

ulm university universität **UUUM**

Libor Market Models with Stochastic Volatility and CMS Spread Option Pricing

Dissertation zur Erlangung des Doktorgrades Dr. rer. nat. der Fakultät für Mathematik und Wirtschaftswissenschaften der Universität Ulm

> vorgelegt von Dipl.-Math. oec., M.Sc. Matthias Lutz aus Wertheim

> > Ulm, im Februar 2011

Amtierender Dekan: Prof. Dr. Paul Wentges 1. Gutachter: Prof. Dr. Rüdiger Kiesel Prof. Dr. Ulrich Stadtmüller 3. Gutachter: Prof. Dr. John Schoenmakers Tag der Promotion: 27.05.2011

To the memory of Iris.

Acknowledgments

First and foremost I would like to express my deepest gratitude to Prof. Dr. Rüdiger Kiesel for his confidence in my work and for giving me the freedom to follow my interests. Moreover, I very much enjoyed being a member of the Institute of Mathematical Finance when he was chair of the institute. He created a pleasant working environment and I always enjoyed teaching under his guidance.

I also would like to express my sincere thanks to Prof. Dr. Ulrich Stadtmüller for being my co-examiner.

To my friends and colleagues, who are too numerous to mention individually: Thank you for many fruitful discussions, your companionship and many great nonuniversity activities. Special thanks go to Christian Hering, Eva Nacca, Andreas Rupp and Dennis Schätz for countless and "highly inspiring" coffee breaks. These will be sorely missed.

I am also deeply indebted to Katrin Jensen and Andreas Rupp for proofreading this thesis and for providing valuable comments.

Ulm, Februrary, 2011 Matthias Lutz

Contents

1	Intr	oduction	L
	1.1	Fixed-Income Instruments	1
	1.2	Objective of the Thesis and Contribution	4
	1.3	Structure of the Thesis	3
2	Inte	erest-Rate Products and Pricing Models	7
	2.1	Probabilistic Framework and Arbitrage-Free Pricing	7
	2.2	The Yield Curve, Forwards & Swaps	9
		2.2.1 Zero-Coupon Bonds and Spot Rates	9
		2.2.2 Forward-Rate Agreements and Forward Rates	0
		2.2.3 The Short Rate and Instantaneous Forward Rates 1	1
		2.2.4 Interest-Rate Swaps 12	2
		2.2.5 Yield-Curve Construction	4
	2.3	Fixed-Income Probability Measures	5
	2.4	Caps, Floors & Swaptions	3
	2.5	Vanilla Models	2
		2.5.1 Black's Model and the Volatility Smile	2
		2.5.2 Local Volatility Models	5
		2.5.3 Stochastic Volatility Models)
	2.6	Other Interest-Rate Options	9
		2.6.1 CMS Swaps, Caps and Floors	9
		2.6.2 CMS Spread Products	7
		2.6.3 More Exotic Products	1
	2.7	Term-Structure Models: From Short-Rate Models to HJM	3
		2.7.1 Short-Rate Models	3
		2.7.2 The HJM Framework	7
3	The	e Libor Market Model 59)
	3.1	Model Set-Up and No-Arbitrage Dynamics	0
		3.1.1 Swap-Rate Dynamics	3
	3.2	A Stochastic-Volatility Extended LMM	5
		3.2.1 Model Description	6
		3.2.2 Pricing European Options	7
	3.3	Parameterization of the LMM	3
		3.3.1 Volatility Structure	9
		3.3.2 Skew Structure	1
	3.4	Calibration	2

		3.4.1 Pre-Calibration: Effective Swaption and Caplet Parameters 7	2
		3.4.2 Main Calibration: Time-Dependent Parameters	3
_			_
4	Effi	cient Pricing of CMS Spread Options 7	5
	4.1	Approximating the Swap-Rate Dynamics	5
	4.2	Iterated Expectations	8
	4.3	The Density of the Integrated Variance	0
		4.3.1 The Branch-Cut Corrected Laplace Transform 8	0
		4.3.2 Calculating the Bromwich Integral	1
		4.3.3 The Optimal Linear Contour	3
	4.4	Calculating CMS Spread Option Prices	6
	4.5	Numerical Results	0
		4.5.1 Time-Dependent Parameters	0
		4.5.2 Constant Parameters	3
	4.6	Conclusion	5
5	Nev	v Correlation Parameterizations 9	7
	5.1	General Considerations	8
		5.1.1 Historical Correlations	8
		5.1.2 Stylized Facts	0
	5.2	Classical Parameterizations	1
	5.3	New Flexible Correlation Parameterizations	3
		5.3.1 Alternative Characterization of the SC-Family 10	4
		5.3.2 Reformulation of the Cholesky Decomposition	5
		5.3.3 New Parametric Forms	7
	5.4	Fitting Historical Correlations	0
	5.5	Conclusion	3
			-
6	DC	T Rank-Reduced Parameterizations 11	.5
	6.1	Existing Methods for Rank-Reducing Correlation Matrices 11	6
	6.2	The DCT Rank Reduction Method	8
	6.3	Numerical Results	0
	6.4	Conclusion	3
7	Ext	racting Correlations from the Market 12	:5
	7.1	Including CMS Spread Options	5
	7.2	Calibration Examples	8
		7.2.1 Data Description	8
		7.2.2 Calibration Results	9
	7.3	Pricing Applications	6
	7.4	Conclusion	0
-			
A	ppen	dices 14	1
	A	Laplace Transform of $V(T)$	1
	В	Time-Dependent Parameter Scenario	7
	С	Standard ρ_{∞} -extension	8
	D	Calibration Results	9

CONTENTS

Abbreviations and Notation	152
Bibliography	162
List of Figures	163
List of Tables	165
Zusammenfassung	167
Index	170

xi

Chapter 1

Introduction

1.1 Fixed-Income Instruments

The interest-rate derivatives market is by far the largest derivatives market in the world. In its 2010 survey¹, the Bank for International Settlements (BIS) reports that the global notional amounts outstanding of over-the-counter (OTC) derivatives² totalled \$583 trillion at the end of June 2010, with interest-rate derivatives accounting for 82.1% of this amount, followed by foreign-exchange contracts (10.8%), credit derivatives (5.4%), equity-linked contracts (1.2%) and commodities contracts $(0.6\%)^3$.

The motivations for using interest-rate derivatives are quite diverse and range from locking in financing costs at a specified fixed rate to pure speculation. While mutual funds or insurance companies may seek to earn superior returns on their investments, corporates and mortgage lenders generally want to hedge their interestrate exposures. According to the International Swaps and Derivatives Association (ISDA)⁴ more than 88% of the world's 500 largest companies use derivatives to manage their interest-rate risks.

Interest-rate derivatives come in different flavors such as swaps, forwards/futures and option-like instruments. An interest-rate swap, for instance, is in its most basic form an agreement between two counterparties to exchange a stream of fixed-rate payments for a stream of floating-rate payments (typically based on a reference rate such as Libor⁵). A pension fund may wish to enter into a receiver swap⁶ to convert the floating-rate coupons, that it is earning on a portfolio of bond-like assets, into a stream of fixed-rate payments to match the expected future pension liabilities, which are often fixed in nature. In contrast, a company with a floating-rate loan

 $^{^1}BIS:$ Triennial and semiannual surveys – Positions in global over-the-counter (OTC) derivatives markets at end-June 2010, Nov. 2010, http://www.bis.org/publ/otc_hy1011.pdf.

²Notional amounts outstanding of exchange traded derivatives reached \$78 trillion, see BIS: Statistical Annex 2010, http://www.bis.org/publ/qtrpdf/r_qa1012.pdf.

³The total gross market value amounts to 25 trillion with 75.0% in interest-rate derivatives, 12.8% in foreign-exchange contracts, 6.9% in credit derivatives, 3.2% in equity-linked contracts and 2.0% in commodities contracts.

⁴International Swaps and Derivatives Association: 2009 ISDA Derivatives Usage Survey, http: //www.isda.org/researchnotes/pdf/ISDA-Research-Notes2.pdf.

⁵London Interbank Offered Rate.

⁶For the party that receives the fixed payments and makes the floating-rate payments the swap is called a receiver swap. Otherwise it is called a payer swap.

may be concerned that future interest rates will possibly rise. Hence, to protect itself from increased financing costs, it may use a payer swap to convert, in effect, the floating-rate loan into a fixed-rate loan. While this strategy eliminates the risk of having to make unexpectedly high future payments, it might be quite "expensive" in a market environment where long-term (fixed) rates are relatively high, compared to short-term (floating) rates, the latter of which the company might expect to stay low or even decrease. In this case the company might be better off staying with the existing loan and buying instead an interest-rate cap. A cap consists of a sequence of interest-rate call options (caplets) and protects the purchaser against a rate rise above a prespecified level. Should market rates move above the cap rate, then the seller pays the purchaser the difference between the market rate and the cap rate. For the company, a cap therefore creates a ceiling on its floating-rate interest costs on the loan. In turn, the company pays the cap seller an upfront premium for this "insurance contract". Interest-rate floors, in contrast to caps, provide holders with protection against drops in market rates. These instruments can be bought, for example, by investors to achieve a minimum return on floating-rate assets or by debt managers to protect against opportunity losses on fixed-rate debt when interest rates fall.

While (standard) swaps, forwards/futures and caps/floors are the most liquidly traded contracts and constitute the largest share of interest-rate derivatives in the market, there exists a great variety of other, more exotic derivatives, that may be structured individually to meet a company's specific risk profile or an investor's market view. Products that were particularly popular in the years 2005 through 2007 were constant maturity swap (CMS) steepener products, with an estimated \$55 billion⁷ of steepener notes issued in 2005.

Steepener (or CMS spread-linked) products allow investors to take a view on the shape of the yield curve. Although the specific characteristics of these products may vary significantly, they are typically sold as a medium-term note, paying high fixed coupons for the first few years, after which the investor receives coupons based on the spread between two CMS rates (often the 10-year and the 2-year rate) multiplied by a leverage factor. Accordingly, these products can be bought by investors who believe that the prevailing yield curve is too flat and will probably steepen. The latter can happen for example if central banks loosen their monetary conditions in response to a weakening economy, or, if short rates and central bank monetary conditions stay the same, but long-term rates go up due to inflationary pressures.

In the past, yield curves in the USA and the euro countries have tended to be upward sloping for most of the time. Between 2004 and 2006, however, yield curves began to flatten significantly and forward curves were even pricing in a future inversion of the yield curve. Many investors thought that the prevailing interest-rate curves were artificially flat and would probably steepen. Due to the low forward spreads, dealers were able to structure products that offered high initial coupons (often 6%–10%) and high leverage factors for the spread coupons. Many of such products were sold to retail investors as well as institutional investors, like pension funds or insurance companies, who were looking at that time for enhanced yields in a falling interest-rate environment.

However, instead of steeping, yield curves continued to flatten in 2006 and 2007,

⁷"Vanilla's the flavour", *Risk*, March 2008, p. 61–63.

leaving owners of steepener notes with mark-to-market losses and below market coupons⁸. In the euro area the situation even worsened in June 2008. The spread between 10-year and 2-year euro swaps had been negative since mid-May, when on June 5 the European Central Bank (ECB) president Jean-Claude Trichet remarked that the bank could possibly raise their interest rates in response to a rising inflation⁹. This caught many dealers and investors on the wrong foot as they were convinced at that time that the credit crises and the weakening economy would force the ECB to cut their rates and that, as a consequence, the curve would remain positive or even steepen¹⁰. After the announcement, the euro curve 10Y–2Y spread underwent an unprecedented drop, which came to a halt at -60.2 basis points (bp) on June 6^{11} , see Figure 1.1. It is interesting to note that the inversion of the yield curve was partially exacerbated by hedging activities around steepener structures¹². In fact, according to Risk magazine¹³, structurers estimated "that probably no more than 30% of the inversion down solely to dynamic hedging activity". Finally, in October 2008,



Figure 1.1: Historical fixings of the EUR 10Y and 2Y swap-rates and their difference.

the ECB started the long awaited cycle of rate cuts¹⁴ and, as a result, CMS spreads came back to more normal levels. With the 10Y-2Y spread back in the 150 bp-200 bp range, investors, who stayed with their steepener notes, were eventually rewarded with high coupon payments (provided the notes had not been previously called by the issuers).

Since 2004/2005, when steepener products emerged in large volumes, also the market in standard ("vanilla") CMS spread options has grown substantially, and prices of such options are nowadays quoted by brokers on a regular basis. Neverthe-

⁸"Getting flattened", *Risk*, February 2006, p. 18–20.

⁹"Bundesbank's inflation warning", *Financial Times*, June 7, 2008.

 $^{^{10}\}ensuremath{^{\prime\prime}}\xspace$ The rates escape", Risk, August 2008, p. 25–27.

¹¹Source: Bloomberg L.P.

 $^{^{12}}$ At that time many steepener notes were structured as CMS spread range accruals, floored at zero. Once the curve breaches the strike price of the digital floors implicitly contained in these structure, the gamma profile of these options suddenly inverts and dealers need to rapidly reallocate their hedge books. See also "The gamma trap", *Risk*, December 2006.

¹³*Risk*, August 2008, p. 26.

¹⁴After it had indeed increased its main refinancing rate by 25bp in July 2008.

less, the correct pricing of CMS spread options is still an area of ongoing research. Indeed, as a recent article by McCloud [McC11] reveals, for much of the year 2009 there were static arbitrage opportunities due to price inconsistencies in the markets for options on CMS rates and CMS spreads. Yet, pricing CMS spread options inline with other market sectors is not the only challenge. Also exotic and plain vanilla CMS spread derivatives must naturally be priced consistently with each other. This is necessary to avoid direct arbitrage opportunities and also to incorporate the true hedging costs of exotic products. Ultimately, the price of an exotic (or any other) derivative security is equal to the sum of hedging costs incurred throughout its life time. With vanilla CMS spread options becoming more and more liquid, these may serve as natural hedging instruments for CMS spread-linked exotic (or other correlation-sensitive products), and hence prices of the vanilla options must be properly taken into account when pricing the exotics.

1.2 Objective of the Thesis and Contribution

In the financial industry, the workhorses for pricing and risk-managing exotic interestrate derivatives are the so-called Libor market models (LMMs). Due to the flexibility of their volatility specification, these models can be calibrated to a wide range of market instruments, and exotic structures can therefore be priced inline with the prevailing market conditions. Moreover, contrary to other models, LMMs allow for rich enough dynamics to capture the decorrelation among rates across the yield curve, which is particularly important for pricing correlation-sensitive products such as CMS spread-linked exotics.

The main objective of the present thesis is to provide efficient methods and tools for calibrating LMMs to market-prices of caps, swaptions and CMS spread options, and in this way extract the volatility and correlation information implicitly contained in these products.

Since the standard (log-normal) LMM, as introduced by Miltersen, Sandmann & Sondermann [MSS97], Brace, Gatarek & Musiela (BGM) [BGM97] and Jamshidian [Jam97] cannot capture the pronounced volatility smiles found in today's interestrate markets, several extensions of the original model have been proposed. Arguably, the most popular of these extensions are based on Heston-type dynamics (see e.g. [AA02], [ABR05] or [Pit05a]). In Chapter 4 we present a new approximation formula for pricing CMS spread options within this model class. At the core of this approximation lies a new method for efficiently evaluating the density of an integrated Cox-Ingersoll-Ross (CIR) process, based on carefully choosing the integration contour for the required Laplace inverse transform. This method is not restricted to the context of CMS spread option pricing, but can be applied whenever a rapid and exact evaluation of the density of an integrated CIR process is required. The problem of pricing CMS spread options within the aforementioned model class is also considered in a recent paper by Antonov & Arneguy [AA09]. We demonstrate that in terms of speed, accuracy and ease of implementation our pricing formula generally outperforms the pricing methods described in their paper.

With an efficient and accurate CMS spread option pricing formula at our disposal, we can include CMS spread options in the general calibration procedure, and in this way back out information about the Libor correlations as implied by the market. This approach is to be contrasted with the common practice of using historically estimated Libor correlations, which is backward-looking in nature and generally does not reflect the prevailing market environment.

No matter whether one adopts the implied or the historical calibration approach, in both cases a parsimonious yet flexible parameterization for the Libor correlations is required. When following the "historical" approach, this is due to the fact that historically estimated Libor correlation matrices are often quite noisy and may contain counterintuitive entries. By fitting a low-parametric functional form to the historical correlations, one tries to obtain a reasonably smooth matrix, which captures only the main features of the historical data. On the other hand, when following the implied route, correlation parameterizations are necessary to avoid overfitting and to obtain stable calibration results. In Chapter 5 we present a new generic method for constructing correlation parameterizations that are always positive definite and derive new flexible low-parametric forms. A study shows that these parametric forms can fit historically estimated correlation matrices better than the popular standard parameterizations.

In practice, the number of driving factors used within a LMM is typically much smaller (say 3–10) than the number of modeled forward rates (often 40–80). Since virtually all existing correlation parameterizations yield full-rank correlation matrices, a PCA-based¹⁵ rank reduction must be applied before the matrices can actually be used within the model. In case of a historical calibration approach, this does not necessarily constitute a serious drawback, as the historically fitted correlation matrix must be rank-reduced only once. However, when calibrating the correlation parameters via the implied route, often thousands of rank reductions must be performed. In this case, the required numerical eigenvalue decompositions may constitute a large share of the total computational costs of the calibration procedure, and may significantly slow it down. In Chapter 6 we therefore develop a new efficient method for rank-reducing parametric forward-rate correlation matrices. The method is based an applying a discrete cosine transform (DCT) to the rows of the Cholesky decomposition of the correlation matrix. As the Cholesky decompositions of our parametric forms are given in closed form, and due to the low computational cost of the DCT rank-reduction method, the combination of the two implicitly generates a new family of low-rank parametric forms.

Finally, in Chapter 7, we put the previously developed methods and tools into action, and calibrate LMMs to real market data. We discuss two possible calibration approaches and demonstrate that with our correlation parameterization generally better market fits can be achieved than with the standard correlation parameterizations. Moreover, one of the main findings of our empirical analysis is that market-implied correlation matrices do not display pronounced upward sloping sub-diagonals – a typical feature of historical correlation matrices, which is sometimes even directly built into correlation parameterizations.

Lastly, we use the calibrated models to provide some pricing examples. In particular, we demonstrate that LMMs with different correlation structures may yield markedly different prices for certain exotic interest-rate products, even though their

¹⁵Principal component analysis.

market fit to a collection of caps, swaptions and CMS spread options is essentially identical.

1.3 Structure of the Thesis

After this introductory chapter, we introduce in Chapter 2 the basic definitions, outline the fundamental concepts of fixed-income modeling and give an overview of the most common interest-rate derivatives. Furthermore, we introduce some of the most important single-rate and full term-structure models. Chapter 3 is devoted to the introduction of one of the most popular families of term-structure models: the Libor market models. Besides the standard (log-normal) LMM we also introduce a class of stochastic-volatility extended LMMs. We present possible parameterizations of the models and discuss pricing and calibration techniques. In Chapter 4 we develop a new formula for efficiently pricing CMS spread options in the aforementioned class of stochastic-volatility extended LMMs. Chapter 5 briefly reviews the standard Libor correlation parameterizations before new flexible parametric forms are developed. The subsequent Chapter 6 introduces a new efficient method for rank-reducing parametric forward-rate correlation matrices. The previously developed methods and tools are finally put into action in Chapter 7, where we calibrate LMMs to real market data, provide pricing examples and discuss empirical findings.

A Word Regarding Computational Implementation and Market Data

All numerical results (including the timing results) are based on C++ implementations, where we made use of the free/open-source QuantLib framework, see http: //quantlib.org. The code was compiled with the Microsoft C/C++ Compiler version 14.0, and run on a standard desktop PC with an Intel Core2Duo 3Ghz processor. Most of the market data used in this thesis was provided by Bloomberg L.P. This data service is available at the Universität Ulm due to the generous support of the Landesbank Baden-Württemberg (LBBW). We thank Dr. Jörg Kienitz for providing the CMS spread options data used for the calibration examples in Chapter 7.

Chapter 2

Interest-Rate Products and Pricing Models

In this chapter we introduce the basic notations to characterize prices and yields of interest-rate products and present a brief review of the literature on interest-rate models.

2.1 Probabilistic Framework and Arbitrage-Free Pricing

Throughout this thesis, we shall always consider an economy with continuous frictionless trading taking place inside some finite time horizon $[0, T^*]$. To model uncertainty and the flow of information we work on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$. Our sample space Ω is equipped with a σ -algebra \mathcal{F} , and "information" is revealed over time according to a filtration $\mathbb{F} = \{\mathcal{F}_t\}_{0 \leq t}$, an increasing family of σ -algebras satisfying $\mathcal{F}_s \subseteq \mathcal{F}_t \subseteq \mathcal{F}$ whenever $0 \leq s \leq t$. In all of the models, that we consider in this thesis, the market is "driven" by some d-dimensional standard Brownian motion (or Wiener process) $W(t) = (W_1(t), \ldots, W_d(t))'$, and the filtration is always the one generated by $W(\cdot)$, i.e., $\mathcal{F}_t = \sigma (\{W(u), 0 \leq u \leq t\})^1$. Prices of assets will generally be modeled via a p-dimensional Itô process $X(t) = (X_1(t), \ldots, X_p(t))'$, characterized by a stochastic differential equation (SDE) of the form²

$$X(t) = X_0 + \int_0^t \mu(s, X(s)) ds + \int_0^t \sigma(s, X(s)) dW(s),$$

or, in differential notation,

$$dX(t) = \mu(t, X(t))dt + \sigma(t, X(t))dW(t), \ X(0) = X_0,$$

with drift $\mu : [0, T^*] \times \mathbb{R}^p \to \mathbb{R}^p$, diffusion coefficient $\sigma : [0, T^*] \times \mathbb{R}^p \to \mathbb{R}^{p \times d}$ and initial condition³ $X_0 \in \mathbb{R}^p$. An introduction to Itô processes, Itô integrals and the related stochastic calculus can be found, e.g. in [KS98] or [Øks03].

