a hybrid model that is integrating commodity and interest rate risk, it is not adequate to equate forward prices with futures prices, as is still common among practitioners. Section ?? describes how to take into account the distinction between futures and forwards when applying the calibration methods proposed in the present section.

The calibration of the model to commodity forward instruments follows the ideas of ?? and is split into two parts. First, a pre-calibration is performed to determine a globally constant speed of mean reversion θ_F and volatility of variance η_F such that the volatility smile given from market option quotes is matched as closely as possible. Second, the volatility term structure λ_F and the volatility skew structure b_F are fitted to option prices.

¹⁶See Equation (11) in ?.

3.1 Step 1 – Calibrating the level of mean reversion and volatility of variance

To obtain an efficient calibration algorithm we follow ? and project the full dynamics of the commodity forwards $F(\cdot, T_n)$ in Equation (??) with time-dependent parameters onto a model with constant parameters using the parameter averaging technique. Formally, the SDE with time-dependent parameters is replaced by an SDE with constant parameters for each maturity, where both have the same marginal distribution. These parameters are called effective parameters, and let $\bar{\lambda}_{F,n}$ denote the effective volatility and $\bar{b}_{F,n}$ denote the effective skew, for all maturity times T_n . The dynamics of $F(t, T_n)$ is then given by

$$dF(t, T_n) = \sqrt{z_F(t)} \left(\bar{b}_{F,n} F(t, T_n) + (1 - \bar{b}_{F,n}) F(0, T_n) \right) \bar{\lambda}_{F,n} dW^{T_n}(t). \quad (15)$$

We assume to have forward processes $F(\cdot, T_1), \dots, F(\cdot, T_N)$ with expiries T_1, \dots, T_N and we further think of T_0 as “now.” Times-to-maturity for an arbitrary calendar time $t \geq 0$ are given by $x_n = T_n - t$ for $n = 0, 1, \dots, N$. For the commodity calibration, market prices for call options¹⁷ on $F(\cdot, T_n)$, with payoff $(F(T_n, T_n) - K_i)^+$, and for several strikes $K_i, i = 1, \dots, k_n$, are assumed to be available and are denoted by $C_{n,i}^{\text{mkt}}$.

Vanilla options on (??) can be calculated efficiently by the Fourier method.¹⁸ We denote the resulting model call prices by $C_{n,i}^{\text{mod}}$.

The calibration problem of the first step is then to find parameters θ_F, η_F and $\bar{b}_{F,n}, \bar{\lambda}_{F,n}$ for $n = 1, \dots, N$, such that

$$\sum_{\substack{n = 1, \dots, N, \\ i = 1, \dots, k_n}} (C_{n,i}^{\text{mod}} - C_{n,i}^{\text{mkt}})^2 \longrightarrow \min.$$

For the global parameters θ_F for mean reversion level of the variance process, and η_F for volatility of variance, this optimisation yields their final values in the calibration. The term structure and skew parameters, $\bar{b}_{F,n}$ and $\bar{\lambda}_{F,n}$ respectively, will be adapted in the next step to fit the market as closely as possible.

3.2 Step 2 – Calibrating the Volatility Term- and Skew-Structure

The calibration of the volatility term structure follows the approach given in ?, but we include the calibration of the skew structure into this procedure.

The term structure of volatility levels is assumed to be piecewise constant for a specified grid of calendar times $t_i = (0, t_1, \dots, t_{n_c})$ and times to maturity $x_j = (x_0, x_1, \dots, x_{n_f})$, which defines a $(n_c \times n_f)$ matrix of volatilities $V = (v_{i,j})_{1 \leq i \leq n_c, 1 \leq j \leq n_f}$. The relation of the d -dimensional model volatility vectors $\lambda(t, T_n)$ (for all n) and matrix V is given by

$$\| \lambda(t, T_n) \| = \sum_{\substack{0 \leq i \leq n_c - 1 \\ 0 \leq j \leq n_f - 1}} \mathbb{1}_{\{t_i \leq t < t_{i+1}, x_j \leq T_n - t < x_{j+1}\}} v_{i,j}. \quad (16)$$

The number of forward times n_f in the volatility matrix do not need to coincide with the number of

¹⁷For notational simplicity, we assume that the option expires at the same time the futures does. In most cases the option expires a few days before the futures expiry. In some cases, like for EUA carbon emission futures, the option can even expire several months before the underlying futures.

¹⁸For details, see e.g. ? (Chapter 9).

traded forwards N , and especially in regions of large forward times a rougher spacing can be chosen for x_j , since volatilities tend to flatten out with increasing forward time. To be able to price options on all of the forwards, the maturity of the longest available forward has to be smaller or equal to the latest calendar time and the longest time to maturity, $T_N \leq \min\{t_{n_c}, x_{n_f}\}$.

We refer to ? for a more detailed discussion on the setup with piecewise constant volatilities, and how to compute total variances efficiently for given calendar and forward times.

In the next section we will use the correlations to obtain a map from the volatility levels $\|\lambda_F(t, T_n)\|$ to the components of the volatility vectors $\lambda_F(t, T_n)$ that are multiplied by the d -dimensional Brownian motion in (??).

In a manner analogous to the volatility levels, we define a $(n_c \times n_f)$ -dimensional matrix $B = (b_{i,j})_{1 \leq i \leq n_c, 1 \leq j \leq n_f}$ for the matrix of piecewise constant skews. To keep notation simple, we use the same grid as for the volatility term structure. The entry $b_{i,j}$ represents the skew corresponding to forward $F(t, T_n)$ with $t_{i-1} \leq t < t_i$ and $x_{j-1} \leq T_n - t < x_j$.

The optimisation of Step 2 is defined with respect to a set of calibration criteria. The first calibration criterion measures the *quality of fit* and is as in Step 1 defined by the sum of squared differences between market and model prices,

$$q = \sum_{\substack{n=1, \dots, N, \\ i=1, \dots, k_n}} (C_{n,i}^{\text{mod}} - C_{n,i}^{\text{mkt}})^2. \quad (17)$$

Since the number of parameters is potentially quite large (and larger than the number of market prices in (??)), we follow ? and specify for the volatility two *smoothness criteria* s_λ given by,

$$s_\lambda = \eta_{\lambda,1} \sum_{j=1}^{n_f} \sum_{i=1}^{n_c-1} (v_{i+1,j} - v_{i,j})^2 + \eta_{\lambda,2} \sum_{i=1}^{n_c} \sum_{j=1}^{n_f-1} (v_{i,j+1} - v_{i,j})^2. \quad (18)$$

The first term measures departures from time-homogeneity; it demands that volatilities with different calendar times but the same time to maturity do not deviate from each other too much. The second term, the forward time smoothness, forces the volatility term structure to be smooth in time to maturity for each fixed calendar time. The larger the weight $\eta_{\lambda,1}$, the more volatility and skew become (calendar) time homogeneous. The larger the weights $\eta_{\lambda,2}$, the flatter the volatility and skew becomes in forward time direction.

We specify an analogous smoothness function for the skew term structure s_b with corresponding weights $\eta_{b,1}$ and $\eta_{b,2}$

$$s_b = \eta_{b,1} \sum_{j=1}^{n_f} \sum_{i=1}^{n_c-1} (b_{i+1,j} - b_{i,j})^2 + \eta_{b,2} \sum_{i=1}^{n_c} \sum_{j=1}^{n_f-1} (b_{i,j+1} - b_{i,j})^2. \quad (19)$$

The smile contribution to the implied volatility coming from the parameters fixed in the first calibration step of Section ?? is unaffected by the smoothness criteria here.

Remark 3.1 *Although the number of parameters is potentially quite large, the optimisation, for instance using a Levenberg–Marquardt approach for minimising the objective value $q + s_\lambda + s_b$, usually gives stable calibration results, since the smoothness criteria force the parameters to a non-parametric but structured form.*

Remark 3.2 *The reason why we have not used the effective parameters $\bar{b}_{F,n}$ and $\bar{\lambda}_{F,n}$ from the global calibration of ?? as target values in our second calibration ??, as for example proposed in ?, is that we here assumed the more complex case of calibrating to options on futures. This requires to compute the convexity correction in each optimisation step, which changes the relation between options prices and the effective volatility and skew parameters. Therefore, the advantage of using pre-computed effective parameters for calibration is not applicable in our general case.*

Remark 3.3 *As pointed out by ?, the degree of freedom is potentially quite large here, and obtaining the volatility and skew term-structure simultaneously is computationally inefficient. One could assume the skew b_F and volatility λ_F to be two almost orthogonal problems (changing the volatility has an minor impact on the skew and vice versa) and solve for them separately. Step 2 can therefore be divided into two parts. First, solve for the skew term structure by fixing the volatility parameters, e.g., to the ones obtained in the first step, $\bar{\lambda}_{F,n}$, together with η and θ . And since the time-dependent skew can be solved for explicitly as in ?, one can target the implied skews from Step 1 instead of option prices and reduce the computational time significantly by avoiding Fourier pricing within each iteration. Then, given the skew term-structure and the parameters η and θ , the volatility term structure is calibrated. However, as mentioned in Remark ??, due to the futures-forward convexity correction we cannot target the implied volatilities directly here.*