In the following we give a brief review of some fundamental results concerning the arbitrage-free pricing of derivative securities (also known as contingent claims). For

¹Augmented to satisfy the "usual conditions", i.e., \mathcal{F}_t must be right-continuous for all t, and \mathcal{F}_0 must contain all \mathbb{P} -null sets of \mathcal{F} .

²For restrictions on μ and σ ensuring existence and uniqueness of the solution to the SDE, see e.g. [Øks03].

³In most cases X_0 will be the currently market observed price/rate of a security and we will omit the initial condition.

a more detailed and rigorous treatment of this topic we refer the reader to [HK00], [MR05] and [BK04].

A trading strategy is a predictable⁴ \mathcal{F}_t -adapted process $\varphi(t, \omega) = (\varphi_1(t, \omega), \ldots, \varphi_p(t, \omega))$, where $\varphi_i(t, \omega)$ denotes the amount of shares of asset X_i held in the portfolio (associated with the trading strategy) at time t. The value process of the trading strategy at time t is given by⁵

$$V_{\varphi}(t) = \varphi(t)' X(t).$$

A trading strategy φ is called *self-financing* if, for any $t \in [0, T^*]$, the value process satisfies

$$V_{\varphi}(t) = V_{\varphi}(0) + \int_0^t \varphi(s)' dX(s).$$

An arbitrage opportunity is a self-financing strategy φ for which $V_{\varphi}(0) = 0$ and for some $t \in [0, T^*]$,

$$V_{\varphi}(t) \geq 0$$
 a.s. and $\mathbb{P}(V_{\varphi}(t) > 0) > 0$.

Absence of arbitrage opportunities in an economy can be characterized via the concept of equivalent martingale measures. Define a *numeraire* to be any a.s. strictly positive non-dividend paying asset and denote it by N. Then we say that a measure \mathbb{Q}^N , equivalent⁶ to \mathbb{P} , is an equivalent martingale measure induced by N if the relative price process

$$X(t)/N(t) = (X_1(t)/N(t), \dots, X_p(t)/N(t))^{n}$$

is a martingale with respect to $(\{\mathcal{F}_t\}, \mathbb{Q}^N)$.

In a nutshell, the absence of arbitrage in a financial market is "equivalent" to the existence of an equivalent martingale measure. There exist various formulations of this statement⁷, and a number of technical conditions are necessary in order to define more precisely in which sense the two concepts are equivalent. For brevity, we ignore these technicalities at this point and refer the reader to the references mentioned at the beginning of this paragraph.

Next, we define a derivative security (or contingent claim) with maturity date T to be an \mathcal{F}_T -measurable random variable H(T). We assume that H(T) has finite variance and say that the derivative security is attainable if there exists an admissible⁸ trading strategy φ , such that $H(T) = \varphi(T)'X(T) = V_{\varphi}(T)$ a.s. We call such a trading strategy φ a replicating strategy for H(T). This means, that holding the portfolio (associated with the replicating strategy) and holding the derivative security are equivalent from a financial point of view. In the absence of arbitrage, the time-0 price H(0) of an attainable derivative security must therefore equal the cost of setting up

 $^{^{4}}$ See [KS98].

⁵For convenience we often drop the dependence on ω when considering random variables or stochastic processes.

⁶Two measures \mathbb{Q} and \mathbb{P} , defined on the same measure space (Ω, \mathcal{F}) , are said to be *equivalent*, if they share the same null-sets, i.e., $\mathbb{Q}(A) = 0 \Leftrightarrow \mathbb{P}(A) = 0, \forall A \in \mathcal{F}$.

⁷Commonly known as the *fundamental theorem of asset pricing*.

⁸A self-financing trading strategy φ is said to be *admissible* if $\int_0^t \varphi(s)' dX(s)$ is a martingale under \mathbb{Q}^N .

the (self-financing) replicating strategy, i.e., $H(0) = V_{\varphi}(0)$. More generally, we must have $H(t) = V_{\varphi}(t) \ \forall t \in [0, T]$ a.s. This observation sets the stage for what is known as arbitrage pricing and allows us to write prices of derivative securities as expectations under an equivalent martingale measure. Specifically, consider a numeraire N and assume the existence of an associated equivalent martingale measure \mathbb{Q}^N . Then, from the martingale property of $V_{\varphi}(t)/N(t)$ and the relation $V_{\varphi}(t) = H(t)$ it immediately follows that

$$\frac{H(t)}{N(t)} = \mathbb{E}_t^N \left[\frac{H(T)}{N(T)} \right],$$
$$H(t) = N(t)\mathbb{E}_t^N \left[\frac{H(T)}{N(T)} \right],$$
(2.1)

or equivalently,

where we used the shorthand notation $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot |\mathcal{F}_t]$ and where $\mathbb{E}^N[\cdot]$ denotes the expectation operator with respect to the measure \mathbb{Q}^N . Equation (2.1) often serves as the starting point for specifying new models. The description of asset-price dynamics is then directly carried out in an equivalent martingale measure⁹, without addressing the real-world evolution of asset prices.

In later sections we will see that it is often convenient to choose a particular numeraire for pricing certain products, as this will simplify computations. For this the so-called *change-of-numeraire technique*¹⁰ is useful: Given an equivalent martingale measure \mathbb{Q}^N induced by a numeraire N, we may define a new measure \mathbb{Q}^M , associated with another numeraire M, via the Radon-Nikodym derivative $\zeta(T^*)$ with respect to \mathbb{Q}^N ,

$$\zeta(t) = \frac{d\mathbb{Q}^M}{d\mathbb{Q}^N}\Big|_{\mathcal{F}_t} = \frac{M(t)N(0)}{N(t)M(0)}, \ 0 \le t \le T^*,$$

such that for an \mathcal{F}_T -measurable random variable H(T) we have

$$H(t) = N(t)\mathbb{E}_t^N \left[\frac{H(T)}{N(T)}\right] = M(t)\mathbb{E}_t^M \left[\frac{H(T)}{M(T)}\right]$$

If all \mathcal{F}_T -measurable random variables (satisfying some integrability conditions) are attainable (i.e., can be replicated or hedged), the market is said to be *complete*. In a complete market all derivatives have unique prices, no matter which equivalent martingale measure we choose.

2.2 The Yield Curve, Forwards & Swaps

2.2.1 Zero-Coupon Bonds and Spot Rates

The most basic fixed-income security is the so-called *zero-coupon bond* (also known as *zero bond* or *discount bond*), which is the building block of all interest rates. A zero-coupon bond maturing at time T > 0 is a contract that guarantees its holder the payment of one unit of currency at time T, with no intermediate payments. We

⁹Often just called *pricing measure*.

¹⁰See Section 2.3 of [BM05] for useful facts and formulae around the change-of-numeraire technique.

will denote the price at time $t \leq T$ of such a bond by P(t,T). Clearly, P(T,T) = 1 for all T > 0.

It is often more convenient to characterize bond prices in terms of interest rates. One such rate is the *continuously compounded spot rate* R(t,T) prevailing at time t for maturity T

$$R(t,T) \coloneqq -\frac{\log P(t,T)}{\tau(t,T)},$$

that is

$$e^{-R(t,T)\tau(t,T)}P(t,T) = 1.$$

Here, the so-called *year fraction* (or *accrual factor*) $\tau(t, T)$ measures the time between t and T (in years). In practice there exists a great variety of *day-count conventions*, which differ according to product type and country. For more details, also on the different business day calendars and date rolling conventions see [ISD06]¹¹. For clearness of exposition we will simply use $\tau(t, T) = T - t$ in the rest of this thesis, although we correctly take into account the respective conventions when dealing with real market data in later chapters.

The continuously compounded spot rate is consistent with the rather idealized assumption of continuously reinvesting any accrued interest at a constant rate. In reality, however, most market quotes are based on simple compounding, where interest accrues proportionally to the time of the investment. Accordingly, we define the *simply-compounded spot rate* L(t,T) as

$$L(t,T) \coloneqq \frac{1}{\tau} \left(\frac{1}{P(t,T)} - 1 \right), \qquad (2.2)$$

that is,

 $(1 + \tau L(t, T)) P(t, T) = 1,$

where $\tau = T - t$ is often called the *tenor* of the spot rate. Examples for such spot rates are *Libor rates* (London Interbank Offered Rate) and *Euribor rates* (Euro Interbank Offered Rate), which are benchmark interest rates at which banks can borrow unsecured funds from other banks in the interbank markets¹². Libor¹³ and Euribor rates are quoted for tenors τ ranging from one week ($\tau = 1/52$)¹⁴ to 12 months ($\tau = 1$) and form the basis for determining the cash flows of other interest-rate derivative securities. One such security is a forward contract known as *forward-rate agreement* (FRA).

2.2.2 Forward-Rate Agreements and Forward Rates

A FRA is a contract between two parties that allows to lock in at the current time t, the interest rate over the future period $[T_1, T_2]$, $t < T_1 < T_2$. More specifically, at

 $^{12}\mathrm{See}$ www.bbalibor.com and www.euribor-ebf.eu.

¹¹For conventions used for EUR market instruments see also [AB09].

 $^{^{13}\}mathrm{All}$ further references to Libor rates in this thesis are equally valid for Euribor rates.

¹⁴For Libor rates the standard day-count convention is "Actual/360", i.e., $\tau = \tau(T_1, T_2)$ is the actual number of days between T_1 and T_2 divided by 360. Consequently, for a 6M Libor rate τ is usually slightly larger than 0.5.

the expiry T_2 of the FRA, a fixed payment based on a fixed rate K (agreed on at the time of initiation of the contract) is exchanged against a floating payment based on the spot (Libor) rate $L(T_1, T_2)$ resetting at time T_1 with tenor $\tau = T_2 - T_1$. Formally, in case of a FRA with unit notional, one pays at time T_2 the amount τK and receives $\tau L(T_1, T_2)$. From the perspective of the fixed-rate payer, the value of the contract at time T_2 is therefore

$$\tau(L(T_1, T_2) - K).$$
 (2.3)

Recalling the definition of L(t,T) in (2.2), this value can be rewritten as

$$\frac{1}{P(T_1, T_2)} - 1 - \tau K. \tag{2.4}$$

The term $1/P(T_1, T_2)$ is an amount of currency held at time T_2 . Multiplying it by $P(T_1, T_2)$ gives its value at time T_1 , which is clearly one. One unit of currency at time T_1 , in turn, is worth $P(t, T_1)$ units of currency at time t. Therefore, the amount $1/P(T_1, T_2)$ at time T_2 is equivalent to an amount of $P(t, T_1)$ units of currency at time t. Discounting also the other two terms in (2.4) back to time t finally yields the time-t value of the FRA

$$V_{\text{FRA}}(t, T_1, T_2, \tau, K) = P(t, T_1) - P(t, T_2) - P(t, T_2)\tau K$$

= $\tau P(t, T_2) \left(\frac{1}{\tau} \left(\frac{P(t, T_1)}{P(t, T_2)} - 1\right) - K\right).$

Most often, at the time of issuance, the fixed rate K of a FRA is chosen such that the cost to enter in the forward contract is zero for either party. The value of K that renders the contract fair at time t, i.e., that makes the FRA have value zero, is given by the *simply-compounded forward rate* $L(t, T_1, T_2)$

$$L(t, T_1, T_2) \coloneqq \frac{1}{\tau} \left(\frac{P(t, T_1)}{P(t, T_2)} - 1 \right).$$
(2.5)

Notice that in order to value a FRA, we simply have to replace the Libor rate $L(T_1, T_2)$ in the payoff (2.3) with the just-defined forward rate $L(t, T_1, T_2)$ and then take the present value of the resulting (deterministic) quantity. The forward rate $L(t, T_1, T_2)$ may thus be interpreted as an estimate of the future (spot) Libor rate $L(T_1, T_2)$, and it is therefore often called forward Libor rate¹⁵.

2.2.3 The Short Rate and Instantaneous Forward Rates

When the maturity S of the forward rate L(t, T, S) collapses towards its expiry T, we obtain the time-t instantaneous forward rate for maturity T > t,

$$f(t,T) \coloneqq \lim_{S \downarrow T} L(t,T,S)$$
$$= -\frac{\partial \log P(t,T)}{\partial T}.$$
(2.6)

¹⁵In fact, we will see below that $L(t, T_1, T_2)$ is the time-t expectation of $L(T_1, T_2)$ under a suitable probability measure.

Clearly, this notation makes only sense if we assume that the zero-bond price function $T \mapsto P(t,T)$ is sufficiently smooth for all T > t. Intuitively, f(t,T) is the time-t forward rate for the infinitesimally small time interval [T, T + dT]. We notice the relationship

$$P(t,T) = \exp\left(-\int_{t}^{T} f(t,u)du\right)$$

between the price of a zero-coupon bond and instantaneous forward rates. Letting the maturity T of the forward rate f(t,T) tend to t, we obtain the so-called *short* rate

$$r(t) \coloneqq f(t,t),$$

which, loosely speaking, can be interpreted as the overnight rate prevailing at time t.

Both, the short rate r(t) and the instantaneous forward rate f(t, T), are of course only theoretical quantities and cannot be directly observed in the market. They form, however, the basis for a number of interest-rate models, some of which we will examine in more detail in Section 2.7.

2.2.4 Interest-Rate Swaps

A swap is a general term for a contract in which two counterparties agree to exchange one stream of cash flows against another stream. These streams are usually called the *legs* of the swap. A *fixed-for-floating interest-rate swap* (also called *plain vanilla* $swap^{16}$ or just swap if there is no danger of confusion) can be considered as a generalization of a FRA, where one leg of the swap is a stream of fixed-rate payments and the other leg is a stream of payments based on floating rates, most often Libor rates.

For concreteness, define a $tenor\ structure,$ i.e., an increasing sequence of maturity times

$$0 \le T_0 < T_1 < \ldots < T_N, \quad \tau_n = T_{n+1} - T_n.$$

At the end of each period $[T_n, T_{n+1}]$, n = 0, ..., N-1, one party (the fixed rate payer) pays the amount

$$\tau_n K$$
,

corresponding to a fixed simple interest rate K, whereas the other party (the floatingrate payer) pays the amount

$$\tau_n L(T_n, T_{n+1}),$$

corresponding to the interest rate $L(T_n, T_{n+1})$ fixing at time T_n . The times T_n and T_{n+1} are normally referred to as the *fixing* and *payment dates* for the *n*-th period, respectively. Notice that the realization of the spot rate $L(T_n, T_{n+1})$ is not known until time T_n .

In practice, payments are of course usually netted, such that cash flows take place in only one direction at each payment date T_{n+1} , $n = 0, \ldots, N-1$. When the fixed leg is paid, the swap is usually called *payer swap*, whereas in the other case we have a *receiver swap*.

 $^{^{16}}$ In finance, the term "plain vanilla" or just "vanilla" is often used to denote the most basic and/or most liquid products from a certain class of derivative securities.

2.2. The Yield Curve, Forwards & Swaps

From the perspective of the fixed-rate payer, the net cash flow at time T_{n+1} is that of a FRA with fixed rate K, fixing date T_n and payment date T_{n+1} ,

$$\tau_n(L(T_n, T_{n+1}) - K).$$

A swap may thus be viewed as a portfolio of $FRAs^{17}$ with time-t value (to the fixed-rate payer)

$$V_{\text{swap}}(t) = \sum_{n=0}^{N-1} V_{\text{FRA}}(t, T_n, T_{n+1}, \tau_n, K)$$
$$= \sum_{n=0}^{N-1} \tau_n P(t, T_{n+1}) (L_n(t) - K),$$

where we have introduced the short-hand notation¹⁸

$$L_n(t) \coloneqq L(t, T_n, T_{n+1}).$$

We have seen above that requiring a FRA to be fair at the time of issuance leads to the definition of forward rates. Similarly, we may require the above fixed-for-floating swap to be fair at time t. The above swap valuation formula can be rewritten as

$$V_{\text{swap}}(t) = \left(\sum_{n=0}^{N-1} \tau_n P(t, T_{n+1})\right) \left(\frac{P(t, T_0) - P(t, T_N)}{\sum_{n=0}^{N-1} \tau_n P(t, T_{n+1})} - K\right)$$
$$= A_{0,N}(t)(S_{0,N}(t) - K), \tag{2.7}$$

where we have defined

$$A_{0,N}(t) \coloneqq \sum_{n=0}^{N-1} \tau_n P(t, T_{n+1})$$
(2.8)

and

$$S_{0,N}(t) \coloneqq \frac{P(t,T_0) - P(t,T_N)}{\sum_{n=0}^{N-1} \tau_n P(t,T_{n+1})}.$$
(2.9)

The quantity $A_{0,N}(t)$ is called the *annuity factor* of the swap (or PVBP for present value of a basis point) and $S_{0,N}(t)$ is the *forward swap rate*. If the swap is *spot-starting*, i.e., $t = T_0$ (today), then $S_{0,N}(t)$ is called *spot swap rate* or just swap rate¹⁹. The (forward) swap rate is also often referred to as the *par* or *break-even rate* of the swap, as it is the value of the fixed rate K that makes the swap have value zero

¹⁷Conversely, a FRA is a one period swap.

¹⁸Observe that $L_n(T_n) = L(T_n, T_n, T_{n+1}) = L(T_n, T_{n+1}).$

¹⁹Normally the schedule of "spot-starting" swaps starts one or two business days after time t, i.e., $T_0 = t + \delta$ where $\delta > 0$ is some contractually specified delay. So, strictly speaking, even spot swap rates are actually forward swap rates.

at time t. An alternative interpretation of the swap rate can be given by considering the following formula

$$S_{0,N}(t) = \sum_{n=0}^{N-1} w_n(t) L_n(t), \qquad (2.10)$$

where the weights $w_n(t)$ are given by

$$w_n(t) \coloneqq \frac{\tau_n P(t, T_{n+1})}{\sum_{k=0}^{N-1} \tau_k P(t, T_{k+1})} \ge 0,$$
$$\sum_{n=0}^{N-1} w_n(t) = 1.$$

The swap rate is therefore a convex combination (and thus a weighted average) of the forward rates associated with the FRAs underlying the swap.

In the above description of the prototypical swap we assumed that fixed-rate payments and floating-rate payments occur at the same dates with the same year fractions. This can of course be easily generalized to the case where the swap legs have different payment dates and day-count conventions (see, e.g. [Sch05], p. 15). In fact, vanilla swaps in the euro market typically have a fixed leg with annual payments and a floating leg with semiannual payments (based on 6M Euribor rates).

2.2.5 Yield-Curve Construction

In the above paragraphs we implicitly assumed that we are given a zero-coupon bond price function (the so-called *discount curve*) for a continuum of maturities in some interval, and we defined the various interest rates and values of interest-rate products in terms of zero-coupon bond prices derived from this curve. This is, however, putting the cart before the horse. In reality at most a few short-dated zero-coupon bonds are directly quoted in the market at any given time, and the discount curve must instead be inferred from a set of liquidly traded fixed-income securities by using an iterative procedure commonly known as *bootstrapping*.

To see how the general bootstrapping procedure works, suppose for simplicity that we can observe in the market at time t a set of N (spot) swap rates $S_{0,i}(t)$, $i = 1, \ldots, N$ defined on a common tenor grid $t = T_0 < T_1 < \ldots < T_N$. Assuming inductively that $P(t, T_i)$ is known for $i = 1, \ldots, n-1$ and using the definition of swap rates in (2.9), we may then compute the zero-coupon bond price for maturity T_n :

$$P(t,T_n) = \frac{1 - S_{0,n}(t) \sum_{i=0}^{n-2} \tau_i P(t,T_{i+1})}{1 + \tau_n S_{0,n}(t)}.$$

Proceeding in this way we can determine, at least in principle, the discount curve at the discrete tenor dates T_i underlying the swap rates. In practice, however, liquid swap rates may only be available for maturities of say 2–5 years and 7, 10, 12, 15, 20, 25 and 30 years.

In order for the above bootstrapping procedure to work, we therefore need to specify an interpolation scheme that allows us to fill the zero-bond prices for the missing tenor dates. Even if we require that all benchmark prices be matched, there is of course an infinite number of ways to interpolate the missing bond prices, let alone the bond prices for maturities that do not belong to the given tenor grid. In recent years, a fairly large body of literature has appeared devoted to the topic of yield curve construction, including [NS87], [Sve94], [Ron00], [And05], [HW06], and [HW08].

Possible interpolation schemes include using linear functions, cubic splines or tension splines applied to either (continuously compounded) spot rates, forward rates or logarithms of zero-coupon bond prices. It should be clear that the interpolation scheme must be carefully chosen as it will determine the regularity of different types of yield curves. Linearly interpolating continuously compounded spot rates, for example, is known to produce saw-tooth shaped forward-rate curves (see, e.g. [AP10a], p. 236). Apart from lacking any economically interpretation, an overly oscillating curve may even produce negative forward or spot rates. For the purpose of this thesis we will use cubic spline interpolation applied to log-discount prices combined with the Hyman cubic monotonic filter [Hym83]. The latter preserves the monotonicity of the input data and removes most of the unpleasant waviness of the resulting forward-rate curves, see also [AB09]. For examples of forward-rate curves that were constructed with this method see Figure 7.2 on page 129 below.