3.3 Construction of the decomposed volatilities $\lambda_F(t, T_n)$ from $\|\lambda_F(t, T_n)\|$

The method for constructing a map from $V = \|\lambda_F(t, T_n)\|$ in Equation (??) to $\lambda_F(t, T_n)$ via volatility factor decomposition or PCA described in ? remains applicable in the present setting, and we provide a brief summary for the reader's convenience. Due to the dependence of the convexity correction for futures contracts on interest rate dynamics (see Sections ?? and ??, below), the volatility factor decomposition has to be included in the calibration process when calibrating to futures and options on futures. As in ?, there is a separate factor decomposition for every calendar time step to be covered by the calibrated volatility function, thus in the following we fix this to an arbitrary t_i . Let v_i denote the i th row (corresponding to calendar time t_i) of V , written as column vector. The matrix C of commodity forward correlations is exogenously given and assumed to be constant over calendar time. For each t_i , the covariance matrix then is calculated by

$$\Sigma = (v_i v_i^\top) \odot C, \quad (20)$$

where \odot means component-wise multiplication (Hadamard product). Decomposing

$$\Sigma = R D^{1/2} (R D^{1/2})^\top$$

results in orthonormal eigenvectors of Σ in the columns of $R = (r_{j,k})_{1 \leq j, k \leq n_f}$ and the corresponding eigenvalues on the diagonal of the matrix $D = (\xi_{j,k})_{1 \leq j, k \leq n_f}$. Choosing the largest d eigenvalues, R and D can be reduced to matrices $R \in \mathbb{R}^{n_f \times d}$ and $D \in \mathbb{R}^{d \times d}$ by retaining only the corresponding d columns in R and $d \times d$ sub-matrix in D . For the factor-reduced volatilities we then have

$$v_{i,j}^2 = \sum_{k=1}^d r_{j,k}^2 \xi_k \quad (j = 1, \dots, n_f). \quad (21)$$

and the stepwise constant volatility function for the forwards is $\lambda_{ijk} = r_{j,k} \sqrt{\xi_k}$.

Remark 3.4 Under the convexity adjustments suggested in Section ??, the forward returns correlation matrix C in (??) is also the correlation matrix of futures returns.

4 Futures/Forward Relation and Convexity Correction

The calibration method in Section ?? is applicable only when forwards and options on forwards are available. This section presents a model-consistent approximate conversion of futures prices to forward prices for all relevant data for calibration, in order to apply the methods of the previous section.

We introduce the notation $G(t, T)$ for a futures price at time t with maturity T , and, as before, $F(t, T)$ will be the corresponding forward price. From no-arbitrage theory we know $F(T, T) = G(T, T)$ and that prices of plain vanilla options on forwards and futures must coincide, whenever the maturities of option, forward and futures are the same. This allows us to use the call prices of options on futures for calibration of forwards, and we only have to assure that the (virtual) forwards have the same maturities as the futures. Due to equation (??) the forward $F(\cdot, T_n)$ is an exponential martingale under the T_n -forward measure.

Denoting by \mathbb{E}^B the expectation under the spot risk-neutral measure \mathbb{Q}^B , futures follow the general relation

$$G(t, T) = \mathbb{E}_t^B [S(T)], \quad (22)$$

see ?, where $S(t)$ is the spot price, which satisfies by no-arbitrage constraints $S(t) = F(t, t) = G(t, t)$ for all t . Integrating (??) and using (??)-(??) gives,

$$\begin{aligned} F(T_n, T_n) &= F(t, T_n) + \int_t^{T_n} \sqrt{z_F(s)} \varphi_F(F(s, T_n)) \lambda_F^\top(s, T_n) dW^{T_n}(s) \\ &= F(t, T_n) + \int_t^{T_n} \sqrt{z_F(s)} \sqrt{z_L(s)} \varphi_F(F(s, T_n)) \lambda_F^\top(s, T_n) \mu_{n-1}(s) ds \\ &\quad + \int_t^{T_n} \sqrt{z_F(s)} \varphi_F(F(s, T_n)) \lambda_F^\top(s, T_n) dW^\mathbb{B}(s). \end{aligned}$$

Putting these relations together and taking the \mathbb{Q}^B -expectation for the futures, as in (??), results in

$$\begin{aligned} G(t, T_n) &= \mathbb{E}_t^B [F(T_n, T_n)] \\ &= F(t, T_n) + \mathbb{E}_t^B \left[\int_t^{T_n} \sqrt{z_F(s)} \sqrt{z_L(s)} \varphi_F(F(s, T_n)) \lambda_F^\top(s, T_n) \mu_{n-1}(s) ds \right] \\ &\quad + \mathbb{E}_t^B \left[\int_t^{T_n} \sqrt{z_F(s)} \varphi_F(F(s, T_n)) \lambda_F^\top(s, T_n) dW^\mathbb{B}(s) \right] \\ &= F(t, T_n) + \mathbb{E}_t^B \left[\int_t^{T_n} \sqrt{z_F(s)} \sqrt{z_L(s)} \varphi_F(F(s, T_n)) \lambda_F^\top(s, T_n) \mu_{n-1}(s) ds \right] \quad (23) \\ &= F(t, T_n) + D(t, T_n), \quad (24) \end{aligned}$$

the third equality follows since the last expectation is that of an Itô integral, further letting $D(t, T_n)$ denote the convexity correction.

From here we consider two alternative ways to proceed.

4.1 Approximation 1 - Freeze all risk factors

The first crude approximation would be to freeze all random variables by setting $\bar{L}_i = L(0, T_i)$, $\bar{F}_j = F(0, T_j)$ and $\bar{z}_L = z_L(t) = 1$, $\bar{z}_F = z_F(t) = 1$ such that the convexity correction $D(t, T_n)$ can be approximated by,

$$D(t, T_n) \approx \int_t^{T_n} \varphi_F(\bar{F}_n) \sum_{j=q(s)}^{n-1} \frac{\tau_j \varphi_L(\bar{L}_j)}{1 + \tau_j \bar{L}_j} \lambda_F^\top(s, T_n) \lambda_L(s, T_j) ds. \quad (25)$$

? demonstrate how the integrals in Equation (??) can be computed when the volatility functions are piecewise constant. Since \bar{F}_n is unknown, it is necessary to solve equation (??) for $F(t, T_n)$ using the approximation (??) above. This is straightforward for displaced diffusions with $\varphi_F(\bar{F}_n)$ piecewise constant subject to the integration variable s .

4.2 Approximation 2 - Freeze LIBORs and Commodity forwards

For the second approximation we choose to freeze the LIBORs and commodity forwards but keep the volatility stochastic. Set $\bar{L}_i = L(0, T_i)$ and $\bar{F}_j = F(0, T_j)$ and use conditioning for the stochastic volatility processes. If the expectation and the integration can be interchanged it follows from the independence of the stochastic volatility processes that the convexity correction $D(t, T_n)$ can be approximated by,

$$D(t, T_n) \approx \int_t^{T_n} \varphi_F(\bar{F}_n) \sum_{j=q(s)}^{n-1} \frac{\tau_j \varphi_L(\bar{L}_j)}{1 + \tau_j \bar{L}_j} \lambda_F^\top(s, T_n) \lambda_L(s, T_j) \mathbb{E}_t^B[\sqrt{z_F(s)}] \mathbb{E}_t^B[\sqrt{z_L(s)}] ds. \quad (26)$$

? show that the first moments of the squared volatility process can be represented as

$$\mathbb{E}_t^B \left[\sqrt{z(T)} \right] = \sqrt{\frac{2e^{-\theta(T-t)}}{n(t, T)}} e^{-\frac{n(t, T)}{2}} \sum_{j=0}^{\infty} \frac{(n(t, T)/2)^j \Gamma(d/2 + j + 1/2)}{j! \Gamma(d/2 + j)}, \quad (27)$$

where $\Gamma(x)$ is the Gamma function, and the parameters n and d are given by

$$n(t, T) = \frac{4\theta e^{-\theta(T-t)}}{\eta^2(1 - e^{-\theta(T-t)})}, \quad d = \frac{4\theta z_0}{\eta^2}. \quad (28)$$

We illustrate the accuracy of the two approximations in Figure ???. Clearly the stochastic volatility factors have too much impact on the futures–forward convexity to be neglected. As demonstrated, Approximation 2 performs exceptionally well and produces an error less than 2 basis points, which is more than acceptable within a calibration routine.

5 Merging Interest Rate and Commodity Calibrations

The calibrations for the interest rate market and the commodity market have so far been considered separately, with exception of the convexity correction in Section ???. This section focuses on linking the calibrations of the two asset classes in order to get a joint commodity and LMM calibration. The relation between these two asset classes is determined by their cross-correlation matrix, which is assumed to be constant over time. Before addressing the problem of merging the calibrations in Section ??, we briefly discuss how to generate the cross-correlation matrix from historical data in Section ??.




Figure 1: Left: Futures prices from Monte Carlo simulation, Approximation 1 and Approximation 2. Right: Relative error in basis points (bps). With $\lambda_L = 0.20$, $\lambda_F = 0.30$, $b_L = 0.9$, $b_F = 0.8$, $\theta_L = \theta_F = 1$, $\eta_L = 0.5$, $\eta_F = 0.4$, $\text{Corr}(dL(t, T_j), dL(t, T_k)) = \text{Corr}(dF(t, T_j), dF(t, T_k)) = 0.9$.

5.1 The Exogenous Cross-Correlations

The cross-correlations between commodity forwards and interest rate forwards are much less pronounced than the correlations between the futures and forwards, respectively, within the asset classes. This makes the estimation for cross-correlations from historical data less stable than for the correlations within the asset classes. For example, the structure of the cross-correlation matrix between Brent Crude Oil futures and USD interest rate forwards in Figure ?? of Section ?? can hardly be explained by obvious rationales. Therefore, one might decide to specify exogenously a very simple cross-correlation structure in practice. The simplest case would be to assign a single value to all cross-correlation in the matrix R_{LF}^{exogen} . If this seems too crude, a linear relationship can be fitted by regressing the empirical cross-correlations for the forward interest rates and commodity futures on a 2-dimensional plane.

5.2 The Cross-Calibration

The quadratic (cross-)covariation process for commodity forwards with maturity T_m and forward interest rates with settlement T_n is given by the dynamics

$$dL(t, T_n)dF(t, T_m) = \left(\sqrt{z_L(t)}\varphi_L(L(t, T_n))\lambda_L(t, T_n) \right)^\top \left(\sqrt{z_F(t)}\varphi_F(F(t, T_m))\lambda_F(t, T_m) \right) dt,$$

where we have used the fact that

$$dW^{T_{n+1}}(t) = dW^{T_m}(t) + \sqrt{z_L(t)}(\mu_n(t) - \mu_{m-1}(t))dt,$$

and $dt dW^{T_m} = 0$. Similar to the pure interest forward rate correlations, the cross-correlations of the increments are given by

$$\text{Corr}(dL(t, T_n), dF(t, T_m)) = \frac{\lambda_L(t, T_n) \lambda_F^\top(t, T_m)}{\|\lambda_L(t, T_n)\| \|\lambda_F(t, T_m)\|}, \quad (29)$$

since only the factorised λ_F and λ_L are (column-)vectors, and all other values in the equation above are scalars. Abbreviate

$$\Lambda_L = (\lambda_L^\top(t, T_n))_{1 \leq n \leq N} = (\lambda_{L;n,k}(t))_{1 \leq n \leq N, 1 \leq k \leq d}, \quad (30)$$

$$\Lambda_F = (\lambda_F^\top(t, T_m))_{1 \leq m \leq M} = (\lambda_{F;m,k}(t))_{1 \leq m \leq M, 1 \leq k \leq d}, \quad (31)$$

$$U_{LF} = (\| \lambda_L(t, T_n) \| \| \lambda_F(t, T_m) \|)_{1 \leq n \leq N, 1 \leq m \leq M}, \quad (32)$$

where m and n are indices for different expiries and forward (settlement) times, respectively, and k is the index relating to the stochastic factor. We skip the calendar time dependence, since it could be considered as fixed throughout the rest of this subsection. Then, the model cross-correlation in (??) can be written as

$$R_{LF}^{\text{model}} = (\Lambda_L \Lambda_F^\top) \circ U_{LF}, \quad (33)$$

where \circ denotes the element-wise division. Merging the individual interest rate and commodity calibrations in order to get a joint calibration means to modify the calibrated matrices Λ_L and Λ_F such that equation (??) is matched (maybe only as closely as possible) for an exogenously given correlation matrix R_{LF}^{exogen} ,

$$R_{LF}^{\text{exogen}} \approx (\Lambda_L \Lambda_F^\top) \circ U_{LF} = R_{LF}^{\text{model}}.$$

Note that the skew parameter, being part of $\varphi_L(\cdot)$ and $\varphi_C(\cdot)$, as well as the smile parameter, being part of the specification of the square root processes $z_F(t)$ and $z_L(t)$, need not to be modified, as they do not determine the cross-correlations, as equation (??) demonstrates.

Clearly, there are many possible approaches to improve the relation $R_{LF}^{\text{exogen}} \approx R_{LF}^{\text{model}}$. The approach we propose in this section will be guided by the following. Firstly, the quality of the individual calibrations to market instruments relating to one asset class only – interest rates or commodities – is more important than the cross-correlation fit. This is because these market instruments are much more liquidly traded, and therefore provide reliable information. Secondly, the adjustment step needs to be carried out in an efficient way, having a common fit criterion for the approximation.

An approach that combines these considerations is to find a rotation matrix Q , that is applied to one of the calibrated matrices, say Λ_F , with the objective (following ? in our choice of metric) to minimise the Frobenius distance between the model and the exogenous covariance matrices. Formally, a matrix Q is sought, such that

$$\| R_{LF}^{\text{exogen}} \circ U_{LF} - \Lambda_L (\Lambda_F Q)^\top \|_F \longrightarrow \min, \quad \text{subject to } QQ^\top = I_d, \quad (34)$$

with I_d the $(d \times d)$ identity matrix and $\| \cdot \|_F$ the Frobenius norm. From the theory of normal distributions it follows that an orthonormal rotation changes the cross-covariances, but not the covariance matrices $\Sigma_L = \Lambda_L \Lambda_L^\top$ and $\Sigma_F = \Lambda_F \Lambda_F^\top$, hence, the individual calibrations remain unaffected.

For the problem of (??) it is quite reasonable to assume $d \leq N$, i.e., the number of stochastic factors has to be equal or less than the number of forward interest rates.¹⁹

The problem of finding the Q satisfying (??) is similar to the so-called ‘‘orthonormal Procrustes’’ problem: For given matrices $A, B \in \mathbb{R}^{m \times p}$, find an orthonormal matrix $Q \in \mathbb{R}^{p \times p}$ that minimises the distance between A and B ,

$$\| A - BQ \|_F \longrightarrow \min, \quad \text{such that } QQ^\top = I_d, \quad (35)$$

¹⁹Note that the roles of Λ_L and Λ_F can be interchanged, and the sufficient assumption actually is $d \leq \max\{M, N\}$. Furnishing the model with more stochastic factors than $\max\{M, N\}$ contributes only spurious complexity to the model. From a practical point of view, the aim is to keep the number of stochastic factors small.

In ?, Section 12.4.1 it is shown that the solution is given by $Q = UV^\top$, where U and V result from the singular value decomposition (SVD) of $B^\top A$, i.e. $B^\top A = UDV^\top$.

Unfortunately, (??) is more complicated than the Procrustes problem in (??), because in our case Q comes under a transposition. As stated in ?, Section 8.3.3, “there seems to be no algebraic solution to the problem” (??), to which they refer as a “scaled orthonormal Procrustes” problem. However, the authors discuss a numerical solution, based on an algorithm of ?, that works well in our case, as we will demonstrate in Section ?? below.

Algorithm 1 Orthogonal scaled procrustes

```

1: procedure OSProcrustes( $Q_0, X_1, S, X_2, \text{TOL}, \text{MAXITER}$ )
2:    $r \leftarrow \|X_1\|_F$ 
3:    $\text{iter} \leftarrow 0$ 
4:    $\delta \leftarrow 2 \cdot \text{TOL}$ 
5:    $Q \leftarrow Q_0$ 
6:   while  $\delta > \text{TOL}$  and  $\text{iter} < \text{MAXITER}$  do
7:      $Z \leftarrow S(X_2^\top X_1 + S^\top Q^\top (rI - X_1^\top X_1))$ 
8:     Compute singular value decomposition  $Z = UDV^\top$ 
9:      $Q_{\text{new}} \leftarrow VU^\top$ 
10:     $\delta \leftarrow | \|X_1 Q S - X_2\|_F - \|X_1 Q_{\text{new}} S - X_2\|_F |$ 
11:     $\text{iter} \leftarrow \text{iter} + 1$ 
12:     $Q \leftarrow Q_{\text{new}}$ 
13:   end while
14:   return  $Q$ 
15: end procedure

```

Proposition 5.1 *The problem (??) is equivalent to the problem of finding a T satisfying*

$$\|D_L T D_F^\top - U_L^\top \Sigma_{LF} U_F\|_F \longrightarrow \min, \quad \text{such that } T T^\top = I_d, \quad (36)$$

where the matrices used are from the singular value decompositions²⁰

$$\begin{aligned} \Lambda_L &= U_L D_L V_L^\top & U_L &\in \mathbb{R}^{N \times N}, D_L \in \mathbb{R}^{N \times d}, V_L \in \mathbb{R}^{d \times d} \\ \Lambda_F &= U_F D_F V_F^\top & U_F &\in \mathbb{R}^{M \times M}, D_F \in \mathbb{R}^{M \times d}, V_F \in \mathbb{R}^{d \times d}, \end{aligned}$$

and

$$\Sigma_{LF} := R_{LF}^{\text{exogen}} \odot U_{LF}.$$

The orthonormal solution Q of (??) is given by

$$Q = (V_L T V_F^\top)^\top \quad (37)$$

Proof: Substituting the SVD yields

$$\begin{aligned} \|\Lambda_L (\Lambda_F Q)^\top - \Sigma_{LF}\|_F &= \|U_L D_L V_L^\top Q^\top V_F D_F^\top U_F^\top - \Sigma_{LF}\|_F \\ &= \|U_L D_L T D_F^\top U_F^\top - \Sigma_{LF}\|_F && \text{abbreviating } T := V_L^\top Q^\top V_F \\ &= \|D_L T D_F^\top U_F^\top - U_L^\top \Sigma_{LF}\|_F && U_L \text{ is orthonormal} \end{aligned}$$

²⁰Note that $\text{rank}(\Lambda_L) = d \leq N$ and $\text{rank}(\Lambda_F) = d \leq M$

$$= \| D_L T D_F^\top - U_L^\top \Sigma_{LF} U_F \|_F \quad U_F \text{ is orthonormal}$$

□

The solution for T in (??) can be numerically approximated by Algorithm 8.1 in ?. See also ? for a nice description of the algorithm with an application in the context of portfolio risk management. For completeness we reproduce the algorithm here (see Algorithm ??).