The choice of the securities used for bootstrapping the yield curve depends on the market under consideration and the liquidity of the available instruments. Typically, different types of interest-rate derivatives are used to infer the short-, medium- and long-term part of the yield curve. A common choice for constructing Libor based yield curves is, for instance, to use Libor deposits for maturities up to 1Y, followed by FRAs and Futures²⁰ covering the window up to 2Y-3Y. Swap rates are then used for the long end of the yield curve up to, say, 30Y or 60Y.

2.3 Fixed-Income Probability Measures

As we have seen in Section 2.1, choosing a martingale (or pricing) measure is essentially a matter of specifying a numeraire, which is a strictly positive asset price process used to renormalize all other security prices. In the following we will list some of the most common numeraires and pricing measures used when pricing fixedincome securities. Throughout this section we will assume that the market under consideration is complete and the corresponding martingale measures exist. Furthermore, we use V(t) to denote the time-t price of a contingent claim making at time T an \mathcal{F}_T -measurable payment of V(T).

²⁰Interest-rate futures are basically exchange-traded equivalents to the over-the-counter (OTC) FRAs. Due to a daily mark-to-market mechanism, the valuation of futures is, however, more involved than the valuation of FRAs, and recovering forward rates from futures rates requires a so-called convexity adjustment. See, e.g. [KN97], [JK05] and [PR06] for more details.

The Risk-Neutral Measure

The short rate r(t) can be used to define the continuously compounded money market account B(t), which satisfies the locally deterministic SDE

$$dB(t) = r(t)B(t)dt, \ B(0) = 1$$

with solution

$$B(t) = e^{\int_0^t r(s)ds}.$$

Taking the money market account as numeraire defines the risk-neutral measure \mathbb{Q} , under which the deflated price process V(t)/B(t) of a derivative security must be a martingale. This yields the standard risk-neutral valuation formula

$$\frac{V(t)}{B(t)} = \mathbb{E}_t^{\mathbb{Q}} \left[\frac{V(T)}{B(T)} \right], \qquad (2.11)$$

or equivalently,

$$V(t) = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-\int_t^T r(s)ds} V(T) \right].$$
(2.12)

In particular, if we set V(T) = 1, we obtain for the time-t price of a zero-coupon bond with maturity T

$$P(t,T) = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-\int_t^T r(s)ds} \right].$$
(2.13)

Observe that in the context of interest-rate derivatives, the payment V(T) will usually depend on realizations of zero-coupon bond prices²¹, which, in turn (through (2.13)) can be related to the evolution of r(t). Hence, all terms inside the expectation in (2.11) depend ultimately on the short rate r(t). It is therefore quite natural to consider the short rate as the basic quantity "driving" the interest-rate market. Directly specifying the dynamics of r(t) leads to so-called short-rate models, which we will examine in more detail in Section 2.7.

The T-Forward Measure

In most simple equity models, including e.g. the Black-Scholes [BS73] model, interest rates are assumed to be deterministic and the discount factor in (2.12) can therefore be pulled out of the expectation. In fixed-income pricing, however, discount factors are not only stochastic (otherwise pricing interest-rate products would be trivial), but they also exhibit a non-trivial dependence on the payoff V(T), making it often difficult to derive analytical formulas for the expectation in (2.12). Therefore, the risk-neutral measure is usually not the first choice when it comes to pricing interestrate derivatives. It is often more convenient, at least for certain simple interest-rate products, to work under the so-called *T*-forward measure \mathbb{Q}^T , associated with taking the *T*-maturity zero-coupon bond as numeraire²². In this case, the time-*t* price of a derivative security simplifies to

$$V(t) = P(t,T)\mathbb{E}_t^T \left[\frac{V(T)}{P(T,T)} \right]$$
$$= P(t,T)\mathbb{E}_t^T \left[V(T) \right],$$

²¹Libor rates, for example, being also just functions of zero-bond prices.

 $^{^{22}}$ Here we implicitly assume that zero-coupon bond prices are always strictly positive.

where in the last equation we have used that P(T,T) = 1.

The reason why the measure \mathbb{Q}^T is called the *T*-forward measure is justified by the following

Lemma 2.3.1. The forward Libor rate L(t, S, T) is a martingale under \mathbb{Q}^T , i.e.,

$$L(t, S, T) = \mathbb{E}_t^T \left[L(S, S, T) \right],$$

for $0 \leq t \leq S < T$.

Proof. By definition (see (2.5))

$$L(t, S, T) = \frac{1}{\tau} \left(\frac{P(t, S)}{P(t, T)} - 1 \right),$$

with the year fraction $\tau = T - S$. As P(t, S)/P(t, T) is a martingale under \mathbb{Q}^T , so is L(t, S, T).

The Spot Measure

When working with a set of forward rates defined on a discrete tenor structure $0 = T_0 < T_1 < T_2 < \ldots < T_N$, it is often convenient to use as numeraire the discrete-time equivalent of the continuously compounded money market account B(t). Specifically, we define

$$B_d(t) = \frac{P(t, T_{\eta(t)})}{\prod_{i=0}^{\eta(t)-1} P(T_i, T_{i+1})} = P(t, T_{\eta(t)}) \prod_{i=0}^{\eta(t)-1} (1 + \tau_i L_i(T_i)),$$

where

$$\eta(t) \coloneqq \min\{n \le N - 1 | t \le T_n\}.$$

$$(2.14)$$

Notice that $\eta(t) + 1$ is the index of the first forward rate $L_{\eta(t)+1}(t)$ that has not expired by time t. Intuitively, $B_d(t)$ is a bank account that is rebalanced only at times T_i in our discrete tenor structure. At t = 0 we start with 1 unit of currency and invest it in zero-coupon bonds with maturity T_1 , returning at time T_1 the amount

$$1/P(0,T_1) = 1 + \tau_0 L(0,0,T_1).$$

This amount is then reinvested ("rolled over") at time T_1 in zero-coupon bonds with maturity T_2 and so forth. The measure associated with taking B_d as numeraire is called the *spot measure* (or *spot Libor measure*), denoted \mathbb{Q}^{B_d} . Accordingly, for the time-t value of a contingent claim we have

$$V(t) = \mathbb{E}_t^{B_d} \left[V(T) \frac{B_d(t)}{B_d(T)} \right],$$

where

$$\frac{B_d(t)}{B_d(T)} = \frac{P(t, T_{\eta(t)})}{P(t, T_{\eta(T)})} \prod_{i=\eta(t)}^{\eta(T)-1} (1 + \tau_i L_i(T_i))^{-1}.$$

Swap Measures

Observe that the annuity $A_{m,n}(t)$ of a (forward) swap rate $S_{m,n}(t)$, spanning the time interval $[T_m, T_n]$, is simply a portfolio of zero-coupon bonds and hence qualifies as a numeraire. The corresponding measure $\mathbb{Q}^{m,n}$ is known as *swap*- or *annuity-measure*. For a security V maturing at time $T \leq T_m$ we obtain from the standard martingale pricing formula

$$V(t) = A_{m,n}(t) \mathbb{E}_t^{m,n} \left[\frac{V(T)}{A_{m,n}(T)} \right]$$

Earlier, we have seen that the Libor rate L(t, S, T) is a martingale under \mathbb{Q}^T . A similar result holds for a forward swap rate under its associated swap measure.

Lemma 2.3.2. The forward swap rate $S_{m,n}(t)$ is a martingale under $\mathbb{Q}^{m,n}$, *i.e.*,

$$S_{m,n}(t) = \mathbb{E}_t^{m,n} \left[S_{m,n}(T) \right].$$

Proof. By definition (see (2.9))

$$S_{m,n}(t) = \frac{P(t, T_m) - P(t, T_n)}{A_{m,n}(t)}$$

As $P(t,T_m)/A_{m,n}(t)$ and $P(t,T_n)/A_{m,n}(t)$ are both prices of tradeable assets expressed in units of the numeraire, they must be martingales. And so must be their difference.

2.4 Caps, Floors & Swaptions

In the following we introduce the two main (option-like) interest-rate derivatives in the fixed-income markets: Interest-rate caps/floors and swaptions.

Caps and Floors

Consider a firm with liabilities funded at the Libor rate, i.e., it has to pay at certain times T_{m+1}, \ldots, T_n the Libor rates resetting at times T_m, \ldots, T_{n-1} , with associated year fractions $\tau_m, \ldots, \tau_{n-1}$. Clearly, such a firm is naturally concerned with the possibility that future interest rates may rise, in which case it has to make higher interest-rate payments. One way to eliminate this risk is to enter into a payer swap, which, in effect, transforms the floating-rate payments into fixed ones, since at times T_{i+1} the firm then has to make the payments

$$\tau_i L_i(T_i) - \tau_i (L_i(T_i) - K) = \tau_i K,$$

where K denotes the fixed rate of the swap. While the firm is now protected against rising interest rates, transforming floating-rate payments into fixed ones also means, however, that the firm does not benefit from a potential future drop of the interestrate level. If the firm wishes to benefit from possibly lower rates in the future, yet wants to "cap" the future payments at a maximum rate K, it may enter into a socalled *interest-rate cap* (or just *cap*). A cap is a portfolio of (European) call options

2.4. Caps, Floors & Swaptions

on successive Libor rates, the so-called *caplets*, which pay off at time T_{i+1} , $i = m, \ldots, n-1$,

$$\tau_i \left(L_i(T_i) - K \right)^+.$$

Note that, by the definition of $L_i(T_i)$, we may write such a payoff as

$$\tau_i(L_i(T_i) - K)^+ = \frac{\tau_i(1+K)}{P(T_i, T_{i+1})} \left(\frac{1}{\tau_i(1+K)} - P(T_i, T_{i+1})\right)^+.$$

Thus, the *i*-th caplet is equivalent to a multiple of a put option with maturity T_i , strike $\tilde{K} \coloneqq \tau_i(1+K)$ and written on the zero-coupon bond with maturity T_{i+1} .

The counterpart of an interest-rate cap is an *interest-rate floor* (or just *floor*), which is a portfolio of put options on successive Libor rates, called *floorlets*, with time- T_{i+1} payoffs

$$\tau_i \left(K - L_i(T_i) \right)^+.$$

An investor with assets earning a floating rate can use a floor to protect herself against low interest-rate scenarios, while still being able to benefit from rising interest rates. Analogously to caplets, floorlets may be interpreted as call options on zero-coupon bonds.

Applying the risk-neutral valuation formula (2.11) yields for the time-t value of a unit-notional cap/floor (covering the time interval $[T_m, T_n]$)

$$V_{\rm cap}(t) = B(t) \sum_{i=m}^{n-1} \tau_i \mathbb{E}_t^{\mathbb{Q}} \left[B(T_{i+1})^{-1} \left(L_i(T_i) - K \right)^+ \right],$$

$$V_{\rm floor}(t) = B(t) \sum_{i=m}^{n-1} \tau_i \mathbb{E}_t^{\mathbb{Q}} \left[B(T_{i+1})^{-1} \left(K - L_i(T_i) \right)^+ \right].$$

The value of the *i*-th caplet/floorlet may be written in a more convenient form by switching to the T_{i+1} -forward measure $\mathbb{Q}^{T_{i+1}}$, defined by using the T_{i+1} -maturity zero-coupon bond as numeraire asset. By performing this measure change we remove the discounting term inside the expectation, and the value of the cap, for instance, then simplifies to

$$V_{\rm cap}(t) = \sum_{i=m}^{n-1} \tau_i P(t, T_{i+1}) \mathbb{E}_t^{T_{i+1}} \left[(L_i(T_i) - K)^+ \right],$$

where the terms inside the expectations now solely depend on the respective Libor rates. In order to price caplets and floorlets it is therefore sufficient to separately model the Libor rates under their respective forward measure. Moreover, recall that by Lemma 2.3.1, $L_i(t)$ must be a martingale under $\mathbb{Q}^{T_{i+1}}$, and thus the drift term of the stochastic process must be zero. Assuming that $L_i(t)$ evolves according to a driftless geometric Brownian motion

$$dL_i(t) = \sigma_i L_i(t) dW^i(t), \qquad (2.15)$$

or equivalently,

$$L_i(t) = L_i(0) \exp\left(-\frac{\sigma_i^2}{2}t + \sigma_i W^i(t)\right), \ 0 \le t \le T_i,$$

where $\sigma_i > 0$ and $W^i(t)$ is a one-dimensional $\mathbb{Q}^{T_{i+1}}$ -Brownian motion, we have that $L_i(T)$ is log-normally distributed. A straightforward evaluation of the corresponding expectation leads to the well-known *Black formula* [Bla76] for the price of the *i*-th caplet

$$V_{\text{caplet}}^{i}(t) = \tau_{i} P(t, T_{i+1}) \text{Bl}\big(K, L_{i}(t), \sigma_{i} \sqrt{T_{i} - t}, 1\big), \ 0 \le t \le T_{i},$$

where, denoting by $\Phi(\cdot)$ the Gaussian cumulative distribution function,

$$Bl(K, L, v, w) = wL\Phi(wd_{+}(K, L, v)) - wK\Phi(wd_{-}(K, L, v)), \qquad (2.16)$$
$$d_{\pm}(K, L, v) = \frac{\log(L/K) \pm v^{2}/2}{v}.$$

Accordingly, the time-t price of the cap/floor is given by a sum of Black formulas

$$V(t) = \sum_{i=m}^{n-1} \tau_i P(t, T_{i+1}) \text{Bl}(K, L_i(t), \sigma_i \sqrt{T_i - t}, w), \ 0 \le t \le T_m,$$

where w = 1 (cap) or w = -1 (floor). It is common market practice to quote the value of a cap or a floor not in terms of its price but rather in terms of a single volatility parameter $\sigma_{m,n}$, the so-called *cap/floor implied volatility*, such that (at time 0)

$$V(0) = \sum_{i=m}^{n-1} \tau_i P(t, T_{i+1}) \text{Bl}(K, L_i(0), \sigma_{m,n} \sqrt{T_i}, w).$$

Caps and floors are among the most liquidly traded interest-rate derivatives in fixed-income markets and the corresponding implied volatilities are quoted in the market for several standard maturities. An example of an at-the-money (ATM)²³ market cap volatility curve is shown in Figure 2.1, where volatilities $\sigma_{m,n}$ are plotted against option maturities T_n^{24} .

Although single caplets/floorlets are not traded, the volatility information for individual forward Libor rates can, at least in principle, be bootstrapped from caps/ floors of different maturities. This so-called volatility bootstrapping is, however, by no means trivial and we refer the reader to [AP10c], Section 16.2, for more details on this topic. Once extracted, the Libor rate volatilities can be used as market inputs when calibrating interest-rate models for pricing other, more complex products.

Remark 2.4.1. Through a deterministic time change (see [AP10a], p. 300) the standard Black formula can be easily extended to allow for time-dependent (deterministic) volatilities $\sigma = \sigma(t)$ of the underlying process. For European vanilla options only the integrated instantaneous variance of the underlying is relevant. Hence, in case of a time-dependent volatility we simply need to replace

$$v = \sigma \sqrt{T - t}$$

²³A cap (resp. floor) with payment times T_{m+1}, \ldots, T_n is said to be at-the-money if $K = S_{m,n}(0)$ (cp. [BM05], p. 18).

²⁴In the euro market T_m is typically equal to 3 months for caps with maturities $T_n=1$, 1.5 and 2 years and with all other T_i 's equally three-month spaced. For caps with maturities larger than 2 years, the first fixing date T_m is typically equal to 6 months with all other T_i 's equally six-month spaced.



Figure 2.1: ATM euro cap implied volatilities $\sigma_{m,n}$ for maturities $T_n = 1$, 1.5, 2, 3, 4, 5, 7, 10, 15, 20 and 30 years.

with

$$v = \sqrt{\int_t^T \sigma^2(s) ds}$$

in the Black formula.

Swaptions

The second class of liquidly traded interest-rate derivatives are *swap options* or more commonly *swaptions*, which are European options on interest-rate swaps. More specifically, the owner of a *payer swaption* has the right (but not the obligation) to enter at the swaption maturity²⁵ into a payer swap with a prespecified fixed rate K. In contrast, a *receiver swaption* is an option to enter into a receiver swap. If the swap, underlying a swaption, has first reset date T_m and last payment date T_n , then the length of the swap $T_n - T_m$ is usually called the tenor of the swaption. If we consider a payer swaption written on a swap covering the interval $[T_m, T_n]$, then we have for the time- T_m payoff of this swaption

$$V_{\rm swptn}(T_m) = (V_{\rm swap}(T_m))^+ = \left(\sum_{k=m}^{n-1} \tau_k P(T_m, T_{k+1})(L_k(T_m) - K)\right)^+$$

Applying the risk-neutral valuation formula yields for the price of the swaption at time $t \leq T_m$

$$V_{\text{swptn}}(t) = B(t)\mathbb{E}_{t}^{\mathbb{Q}}\left[B(T_{m})^{-1}\left(\sum_{k=m}^{n-1}\tau_{k}P(T_{m},T_{k+1})(L_{k}(T_{m})-K)\right)^{+}\right].$$

Observe that contrary to the cap/floor case, this payoff cannot be decomposed into a sum of simpler products. In particular, the payoff does not only depend on the evolution of the individual Libor rates (as was the case with caplets/floorlets),

 $^{^{25}}$ We will consider in the following only the most common case, where the maturity of the option coincides with the first reset date of the underlying swap.

but it also depends on the *joint* behavior of the rates. We will return to this topic in later chapters.

Using formula (2.7) for the swap value at time T_m , we obtain the following alternative formulation for the time-t price of the (payer) swaption

$$V_{\text{swptn}}(t) = B(t)\mathbb{E}_{t}^{\mathbb{Q}}\left[B(T_{m})^{-1}A_{m,n}(T_{m})\left(S_{m,n}(T_{m})-K\right)^{+}\right].$$

Finally, switching to the swap measure we can rewrite this formula in the more compact form

$$V_{\text{swptn}}(t) = A_{m,n}(t)\mathbb{E}_{t}^{m,n}\left[\left(S_{m,n}(T_{m}) - K\right)^{+}\right].$$
(2.17)

From this it is easy to see that (under the right measure) a payer swaption is basically a call option on the forward swap rate, with the strike K being equal to the fixed rate of the swap. Similarly, a receiver swaption can be interpreted as a put option on the forward swap rate.

From Section 2.3 we have that $S_{m,n}(t)$ must be a martingale under the swap measure $\mathbb{Q}^{m,n}$. Assuming that the swap rate follows a geometric Brownian motion yields again a Black-type formula

$$V_{\text{swptn}}(t) = A_{m,n}(t) \text{Bl}(K, S_{m,n}(t), \sigma_{m,n}\sqrt{T_m - t}, w), \ 0 \le t \le T_m,$$
(2.18)

where $\sigma_{m,n}$ now denotes the swap-rate volatility and where w = 1 for a payer swaption and w = -1 for a receiver swaption. As with caps/floors it is common market practice to express market prices of swaptions in terms of Black-implied volatilities, i.e., volatilities $\sigma_{m,n}$ that recover the market price when plugged into the Black formula²⁶. Swaptions are very liquid for a large number of maturity-tenor-combinations. Figure 2.2 shows implied ATM²⁷ volatilities as quoted in the euro market²⁸ on 1/14/2008.

2.5 Vanilla Models

2.5.1 Black's Model and the Volatility Smile

The options introduced in the last section depend only on a single forward rate, and, provided we choose the *natural martingale measure*, it is therefore sufficient to have a model for the evolution of this particular rate. We adopt the terminology of [AP10a] and shall call models of this type *vanilla models* to distinguish them from full termstructure models, which we will encounter in later sections. The most prominent example for a vanilla model is the Black model, which assumes that the forward rate under consideration, say S(t), follows a geometric Brownian motion. Accordingly, for some fixed time T, S(T) is log-normally distributed and prices for European call and put options are given by so-called Black formulas. Although it is standard market

 $^{^{26}}$ Note, that this can be done independently of whether or not one believes that the Black model is a realistic model. One simply uses the Black formula to express option prices in a different (often more intuitive) coordinate system.

²⁷A swaption at time t = 0 (today) is said to be at-the-money if the strike price equals the forward swap rate, i.e., if $K = S_{m,n}(0)$.

²⁸In the euro market swaptions are typically cash-settled and the annuity factor $A_{m,n}(t)$ in (2.18) must be replaced with a so-called cash-settled annuity factor to obtain the corresponding standard market formula, see [Mer08].



Figure 2.2: ATM euro swaption volatilities as quoted on 1/14/2008.

practice to express option prices in terms of Black implied volatilities, the Black model (or the Black-Scholes model [BS73] in an equity context) is not consistent with market prices in the sense that for a given maturity, vanilla options written on the same underlying but with different strike prices require different volatilities to be plugged into the Black formula. In Figure 2.3 we present the implied volatilities of swaptions with a maturity of 10 years and a tenor of 5 year (usually simply called "10y-into-5y" or "10x5" swaptions) for different strike prices K.



Figure 2.3: Implied volatility smile for 10x5 EUR swaptions on 4/26/2010. The ATM-strike is $K_{ATM} = 4.45\%$.