The objective is to find a Q that minimises

$$\| X_1 Q S - X_2 \|_F, \quad \text{subject to } Q Q^\top = I,$$

where S has entries on its diagonal only.

The stopping of the numerical procedure is controlled by the maximal number of iterations (MAXITER) and the tolerance in change of the Frobenius norm (TOL) in the objective above.

As initial guess the trivial transformation $Q_0 = I$ can be used.

5.3 Cross-Calibration Computational Complexity

We conclude the section with a note on the computational effort required by the proposed algorithm. From a theoretical point of view, one would expect the Procrustes method to be faster than the Levenberg–Marquardt optimisation as used in ?, because the number of parameters to be optimised is given by $n_{\text{opt}} = d(d-1)/2$, which means the dimension of the optimisation problem grows quadratically in the number of stochastic factors d .

The Levenberg–Marquardt optimisation requires at least $O(n^3)$ or $O(mn^2)$ operations in each iteration step (see ?, Section 10.3), where n is the number of optimisation parameters, and m is the number of residuals in the objective function. In our case we get $O(n_{\text{opt}}^3) = O(d^6)$ operations per iteration for the unconstrained optimisation problem. The constraint $Q Q^\top = I$ is non-linear, hence cannot be represented by simple bounds for the parameters, though it is possible to re-parametrise the problem expressing the orthonormal transformation Q by n_{opt} successive rotations, where the angles are the new parameters in an unconstrained Levenberg–Marquardt optimisation, see ? for a detailed description.

In contrast, the Procrustes problem requires only $O(d^3)$ operations per iteration for the matrix multiplications and the singular value decomposition.

We applied the Procrustes method to several randomly chosen 100×100 cross-correlation matrices, and it converged in all cases within a tolerance of less than 10^{-5} after 200 seconds and about 120,000 iterations. Our results are in line with the numerical experiments in ?, in which the Procrustes problem turned out to be much faster than the least-squares optimisation using Levenberg–Marquardt. That paper reports that the latter method failed to optimise a 65×65 transformation even “after several hours,” whereas the Procrustes method found an approximate solution in “about 5 minutes.”

The transformation problems we encounter in the model calibration might be of much smaller size than dimension 100, typically somewhere between 4 – 20. However, in the case of calibrating the model to commodity futures, the computation of the optimal transformation matrix has to be iterated, which led us to prefer the much more efficient Procrustes algorithm over the “brute force” Levenberg–Marquardt optimisation used in ?.

6 Applying the calibration to market data

6.1 Summary of the Calibration Procedure

The following scheme summarises the steps of the calibration procedure, bringing together the steps discussed in the previous sections and assuming the calibration of the (domestic) interest rate LMM has already been carried out.²¹

I. Preliminary calculations applied to the LMM calibration outcome.

1. For each calendar time t_i ($1 \leq i \leq n_L$):
 - (a) Compute the covariance matrix Σ_i^L as in (??).
 - (b) Decompose Σ_i^L into λ_i^L using PCA as described in Section ??.

II. Hybrid Calibration.

1. Calibrate on all commodity data to obtain the global θ_F , and η_F .
2. Calibrate on all data to obtain the skew term-structure B_F , see Remark ??.
3. For each calendar time t_i ($1 \leq i \leq n_F$):
 - (a) Compute the covariance matrix Σ_i^F as in (??).
 - (b) Decompose Σ_i^F into U_i^F using PCA as described in Section ??.
 - (c) Compute the rotation matrix Q using Algorithm ??
4. Compute forward prices from futures prices using (??).
5. Compute model options prices on forwards.
6. Continue with Step ?? until a sufficiently close fit to commodity market instruments and assumed cross-correlations is reached.

The dynamics of the thus calibrated hybrid Commodity LMM with SLV can then be written as follows. Let W be a d -dimensional Brownian motion and denote by $\lambda_{i,j}^L$ and $\lambda_{i,j}^F$ the d -dimensional vectors in Λ_L and Λ_F of volatilities for calendar times $t \in [t_{i-1}, t_i)$ and times to maturity $x \in [x_{j-1}, x_j)$. Then, the dynamics of the forward interest rates $L(t, T)$ and the dynamics of the commodity forwards $F(t, T)$ are

$$\begin{aligned} dL(t, T_k) &= \sqrt{z_L(t)} \varphi_L(L(t, T)) \lambda_{i,j}^L dW^{T_{k+1}}(t), \\ dF(t, T_k) &= \sqrt{z_F(t)} \varphi_F(F(t, T)) \lambda_{i,j}^F dW^{T_k}(t), \end{aligned}$$

for all maturity times satisfying $T_k = t_i + x_j$ (for some $1 \leq i \leq n_c$ and $1 \leq j \leq m_f$ or $1 \leq j \leq n_f$, respectively) and all calendar times $t_{i-1} \leq t < t_i$.

²¹Note that we must iterate over repeated calibration to the commodity market and to the cross-correlations, as the conversion of commodity futures into forwards depends on cross-correlations.




Figure 2: Left: Historical Brent Crude Oil 1M, 1Y, 2Y and 4Y Futures contracts between January 2014 and March 2015. Historical Brent Crude Oil 1M, 1Y, 2Y and 4Y at-the-money implied volatilities between January 2014 and March 2015

6.2 Calibration Setup

The Brent Crude Oil is selected as commodity, and the USD forward rate as interest rate.

The calibration date is January 13th, 2015. Figure ?? illustrates the historical Brent crude oil futures and implied volatilities for a selected number of contracts. As one can observe, the futures price (implied volatility) are close to their lowest (highest) levels on the selected calibration date; thus we are using a somewhat “stressed” market scenario to best illustrate the effectiveness of the calibration.

We assume that both the instantaneous volatility functions and skew functions are piecewise constant. The LIBOR forward rates and Brent crude oil futures correlation matrix as used for calibration is historically estimated from the time series of LIBOR forward rates and Brent futures prices covering 3 months before the calibration date.

For the correlations and cross-correlations we use a parametric form given for the correlation of the underlying (LIBOR, oil futures) with fixings/maturities T_k and T_j , which is given by²²

$$\text{Corr}(dL(t, T_j), dL(t, T_k)) = q(T_j - t, T_k - t),$$

where

$$q(x, y) = \rho_\infty + (1 - \rho_\infty) \exp(-a(\min(x, y))|x - y|), \quad (38)$$

$$a(z) = a_\infty + (a_0 - a_\infty) e^{-\kappa z}, \quad (39)$$

and subject to $0 \leq \rho_\infty \leq 1$, $a_0, a_\infty, \kappa \geq 0$. The two calibrated correlation matrices for the LIBOR forward rates and the Brent crude oil futures are illustrated in Figure ??.

6.3 The Interest Rate Calibration

The calibration of the SLV–LMM for the USD interest rate market was conducted based on ? by first performing a pre–calibration to fit the SLV–LMM, with constant parameters for each tenor and expiry, to the swaption cube. Figure ?? illustrates the obtained effective at–the–money volatility $\bar{\lambda}_{n,m}^{\text{mkt}}$ and effective

²²The origins of this parametric form can be traced back to ?; see Chapter 22 of ? for a detailed discussion of the rationale behind a parametric form of this type. Our notation follows ?.

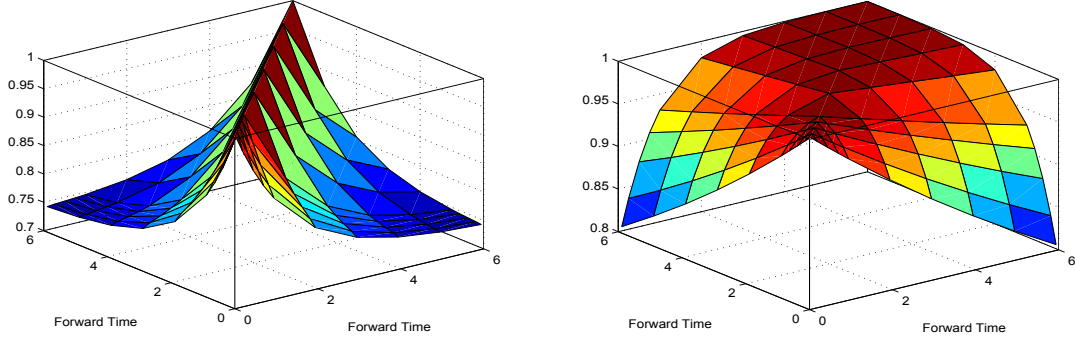


Figure 3: Left: The historically estimated LIBOR forward rate correlation matrix. Right: The historically estimated Brent crude oil futures correlation matrix.




Figure 4: Left: The calibrated implied volatility term structure for the swaption cube. Right: The calibrated implied skew term structure for the swaption cube.

skew $\bar{b}_{n,m}^{\text{mkt}}$ for each tenor T_n and expiry T_m .

These obtained model parameters then serve as target values in the main calibration when calibrating the volatility and skew term-structure as described in ?. Figure ?? shows the resulting volatility and skew term-structure obtained by the main calibration. Figure ?? shows the quality of fit of the effective at-the-money volatility $\bar{\lambda}_{n,m}^{\text{mkt}}$ and effective skew $\bar{b}_{n,m}^{\text{mkt}}$. The overall model fit to the implied volatility is very good as demonstrated by Figure ?? for selected expiries and tenors. Moreover, we chose $d_L = 4$ factors, which typically explain about 99.99% of the overall variance.

6.4 The Cross-Calibration

The calibration of the Brent Crude Oil futures was achieved by the method described in Section ?. The market instruments are futures and options on futures. Figure ?? shows the calibrated volatility and skew surface. Calendar and forward times go out to 3 years, and although on the exchange futures with expiries in every month are traded, we chose the calendar and forward time vectors to be unequally spaced (while still calibrating to all traded instruments), $[0, 1M, 3M, 6M, 9M, 1Y, 2Y, 3Y]$. This setup speeds up the calibration without losing much structure in the volatility surface, since the market views futures with long maturity to have almost flat volatility. For weighting in the calibration objective function we have chosen $\eta_{\lambda,1} = \eta_{b,1} = 0.1$ (time homogeneity, i.e. smoothness in calendar time direction),

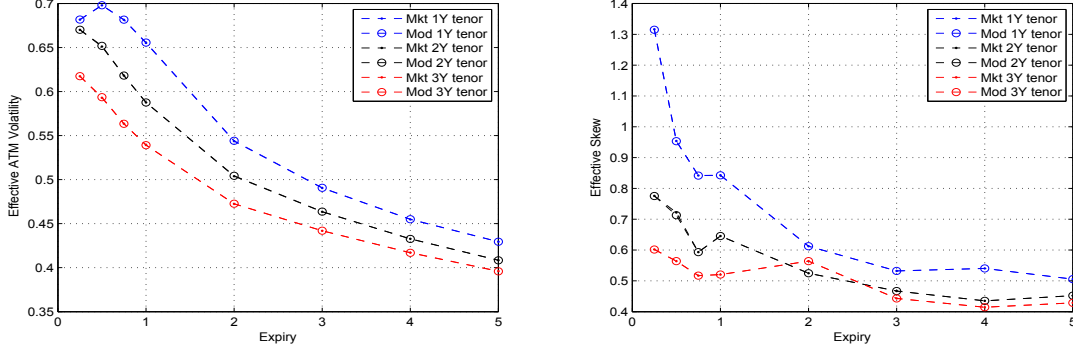


Figure 5: Left: The calibrated effective volatility $\bar{\lambda}_{n,m}^{\text{mkt}}$ for the swaption cube. Right: The calibrated effective skew $\bar{b}_{n,m}^{\text{mkt}}$ for the swaption cube

Number of Factors	1	2	3	4	5	...	19
Percentage of Overall Variance	98.14	99.84	99.96	99.98	99.99	...	100

Table 1: The percentage of overall variance that can be generated by the first i factors, obtained from PCA of the commodity forward covariance matrix for the first calendar time t_1 .

$\eta_{\lambda,2} = \eta_{b,2} = 0.01$ (smoothness in forward time direction).

The exogenously given target cross-correlation matrix (calculated from the historical time series covering 3 months before the calibration time) and the estimated cross-correlation from historical futures returns and the cross-correlation matrix for the calendar time that fitted worst are illustrated in Figure ??.

Finally, we link both separately calibrated volatility matrices to one set of stochastic factors. Table ?? shows how much of the overall variance, i.e. of the sum of variances over all factors, can be explained by the leading factors, when the factors are sorted according to decreasing contribution to total variance of the commodity forwards. The first two factors already account for more than 99% of the overall variance. Note that if it is necessary to interpolate the forward interest rate volatility matrix in order to match the calendar times of the commodity volatility matrix, the forward rate covariance matrix will change and, hence, eigenvalue decompositions of the calendar time adjusted covariance matrices yield different results than an eigenvalue decompositions of the original covariance matrices as used for calibration. However, these differences should not be substantial as long as the calibrated volatility matrix is sufficiently smooth in calendar time.

The model fit to the commodity implied volatility is illustrated in Figure ??.

6.5 Pricing Performance and Accuracy

The calibration routine was written in MATLAB and timing the execution was performed on a Intel Core i7-2600K 3.40GHz.

The interest rate volatility and skew term-structure calibration takes roughly 15 seconds, and the cross-correlation calibration (including fitting the commodity volatility) takes just roughly 10 seconds, where the cross-calibration takes about roughly 2 seconds. This would be fast enough for production purposes

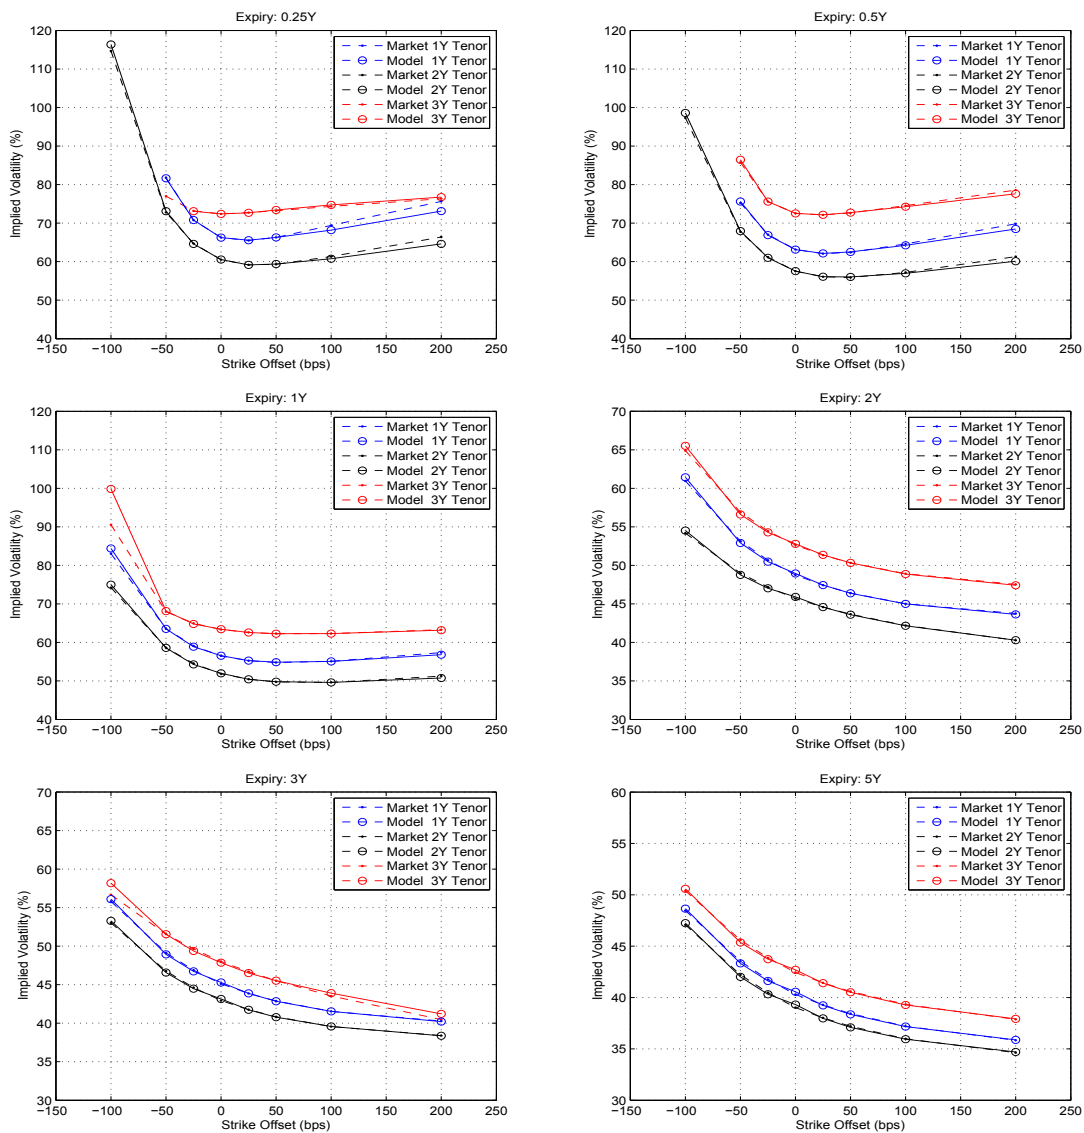


Figure 6: The calibrated interest smile with maturity 0.25, 0.5, 1, 2, 3, 5 years and tenors 1, 2, 3 years.

in most instances, though the fact that this prototype implementation runs in a serial, interpreted language environment leaves scope for substantial speed-ups through through parallelisation and the use of a compiled programming language (such as C++).