The mapping $K \mapsto \sigma_{imp}(K)$ is commonly known as the volatility smile, or, in case the smile is predominantly downward sloping, as volatility skew. Clearly, if market prices were consistent with Black's model, then the smile would be flat, since the implied volatility should not depend on the strike K. There are several conventional explanations²⁹ for the volatility smile/skew and some key words are supply and demand effects, big downward jumps, the leverage effect or anticorrelated volatility moves; see e.g. [KG10]. We will simply interpret the volatility smile in

²⁹Some of these explanations originate from the equity markets and cannot be directly carried over to an interest-rate setting. The "leverage effect", for example, explains an increase of volatility for falling stock prices by an increase of the debt-to-equity ratio, i.e., the leverage of the firm.

terms of the market implied terminal distribution of the forward rate. To this end let C(t, S(t), T, K) denote the undiscounted time-t price of a European call option with maturity T and underlying S(t)

$$C(t, S(t), T, K) = \mathbb{E}_t \left[(S(T) - K)^+ \right].$$

Assuming that the density f_S of S(T) given S(t) exists, this may be rewritten as

$$C(t, S(t), T, K) = \int_{K}^{\infty} (x - K) f_S(x, T; t, S(t)) dx,$$

and differentiating twice with respect to K finally gives

$$\frac{\partial^2}{\partial K^2} C(t, S(t), T, K) \Big|_{K=x} = f_S(x, T; t, S(t)).$$
(2.19)

This classical result, which is due to Breeden & Litzenberger [BL78], allows us (at least theoretically) to recover the market-implied density from a continuum of call/put prices. If market prices are not compatible with a single volatility parameter, then this simply means that the market-implied density is not a log-normal density as postulated by Black's model. In case of a proper (i.e., U-shaped) volatility smile, where out-of-the-money (OTM)³⁰ implied volatilities are larger than the ATM volatility, the market-implied distribution for $\log(S(T))$ has fatter tails (or a higher kurtosis) than a normal distribution. A predominantly downward sloping volatility smile (or skew), on the other hand, corresponds to a left-skewed distribution for $\log(S(T))$.

In general, market observed implied volatilities do not only depend on the strike but also on the maturity of the option, and the mapping $(T, K) \mapsto \sigma(T, K)$ is commonly referred to as the *(implied) volatility surface*. In this regard, however, it is important to differentiate between equity and fixed-income markets. In the equity case one usually means by "volatility surface" implied volatilities corresponding to options with different strike prices and different maturities, but written on the same stock (and therefore connected to the same stochastic process). Also in fixed-income markets one can observe, for instance, implied caplet volatilities for different strikes and maturities, and one therefore often speaks of the *caplet volatility surface*³¹. Note, however, that caplets with different maturities are options written on different forward rates, i.e., written on different underlyings. Put differently, for a given forward Libor or swap rate we can usually observe at most one single volatility smile in the market.

The obvious question is now of course: Are there alternative models whose dynamics are compatible with the market implied distributions? In the following we briefly review some of the most popular diffusive models that can account for volatility smiles or skews.

 $^{^{30}}$ A call option is said to be out of the money (in-the-money (ITM)) if the forward price is less (greater) than the strike price. Conversely, a put option is said to be out-of-the-money (in-the-money) if the forward price is greater (less) than the strike price.

³¹In case of swaptions, implied volatilities are not only indexed by strike and maturity, but also by the tenor of the underlying swap and one therefore speaks of the *swaption volatility cube*.
2.5.2 Local Volatility Models

In this section we consider models where the forward rate S(t) under its natural martingale measure is assumed to satisfy a one-dimensional SDE of the form

$$dS(t) = \varphi(t, S(t))dW(t), \qquad (2.20)$$

with a state-dependent diffusion coefficient $\varphi:[0,\infty)\times\mathbb{R}\to\mathbb{R}$ known as *local volatility* function. Models that fit into this framework were first³² proposed by Cox [Cox96], Cox & Ross [CR76] and Rubinstein [Rub83], although it were the works³³ by Dupire [Dup94] and Derman & Kani [DK94] that coined the term "local volatility".

Dupire noted that there is a unique diffusion coefficient $\sigma(t, s)$ that is consistent with a continuum (in maturity and strike dimension) of European call option prices observed at time t = 0, namely

$$\varphi^2(T,K) = \frac{\frac{\partial}{\partial T}C(0,S(0),T,K)}{\frac{1}{2}K^2\frac{\partial^2}{\partial K^2}C(0,S(0),T,K)}.$$

Computing this local volatility function in practice, however, is by no means trivial and requires a smooth and arbitrage-free interpolation scheme for interpolating the finitely many (often noisy) implied volatilities observed in the market. Apart from that, local volatility functions calculated in this way usually come out as being strongly level-dependent and non-monotonic in S, and the resulting model will imply non-stationary volatility smile dynamics, which is at odds with typical market behavior; see the discussions in Chapter 12 of [Reb04] and Section 7.1 of [AP10a]. It is therefore advisable to use only local volatility models with relatively simple and monotonic functions $\varphi(t, s)$, even though it is then usually not possible to perfectly fit observed market volatility smiles. Two of the most popular monotonic choices for $\varphi(t, s)$ will be discussed in the following.

The CEV Model

There is empirical evidence³⁴ that the *absolute* volatility level of interest rates (under the real measure) generally scales less than proportionally with the level of interest rates. Consequently, when interest rates fall, *percentage* volatilities tend to increase. This observation can, at least heuristically, be taken as a possible explanation for the existence of volatility skews in the interest-rate markets, with high implied volatilities for low-strike options. In order to obtain a less than proportional volatility scaling, one can choose the following local volatility function

$$\varphi(t,s) = \lambda s^{\beta},$$

where $0 < \beta < 1$ and $\lambda > 0$. The corresponding model is known as the *constant*elasticity-of-variance (CEV) model and was first considered by Cox [Cox96].

³²Note, however, that also the Black model trivially fits into this framework with $\sigma(t, x) = \sigma x$.

 $^{^{33}\}mathrm{Dupire}$ developed the continuous-time theory while Derman & Kani considered a discrete-time binomial tree setting.

 $^{^{34}}$ See [Reb03] and Section 8.5 of [RMW09].

Compared to Black's model, the CEV model clearly has one additional parameter β , which can be used to obtain a better fit to the observed market prices of options. Observe, that in the limiting case $\beta = 1$, $S(\cdot)$ is a log-normal process as in the Black model, and hence the corresponding volatility smile is perfectly flat. Contrary, if $\beta = 0$ we have that $S(\cdot)$ is a Gaussian process, and the model is then called the *Normal* or *Bachelier model*. The volatility smile generated by this model will be steeply downward sloping, as the process has a constant (absolute) volatility term. Undiscounted call option prices³⁵ within the Bachelier model are straightforward to compute and given by the following formula

$$C_{\text{Bach}}(t, S(t), K, T) \coloneqq \mathbb{E}_t \left[(S(T) - K)^+ \right]$$
$$= (S(t) - K) \Phi(d) + \lambda \sqrt{T - t} \phi(d), \qquad (2.21)$$

with

$$d = \frac{S(t) - K}{\lambda \sqrt{T - t}},$$

and where $\Phi(\cdot)$ and $\phi(\cdot)$ denote the standard normal cumulative distribution function (CDF) and the probability density function (PDF), respectively.

For values of β between zero and one we have the following proposition (cp. [AA00]):

Proposition 2.5.1. For the stochastic differential equation

$$dS(t) = \lambda S(t)^{\beta} dW(t), \qquad (2.22)$$

with $0 < \beta < 1$ and $\lambda > 0$ the following holds:

- 1. All solutions to (2.22) are non-explosive.
- 2. For $\beta \geq 1/2$ the SDE (2.22) has a unique solution.
- 3. S = 0 is an attainable boundary for the process (2.22).
- 4. For $\beta < 1/2$ the SDE (2.22) does not have a unique solution, unless a separate boundary condition is specified for the boundary behavior in S = 0.

In order to obtain a unique solution for the case $\beta < 1/2$, S = 0 is usually defined to be an absorbing barrier, i.e., if the solution ever hits zero it stays there. This condition is not only consistent with the case $\beta \ge 1/2$ (for which S = 0 is a natural absorbing boundary), but it is also the only boundary condition that allows no arbitrage opportunities.

The CEV model is quite popular in mathematical finance since it is complete and analytically tractable; see [Sch89], [DL01] and [DS02]. In particular, the transition density can be expressed explicitly in terms of modified Bessel functions. Consequently, prices of European-type option can be computed, at least in principle, by integrating the payoff function against this transition density. In case of vanilla options, this results in pricing formulas involving the cumulative distribution functions of non-central χ^2 -distributed random variables.

³⁵As usual, put option prices can be obtained via put-call parity.

2.5. Vanilla Models

The above specification of the CEV model can be easily extended to allow for a time-dependent volatility parameter $\lambda = \lambda(t)$; see [AA00]. One drawback of the CEV model is that S(t) has a positive probability of being absorbed at zero, which is at least from an empirical point of view not desirable.

Displaced Diffusions

The displaced-diffusion process (or shifted log-normal process) is obtained by setting

$$\varphi(t,s) = \lambda s + \alpha, \ \alpha \in \mathbb{R}, \ \lambda > 0, \tag{2.23}$$

and was first introduced by Rubinstein [Rub83]. Throughout this thesis we will use the following, often more convenient form

$$\varphi(t,s) = \lambda(\beta s + (1-\beta)S(0)), \ \beta \in [0,1], \ \lambda > 0,$$
(2.24)

which is obtained from (2.23) by a simple reparameterization. In order to get an intuitive feel for the properties of displaced diffusions, consider the corresponding SDE

$$dS(t) = \lambda(\beta S(t) + (1 - \beta)S(0))dW(t).$$
(2.25)

For large values of S(t), the diffusion coefficient scales approximately proportionally with the level of S(t); we are in a log-normal setting. For small (absolute) values of S(t), we have

$$dS(t) \approx \lambda (1 - \beta) S(0) dW(t),$$

that is, the process is approximately normal. Roughly speaking, we can therefore consider a displaced diffusion as a (dynamic) mixture of a log-normal and a normal process. As with the CEV process, normal ($\beta = 0$) and log-normal ($\beta = 1$) processes may be seen as special cases of displaced-diffusion processes.

Another interpretation of displaced-diffusion processes can be given in terms of a transformation of the process. For $\beta \in (0, 1]$ we may define $X(t) := (\beta S(t) + (1 - \beta)S(0))$ and obtain by Itô's formula

$$dX(t) = \beta \lambda X(t) dW(t),$$

that is, a linear transform of S(t) rather than S(t) itself follows a log-normal process. Call option prices in the displaced-diffusion model are given by the following

Proposition 2.5.2. Consider the displaced diffusion

$$dS(t) = \lambda(\beta S(t) + (1 - \beta)S(0))dW(t), \ S(0) > 0,$$

where $\lambda > 0$ and $\beta \in (0,1]$. Define $\alpha := (1 - \beta)S(0)/\beta$ and assume $S(t), K > -\alpha$. Then we have

$$C_{DD}(t, S(t), K, T) \coloneqq \mathbb{E}_t \left[(S(T) - K)^+ \right]$$

= Bl(K + \alpha, S(t) + \alpha, \lambda\beta\sqrt{T-t}, 1), (2.26)

where Bl(K, S, v, 1) is given in (2.16).

Proof. Follows directly from the standard Black formula after recalling from above that $X(t) = \beta(S(t) + \alpha)$ is a log-normal process.

Using the parameterization (2.24) instead of (2.23) has the advantage that the parameters act almost "orthogonal" on the shape of the volatility smile. While β provides control over the skew, i.e., the slope of the volatility smile, σ determines the overall volatility level. Possible shapes of volatility smiles that can be generated with a displaced-diffusion model are depicted in Figure (2.4), where we have plotted volatility smiles for β values ranging from 0.01 to 1. As noted above, setting $\beta = 1$ yields the Black model, in which case the volatility smile is perfectly flat.



Figure 2.4: Left: Displaced-diffusion volatility smiles for $\beta = 0.01, 0.2, 0.4, 0.6, 0.8, 1, S(0) = 4.45\%, T = 5 and <math>\sigma = 15.4\%$. Right: CEV and displaced-diffusion volatility smiles for $\beta = 0.4, \sigma_{DD} = 15.4\%, \sigma_{CEV} = 2.38\%$. S(0) and T as before.

Option prices produced by CEV and displaced-diffusion dynamics are often remarkably similar. This can be seen on the right hand side of Figure (2.4), which displays volatility smiles for $\beta_{\text{CEV}} = \beta_{\text{DD}} = 0.4$. The volatility parameter of the CEV process was chosen such that the ATM volatilities match. A partial explanation for the similarity between the CEV and the displaced-diffusion model can be given by expanding the CEV local volatility function $\varphi(s) = \lambda s^{\beta}$ about the ATM level s = S(0), which yields

$$\varphi(S(t)) \approx \varphi(S(0)) + \varphi'(S(t))(S(t) - S(0))$$
$$= \lambda S(0)^{\beta - 1} (\beta S(t) + (1 - \beta)S(0)).$$

Hence, choosing $\lambda_{\text{DD}} = \lambda_{\text{CEV}} S(0)^{\beta-1}$, the displaced-diffusion process can be seen as an approximation to the CEV process (with the same skew parameter β). For more details on the quality of this approximation and for a quantification of the similarity of option prices associated with the two dynamics see [SG09].

One clear drawback of displaced diffusions is that for $\beta < 1$ (which is typically the case) S(t) can become negative. More precisely, for $\beta \in (0, 1)$ the process S(t)can assume values in $(-\alpha, \infty)^{36}$. For not too low forward-rate levels, the probability of negative rates is often small enough to ignore³⁷, and prices of vanilla options are

³⁶With $\alpha > 0$ given in Proposition 2.5.2.

 $^{^{37}}$ For the parameter scenarios considered in Figure 2.4, the probability is in all cases less than 0.5%.

remarkably close to prices generated by a CEV model, as we have seen above. Nevertheless, as argued in Rogers [Rog96], there exist certain interest-rate products which are highly sensitive to even very low probabilities of negative rates. For pricing such products one should obviously avoid models that allow negative rates. Additionally, for very low interest-rate environments and/or high percentage volatilities, the probability of negative rates may become large enough to even have an impact on vanilla option prices. Also in this case, displaced-diffusion models should be used with care.

Despite the drawback of allowing for negative rates, the displaced-diffusion model is computationally far more tractable than the CEV model, since the standard "lognormal calculus" known from the Black model can be applied³⁸. And, as noted above, for not-too-extreme interest-rate- and volatility-scenarios, displaced diffusions are (very) accurate approximations of CEV processes. So even if one ultimately wants to use the financially more appealing CEV model³⁹, the displaced-diffusion approximation can nevertheless be used as a tool for efficiently calibrating such a model.

In the following we will consider an extension of the displaced-diffusion model, where we let the local volatility function φ explicitly depend on calender time t,

$$\varphi(t,s) = \lambda(t) \big(\beta(t)s + (1 - \beta(t))S(0)\big). \tag{2.27}$$

By itself, European option pricing generally does not require time-dependent skew or volatility parameters, as only the terminal distribution of the underlying is relevant. However, time-dependent coefficients will later often arise when we model the dynamics of interest rates in full term-structure models.

First assume, that only the volatility coefficient $\lambda = \lambda(t)$ depends on time. In this case we may simply invoke again the time-change argument mentioned in Remark 2.4.1 to obtain a call option pricing formula that is identical to (2.26), except that we must replace $v = \lambda \beta \sqrt{T-t}$

with

$$v = \beta \sqrt{\int_t^T \lambda^2(s) ds}.$$

Matters become more involved if we also let the skew parameter β depend on time. For a general skew function $\beta(t)$ there do not exist closed-form option pricing formulas. Although option prices could in principle be computed by using PDE⁴⁰ methods without much difficulty, this is generally too slow or too inaccurate for calibration applications. In the following we therefore give a brief review of the so-called parameter-averaging technique, pioneered by Piterbarg [Pit05a], [Pit05b].

The basic idea of parameter averaging is to find a model with *time-independent* parameters, that yields European option prices approximately matching the prices from the time-dependent model. We have already seen a form of parameter averaging when we considered the Black model or the displaced-diffusion model with

³⁸Calculating vanilla option prices in the CEV model, on the other hand, requires the evaluation of a non-central χ^2 -distribution, which is computationally fairly demanding.

³⁹In the sense that it guarantees S(t) to remain non-negative.

⁴⁰Partial differential equation.

time-dependent volatility functions. From above we have, that prices of T-maturity European options in the model

$$dS(t) = \lambda(t) \big(\beta S(t) + (1 - \beta)S(0)\big) dW(t)$$

are the same as in the model

$$dS(t) = \bar{\lambda} \big(\beta S(t) + (1 - \beta) S(0)\big) dW(t), \qquad (2.28)$$

where $\bar{\lambda}$ is given by

$$\bar{\lambda}^2 = \frac{1}{T} \int_0^T \lambda(s)^2 ds.$$
(2.29)

Observe that in this case the option pricing formula based on the *effective volatility* $\bar{\lambda}$ (for maturity T) is not an approximation, but yields the exact price.

We now focus on finding an *effective skew* $\bar{\beta}$. First recall that European prices depend only on the terminal distribution of the underlying. So if we consider the two processes

$$dS(t) = \lambda(t) \big(\beta(t)S(t) + (1 - \beta(t))S(0)\big) dW(t),$$

$$d\bar{S}(t) = \lambda(t) \big(\bar{\beta}S(t) + (1 - \bar{\beta})S(0)\big) dW(t),$$

we could try to find $\bar{\beta}$ such that for a given maturity T

$$\mathbb{E}\left[(S(T) - \bar{S}(T))^2\right] \to \min.$$

While the exact solution to this problem is generally not known in closed analytical form, Piterbarg [Pit05a] uses asymptotic expansion techniques to find the following approximate solution for the effective skew (over the time horizon [0, T])

$$\bar{\beta} = \int_0^T \beta(t) w_T(t) dt, \qquad (2.30)$$

with weight function $w_T(\cdot)$ given by

$$w_T(t) = \frac{v(t)^2 \lambda(t)^2}{\int_0^T v(t)^2 \lambda(t)^2 dt}, \quad v(t)^2 = \int_0^t \lambda(s)^2 ds.$$
(2.31)

With the effective volatility and the effective skew, vanilla option prices can now be computed efficiently by using the constant-parameter formula given in (2.26). Test results, which demonstrate that the above approximation is very accurate even for quite long maturities, can be found in [Pit05a].

2.5.3 Stochastic Volatility Models

In the previous section we introduced vanilla models where the instantaneous volatility is a deterministic function of the underlying interest rate. With such models it is possible to account for deviations from a log-normal behavior of interest rates, and generate downward sloping volatility skews. Comparing, however, the volatility smiles shown in Figure 2.4 and 2.3, we find that the smiles implied by the CEV or displaced-diffusion model are generally too "linear" compared to the more convex market-implied volatility smile. It therefore seems necessary to somehow enrich our model setup to be able to achieve better market fits.

A number of empirical studies provide evidence that besides the level-dependence of interest-rate volatility, there are additional sources of randomness affecting volatilities in fixed-income markets; see, e.g. [CDG02], [HW03], [CS04] and [LZ06] to just name a few. It therefore seems natural to allow the volatility to be driven by an additional Brownian motion. The resulting models are known as *stochastic volatility models* and generally not only provide more realistic dynamics⁴¹, but also give better fits to market-observed implied volatilities. These advantages are bought at the cost of completeness: Stochastic volatility models are generally not complete anymore, that is, contingent claims cannot be replicated or hedged by trading in the underlying alone. In this case, either additional hedging instruments must be added to the hedging portfolio or one must resort to minimum variance hedging or similar techniques⁴².

As in the previous section we consider in what follows a forward swap or Libor rate S(t) under its natural martingale measure, say \mathbb{Q} . Further, we let W(t) and Z(t) be two one-dimensional Brownian motions under \mathbb{Q} with correlation $\rho \in (0, 1)$

$$\langle dZ(t), dW(t) \rangle = \rho dt.$$

A fairly general family of stochastic volatility models is obtained by specifying the following dynamics

$$dS(t) = \lambda \varphi(S(t))\psi(V(t))dW(t), \qquad (2.32)$$

$$dV(t) = \kappa(\theta - V(t))dt + \xi \vartheta(V(t))dZ(t), V(0) = v_0,$$

where $\lambda, \kappa, \theta, \xi$ are positive constants and $\varphi(\cdot), \psi(\cdot)$ and $\vartheta(\cdot)$ are smooth deterministic functions.

Classical references investigating specific models that fit into this framework are for example Hull & White [HW87], Wiggins [Wig87], Stein & Stein [SS91], Chesney & Scott [CS89] or Hagan et. al [HKLW02].

The Heston Model

One of the most widely used stochastic volatility model is due to Heston [Hes93] and corresponds to choosing

$$dS(t) = S(t)\sqrt{V(t)}dW(t), \qquad (2.33)$$

$$dV(t) = \kappa(\theta - V(t))dt + \xi\sqrt{V(t)}dZ(t), \qquad (2.34)$$

with $\kappa, \theta, \xi > 0$. Here, the stochastic-variance process V(t) follows a mean reverting square-root diffusion, also known as Cox-Ingersoll-Ross (CIR) process, see [CIR85].

⁴¹This is important for hedging purposes. Note, however, that we must be careful when we interpret properties under the real measure and then move to a pricing measure. Measure changes generally affect the dynamic properties of processes, see also Chapter 3. 42 Conter 4. [DEUE00]

⁴²See, e.g. [PSHE09].