We conclude this section by studying the accuracy of the pricing derived in this paper and used within the calibration routine. We focus only on the overall hybrid calibration since ? already performed extensive numerical tests on the accuracy of the parameter averaging approximation for the interest rate calibration itself, which is input to our routine. To summarize ?, in terms of (interest rate) Black volatility, the author reported a difference between the Monte Carlo simulation and the semi-analytic approximations for interest rate swaptions, which was within 15 basis points for at-the-money swaptions, and about 20 basis points for in-the-money and out-of-the-money swaptions, for various tenors, expiries and strikes.

For the overall hybrid calibration, we let *Market* refer to the market implied commodity volatilities, *MC* to the implied commodity volatilities from the Monte Carlo simulation of the calibrated model




Figure 7: Left: The calibrated commodity volatility surface. Right: The calibrated commodity skew surface.




Figure 8: Left: The target cross-correlation matrix (green) estimated from historical futures returns, the rotated cross correlation (blue) and the un-rotated cross correlation (grey) matrix, for the first calendar time. Right: The absolute differences between target and rotated cross-correlations.

and *Approx* to the implied commodity volatilities obtained from the semi-analytic approximations used within the calibration.

Table ?? and Figure ?? provide full a comparison of (correlated hybrid) commodity implied Black volatilities, for the market volatilities, Monte Carlo simulated commodity implied Black volatilities and those obtained from the semi-analytic approximations (used within the calibration), for different expiries and strike offsets. For each expiry 0.16, 0.24, 0.32, 0.58, 0.82, 1.79, 2.28 and 3.79 years,²³ we price the options written on commodity futures with strikes offsets -15%, -7.5%, 0%, +7.5% and +15% (the +15% implied volatility was not available for the 0.16Y option). The model parameters are the ones obtained from the calibration in this section. For the Monte Carlo simulation we use 1,000,000 paths with monthly time steps.

The semi-analytical approximation demonstrates similar accuracy as in ?, this can be seen by the reported *MC-Market* error. All reported errors are below 25 bps of Black volatilities. For options with expiry longer than 1Y, the approximation error is about or below 10 bps. As in ? the approximation quality seems to improve for options with longer maturity. This is an effect due to the Monte Carlo discretisation

²³These are a selection of current expiries, i.e., the time between the calibration date 13 January 2015 and the various expiry dates, for the standardised commodity options.

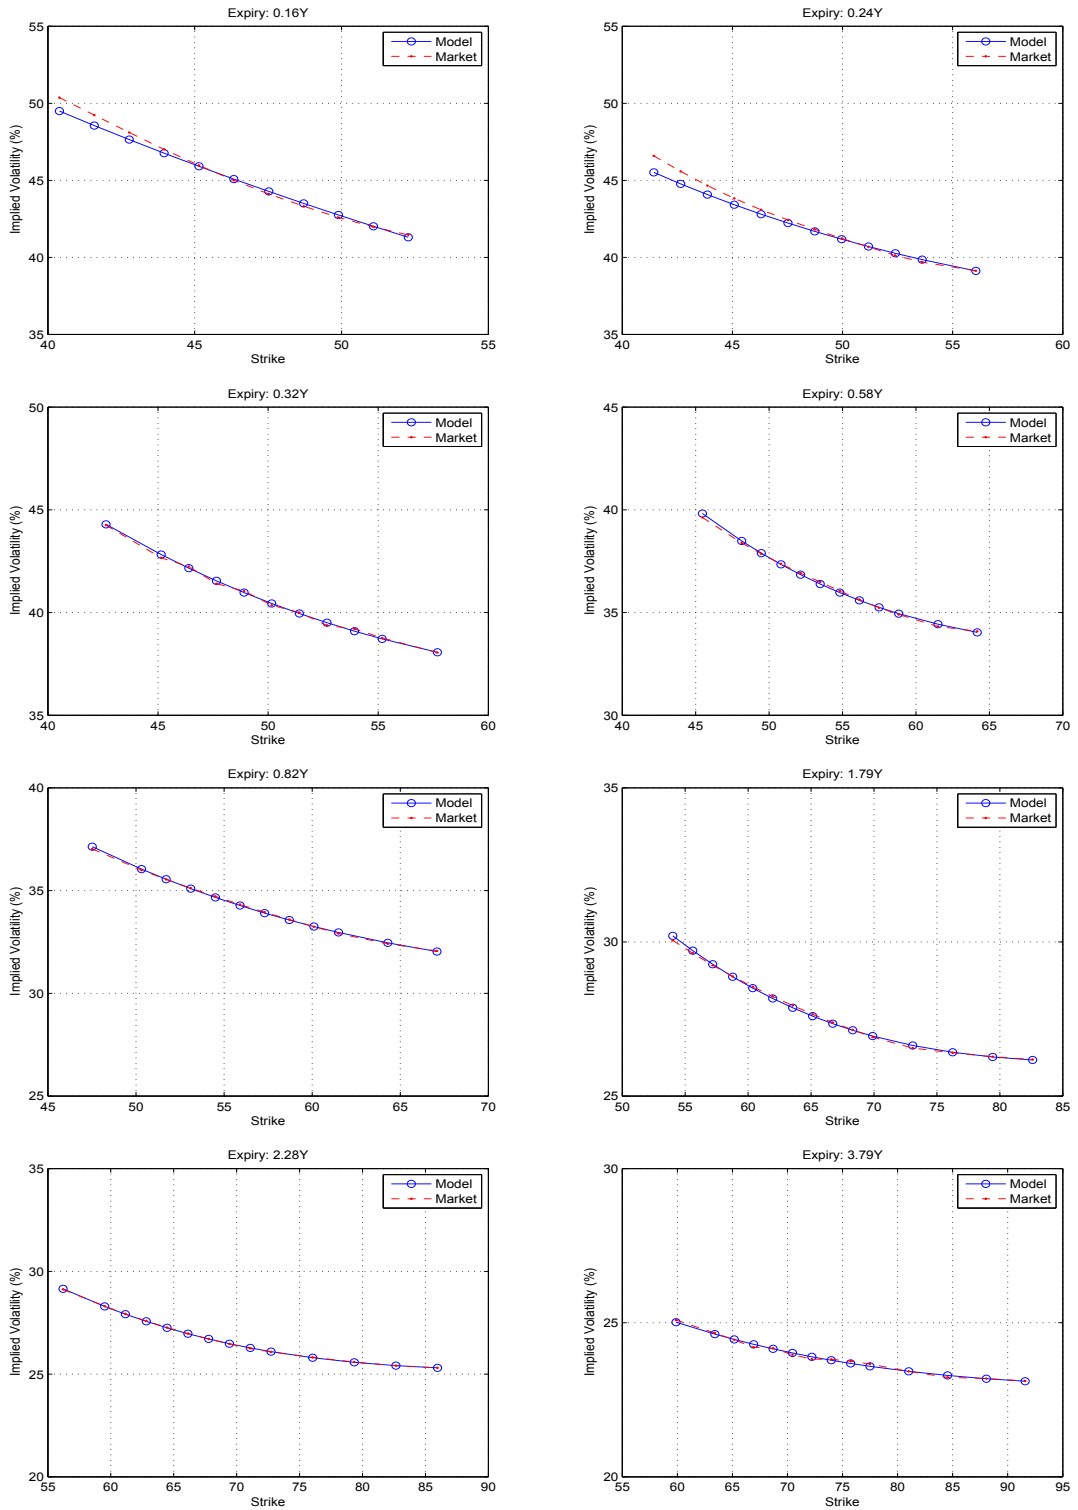


Figure 9: The calibrated commodity smile.

error rather than an approximation error, due to the low number of time steps used within the simulation. We noticed that by increasing the number of steps per year, this error decreased.

MC-Market, illustrates the differences in Black volatility between the given market implied volatilities and the ones obtained by the Monte Carlo simulation. This numerical test is performed to demonstrate that given the full calibrated model, which was calibrated using the semi-analytical approximation, we are able to regenerate the market implied volatilities without loss of essential information. The biggest differences are for the short dated options where we observed a difference of 40 bps for the 0.16Y ATM option. As the expiry increases, this error, for all strikes, decreases, where we obtain an error of 20-30 bps for the 0.24Y and 0.32Y expiries, an error of 10-20 bps for the 0.58Y and 0.82Y expiries. For the longer expiries this error decreases to about or below 10 bps.

7 Conclusion

As the market data example in the previous section demonstrates, the LMM approach to term structure modelling remains is one of the most flexible for good calibration of the model to market data, even when it is extended to allow for market quotes across multiple strikes (volatility “skews” and “smiles”) and the integration of (and correlation between) multiple sources of risk — commodity and interest rate risk in our present example. This is due to actual market observables (in particular forward LIBORs) being modelled directly and model prices for calibration instruments (e.g. caps/floors, swaptions, commodity futures and options) being available either in exact or accurate approximate closed form.

The dynamics of all market variables can be expressed in terms of the same, vector-valued Brownian motion and correlation between market variables is obtained via the sum products of the respective vector-valued volatilities. As a consequence, the calibration across multiple sources of risk can be broken down into stages, simplifying the high-dimensional optimisation problems to be solved at each stage. The interest rate market can be calibrated separately using well established procedures. We chose to base our interest rate volatility calibration on the robust and widely used method of χ^2 , but this could easily be replaced by a different method without impacting the remainder of our calibration approach. The model is fitted to an exogenously given correlation structure (typically estimated statistically from historical data) between forward interest rates and commodity prices (cross-correlation) — without impacting the interest rate volatility calibration already obtained — via a modified version of the “orthonormal Procrustes” problem in linear algebra, for which an efficient numerical solution exists.

Finally, it is worth noting that a model which is calibrated to the term structures of commodity futures and options will implicitly reflect any seasonality and/or mean reversion of commodity prices. If seasonality is present and priced by the market, this information will be contained in the initial term structure of commodity futures prices, and integrated into the model by the fact that it fits the initial observed term structure by construction. If mean reversion is present and priced by the market, this information will be contained in the term structure of commodity volatilities, or, more completely, the commodity option implied volatility surface, to which the model is calibrated.²⁴

BIBLIOGRAPHY

Andersen, L. and Brotherton-Ratcliffe, R. (2005), ‘Extended Libor Market Models with Stochastic Volatility’, *Journal of Computational Finance* **9**, 1–40.

²⁴Specifically, mean reversion would manifest itself in the market as a downward sloping term structure of commodity option implied volatilities.

Expiry	Offset	Market	Approx	MC	Approx-Market	MC-Approx	MC-Market
0.16y	-15%	50.37	50.11	50.36	-0.26	0.25	-0.01
	-7.5%	47.01	47.05	47.30	0.04	0.25	0.29
	0%	44.12	44.27	44.52	0.15	0.25	0.40
	7.5%	41.98	41.72	41.97	-0.26	0.25	-0.01
0.24y	-15%	46.59	46.10	46.35	-0.49	0.25	-0.24
	-7.5%	43.83	43.93	44.18	0.10	0.25	0.35
	0%	41.82	41.97	42.21	0.15	0.24	0.39
	7.5%	40.13	40.20	40.42	0.07	0.22	0.29
	15%	39.13	38.57	38.77	-0.56	0.20	-0.35
0.32y	-15%	44.25	44.22	44.42	-0.04	0.20	0.16
	-7.5%	42.24	42.11	42.32	-0.12	0.21	0.09
	0%	40.35	40.42	40.64	0.07	0.22	0.29
	7.5%	39.21	39.10	39.32	-0.11	0.22	0.12
	15%	38.05	38.09	38.32	0.04	0.23	0.27
0.58y	-15%	39.63	39.60	39.78	-0.03	0.18	0.15
	-7.5%	37.86	37.88	38.04	0.02	0.16	0.19
	0%	36.48	36.45	36.61	-0.04	0.16	0.13
	7.5%	35.23	35.27	35.43	0.04	0.16	0.20
	15%	34.32	34.30	34.46	-0.02	0.16	0.14
0.82y	-15%	37.03	37.02	37.18	0.00	0.16	0.15
	-7.5%	35.54	35.55	35.71	0.01	0.16	0.17
	0%	34.32	34.31	34.47	-0.02	0.16	0.15
	7.5%	33.24	33.26	33.42	0.02	0.16	0.18
	15%	32.41	32.38	32.53	-0.03	0.15	0.11
1.79y	-15%	30.05	30.05	30.17	0.00	0.12	0.12
	-7.5%	28.88	28.88	28.99	0.00	0.11	0.11
	0%	27.93	27.93	28.03	-0.01	0.10	0.09
	7.5%	27.15	27.15	27.26	0.01	0.10	0.11
	15%	26.56	26.54	26.64	-0.02	0.10	0.08
2.28y	-15%	29.12	29.15	29.23	0.03	0.08	0.11
	-7.5%	27.93	27.91	27.99	-0.01	0.08	0.06
	0%	26.97	26.97	27.04	0.00	0.07	0.07
	7.5%	26.25	26.27	26.34	0.02	0.07	0.09
	15%	25.82	25.79	25.86	-0.03	0.07	0.04
3.79y	-15%	25.08	25.07	25.10	-0.02	0.04	0.02
	-7.5%	24.44	24.45	24.48	0.01	0.03	0.04
	0%	23.96	24.00	24.03	0.03	0.03	0.06
	7.5%	23.75	23.68	23.71	-0.07	0.03	-0.04
	15%	23.42	23.49	23.51	0.07	0.03	0.09

Table 2: Pricing results for implied Black commodity volatilities, for commodity options with expiries 0.16, 0.24, 0.32, 0.58, 0.82, 1.79, 2.28 and 3.79 years, and strikes offsets -15%, -7.5%, 0%, +7.5% and +15% . The values are reported in percent (%).