The mean-reversion level θ determines the long-term mean of the variance process, while the speed of mean reversion κ represents the rate at which V(t) is pulled back to this long-term mean. The parameter ξ is the volatility of variance.

Observe that the diffusion term of V(t) in (2.34) is the same as that of a CEV process with exponent $\beta = 1/2$. Opposed to a CEV process, however, the drift term in (2.34) prevents the CIR process from being absorbed whenever it hits zero. If the upward drift is "sufficiently large", then the origin is even inaccessible, as the following well-known result shows

Proposition 2.5.3. The SDE (2.34) has a unique solution. If $2\kappa\theta \ge \xi^2$, i.e., the socalled Feller condition holds, then V = 0 is unattainable. If $2\kappa\theta < \xi^2$ then V = 0 is an attainable boundary but it is strongly reflecting.

Proof. See, e.g. [KT81].

Similar to the CEV case, the conditional distribution of V(t) is linked to a noncentral χ^2 -distribution:

Proposition 2.5.4. Let $0 \le s < t$ and define

$$c = \frac{2\kappa}{\xi^2 \left(1 - e^{-\kappa(t-s)}\right)}, \ u = cV(s)e^{-\kappa(t-s)}, \ q = \frac{2\kappa\theta}{\xi^2} - 1.$$

Then, the transition density respectively the CDF of V(t) given V(s) are given by

$$f(x,t;V(s),s) = ce^{-u-cx} \left(\frac{cx}{u}\right)^{q/2} I_q\left(2\sqrt{ucx}\right)$$

and

$$\mathbb{P}(V(t) < x | V(s)) = \chi^2(2cx; 2q + 2, 2u),$$

where $I_q(\cdot)$ denotes the modified Bessel function of the first kind of order q and $\chi^2(z;\nu,\gamma)$ is the CDF of a non-central χ^2 distribution with ν degrees of freedom and non-centrality parameter γ ,

$$\chi^2(z;\nu,\gamma) = e^{-\gamma/2} \sum_{j=0}^{\infty} \frac{(\gamma/2)^j}{j! \, 2^{j+\nu/2} \Gamma(j+\nu/2)} \int_0^z y^{j+\nu/2-1} e^{-y/2} dy.$$

Proof. See [CIR85].

The effects that the stochastic-variance parameters have on the model-implied volatility smile can be summarized as follows. The mean reversion level θ is responsible for the overall level of the implied volatility smile, while κ and ξ provide control over the curvature or convexity of the smile. The skewness of the volatility smile is determined by the correlation ρ . In principle, the initial value v_0 of the unobservable stochastic-variance process represents another "free parameter", which we could use to improve the fit to market prices. In most cases, however, one is best advised to simply set $v_0 = \theta$; see Section 13.7 of [Reb04] for a discussion of this topic.

One reason why the Heston model is so popular among academics and practitioners alike is its analytical tractability. In particular, the characteristic function for the logarithm of the underlying is known in closed form (see [AMST07] and Lee [Lee04]):

Proposition 2.5.5. The characteristic function of $\log S(t)$ is given by

$$\phi(u;t) \coloneqq \mathbb{E}\left[e^{iu\log S(t)}\right] = \exp\left\{iu\log S(0) + A(u,t) + B(u,t)v_0\right\},\tag{2.35}$$

where

$$\begin{split} A(u,t) &= \frac{\kappa\theta}{\xi^2} \left((\kappa - \rho\xi iu - \gamma)t - 2\log\left(\frac{1 - ge^{-\gamma t}}{1 - g}\right) \right), \\ B(u,t) &= \frac{\kappa - \rho\xi iu - \gamma}{\xi^2} \left(\frac{1 - e^{-\gamma t}}{1 - ge^{-\gamma t}}\right), \\ g &= g(u) = \frac{\kappa - \rho\xi iu - \gamma}{\kappa - \rho\xi iu + \gamma}, \\ \gamma &= \gamma(u) = \sqrt{(\rho\xi iu - \kappa)^2 + \xi^2(iu + u^2)}. \end{split}$$

The domain⁴³ of $\phi(u;t)$ is the strip $\{u \in \mathbb{C} \mid \text{Im}(u) \in (\alpha_-, \alpha_+)\}$, where $\alpha_- < -1$ and $\alpha_+ > 0$ and solve

 $g(i\alpha)\exp\{d(i\alpha)t\} = 1.$

With the characteristic function at our disposal, we can take advantage of the general Fourier-option pricing theory presented for instance by Carr & Madan [CM99], Lewis [Lew01] or Lee [Lee04]. In particular, the following result allows for the efficient computation of vanilla option prices

Proposition 2.5.6. Let X be a random variable and let $\psi(u)$ denote its characteristic function

$$\phi(u) = \mathbb{E}\left[e^{iuX}\right].$$

Then, for $k \in \mathbb{R}$,

$$\mathbb{E}\left[\left(e^{X}-e^{k}\right)^{+}\right] = \phi(-i) - \frac{K}{2\pi} \int_{i\alpha-\infty}^{i\alpha+\infty} \frac{e^{i\omega k}\phi(-\omega)}{\omega^{2}-i\omega} d\omega \qquad (2.36)$$

for any $\alpha \in (0,1)$ for which the right-hand side exists.

Proof. See [Lew01].

Performing the Fourier inversion in (2.36) via standard numerical integration procedures is much faster (or more accurate) than relying on Monte Carlo (MC) or PDE methods, and in particular allows the efficient calibration of the Heston model to observed market prices.

Remark 2.5.1. There exist various formulas for pricing vanilla options via Fourier transforms depending, for example, on whether the Fourier transform is performed in log-strike (Carr & Madan [CM99]) or log-forward (Lewis [Lew01]) coordinates. See [Sch10] for a good overview of the various approaches.

⁴³Also known as the strip of convergence. This is directly related to the existence of moments.

Remark 2.5.2. The constant α in the above proposition determines an integration path in the complex plane. Choosing $\alpha = 1/2$ corresponds to the so-called Lewis-Lipton-formula [Lip02], which is common in practice and gives stable numerical results in most cases. The particular choice for α can be optimized to improve the numerical properties of the integral, see e.g. [LK07]. In Chapter 4 we shall consider again the topic of choosing optimal integration contours in the complex plane.

Remark 2.5.3. Integrating functions in the complex plane requires some care, especially when dealing with multi-valued functions such as the complex logarithm. Restricting the logarithm to its principal branch (as is done in most software packages), the integrand may become discontinuous when the integration contour should cross a branch cut, leading to wrong option prices. The original formulation of the Heston characteristic function as given in [Hes93] is not free of such problems and a so-called *branch-cut tracking* or *rotation-count algorithm* must be applied, see e.g. [KJ05]. In contrast, the formulation given in Proposition 2.5.5 (which is algebraically equivalent to the original formulation) does not require branch-cut tracking; see [AMST07] for a detailed discussion.

The Displaced Heston Model

As we have noted above, the skewness of the volatility smile implied by the Heston model can be controlled via the correlation between the underlying and the stochastic variance. Most often market-observed volatility smiles are skewed to the left⁴⁴, which requires a negative correlation. If S(t) represents the price of a stock, then $\rho < 0$ is inline with the observation that volatility tends to go up if prices go down - a feature commonly known as the "leverage effect". In the interest-rate markets, however, the skew effect is mainly due to the strong dependence of the percentage volatilities on the level of the underlying rates. After accounting for this fact, the evolution of the rates will often be (almost) uncorrelated with the dynamics of the stochastic volatility or variance⁴⁵. In the following we will therefore always set $\rho = 0$ and control the skewness via the local volatility function $\varphi(\cdot)$ in (2.32). In terms of fitting capabilities this is not a limitation: With the right function $\varphi(\cdot)$, we can generate volatility smiles very similar to those generated by a non-zero correlation parameter ρ . Also from a hedging perspective it is not overly important⁴⁶, whether we generate the skew via the correlation ρ or the function $\varphi(\cdot)$ as long as the right hedge ratios are used; see Section 8.9 of [AP10a]. Lastly, let us note that setting the correlation to zero is also appealing from a practical point of view, since, as we will see in the next chapter, this will simplify matters when performing common measure changes in full term-structure models.

When it comes to choosing a concrete local volatility function $\varphi(\cdot)$, we may reconsider the two examples from the last section and use either a CEV- or a displaceddiffusion-like specification. Due to the better tractability (and the often close resemblance) we shall concentrate in the rest of this thesis on the latter specification, which

⁴⁴At least in the interest-rate and equities markets. FX (Foreign Exchange) smiles are typically symmetric, while smiles in emerging markets may be right-skewed, see e.g. [Tan10].

 $^{^{45}}$ See [Reb04], Chapter 23 and [CS04].

 $^{^{46}\}mathrm{At}$ least for vanilla-like options.

results in the so-called displaced Heston model,

$$dS(t) = \lambda \left(\beta S(t) + (1 - \beta)S(0)\right) \sqrt{V(t)} dW(t), \qquad (2.37)$$

$$dV(t) = \kappa(\theta - V(t))dt + \xi\sqrt{V(t)}dZ(t), \qquad (2.38)$$

with $\kappa, \theta, \xi, \lambda > 0$, $\beta \in (0, 1]$ and where we will always assume that $\langle dZ(t), dW(t) \rangle = 0dt$ and $v_0 = \theta$. Observe that in the above specification λ and θ are redundant parameters in the sense that we may fix one of these parameters without restricting the model. For reasons that will become apparent in the next chapter, we shall always set $\theta = 1$.

Following Andersen and Andreasen [AA02], we may use the same transformation as in the displaced-diffusion case and define for $\beta \in (0, 1]$ the process $X(t) := (\beta S(t) + (1 - \beta)S(0))$. By Itô's formula we obtain

$$dX(t) = \beta \lambda X(t) \sqrt{V(t)} dW(t),$$

$$dV(t) = \kappa(\theta - V(t)) dt + \xi \sqrt{V(t)} dZ(t).$$

Introducing $\tilde{\theta} := (\beta \lambda)^2 \theta$, $\tilde{\xi} := \beta \lambda \theta$ and $\tilde{V}(t) := (\beta \lambda)^2 V(t)$, the above system of SDEs may be rewritten as

$$dX(t) = X(t)\sqrt{\tilde{V}(t)}dW(t),$$

$$d\tilde{V}(t) = \kappa(\tilde{\theta} - \tilde{V}(t))dt + \tilde{\xi}\sqrt{\tilde{V}(t)}dZ(t)$$

which is of standard Heston form (2.33)–(2.34). Combining this with the closed-form characteristic function of the Heston model and Proposition 2.5.6, we obtain, after fixing $\alpha = 1/2$, the following call price formula:

Proposition 2.5.7. Call option prices in the displaced Heston model (2.37)–(2.38) are given by

$$C_{DDH}(0, S(0), K, T) \coloneqq \mathbb{E}\left[(S(T) - K)^+ \right] \\ = \frac{S(0)}{\beta} - \frac{\tilde{K}}{2\pi\beta} \int_{-\infty}^{\infty} \frac{e^{(1/2 - i\omega)k}\phi(-(\omega + i/2); T)}{\omega^2 + 1/4} d\omega, \qquad (2.39)$$

where we have defined

$$\tilde{K} = \beta K + (1 - \beta)S(0), \quad k = \log(S(0)/\tilde{K}),$$

and where

$$\begin{split} \phi(u;T) &= \exp\left\{A(u,T) + B(u,T)\right\},\\ A(u,T) &= \frac{\kappa}{\xi^2} \left((\kappa - \gamma)T - 2\log\left(\frac{1 - ge^{-\gamma T}}{1 - g}\right)\right),\\ B(u,T) &= \frac{\kappa - \gamma}{\xi^2} \left(\frac{1 - e^{-\gamma T}}{1 - ge^{-\gamma T}}\right),\\ g &= g(u) = \frac{\kappa - \gamma}{\kappa + \gamma},\\ \gamma &= \gamma(u) = \sqrt{\kappa^2 + (\lambda\beta\xi)^2(iu + u^2)}. \end{split}$$

After having presented a pricing formula for European options in the displaced Heston model, let us now consider some volatility smiles generated by this model. As can be seen in Figure 2.5, the displaced Heston model is capable of fitting the market smile from Figure (2.3) quite well. The main effect of the volatility of variance parameter ξ is adding curvature to the implied volatility smile, or equivalently generating implied marginal distributions with fatter tails. Observe that the smile corresponding to $\xi = 0$ is in effect that of a displaced-diffusion model. The speed of mean reversion κ has a similar effect on the volatility smile as ξ . Both together can be used to control how fast the convexity of the smile decays in *T*-direction.



Figure 2.5: Market volatility smile (cp. Fig. (2.3)) and fitted displaced Heston smile with parameters S(0) = 4.45%, T = 5, $\beta = 0.1$, $\lambda = 0.16$, $\xi = 0.89$, $\kappa = 0.10$. The other smiles correspond to different values of ξ as indicated in the graph, with λ being chosen such that the ATM volatility levels match (all other parameters being the same as before).

The above model can be extended by allowing for time-dependent parameters. Besides increasing the flexibility of the vanilla model, time-dependent parameters will naturally emerge when we will consider full term-structure models in later chapters. Let us consider a model with time-dependent volatility $\lambda = \lambda(t)$ and skew $\beta = \beta(t)$:

$$dS(t) = \lambda(t) (\beta(t)S(t) + (1 - \beta(t))S(0)) \sqrt{V(t)} dW(t), \qquad (2.40)$$

$$dV(t) = \kappa (1 - V(t))dt + \xi \sqrt{V(t)}dZ(t), \ V(0) = 1.$$
(2.41)

As with the displaced-diffusion model, we can use parameter-averaging techniques to derive (constant) effective parameters, which allow us to efficiently price vanilla options by reusing the formula from Proposition 2.5.7. The effective skew $\bar{\beta}$ over a time horizon [0, T] in the model (2.40)–(2.41) is obtained similarly as in the displaceddiffusion case and is given by (see [Pit05a])

$$\bar{\beta} = \int_0^T \beta(t) w_T(t) dt, \qquad (2.42)$$

with weight function

$$w_T(t) = \frac{v(t)^2 \lambda(t)^2}{\int_0^T v(s)^2 \lambda(s)^2 ds},$$
(2.43)

$$v(t)^{2} = \int_{0}^{t} \lambda(s)^{2} ds + \xi^{2} e^{-\kappa t} \int_{0}^{t} \lambda(s)^{2} \frac{e^{\kappa s} - e^{-\kappa s}}{2\kappa} ds.$$
(2.44)

In order to also reduce the time-dependent volatility to an "averaged" effective volatility, we follow Piterbarg [Pit05a] and consider the price of an (undiscounted) ATM option $(K = S_0)$

$$\mathbb{E}\left[(S(T) - S_0)^+ \right] = \mathbb{E}\left[\mathbb{E}\left[(S(T) - S_0)^+ \right] | \{V(t)\}_{t=0}^T \right] \right].$$

Due to the independence of the Brownian motions driving S(t) and V(t), we have that given a particular path $\{V(t)\}_{t=0}^{T}$, the forward-rate process S(t) follows an ordinary displaced-diffusion process with a (conditionally deterministic) time-dependent volatility function. In particular (see (2.28)–(2.29)), the distribution of S(T) depends only on the total integrated variance

$$\bar{V}(T) \coloneqq \int_0^T \lambda(t)^2 V(t) dt,$$

such that we may write

$$\mathbb{E}\left[\left(S(T) - S(0)\right)^{+}\right] = \mathbb{E}\left[g(\bar{V}(T))\right],$$

where $g(\cdot)$ denotes the call price formula in the displaced-diffusion model, see Proposition 2.5.2. The problem of finding an effective volatility can now be represented as finding $\bar{\lambda}$ such that

$$\mathbb{E}\left[g\left(\int_0^T \lambda(t)^2 V(t)dt\right)\right] = \mathbb{E}\left[g\left(\bar{\lambda}^2 \int_0^T V(t)dt\right)\right)\right].$$
(2.45)

Neither of the expectations in (2.45) is easy to compute. However, the Laplace transform of $\int_0^T V(t)dt$ is given in closed form, while the Laplace transform of $\bar{V}(T)$ is easy to compute numerically, see [AA02]. This suggests approximating g(v) by a function of the form

$$g(v) \approx a + be^{-cv}$$

around the mean of $\overline{V}(T)$

$$\zeta = \mathbb{E}\left[\bar{V}(T)\right] = \int_0^T \lambda(t)^2 dt$$

With this approximation, the problem (2.45) can now be represented as finding $\bar{\lambda}$ such that

$$\psi_0\left(-\frac{g''(\zeta)}{g'(\zeta)}\bar{\lambda}^2\right) = \psi\left(-\frac{g''(\zeta)}{g'(\zeta)}\right),\tag{2.46}$$

with

$$\psi_0(\mu) = \mathbb{E}\left[e^{-\mu\int_0^T V(t)dt}\right]$$

and

$$\psi(\mu) = \mathbb{E}\left[e^{-\mu \bar{V}(T)}\right]$$

being the Laplace transforms of $\int_0^T V(t) dt$ and $\overline{V}(T)$, respectively.

By means of the parameter-averaging technique we have essentially projected the time-dependent model (2.40)-(2.41) onto the model

$$dS(t) = \bar{\lambda} \big(\bar{\beta} S(t) + (1 - \bar{\beta}) S(0) \big) \sqrt{V(t)} dW(t), \qquad (2.47)$$

$$dV(t) = \kappa (1 - V(t))dt + \xi \sqrt{V(t)}dZ(t), \ V(0) = 1,$$
(2.48)

such that both models produce approximately the same T-maturity vanilla option prices. This allows us again to make use of the efficient Fourier pricing formula given in Proposition 2.5.7.

Let us make a final remark regarding the particular choice of a model. So far, our considerations were mainly driven by the objective of finding a model that matches market-observed vanilla option prices, or equivalently implied volatilities, as closely as possible (while being at the same time computationally tractable). The rationale behind this is that vanilla options are typically (sometimes the only) liquidly traded instruments from which we can extract market information about dynamics of the underlying. Moreover, when a model is used for pricing/hedging exotic products, vanilla options are typically used as additional hedging instruments (besides the underlying) and the argument is then, that the model should at least be able to correctly price these hedging instruments. As we have noted several times before, knowing the prices of European options is equivalent to knowing the distribution of the underlying at fixed times, conditional on its current value. It is important to notice, however, that the underlying stochastic process carries much more information than the conditional laws. The dynamic behavior and the hedging implications of two different models may therefore be quite different, even though they perfectly agree on a set of European option prices. In case of vanilla-like options, Poulsen et al. [PSHE09] demonstrate that as long as the modeling is done sensibly and the right hedging strategies are used, the particular flavor of (stochastic volatility) model matters little for hedging these options. The situation may change though, if we consider the pricing and hedging of more complex (i.e., exotic) derivatives. Schoutens et al. [SST04] and Ayache et al. [AHNW04] among others show, that models that produce similar vanilla option prices may very well give markedly different prices of exotic options (cp. Section 7.3). This must be taken into account when a model is to be chosen for pricing a certain product. Ultimately, to put it in the words of Lipton [Lip02], "the right criterion, as advocated by a number of practitioners and academics, is to choose a model that produces hedging strategies for both vanilla and exotic options resulting in profit and loss distributions that are sharply peaked at zero."

38

2.6 Other Interest-Rate Options

In Section 2.4 we introduced caps/floors and swaptions, which are European options on Libor and swap rates and are among the most liquidly traded interest-rate derivatives in the market. In the following we will extend this set of derivative interest-rate securities by considering more general examples of market payoffs. In all cases we will assume that we are given a discrete tenor structure $0 \leq T_0 < T_1 < \ldots < T_N$.

2.6.1 CMS Swaps, Caps and Floors

Most of the time, it can be observed that short-term interest rates (e.g. 6M or 1Y rates) are lower than long-term interest rates (e.g. 20Y rates), and the corresponding yield curves are hence upward sloping⁴⁷. So even if today's long-term swap or forward Libor rate for a given maturity T is relatively high (say 6% or 7%), it is quite likely that once we reach time T, the prevailing 6M spot Libor rate at that time is (again) only between 2% to 3%. For an investor it is therefore often more attractive to receive a fixed coupon rather than a floating $coupon^{48}$. But what if an investor thinks that interest rates (across the yield curve) are currently too low and will increase in the future? In this case, the holder of a fixed-rate bond will end up receiving a low fixed coupon over the lifetime of the bond (e.g. 10 years), even if interest rates go up. The solution for our investor might be a bond with coupons based on a so-called constant-maturity swap (CMS) rate. As the market for plain vanilla swaps is very liquid, market quotes for swap rates can be used themselves for defining the payoffs of other securities. A CMS rate is defined to be the break-even swap rate (see (2.9)) of a standard fixed-for-floating swap of a fixed maturity, e.g. 10 years or 30 years. Accordingly, a coupon bond based on the 10Y CMS rate pays on each coupon date (e.g. every 6 months) the currently prevailing 10Y swap rate. With such a security, an investor can take advantage of an upward sloping yield curve (as the payments are based on the long end of the yield curve), but still benefits if interest rates go up (because the coupons are floating rather than fixed). CMS linked products are also used, for example, by mortgage lenders or insurance companies to hedge risks connected to movements in, say, the 10Y or 30Y point of the yield curve.

CMS Swaps

A *CMS* swap is a fixed-for-floating swap, but in contrast to a plain vanilla swap, the floating leg payments are based on CMS rather than Libor rates. More specifically, let $S_{k,k+n}(\cdot)$ denote the *n*-period swap rate with first fixing date T_k , as defined in (2.9). Then, the holder of a (payer) CMS swap with fixed rate K, based on the *n*-period CMS rate, receives at a payment time T_{k+1} the *n*-period swap rate⁴⁹ fixed at time T_k and pays the fixed rate K^{49} . Thus, the value of a CMS swap is given by

$$V_{\rm swap}^{\rm CMS}(t) = B(t) \sum_{k=0}^{N-1} \tau_k \mathbb{E}_t^{\mathbb{Q}} \left[B(T_{k+1})^{-1} (S_{k,k+n}(T_k) - K) \right], \qquad (2.49)$$

⁴⁷This can be readily observed in Figure 5.1 on p. 99, which shows the evolution of the EUR forward-rate curve for the period 1/14/2008 - 4/26/2010.

⁴⁸Put in swap terminology, it is usually better to receive the fixed leg and pay the floating leg.
⁴⁹Multiplied by an accrual factor.

or, using the T_{k+1} -forward measure in each period,

$$V_{\rm swap}^{\rm CMS}(t) = \sum_{k=0}^{N-1} \tau_k P(t, T_{k+1}) \mathbb{E}_t^{T_{k+1}} \left[S_{k,k+n}(T_k) - K \right].$$
(2.50)

It is standard market practice to quote prices of such swaps as a spread over Libor, i.e., in (2.50) the fixed leg is replaced by a floating Libor leg, which pays $L_k(T_k) + s$ at payment time T_{k+1} . The (constant) fixed spread s, that makes this swap fair, i.e., have value zero today, is then the quoted "price" of the CMS swap.

While plain vanilla swaps can be valued by simple no-arbitrage arguments and by just knowing today's yield curve, the valuation of CMS swaps requires a termstructure model (or at least a vanilla model), as we will see in the following. For notational simplicity let us focus on a particular CMS cash flow from (2.50)

$$V_{\text{CMS}}(0) \coloneqq P(t, T_{k+1}) \mathbb{E}^{T_{k+1}} \left[S_{k,k+n}(T_k) \right],$$

and write this as

$$V_{\text{CMS}}(0) = P(t, T + \delta) \mathbb{E}^{T+\delta} \left[S(T) \right], \qquad (2.51)$$

where the notation should be obvious. We recall that by Lemma 2.3.2, the swap rate S(t) is a martingale under its associated swap measure \mathbb{Q}^A , where A(t) denotes the the annuity factor $A(t) = A_{k,k+n}(t)$ of S(t). In particular, we have

$$S(0) = \mathbb{E}^A \left[S(T) \right],$$

i.e., today's forward swap rate is simply the expectation of the (spot) swap rate at the fixing time T under the natural swap measure. This does not hold⁵⁰ under the $T + \delta$ -forward measure $\mathbb{Q}^{T+\delta}$. Note that by changing to the swap measure, we can rewrite the expectation of S(T) under $\mathbb{Q}^{T+\delta}$ as

$$CMS(0) \coloneqq \mathbb{E}^{T+\delta} [S(T)]$$

= $\frac{A(0)}{P(0, T+\delta)} \mathbb{E}^A \left[\frac{P(T, T+\delta)}{A(T)} S(T) \right].$ (2.52)

The quantity CMS(0) is commonly known as the forward CMS rate, whereas

$$CA(0) \coloneqq CMS(0) - S(0) \tag{2.53}$$

is called the *CMS convexity adjustment* (applied to S(0)). In order to compute the expectation in (2.52), we generally need a term-structure model, as the multiplier $P(T, T + \delta)/A(T)$ depends on the joint distribution of a whole set of interest rates. Note that we may always write (using the tower property of expectations)

$$\mathbb{E}^{A}\left[\frac{P(T,T+\delta)}{A(T)}S(T)\right] = \mathbb{E}^{A}\left[\mathbb{E}^{A}\left[\frac{P(T,T+\delta)}{A(T)} \mid S(T)\right]S(T)\right]$$
$$= \mathbb{E}^{A}\left[g(S(T))S(T)\right]$$
(2.54)

 $^{^{50}}$ If we neglect the degenerate case of a one-period swap rate.

2.6. Other Interest-Rate Options

with a deterministic function $g(\cdot)$ given by

$$g(s) = \mathbb{E}^A \left[\frac{P(T, T+\delta)}{A(T)} \mid S(T) = s \right].$$
(2.55)

While for computing the expectation in Equation (2.54) a Vanilla model (which models the evolution of S(t)) is sufficient, determining the function $g(\cdot)$ still requires a full term-structure model. However, convexity adjustments become sizeable only for longer maturities, and the long end of the yield curve is usually rather flat and often tends to move in parallel. This may serve as a justification for approximating the "real" function $g(\cdot)$ by a rather simple function of the swap rate S(T).

Often such approximations are inspired by either real term-structure models (see also Chapter 4) or simple "bond mathematics". A representative for the latter approach is the following popular choice for $g(\cdot)$ (see [Hag03])

$$\frac{P(T, T+\delta)}{A(T)} \approx \frac{(1+\tau S(T))^{-\delta/\tau}}{\sum_{j=1}^{n} \tau (1+\tau S(T))^{-j}} \\
= \frac{S(T)}{(1+\tau S(T))^{\delta/\tau}} \cdot \frac{1}{1-(1+\tau S(T))^{-n}} \\
=: g(S(T)),$$
(2.56)

where it is assumed that the accrual factors are approximately equal, i.e., $\tau_i \approx \tau$. Another popular approximation is based on the linear swap-rate model of Hunt and Kennedy [HK00]; see also [BS03] and [Pel03] for a comparison with true (Monte Carlo simulated) convexity adjustments from a Libor market model. For more sophisticated specifications of $g(\cdot)$ see [Hag03] or Chapter 16 of [AP10c].

After a particular function $g(\cdot)$ has been selected, the value of the convexity adjustment (2.53), or equivalently the value of the CMS cash flow (2.51), is fully determined by the distribution of S(T) under the swap measure \mathbb{Q}^A and the expectation (2.54) can be computed by using, for instance, a suitably calibrated Vanilla model.

Alternatively, we can use the Breeden-Litzenberger formula (2.19) to infer the market-implied density of S(T) from a continuum of market prices of payer/receiver swaptions and then compute the expectation by integrating against this density. This approach is appealing, in that it yields model-independent (up to the choice of $g(\cdot)$) and market-consistent convexity adjustments, but requires suitable methods for interand extrapolating the (often few) market-observable swaption prices.

Similar in spirit to the Breeden-Litzenberger formula is the following static replication formula due to Carr & Madan [CM01]:

Proposition 2.6.1. Let p(0, S(0); T, K) respectively c(0, S(0); T, K) denote the undiscounted time-0 prices of European put and call options with strike K and expiry T. Then, for any twice-continuously differentiable function $f(\cdot)$, the undiscounted value of a European option with payoff function $f(\cdot)$ and expiry T may be written as

$$\mathbb{E}[f(S(T))] = f(K^*) + f'(K^*)(S(0) - K^*) + \int_0^{K^*} f''(K)p(0, S(0); T, K)dK + \int_{K^*}^{\infty} f''(K)c(0, S(0); T, K)dK$$
(2.57)

for any $K^* \in \mathbb{R}_+$.

Notice that this formula not only allows us to express the value of a European option in terms of put and call options in a totally model-independent way, but it also provides us with a static⁵¹ hedging portfolio, where the portfolio weights for the strike-K put and calls are equal to f''(K)dK. Of course, in reality one does not have an infinite number of options to construct this replication strategy. To account for this fact one will usually choose a collection of (tradeable) strikes $\{K_i\}$ and then use a formula of the form

$$\mathbb{E}[f(S(T))] \approx f(K^*) + f'(K^*)(S(0) - K^*) + \sum_i w_p(i)p(0, S(0); T, K_i) + \sum_i w_c(i)c(0, S(0); T, K_i), \quad (2.58)$$

where the weights $w_p(i)$ and $w_c(i)$ are chosen such that the sums in (2.58) approximate⁵² the integrals in (2.57).

Using the static replication formula (with $K^* = S(0)$), we may write the value of the CMS cash flow (2.51) as

$$V_{\text{CMS}}(0) = A(0)\mathbb{E}^{A} \left[g(S(T))S(T)\right]$$

= $A(0)S(0)g(S(0))$
+ $\int_{0}^{S(0)} w(K)V_{\text{rec}}(0,K)dK + \int_{S(0)}^{\infty} w(K)V_{\text{pay}}(0,K)dK,$ (2.59)

with

$$w(K) = \frac{\partial^2}{\partial K^2} (g(K)K)$$

and where $V_{\rm rec}(0, K)$ and $V_{\rm pay}(0, K)$ are the time-0 prices of receiver and payer swaptions, respectively:

$$V_{\rm rec}(0,K) = A(0)\mathbb{E}^{A} \left[(K - S(T))^{+} \right],$$

$$V_{\rm pay}(0,K) = A(0)\mathbb{E}^{A} \left[(S(T) - K)^{+} \right].$$

As before, the swaption prices in (2.59) can be either computed by using a calibrated model⁵³ or can be directly observed in the market.

Depending on the payment lag δ and the tenor *n* of the CMS rate, the function $g(\cdot)$ is often observed to be slowly varying and almost linear⁵⁴, regardless of the model used to obtain it. We may therefore approximate $g(\cdot)$ by a linear function and in this way avoid the necessity of having to perform numerical integrations. Linearizing for example the function $g(\cdot)$ as given in (2.56) about the initial value S(0)

$$g(S(T)) \approx g(S(0)) + g'(S(0))(S(T) - S(0)),$$

⁵¹I.e., the portfolio never needs to be adjusted throughout the life time of the option.

⁵²Alternatively, the weights may be also chosen to super- or sup-replicate the payoff.

⁵³In practice, often the SABR model of Hagan et. al [HKLW02] is used for this task. See also [MP06].

⁵⁴See, e.g. the appendix of [ZK10].

2.6. Other Interest-Rate Options

and assuming that $A(0)/P(0, T + \delta) \approx g(S(0))$, yields the following expression for the convexity adjustment

$$CA(0) = S(0) \frac{1 + \tau S(0) \left(1 - \delta - \frac{n}{(1 + \tau S(0))^n - 1}\right)}{1 + \tau S(0)} \left(\frac{\mathbb{E}^A[S(T)^2]}{S(0)^2} - 1\right).$$
(2.60)

Computing the actual value of the convexity adjustment is now only a matter of computing the second moment of S(T), which is known in closed form for quite a wide range of models. In the flat-smile case, i.e., when S(T) is log-normal with volatility σ , we have

$$\mathbb{E}^{A}[S(T)^{2}] = S(0)^{2} e^{\sigma^{2}T}, \qquad (2.61)$$

and only a single volatility parameter is needed for computing the convexity adjustment.

Figure 2.6 illustrates the impact of the (Black) volatility σ and the maturity Ton the size of the convexity adjustment (2.60). We assumed that interest rates are flat at 5% (continuously compounded) for all maturities. As can be clearly seen, the convexity adjustment is increasing in maturity T and volatility σ^{55} . In the flat-smile case this is quite obvious from the exponential factor in (2.61). However, the same observations generally also hold for other models and convexity-adjustment formulas. For comparison, we have also plotted the convexity adjustments as obtained from formula (2.60), if the second moment is taken from a displaced Heston model, where the parameters are chosen so as to match an ATM volatility level of 15% for each maturity T. The skewness (and the fatter tails) of the displaced Heston distribution obviously does not have a huge impact on the size of the second moment (and therefore on the convexity adjustment), as long as the "average" volatility level is the same as in the corresponding Black model. Therefore, even in the presence of a true volatility smile/skew, the "Black implied" convexity adjustment⁵⁶ with the market-observed ATM volatility may be used as a first rough approximation⁵⁷.

CMS Caps and Floors

Options on CMS rates are traded in the market in the form of *CMS caps and floors*, which, similar to ordinary (Libor) caps and floors, consist of sequences of European options:

$$V_{\rm cap}^{\rm CMS}(t) = \sum_{k=0}^{N-1} \tau_k P(t, T_{k+1}) \mathbb{E}_t^{T_{k+1}} \left[(S_{k,k+n}(T_k) - K)^+ \right],$$

$$V_{\rm floor}^{\rm CMS}(t) = \sum_{k=0}^{N-1} \tau_k P(t, T_{k+1}) \mathbb{E}_t^{T_{k+1}} \left[(K - S_{k,k+n}(T_k))^+ \right].$$

⁵⁵Furthermore, it is generally also increasing in the length of the CMS tenor. Notice that a oneperiod (forward) swap rate is actually a (forward) Libor rate and accordingly, the necessary convexity adjustment would be zero.

⁵⁶That is, combining (2.60) and (2.61).

⁵⁷From (2.59) it can be seen that low- and high-strike options "equally" enter the valuation formula for the convexity adjustment. Therefore, using the ATM volatility may be a sufficiently accurate approximation. This situation will change in case of CMS options, see below.



Figure 2.6: CMS convexity adjustments in basis points for n = 10, $\delta = 0.5, \tau = 1$ as implied by formula (2.60) for the Black and the displaced Heston model. Black volatilities σ as indicated in the graph. Displaced Heston parameters: $\kappa = 0.1, \xi = 0.9, \beta = 0.4$. The parameter λ in the displaced Heston model is chosen such that for each maturity the implied ATM level of the smile is 15%.

As before, let us focus on a single CMS caplet with initial value

$$V_{\text{caplet}}^{\text{CMS}}(0) = P(t, T + \delta) \mathbb{E}^{T+\delta} \left[(S(T) - K)^+ \right]$$

Introducing again a suitable function $g(\cdot)$, this may be written as

$$V_{\text{caplet}}^{\text{CMS}}(0) = A(0)\mathbb{E}^{A}\left[g(S(T))(S(T) - K)^{+}\right].$$
 (2.62)

As with the "plain" CMS cash flow, the expectation on the right hand side of (2.62) can be either computed with a vanilla model or inferred from market prices of swaptions via a replication argument. While the replication formula from Proposition 2.6.1 is not directly applicable, as the payoff function

$$f(s) \coloneqq g(s)(s-K)^+$$

is not differentiable at s = K, it can nevertheless be easily generalized to also cover this case; see e.g. [Hag03]. The replication formula for the CMS caplet then reads

$$V_{\text{caplet}}^{\text{CMS}}(0) = h'(K)V_{\text{pay}}(0,K) + \int_{K}^{\infty} h''(k)V_{\text{pay}}(0,k)dk, \qquad (2.63)$$

where h(s) = g(s)(s - K). We emphasize again that this formula not only allows us to value the CMS caplet, but also provides us with a model-independent static hedging portfolio consisting of payer swaptions. Notice that, in contrast to the plain CMS cash flow from above, now only swaptions with strikes greater than or equal to the caplet strike K enter the valuation formula. An analogous result holds for CMS floorlets

$$V_{\text{floorlet}}^{\text{CMS}}(0) = h'(K)V_{\text{rec}}(0,K) + \int_0^K h''(k)V_{\text{rec}}(0,k)dk, \qquad (2.64)$$

where h(s) = g(s)(K - s). In this case only (receiver) swaptions with strikes less than or equal to the floorlet strike K enter the valuation formula.

Remark 2.6.1. From formulas (2.59), (2.63) and (2.64) it can be easily seen that the values of CMS-linked products depend on prices (or implied volatilities) of swaptions with far-from-the-money strikes, which may not be traded in the market. Provided that CMS swaps/floor/caps are traded liquidly enough, we can therefore use these instruments in combination with the aforementioned formulas to infer (at least theoretically) market-implied information about the asymptotic behavior of the volatility smile.

Remark 2.6.2. CMS floors are often implicitly contained in bonds paying CMS coupons. Most of these bonds include a floor in order to limit the investor's risk of receiving very low (or negative) coupons.

CMS-Linked Payoffs and Moment Explosions

In the previous sections we have seen that switching from the forward measure to the natural swap measure, under which the modeling typically takes place, introduces a certain function $g(\cdot)$. "Payoffs" of CMS-linked cash flows under \mathbb{Q}^A are hence of the form

$$g(s)s \text{ or } g(s)(s-K)^+.$$
 (2.65)

While the original payoff functions are linear (at least piecewise), multiplication by g(s) generally results in payoff functions that grow super-linearly. Therefore, convexity adjustments, respectively prices of CMS-linked cash flows, implicitly or explicitly⁵⁸ depend on higher order moments of S(T) under \mathbb{Q}^A . In the Black model this is not a matter of concern as all moments of S(T) exist. This statement also holds for the displaced-diffusion model, being essentially just a transformed Black model, as well as for the CEV model, whose local-volatility function grows only sub-linearly. Stochastic volatility models, however, which typically have fat-tailed distributions, may suffer from so-called *moment explosions*, that is, certain higher-order moments either may not exist or exist only up to some finite *explosion time* T^* .

In the particular case of the displaced Heston model, we have the following conditions on the finiteness of moments (cp. [AP07]).

Proposition 2.6.2. Consider the displaced Heston model (2.37)–(2.38). For a given $\nu > 1$, set $\mu = (\lambda \beta)^2 \nu (\nu - 1) > 0$ and define

$$b = \frac{2\mu}{\xi^2} > 0, \ a = -\frac{2\kappa}{\xi^2} < 0, \ D = a^2 - 4b.$$

The moment $\mathbb{E}[S(T)^{\nu}]$ will be finite for all T > 0 if $D \ge 0$. If D < 0, then $\mathbb{E}[S(T)^{\nu}]$ will be finite for $T < T^*$ and infinite for $T^* \ge T$, where T^* is given by

$$T^* = \frac{2}{\eta\xi^2} \left(\pi + \arctan\left(2\eta/a\right) \right), \ \eta \coloneqq \frac{1}{2} \sqrt{-D}.$$

For CMS linked cash flows, the second-order moment is of particular importance. In case of the convexity approximation given in (2.60) this is quite obvious. But

 $^{^{58}}$ As is evident for example from Formula (2.60).

notice, that also in case of the generic expression for $g(\cdot)$ in (2.55) we have

$$g(s) = \mathbb{E}^{A} \left[\frac{P(T, T+\delta)}{A(T)} \mid S(T) = s \right]$$
$$= s \mathbb{E}^{A} \left[\frac{P(T, T+\delta)}{1 - P(T, T_{k+n})} \mid S(T) = s \right]$$
$$\leq s \mathbb{E}^{A} \left[\frac{1}{1 - P(T, T_{k+n})} \mid S(T) = s \right].$$

All rational interest-rate models would have the last conditional expectation decay exponentially fast to 1 for large s, which leaves us again (see (2.65)), to leading order, with the evaluation of the second-order moment $\mathbb{E}[S(T)^2]$.

For illustration, Figure 2.7 shows the dependence of the explosion time T^* on β and λ . The stochastic-volatility parameters κ and ξ were set to realistic values (cp. Section 7.2). As is evident from the figure, lowering the skew parameter β has a dampening effect, pushing T^* further into the future. This is intuitively clear, since lowering β essentially redistributes probability mass from the right tail of the probability distribution to the left tail and hence produces more left-skewed distributions.



Figure 2.7: Critical explosion time T^* for the second-order moment of S(T) in the displaced Heston model. Stochastic volatility parameters: $\kappa = 0.06, \xi = 0.9$.

With (ATM) volatilities of long-maturity swaptions typically in the 10%–20% range and with skews β typically less than 0.3, moment explosion are most often not an issue. However, in high-volatility regimes or when long-dated options are to be priced, the finiteness of the relevant moments should always be checked. This is particularly important, since numerical pricing routines usually have a built-in dampening effect. It might therefore not be obvious that a stochastic-volatility model suffers from moment explosions. For example, when computing an expectation by numerical integration, the integral is typically truncated at some large finite value. If the integral value initially grows only very slowly in the upper bound, then one might be lead to the conclusion that the expectation exists, even if, in fact, it does not. Similarly, a user who is performing Monte Carlo simulations might see reasonable looking prices, albeit that these prices will not converge when increasing the number of Monte Carlo paths.

2.6.2 CMS Spread Products

CMS spread products are linked to the difference between two CMS rates, where typical examples are the 10Y vs. 2Y or the 30Y vs. 2Y spread. In its simplest form, spread-linked cash flows can be generated by going long a CMS swap based on, say, the 10Y rate and going short a CMS swap based on the 2Y rate. Neglecting the fixed leg, this strategy pays off at a payment time $T_k + \delta_k$, $\delta_k \ge 0$ the amount

$$S_{k,k+n}(T_k) - S_{k,k+m}(T_k),$$

where in case of an underlying tenor structure based on semiannual (quarterly) fixings we would have n = 20(40) and m = 4(8) to cover 10 years and 2 years, respectively. More general payoffs linked to CMS spreads can be found in the market in the form of CMS spread-leveraged notes or exotic swaps paying coupons⁵⁹ of the form

$$C_{k} = \max\left\{\min\left\{g \cdot (S_{k,k+n}(T_{k}) - S_{k,k+m}(T_{k})) + s, c\right\}, f\right\},$$
(2.66)

with a gearing factor g, spread s, cap c and floor f. CMS spread-linked structures are not affected by shifts in the overall level of the yield curve, but rather depend on the slope of the yield curve. By buying such structures, investors can express their view on (or hedge risks associated with) the relationship between the long and the short end of the yield curve.