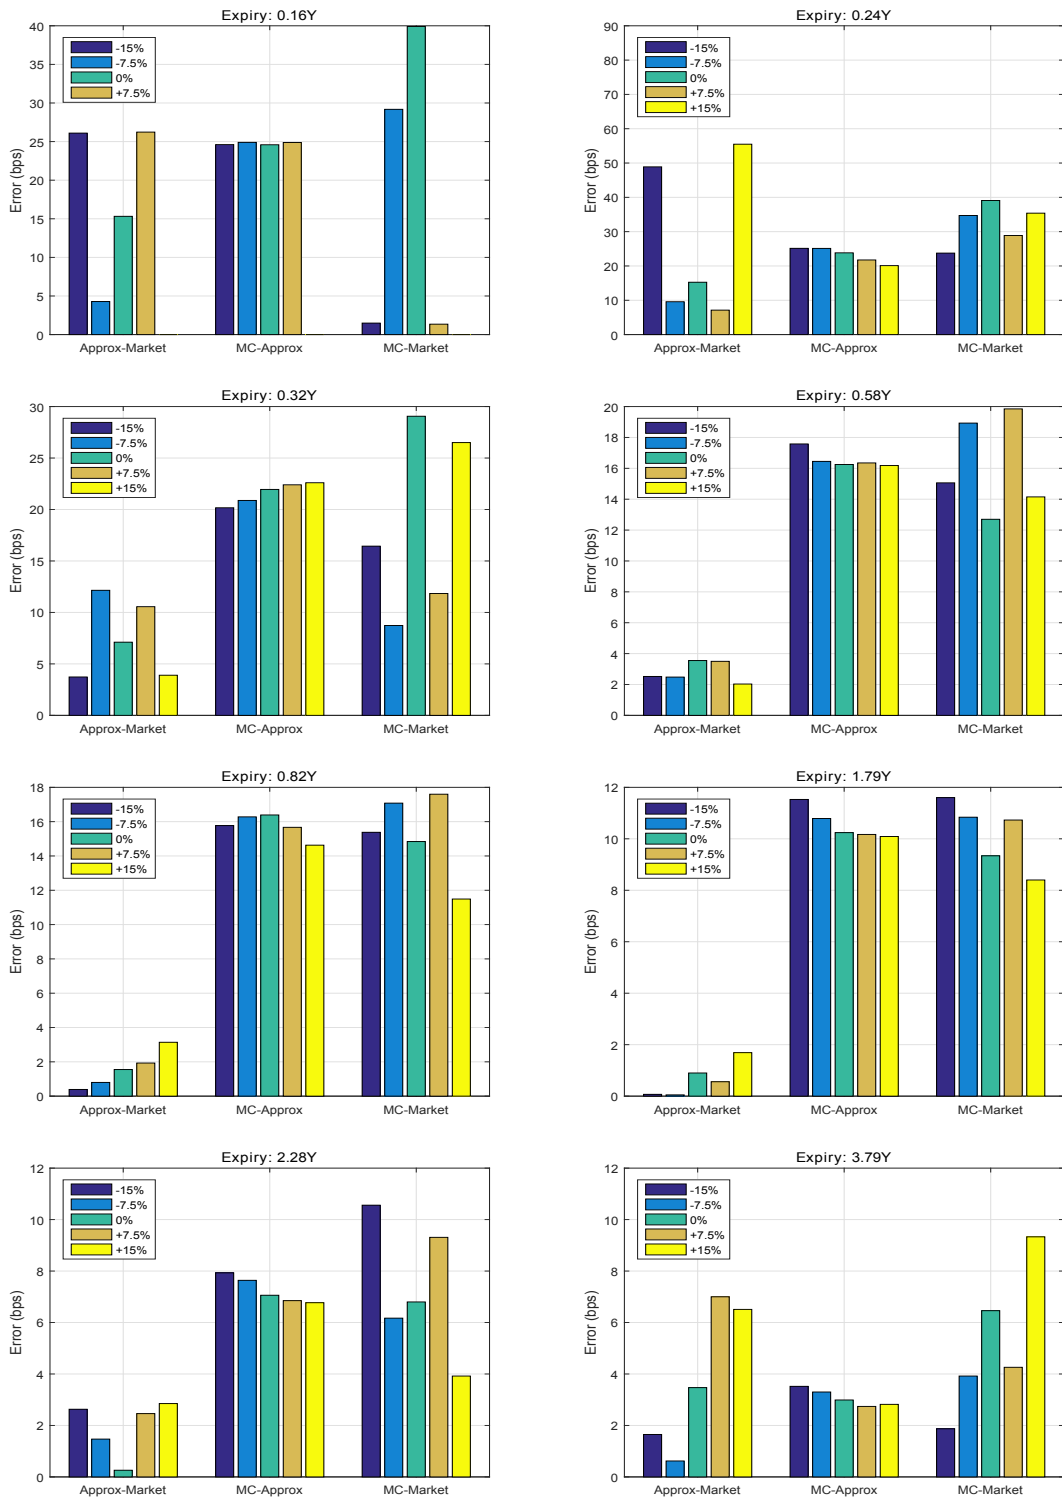


Figure 10: Reported errors between the volatilities implied from the Market, the Monte Carlo (MC) simulation and the semi-analytic approximations (Approx) used in within the calibration. The values are reported in basis points (bps).

- Andersen, L.B and Piterbarg, V. (2010), *Interest Rate Modeling: Models, Products, Risk Management*. In three volumes. Atlantic Financial Press.
- Anderson, G., Goldberg, L., Kercheval, A.N., Miller, G and Sorge, K. (2005), ‘On the Aggregation of Local Risk Models for Global Risk Management’, *Journal of Risk* **8(1)**, 25–40.
- Black, F. and Scholes, M. (1973), ‘The Pricing of Options and Corporate Liabilities’, *Journal of Political Economy* **81(3)**, 637–654.
- Black, F. (1976), ‘The Pricing of Commodity Contracts’, *Journal of Financial Economics* **3(1–2)**, 167–179.
- Brace, A., Gatarek, D. and Musiela, M. (1997), ‘The Market Model of Interest Rate Dynamics’, *Mathematical Finance* **7**, 127–154.
- Cheng, B. Nikitopoulos Sklibosios, C. and Schlögl, E. (2016a), ‘Pricing of Long-Dated Commodity Derivatives with Stochastic Volatility and Stochastic Interest Rates’, Available at SSRN: <http://ssrn.com/abstract=2712025>.
- Cheng, B. Nikitopoulos Sklibosios, C. and Schlögl, E. (2016b), ‘Empirical Pricing Performance on Long-Dated Crude Oil Derivatives: Do Models with Stochastic Interest Rates Matter?’, Available at SSRN: <http://ssrn.com/abstract=2721716>.
- Cheng, B. Nikitopoulos Sklibosios, C. and Schlögl, E. (2016c), ‘Hedging Futures Options with Stochastic Interest Rates’, Available at SSRN: <http://ssrn.com/abstract=2840635>.
- Cheng, B. Nikitopoulos Sklibosios, C. and Schlögl, E. (2016d), ‘Empirical Hedging Performance on Long-Dated Crude Oil Derivatives’, Available at SSRN: <http://ssrn.com/abstract=2840622>.
- Choy, B., Dun, T. and Schlögl, E. (2004), ‘Correlating Market Models’, *Risk*, **17**, 124–129.
- Cortazar, G. and Schwartz, E. S. (1994), ‘The Valuation of Commodity Contingent Claims’, *Journal of Derivatives* **1**, 27–39.
- Cox, J.C., Ingersoll, J.E. and Ross, S.A. (1981), ‘The Relation between Forward Prices and Futures Prices’, *Journal of Financial Econometrics* **9**, 321–346.
- Derman, Emanuel (2003), ‘Laughter in the Dark — The Problem of the Volatility Smile’, May 26, 2003.
- Gibson, R. and Schwartz, E. S. (1990), ‘Stochastic Convenience Yield and the Pricing of Oil Contingent Claims’, *Journal of Finance* **45**, 959–976.
- Golub, G. and Van Loan, C. (1996), *Matrix Computations*, 3rd ed., Johns Hopkins.
- Gower, J.C. and Dijkstraerhuis, G.B. (2004), *Procrustes Problems*, Oxford University Press.
- Grzelak, L. and Oosterlee, C. (2011a), ‘On Cross-Currency Models with Stochastic Volatility and Correlated Interest Rates’, *Applied Mathematical Finance*, 1–35.
- Grzelak, L. and Oosterlee, C. (2011b), ‘On the Heston Model with Stochastic Interest Rates’, *SIAM J. Financial Math.* **2**, 255–286.
- Hagan, P., Kumar, D., Lesniewski, A. and D. Woodward. (2002), ‘Managing Smile Risk’, *Wilmott Magazine* **September**, 84–108.
- Heath, D., Jarrow, R. and A. Morton. (1992), ‘Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation’, *Econometrica* **60**, 77–105.

- Jamshidian, F. (1997), 'LIBOR and Swap Market Models and Measures', *Finance and Stochastics* **1**, 293–330.
- Joshi, M. and R. Rebonato (2003), 'A Displaced–Diffusion Stochastic Volatility LIBOR Market Model: Motivation, Definition and Implementation', *Quantitative Finance* **3**(6), 458–469.
- Kercheval, A.N. (2006), 'Optimal Covariances in Risk Model Aggregation', *Proceedings of the Third IASTED International Conference on Financial Engineering and Applications*, ACTA Press, 30–35
- Koschat, M.A. and D.F. Swayne (1991), 'A weighted Procrustes criterion', *Psychometrika* **56**, 229–223.
- Miltersen, K. (2003), 'Commodity Price Modelling that Matches Current Observables: A New Approach', *Quantitative Finance* **3**, 51–58.
- Miltersen, K., Sandmann, K. and Sondermann, D. (1997), 'Closed Form Solutions for Term Structure Derivatives with Log-Normal Interest Rates', *Journal of Finance* **52**, 409–430.
- Miltersen, K. and Schwartz, E. (1998), 'Pricing of Options on Commodity Futures with Stochastic Term Structures of Convenience Yield and Interest Rates', *Journal of Financial and Quantitative Analysis* **33**, 33–59.
- Musiela, M. and Rutkowski, M. (1997a), 'Continuous–time Term Structure Models: Forward Measure Approach', *Finance and Stochastics* **1**, 261–291.
- Musiela, M. and Rutkowski, M. (1997b), *Martingale Methods in Financial Modelling*, Springer Verlag.
- Nocedal, J. and Wright, S.J. (2006), *Numerical Optimization*, Springer Verlag.
- Pedersen, M.B. (1998), 'Calibrating Libor Market Models', *SimCorp Financial Research Working Paper*.
- Pilz, K.F. and Schlögl, E. (2013), 'A hybrid commodity and interest rate market model', *Quantitative Finance*, **13**, 543–560.
- Piterbarg, V. (2005a), 'Time to smile', *Risk Magazine*, **18**, 71–75.
- Piterbarg, V. (2005b), 'Stochastic volatility model with time dependent skew', *Applied Mathematical Finance* **12**, 147–185.
- Rebonato, R. (1999), 'Calibrating the BGM Model', *Risk*, March, 88–94.
- Rebonato, R. (2004), *Volatility and Correlation: The Perfect Hedger and the Fox*, 2nd edition, Wiley, Chichester.
- Rebonato, R. and P. Jäckel (2000), 'The Most General Methodology to Create a Valid Correlation Matrix for Risk Management and Option Pricing Purposes', *Journal of Risk* **2**, 17–27.
- Rebonato, R., McKay, K. and R. White (2009), *The SABR/LIBOR Market Model*, Wiley.
- Schlögl, E. (2002a), 'Arbitrage-Free Interpolation in Models of Market Observable Interest Rates', In K. Sandmann and P. Schonbucher (Eds.), *Advances in Finance and Stochastics*, Springer, Heidelberg.
- Schlögl, E. (2002b), 'A Multicurrency Extension of the Lognormal Interest Rate Market Models' *Finance and Stochastics* **6**, 173–196.
- Schwartz, E. S. (1982), 'The Pricing of Commodity–Linked Bonds', *Journal of Finance* **37**, 525–539.
- Schwartz, E. S. (1997), 'The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging', *Journal of Finance* **52**, 923–973.