As discussed earlier, yield curves tend to be upward sloping most of the time⁶⁰. Moreover, the long end is often observed to be relatively flat, which implies that future realized 10 year and 2 year rates (as seen today) will be approximately equal. Therefore, if an investor believes that either the current yield curve will steepen further⁶¹ or that future realized yield curves will be again upward sloping, and not as flat as projected by today's long end of the yield curve, she might be interested in buying a CMS spread-linked note⁶². Depending on whether a certain product benefits from a curve steepening or flattening, such structures are often also called (curve) steepeners or flatteners.

A standard CMS call spread option or CMS spread cap pays at a sequence of payment times $T_k + \delta_k$ the amounts

$$(S_{k,k+n}(T_k) - S_{k,k+m}(T_k) - K)^+.$$
(2.67)

That is, similar to Libor and CMS caps, CMS spread caps consist of a sequence of caplets (European call options). There exists a relatively liquid broker market for such spread options on EUR and US\$ CMS rates and prices are quoted on a regular basis.

CMS spread-linked products typically cannot be valued in terms of vanilla products by a replication argument. Observe that the value of the payoff (2.66) depends

⁵⁹In case of a swap, the exotic leg pays the structured coupons, while the other leg may receive fixed rates or floating Libor (plus spread).

 $^{^{60}\}mathrm{At}$ least in the USA, Japan and the euro countries.

⁶¹As remarked in the introduction, this might be the case, for instance, if a central bank is expected to loosen monetary policy due to a weak economy. Alternatively, long rates may be expected to increase due to inflationary pressure.

⁶²With n > m and g > 0.

on the joint distribution of the swap rates under a common pricing measure. For notational simplicity let us assume that we have fixed k, n and m and denote the two swap rates by $S_i(T)$, i = 1, 2 and the payment time by $T+\delta$. Using the $T+\delta$ -forward measure, we obtain for the (undiscounted) value of the payoff (2.67)

$$V(0;T,K) = \mathbb{E}^{T+\delta} \left[(S_1(T) - S_2(T) - K)^+ \right].$$
(2.68)

Note that generally there does not exist a common measure under which both $S_1(t)$ and $S_2(t)$ are martingales. Although, in principle, we could switch to one of the natural swap measures⁶³ \mathbb{Q}^{A_i} , the $T + \delta$ -forward measure is often more natural and more convenient to work with. In particular, when working under the forward measure, both rates are "equally" affected by certain measure-change related adjustments, and possible approximation errors often tend to at least partially cancel.

Similar to single-rate vanilla options, market prices of CMS spread options are often quoted in terms of implied volatilities. More specifically, the *implied Normal* (also known as *basis-point*) spread volatility $\sigma(T, K)$ for strike K and maturity T is defined by equating the undiscounted market price of a spread option to the option price formula in the Bachelier (or Normal) model, i.e.,

$$V_{\rm mkt}(0;T,K) = C_{\rm Bach}(0,S(0),K,T,\sigma(T,K))$$
(2.69)

where $C_{\text{Bach}}(t, s, K, T, \lambda)$ is the call option pricing formula in the Bachelier model with volatility λ as defined in (2.21). In the above equation, S(0) denotes the *convexityadjusted forward (CMS) spread* $S(0) := \mathbb{E}^{T+\delta} [S_1(T) - S_2(T)]$, which can be computed by using the techniques discussed in Section 2.6.1. Notice that the Bachelier rather than the Black model is used for quoting implied volatilities, as the spread $S_1(T) - S_2(T)$ can certainly become negative.

The Bachelier model provides a convenient way of quoting spread-option prices, but similar to single-rate vanillas, Normal spread-implied volatilities will generally depend on strike and maturity, i.e., there will be again volatility smiles reflecting the fact that market prices are not consistent with assuming a Normal model for the spread⁶⁴. Furthermore, the Normal model provides no link to the distributions (including swaption volatilities) of the underlying swap rates, which makes the model of limited use for risk-management and hedging purposes.

In order to price CMS spread options in a really consistent and arbitrage-free way, a proper term-structure model is generally needed, and we will devote Chapter 4 to the topic of pricing CMS spread options in a certain class of term-structure models. Before moving on, we shall briefly review an approach, that, despite some obvious drawbacks, has become the de facto market standard for pricing "vanilla"⁶⁵ European CMS spread options.

The Copula Approach

In Sections 2.5.1–2.6.1 we have seen that for pricing single-rate derivatives it is often sufficient to use vanilla models, which model only the evolution of one particular rate.

 $^{^{63}\}mathrm{In}$ which case at least on rate would be a martingale.

 $^{^{64}}$ For market-implied spread volatility surfaces see Figures 7.4 and 7.10 in Section 7.2.

 $^{^{65}}$ The market for CMS spread options of the most basic form (2.67) has become relatively liquid in recent years, so that one can almost speak of a vanilla (spread) options.

This is not only more convenient and easier from a mathematical point of view, but it is usually also faster and more accurate than using a full term-structure model. "More accurate" is to be understood in the sense that it is generally more difficult to fit the market-implied distribution of one particular rate with a model that simultaneously specifies the dynamics of the whole yield curve. Given the high traded volumes in many derivatives markets, there is often not much room for pricing errors due to not being able to fit market-observable prices.

The same carries over to vanilla CMS spread options, which have become fairly liquid in recent years. Essentially, what we would like to have is a fast and flexible (two-rate) vanilla model, that specifies the joint dynamics of only the two underlying swap rates. In order to obtain such a model, one usually proceeds along the following lines. First recall, that we can deduce the distribution of each swap rate $S_i(T)$ under its swap measure \mathbb{Q}^{A_i} from market prices of swaptions across strikes. Further, by specifying functions $g_i(\cdot)$ we can translate these swap-rate distributions into the corresponding distributions under the $T + \delta$ -forward measure. More precisely, we have for the density of $S_i(T)$ under $\mathbb{Q}^{T+\delta}$ (see Equations (2.52) and (2.54))

$$\mathbb{Q}^{T+\delta}\left(S_i(T) \in ds\right) = \frac{A_i(0)}{P(0, T+\delta)} g_i(s) \mathbb{Q}^{A_i}(S_i(T) \in ds)$$
(2.70)

with

$$g_i(s) = \mathbb{E}^{A_i} \left[\frac{P(T, T+\delta)}{A_i(T)} \mid S_i(T) = s \right].$$
(2.71)

Formulas (2.70)–(2.71) are exact as written, but $g_i(\cdot)$ would virtually always be approximated by a simple function of $S_i(T)$ as was demonstrated in Section 2.6.1. In a two-rate setting this approach will lead to some inconsistencies, since each $P(T, T + \delta)/A_i(T)$, i = 1, 2 will generally depend on both swap rates $S_1(T)$ and $S_2(T)$, and the calculation of $g_i(s)$ should therefore incorporate the dependence structure of both rates. Nevertheless, for tractability reasons the measure change related calculations are most often done independently from the dependence structure modeling.

Having determined the marginal distributions of $S_i(T)$ under the forward measure, the joint distribution of $(S_1(T), S_2(T))$ can now be obtained by linking the margins with a so-called *copula*. A two-dimensional⁶⁶ copula $C : [0, 1]^2 \to [0, 1]$ is a cumulative distribution function on $[0, 1]^2$ with standard uniform margins, i.e.,

$$C(u, 1) = u$$
 and $C(1, v) = v$.

For given marginal CDFs $F_i(\cdot)$, i = 1, 2, a copula can be used to construct a bivariate CDF

$$F_C(x_1, x_2) \coloneqq C(F_1(x_1), F_2(x_2)) \tag{2.72}$$

having margins $F_1(\cdot)$ and $F_2(\cdot)$. Conversely, if $F(\cdot, \cdot)$ is a joint CDF with margins $F_1(\cdot)$ and $F_2(\cdot)$, then there exists a copula⁶⁷ $C(\cdot, \cdot)$, such that

$$F(x_1, x_2) = C(F_1(x_1), F_2(x_2)).$$

 $^{^{66}\}mathrm{The}$ generalization to d dimensions is straightforward.

⁶⁷If the margins are continuous, then it is unique.

Using a somewhat loose notation, we may therefore write

joint distribution = margins
$$\oplus$$
 copula.

For a more thorough introduction to copulas and their use in finance, we refer the reader to [CLV04], [MFE05] and [Nel06].

Provided that the densities of the margins and of the copula exist, the density of the CDF $F_C(\cdot, \cdot)$ exists and is given by

$$f_C(x_1, x_2) = c(F_1(x_1), F_2(x_2))f_1(x_1)f_2(x_2),$$

with $f_i(x)$, i = 1, 2 denoting the marginal densities and $c(\cdot, \cdot)$ the copula density. Calculating spread option prices as given in (2.68) is now only a matter of calculating a two-dimensional integral

$$V(0;T,K) = \iint (x_1 - x_2 - K)^+ f_C(x_1, x_2) dx_1 dx_2.$$

One of the most commonly used copulas is the so-called *Gaussian copula*, which represents the dependence structure of a bivariate Normal distribution and is defined by

$$C_{\text{Gauss}}(x_1, x_2) = \Psi \left(\Phi^{-1}(x_1), \Phi^{-1}(x_2); \rho \right),$$

where $\Psi(\cdot, \cdot; \rho)$ is the bivariate standard Normal CDF with correlation $\rho \in (-1, 1)$ and $\Phi(\cdot)$ is the (one-dimensional) standard Normal CDF. In practice, however, it will generally not be possible to fit market-observed spread-option prices across strikes with only one correlation parameter⁶⁸ ρ and other, more flexible copulas must therefore be used instead. For examples of parametric copula families, that are possibly better suited for pricing CMS spread options, we refer the reader to Chapter 17 of [AP10c].

Arguably the main reason why the copula approach is so popular in practice, is the ease with which the joint distribution of $S_1(T)$ and $S_2(T)$ can be parameterized and manipulated, while perfectly matching the market-implied marginal distributions. Market-observable prices of European spread options can therefore be easily inter- and extrapolated in an arbitrage-free way. However, using copulas is a totally static approach, where only the joint terminal distribution is of importance, while the dynamic behavior of the underlying rates and the spread process is neglected. Moreover, the copula parameters often lack any economic interpretation and may change in an unexpected way when the underlying rates move. It is therefore questionable whether a copula based model can be reliably used for hedging purposes.

Since we are mainly interested in pricing exotic CMS spread options and extracting market implied information about the dynamic behavior of interest rates, we will not consider the copula approach any further.

Remark 2.6.3. An interesting application of copulas is the derivation of certain noarbitrage bounds. More specifically, for a given set of single-rate European option prices, the so-called Fréchet-Hoeffding copula bounds can be used to derive lower and upper bounds for prices of spread- and other two-rate options. Moreover, superreplicating strategies can be constructed to exploit possible arbitrage opportunities. For more details see [CLV04] and [McC11].

 $^{^{68}\}text{Using different }\rho's$ for different strikes will lead to a *correlation smile* or *frown*.

2.6.3 More Exotic Products

The various interest-rate derivatives that we have introduced in the last sections were all of European type and their payoffs depended only on one or two rates. In the following we give a brief overview of more sophisticated (exotic) interest-rate products.

Bermudan Swaptions

A financial product is called *Bermudan* if it has multiple exercise dates, i.e., at prespecified time points T_i , the holder of such a product can choose between different payments or financial products. In this sense, Bermudan options line up in between European options, which can be "exercised" only at expiry T, and American options, which can be exercised at (almost) any time until expiry⁶⁹.

A Bermudan swaption is an option to enter into a vanilla fixed-for-floating swap at any (or any from a subset) of the swap fixing times, say $\{T_k\}_{k=0}^{N-1}$. If the option is exercised at time T_n , then the option expires and the holder enters into a swap with first fixing time T_n , last payment time T_N and fixed rate K, which is the strike of the Bermudan swaption. At time T_n , the value of the Bermudan option will therefore be the maximum of the swap value and a Bermudan swaption with exercise dates $\{T_k\}_{k=n+1}^{N-1}$. Bermudan swaptions are, by far, the most liquid exotic interestrate derivatives and are used, for example, by mortgage lenders to hedge against prepayment risks associated with home mortgage financing.

Exotic Coupons

In an *exotic swap* a regular floating Libor leg is swapped against a leg paying structured coupons $C_k = C_k(X_k(T_k))$, that are allowed to be arbitrary functions of observed interest rates $X_k(T_k)$, such as Libor or CMS rates. Some examples are:

• Capped and floored *floaters*,

$$C_k(x) = \max \{ \min \{ g \cdot x + s, c \}, f \},\$$

with gearing g, spread s, cap c and floor f.

• Capped and floored *inverse floaters*,

$$C_k(x) = \max \{ \min \{ s - g \cdot x, c \}, f \},\$$

with gearing g, spread s, cap c and floor f.

• Digitals,

$$C_k(x) = R \mathbb{1}_{[K,\infty)}(x)$$
 or $C_k(x) = R \mathbb{1}_{(-\infty,K]}(x),$

with coupon rate R and strike K.

⁶⁹Most interest-rate products are closely linked to a certain discrete tenor structure, on which the involved rates and fixing/payment times are defined. In a fixed-income setting Bermudan style options are therefore more natural and more common than American options.

• Range accruals,

$$C_{k} = R \frac{1}{N_{k}} \sum_{i=1}^{N_{k}} \mathbb{1}_{[l,u]} (X_{i}(T_{i})),$$

with coupon rate R, lower bound l and upper bound u. That is, the coupon is proportional to the number of days (in the coupon period) on which a certain reference rate stayed within a given range or corridor.

• *Multi-rate* coupons

$$C_k = gX_k^1(T_k)\mathbb{1}_{[l,\infty)}\left(X_k^2(T_k)\right)$$

with gearing g and level l, that is, a coupon proportional to one rate is paid if another rate is above some level l.

• Coupons with a *snowball* (also called *ratchet* or *ladder*) feature

$$C_k = \max \{\min \{C_{k-1} + Y_k, c\}, f\}$$

where C_{k-1} is the previous coupon and Y_k may be any of the coupons introduced above. Coupons of this form are called *path-dependent*, since they depend not only on current interest rates, but also on rate observations from previous coupon periods.

All coupons defined above can certainly also be linked to CMS spreads, rather than Libor or CMS rates. Notice that the first four coupon definitions involve only Europeantype options, which can be statically replicated⁷⁰ with vanilla options (see Prop. 2.6.1). Hence, by themselves, exotic swaps paying such coupons would not require the use of term-structure models for valuation or risk-management. This changes though, if these swaps are equipped for example with so-called callability-features, which we will introduce momentarily.

Before doing so, let us note that exotic swaps often emerge as part of bonds or notes, sold by banks to investors. Consider for example an investor, who invests a certain *principal amount* (e.g., 10 mil. euro) into a structured note, paying any of the above coupons. The issuer of the note receives the principal amount and invests it into a money market account, which pays the Libor rate plus or minus a spread. At maturity, the investor gets back his principal amount. From the perspective of the issuer, who pays the structured coupons and receives floating Libor, the net cash-flows of the note are equivalent to those of an exotic swap.

Callable Exotics

Structured notes are often made callable, that is, the issuer of the note has the right to cancel, or *call*, the note on a prespecified set of dates, which most often coincide with the fixing/payment times of the coupons⁷¹. If the note is called, the issuer returns the principal to the investor and no future coupon payments are made. Obviously, the issuer will only call the note if it is optimal from his point of view and suboptimal

52

 $^{^{70}\}mathrm{At}$ least in the single-rate case.

 $^{^{71}}$ In general there will be an initial *lockout* or *no-call* period (e.g., the first two years), where the issuer cannot call the note.

from the investor's point of view. As a compensation⁷², the investor can be offered, for instance, very high fixed initial coupons (cp. Section 7.3). Selling optionality to the issuer is a so-called *yield enhancement* strategy.

From the perspective of the issuer, a callable note is equivalent to a straight exotic swap plus a Bermudan-style option⁷³ to enter a swap with opposite cash-flows, such that all cash-flows cancel if the option is exercised.

The above list of exotic interest-rate derivatives is of course far from being exhaustive and we refer the reader to [BM05], [Fri07] and [AP10a] for other possible coupon and product specifications. Further examples of structured notes and also motivations for when and why such products might be interesting to investors can be found, e.g. in [Nef08] and [Tan10]. Some concrete examples will also be provided in Chapter 7.

2.7 Term-Structure Models: From Short-Rate Models to HJM

So far, our main focus has been on vanilla models, which specify the dynamics of only one or (with some tweaks) two interest rates. These models are sufficient for pricing relatively simple products such as for example European-type options on a single Libor or swap rate.

Many practically relevant securities, however, which involve for example pathdependent payoffs or early-exercise features, depend on the dynamics of the entire yield curve and not just on one or two points on it. Pricing and risk-managing such products therefore necessitates full term-structure models, which specify the joint arbitrage-free evolution of all points on the yield curve.

In Section 2.2 we have seen that the different types of interest rates are related through certain formulas involving zero-bond prices and are hence in some sense equivalent. Therefore, there is a certain arbitrariness in choosing a particular "coordinate system" for describing the arbitrage-free evolution of the yield curve, and we could take for example the short rate, instantaneous forward rates or a set of Libor rates as our basic modeling quantities. In fact, all these choices correspond to different classes of interest-rate models. Despite the conceptual equivalence, it should be clear that choosing a particular coordinate system will have implications for how easily a model can be calibrated to market-observable prices.

2.7.1 Short-Rate Models

Historically, the first dynamic term-structure models were *short-rate models* with arguably the most prominent examples being the Vasicek [Vas77] and the Cox-Ingersoll-Ross [CIR85] model. In this model class, the dynamics of the short-rate under the

 $^{^{72}\}mathrm{Note}$ that the investor implicitly sells the issuer an option.

⁷³Essentially, it is a Bermudan swaption, where the underlying of the swaption is not a vanilla fixed-for-floating swap but an exotic swap.

risk-neutral measures 74 \mathbb{Q} are generally assumed to be of the form

$$dr(t) = \mu(t, r(t))dt + \sigma(t, r(t))dW(t),$$

where $\mu(t, r)$ and $\sigma(t, r)$ are well-behaved functions and W(t) is a one-dimensional Brownian motion under \mathbb{Q} . In Table 2.1 we give some classical examples of short-rate models.

Once the short-rate dynamics have been specified, discount bond prices are given in terms of the fundamental formula

$$P(t,T) = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-\int_t^T r(s)ds} \right], \qquad (2.73)$$

which we have already encountered in Section 2.3. Short-rate models are generally quite tractable and for many models listed in Table 2.1, the above expectation can be calculated analytically. Depending on the particular model, also prices of caplets and swaptions may be given in closed or semi-closed form, which simplifies the calibration to market-observable prices.

Model	Dynamics
Vasicek (1977)	$dr(t) = \kappa(\theta - r(t))dt + \sigma dW(t)$
Dothan (1978)	$dr(t) = \kappa r(t)dt + \sigma r(t)dW(t)$
Brennan & Schwartz (1979)	$dr(t) = \kappa(\theta - r(t))dt + \sigma r(t)dW(t)$
Cox, Ingersoll & Ross (1985)	$dr(t) = \kappa(\theta - r(t))dt + \sigma\sqrt{r(t)}dW(t)$
Ho & Lee (1986)	$dr(t) = \theta(t)dt + \sigma dW(t)$
Hull & White (1990)	$dr(t) = \kappa(t)(\theta(t) - r(t))dt + \sigma(t)dW(t)$
	$dr(t) = \kappa(t)(\theta(t) - r(t))dt + \sigma(t)\sqrt{r(t)}dW(t)$
Black & Karasinski (1991)	$dr(t) = r(t)(\theta(t) + \sigma^2/2 - \kappa \log r(t))dt + \sigma r(t)dW(t)$
Pearson & Sun (1994)	$dr(t) = \kappa(\theta - r(t))dt + \sigma\sqrt{r(t) - \beta} dW(t)$

Table 2.1: Classical (one-factor) short-rate models. Here, κ , θ , σ and β denote scalar (possibly deterministically time-dependent) model parameters.

One of the main drawbacks of the earlier (time-homogeneous) models is, that they are not capable of fitting the initial yield curve exactly. Clearly, using just three parameters, say κ , θ , σ , it will generally be impossible to obtain a perfect match between the (theoretical) model prices $P_{\text{mod}}(0, T_i)$ and the observed market prices $P_{\text{mkt}}(0, T_i)$, $i = 1, \ldots, N$, where typically $N \gg 3$. Models of this type are therefore also referred to as "endogenous term-structure models", meaning that the initial yield curve is a model output rather than an input. This situation is somewhat unsatisfactory in that it gives rise to arbitrage opportunities when such a model is used for pricing and hedging derivatives (even though internally the model is consistent and arbitrage-free). One way to solve this problem is to increase the degrees of freedom of the model by allowing for time-dependent parameters. Famous examples of this

⁷⁴In the original works, model dynamics were often derived via equilibrium considerations or by first specifying the dynamics under the real-world measure, and then changing to the risk-neutral measure via a certain market price of risk. We instead follow the more modern approach and start by specifying the dynamics under the risk-neutral measure directly.

approach are the Hull & White [HW90] extended versions of the Vasicek and the CIR model. If the only concern is matching the initially observed bond prices, then it is usually sufficient to make just one parameter time-dependent. In most cases this will be the "level parameter" θ , which will then be a function directly given in terms of the initial discount curve⁷⁵; see also Equation (2.78) below. Nevertheless, often also the volatility parameter σ is allowed to depend on time, which helps improving the fit of the model to market-observable prices of caps and swaptions.

Before the reader starts assuming that introducing time-dependent parameters is a panacea, we should note that this approach has its limitations and should not be taken too far. Trying to perfectly match a given set of market prices with timedependent parameters may result in a strongly time-inhomogeneous model, implying a rather unrealistic evolution of future forward rates and volatilities⁷⁶. As a consequence, model-implied prices and hedging strategies for exotic products might not be reliable.

Gaussian short-rate models such as the Vasicek model or its popular Hull-White extension are generally fairly tractable. A drawback of these models is, however, that they allow interest rates to become negative. Given the prevalence of log-normal distributions in derivatives pricing theory, it should therefore come as no surprise that many authors have attempted to introduce short-rate models with log-normallike behavior, i.e., models where the dynamics of the short-rate r(t) are of the form

$$dr(t) = O(dt) + \sigma(t)r(t)dW(t)$$

for some deterministic function $\sigma(t)$. Well-known examples are the models due to Dothan [Dot78] and Black & Karasinski [BK91], see also Table 2.1. Although lognormal models guarantee interest rates to stay positive, they generally do not have closed-form pricing formulas for zero-coupon bonds or interest-rate options⁷⁷, which substantially reduces their tractability. An even more severe drawback, shared by all log-normal short-rate models, is the fact that the expected return of investing in the continuously compounded money market account over any positive time interval

$$\mathbb{E}_{t}^{\mathbb{Q}}\left[\frac{B(t+\Delta t)}{B(t)}\right] = \mathbb{E}_{t}^{\mathbb{Q}}\left[e^{\int_{t}^{t+\Delta t} r(s)ds}\right]$$
(2.74)

will be infinite⁷⁸, see [SS97]. This disqualifies log-normal short-rate models from being used in many applications.

Multi-Factor Short-Rate Models

The short-rate models that we have considered so far are all *one-factor models*, driven by only a single Brownian motion. The instantaneous correlations between

 $^{^{75}}$ The yield curve is now an input rather than an output and we therefore have an *exogenous term-structure model*.

⁷⁶See also the discussions in Chapter 3 of [BM05] and Chapter 10 of [AP10b].

⁷⁷The Dothan model is the only log-normal short-model with an analytical formula for zero-coupon bonds. However, the formula involves a double integral with modified Bessel functions, so the advantage of having an "explicit" formula is dramatically reduced.

⁷⁸Roughly speaking, the problem is that if r(t) is log-normal then we have a "double exponential" expression inside the expectation on the right hand side of (2.74).

forward rates of different maturities are hence all equal to one and the points along the forward-rate curve always move in the same direction. This is contrary to what can be observed in reality⁷⁹ and precludes these models from being used for pricing and risk-managing products, that strongly depend on correlations across the term structure of forward rates. Accordingly, in order to properly deal with securities that depend on "non-parallel" moves of the forward-rate curve, we need to extent the above models to allow for multiple Brownian drivers. Often this can be done by following an additive approach, where the short-rate is assumed to be the sum of a collection of one-dimensional processes. One very popular example is the two-factor additive Gaussian model (also known as G2++), which is a generalization of the Vasicek model. Here, the short-rate under the risk-neutral measure is modeled as (see [BM05], p. 143)

$$r(t) = x(t) + y(t) + \varphi(t),$$

where the processes x(t) and y(t) satisfy

$$dx(t) = -ax(t)dt + \sigma dW_1(t), \ x(0) = 0,$$

$$dy(t) = -by(t)dt + \eta dW_2(t), \ y(0) = 0,$$

where $(W_1(t), W_2(t))$ is a two-dimensional Brownian motion with instantaneous correlation $\rho \in (-1, 1)$

$$dW_1(t)dW_2(t) = \rho dt,$$

and where a, b, σ, η are positive constants. The deterministic function $\varphi(t)$ would then be chosen such that the model matches the initially observed yield curve. For more general formulations of d-factor short-rate models see Chapter 12 of [AP10b].

Multi-factor models are not only more realistic in capturing the main characteristics of yield-curve movements, but they are also more flexible in that they can generally better fit market-observed prices of caplet and swaption volatilities, without relying too heavily on time-dependent parameters. On the other hand, however, multi-factor short-rate models are numerically substantially more demanding than their one-factor counterparts. Analytical tractability is therefore often key to making such models applicable in practice. This is the main reason why Gaussian models like the G2++ model are so popular, despite their unpleasant feature of allowing for negative rates. A CIR2++ model, which would guarantee rates to stay positive, would be distinctively more difficult to handle numerically.

A major disadvantage of short-rate models in general is that these models are formulated in terms of an artificial non-observable quantity. Dynamic properties of, say, Libor and swap rates can be manipulated only indirectly through the short-rate parameters. In particular, it is often difficult to precisely control the decorrelation among different points along the forward-rate curve (no matter how many factors are used), which makes short-rate models only of limited use for pricing and riskmanaging CMS spread-linked securities. We therefore stop our review of short-rate models at this point and refer the reader to [BM05] and [AP10b] for more details on this model class.

⁷⁹See Chapter 7.

2.7.2 The HJM Framework

As we have noted in the introduction of this section, specifying the arbitrage-free evolution of the yield curve can be done in various ways, of which using the short-rate process is only one example. Heath, Jarrow and Morton (HJM) [HJM92] take the instantaneous forward rates⁸⁰

$$f(t,T) = -\frac{\partial \log P(t,T)}{\partial T},$$

as the fundamental quantities to model and assume, that under the risk-neutral measure \mathbb{Q} , the forward rate $f(\cdot, T)$ for each $T \in [0, T^*]$ evolves according to

$$df(t,T) = \alpha(t,T)dt + \sigma_f(t,T)'dW(t), \ 0 \le t \le T,$$
(2.75)

where W(t) is a d-dimensional Q-Brownian motion and $\alpha(\cdot, T)$ resp. $\sigma_f(\cdot, T)$ are oneresp. d-dimensional adapted processes. Observe that here the entire instantaneous forward-rate curve is used as the underlying "state variable" and Equation (2.75) indeed represents an infinite dimensional system of SDEs. The main achievement of Heath, Jarrow and Morton [HJM92] is usually considered to be the so-called *HJM drift condition*:

Proposition 2.7.1. Assume that under the risk-neutral measure \mathbb{Q} the forward-rate dynamics are given by (2.75), then we have

$$\alpha(t,T) = \sigma_f(t,T)' \int_t^T \sigma_f(t,s) ds, \ 0 \le t \le T \le T^*, \ \mathbb{Q} - a.s.,$$
(2.76)

so that the integrated dynamics of f(t,T) under the risk-neutral measure are

$$f(t,T) = f(0,T) + \int_0^t \sigma_f(u,T)' \int_u^T \sigma_f(u,s) ds du + \int_0^t \sigma_f(u,T)' dW(u).$$
(2.77)

This demonstrates that, contrary to short-rate models, where we are free⁸¹ to choose the diffusion and the drift coefficient, an HJM model is fully specified once the diffusion coefficients $\sigma_f(t,T)$ have been specified for all t and T. A clear advantage of these models is that they take the initial forward-rate curve f(0,T), $0 \leq T \leq T^*$ as an exogenous input and therefore perfectly fit the initial yield curve by construction (i.e., we have an exogenous model). A drawback, however, is the sheer dimensionality of the model: In order to describe the underlying state variable at time t, we need to keep track of a continuum of instantaneous forward rates $\{f(t,T), t \leq T \leq T^*\}$. Note also, that Equation (2.77) constitutes an arbitrage-free framework rather than a specific model, as we can almost arbitrarily⁸¹ choose the diffusion coefficients $\sigma(t,T)$. In fact, any diffusive arbitrage-free interest-rate model can be derived as a special case of the HJM framework. This includes the short-rate models from above just as well as the so-called market models that we will consider in the next chapter.

⁸⁰See also Equation (2.6).

⁸¹Subject to regularity conditions.

Note, for example, that from Equation (2.77) it follows that the short rate r(t) under \mathbb{Q} is given by

$$r(t) = f(t,t) = f(0,t) + \int_0^t \sigma_f(u,t)' \int_u^t \sigma_f(u,s) ds du + \int_0^t \sigma_f(u,t)' dW(u).$$

Suppose now that d = 1 and that the volatilities of the instantaneous forward rates are given by

$$\sigma_f(t,T) = \sigma e^{-\kappa(T-t)}.$$

Then, straightforward calculations show⁸² that the dynamics of the short rate under \mathbb{Q} are given by

$$dr(t) = \kappa(\theta(t) - r(t))dt + \sigma dW(t),$$

with

$$\theta(t) = \frac{1}{\kappa} \frac{\partial}{\partial T} f(0,T) \Big|_{T=t} + f(0,t) + \frac{\sigma^2}{2\kappa^2} \left(1 - e^{-2\kappa t}\right).$$
(2.78)

That is, we have obtained a version of the Hull-White extended Vasicek model, where the time-dependent level parameter $\theta(t)$ ensures that the initial yield curve is recovered.

Even though we can, in principle, derive any arbitrage-free interest-rate model⁸³ by appropriately specifying the volatility coefficients $\sigma(t, T)$ in the HJM framework, this is generally not the most natural way to derive a specific model. Moreover, without any further restrictions on the (possibly state-dependent) coefficients $\sigma(t, T)$, a particular choice will generally lead to a model that is not Markovian in a finite number of state variables, i.e., we have to keep track of a continuum of processes. There exist however, certain conditions on the volatility structure of forward rates under which the short rate r(t) is either outright Markov or at least can be written in terms of a finite-dimensional Markovian vector of state variables, see e.g., Carverhill [Car94] and Ritchken & Sankarasubramanian [RS95]. This is important for practical applications, as the models can then be efficiently implemented using, for instance, lattice-based methods.

As a final remark let us note that in order to avoid negative forward rates one might again be tempted to consider volatility specifications of the form

$$\sigma_f(t,T) = \sigma(t,T)f(t,T)$$

for some deterministic function $\sigma(t, T)$, in which case the forward rate f(t, T) would be log-normally distributed under the *T*-forward measure. However, similar to lognormal short-rate models, this specification suffers from severe technical problems. More specifically, forward rates will explode to infinity with positive probability and zero-coupon bond prices are hence all equal to zero, leading to obvious arbitrage opportunities. Sandmann and Sondermann [SS97] observed that these log-normal explosion problems can be avoided altogether by shifting from continuously-compounded to simply-compounded interest rates. This ultimately laid the foundations for the so-called Libor market models, which we will consider in the next chapter.

⁸²See e.g. [Fri07], p. 359.

⁸³As long as it is driven by Brownian motions.

Chapter 3

The Libor Market Model

As we have argued earlier, working with interest rates based on continuous compounding, such as the short-rate r(t) or forward rates f(t,T), is not particularly attractive in practical applications. For one, these rates are neither directly observable in the market, nor do they directly specify the payoff of any traded derivative contracts. This renders the calibration to prices of market-observable instruments often cumbersome and does not allow for an intuitive interpretation of the corresponding model parameters. Moreover, none of the short-rate or low-factor Markovian HJM models is compatible with the standard market-practice of using Black-type formulas for pricing vanilla interest-rate options.

In the late 1990's the focus therefore shifted from the unobservable instantaneous rates to quoted market rates and a new class of so-called *market models* emerged. While the pioneering works by Miltersen, Sandmann & Sondermann [MSS97] and Brace, Gątarek & Musiela (BGM) [BGM97] focused on Libor rates, Jamshidian [Jam97] later extended the approach also to a swap-rate context. One of the most striking feature of market models, and eventually the main reason for their success, is their consistency with the market practice of pricing caps, floors and swaptions by means of Black-type formulas, while at the same time being proper arbitrage-free term-structure models¹. Accordingly, these models can be calibrated to a set of plain-vanilla options "virtually by inspection"².

In the following section we will introduce the log-normal *Libor market model* (LMM), which is usually considered as *the* LMM³. Several approaches have appeared in the literature for deriving the no-arbitrage dynamics of forward Libor rates: See, for instance, the original articles mentioned above as well as Musiela & Rutkowski [MR97] or Rutkowski [Rut99]. While the various derivations are by and large equivalent, they nevertheless display significant differences⁴. In the earlier works, Libor-rate dynamics were often derived within the HJM framework by explicitly specifying (a continuum of) bond-price or instantaneous forward-rate volatilities. Showing that the LMM falls within the HJM model class, however, represents an unnecessary extra burden. We shall therefore follow the more modern *forward-measure approach*

²See [Reb02].

¹At that time this was particularly remarkable in the light of the problems one was experiencing with log-normal short-rate and forward-rate models.

³Sometimes also referred to as BGM model.

⁴For a nice overview of the different approaches see the survey article by Rutkowski [Rut01].

as presented e.g. in [Rut99] and work only with a finite collection of forward Libor rates.

3.1 Model Set-Up and No-Arbitrage Dynamics

In the following let $0 \leq T_0 < T_1 < \ldots < T_N$ be a discrete tenor structure together with a sequence $\tau_n \coloneqq T_{n+1} - T_n$, $n = 0, \ldots, N-1$ of year fractions. Typically the τ_n 's will be either set to (roughly) 0.25 or 0.5 (corresponding to 3 or 6 months), depending on the accrual period of the associated observable Libor rates. Rather than keeping track of the entire yield curve, we focus only on the evolution of a finite collection of Libor forward rates

$$L_n(t) \coloneqq \frac{1}{\tau_n} \left(\frac{P(t, T_n)}{P(t, T_{n+1})} - 1 \right), 0 \le t \le T_n, \quad 0 \le n \le N - 1.$$

Notice that the forward Libor rate $L_n(t)$ expires at time T_n , so that at a given time instant $t \in [0, T_{N-1}]$ only the forward Libor rates with indices $n \geq \eta(t)^5$ are still "alive".

Further, let $\{\sigma_n(t)\}_{n=0}^{N-1}$ be a collection of deterministic (bounded) *d*-dimensional instantaneous volatility functions, for some fixed $d \in \{1, \ldots, N\}$. As shown in Lemma 2.3.1, $L_n(t)$ must be a martingale under the T_{n+1} -forward measure \mathbb{Q}^{T_n+1} . Therefore, we may postulate that for $n \in \{0, \ldots, N-1\}$, the forward Libor rate $L_n(t)$ has the following driftless dynamics under the respective forward measure \mathbb{Q}^{T_n+1} :

$$dL_n(t) = L_n(t)\sigma_n(t)'dW^{n+1}(t), \ 0 \le t \le T_n,$$
(3.1)

where $W^{n+1}(t) := W^{T_{n+1}}(t)$ is a *d*-dimensional standard Brownian motion under $\mathbb{Q}^{T_{n+1}}$. The unique strong solution of the SDE (3.1) is then easily seen to be

$$L_n(t) = L_n(0) \exp\left(-\frac{1}{2}\int_0^t \|\sigma_n(s)\|^2 ds + \int_0^t \sigma_n(s)W^{n+1}(s)\right), \ 0 \le t \le T_n, \quad (3.2)$$

and hence $L_n(t)$ is log-normally distributed for all $0 \le t \le T_n$.

In order to fully specify the model (and to establish existence of the model), we need to find the joint dynamics of all forward Libor rates under one common measure:

Proposition 3.1.1. Let $k \in \{1, ..., N\}$ be fixed and assume that $L_n(t)$, $0 \le t \le T_n$, n = 0, ..., N-1 satisfy (3.1). Then the following relations for the Libor dynamics under the forward measure \mathbb{Q}^{T_k} hold:

$$n > k - 1: \quad dL_n(t) = L_n(t)\sigma_n(t)' \left(\sum_{i=k}^n \frac{\tau_i L_i(t)}{1 + \tau_i L_i(t)}\sigma_i(t)dt + dW^k(t)\right), \tag{3.3}$$

$$n < k - 1: \quad dL_n(t) = L_n(t)\sigma_n(t)' \left(-\sum_{i=n+1}^{k-1} \frac{\tau_i L_i(t)}{1 + \tau_i L_i(t)} \sigma_i(t) dt + dW^k(t) \right), \quad (3.4)$$

where $0 \leq t \leq \min\{T_n, T_k\}$ and $W^k(t)$ is a d-dimensional standard Brownian motion under \mathbb{Q}^{T_k} . The above system of SDEs admits a unique strong solution.

⁵See Equation (2.14) for the definition of the index function $\eta(t)$.
Proof. First we consider the case n > k - 1. In order to derive the Radon-Nikodym derivative of \mathbb{Q}^{T_n} with respect to $\mathbb{Q}^{T_{n-1}}$, we employ the change-of-numeraire technique (see p. 9)

$$\begin{aligned} \zeta(t) &= \frac{d\mathbb{Q}^{T_{n-1}}}{d\mathbb{Q}^{T_n}} \Big|_{\mathcal{F}_t} \\ &= \frac{P(t, T_{n-1})P(0, T_n)}{P(t, T_n)P(0, T_{n-1})}, \\ &= \frac{1 + \tau_{n-1}L_{n-1}(t)}{1 + \tau_{n-1}L_{n-1}(0)}, \ 0 \le t \le T_{n-1}. \end{aligned}$$

From this formula we can easily compute the dynamics of $\zeta(t)$ under \mathbb{Q}^{T_n}

$$\begin{split} d\zeta(t) &= \frac{\tau_{n-1}}{1 + \tau_{n-1}(0)} dL_{n-1}(t) \\ &= \frac{\tau_{n-1}L_{n-1}(t)}{1 + \tau_{n-1}L_{n-1}(0)} \sigma_{n-1}(t)' dW^n(t) \\ &= \zeta(t) \frac{\tau_{n-1}L_{n-1}(t)}{1 + \tau_{n-1}L_{n-1}(t)} \sigma_{n-1}(t)' dW^n(t). \end{split}$$

Observe that the Girsanov kernel

$$\varphi(t) \coloneqq \frac{\tau_{n-1}L_{n-1}(t)}{1 + \tau_{n-1}L_{n-1}(t)}\sigma_{n-1}(t)$$

is bounded and hence satisfies the Novikov condition (cp. [Bjö09], p. 167)

$$\mathbb{E}^{T_n}\left[\exp\left(\frac{1}{2}\int_0^{T_{n-1}}\|\varphi(t)\|^2dt\right)\right]<\infty.$$

From the Girsanov theorem (cp. [Bjö09], p. 164) it then follows that

$$dW^{n}(t) = \frac{\tau_{n-1}L_{n-1}(t)}{1 + \tau_{n-1}L_{n-1}(t)}\sigma_{n-1}(t)dt + dW^{n-1}(t),$$

where $W^{n-1}(t)$ is a *d*-dimensional standard Brownian motion under $\mathbb{Q}^{T_{n-1}}$. Applying this inductively we obtain

$$dW^{n+1}(t) = \sum_{i=k}^{n} \frac{\tau_i L_i(t)}{1 + \tau_i L_i(t)} \sigma_i(t) dt + dW^k(t)$$

and hence

$$dL_n(t) = L_n(t)\sigma_n(t)'dW^{n+1}(t)$$

= $L_n(t)\sigma_n(t)'\left(\sum_{i=k}^n \frac{\tau_i L_i(t)}{1 + \tau_i L_i(t)}\sigma_i(t)dt + dW^k(t)\right),$

which is (3.3). The corresponding result for n < k - 1 follows similarly.

Finally, by applying Itô's formula we obtain

$$n > k - 1: \ d \log L_n(t) = \sigma_n(t)' \left(\sum_{i=k}^n \frac{\tau_i L_i(t)}{1 + \tau_i L_i(t)} \sigma_i(t) dt - \frac{1}{2} \sigma_n(t) dt + dW^k(t) \right),$$
$$n < k - 1: \ d \log L_n(t) = \sigma_n(t)' \left(-\sum_{i=n+1}^{k-1} \frac{\tau_i L_i(t)}{1 + \tau_i L_i(t)} \sigma_i(t) dt - \frac{1}{2} \sigma_n(t) dt + dW^k(t) \right).$$

Existence and uniqueness of a strong solution to this system of SDEs is ensured (see e.g. $[\emptyset ks03]$) by the fact that the diffusion coefficients are deterministic and bounded, and the drift coefficients are bounded and continuous functions of the forward Libor rates.

Corollary 3.1.1. Under the assumptions of the above proposition and for n > k, the following relationship between the Brownian motions $W^n(t)$ and $W^k(t)$ under the respective forward measures \mathbb{Q}^{T_n} and \mathbb{Q}^{T_k} holds:

$$dW^{n}(t) = \sum_{i=k}^{n-1} \frac{\tau_{i}L_{i}(t)}{1 + \tau_{i}L_{i}(t)} \sigma_{i}(t)dt + dW^{k}(t).$$

Observe that the numeraire process $P(\cdot, T_k)$ associated with the T_k -forward measure is only alive up to time T_k . Accordingly, the Libor rates can be evolved under this measure only up to time $t = T_k$. In most cases one therefore works either under the so-called *terminal measure*, i.e., the forward measure \mathbb{Q}^{T_N} , associated with the last zero bond defined on the tenor structure, or under the spot Libor measure \mathbb{Q}^{B_d} (see Section 2.3), involving repeatedly rolling over in the bond with the shortest time to maturity available. In both cases the numeraire process remains alive throughout the time span of the tenor structure $\{T_n\}_{n=0}^N$. This is necessary for the evaluation of derivative securities that may involve random payoffs at any date in the tenor structure. The forward Libor dynamics under the spot Libor measure are given by the following

Proposition 3.1.2. Let $L_n(t)$, $0 \le t \le T_n$, n = 0, ..., N-1 satisfy (3.1). The forward Libor dynamics under the spot Libor measure \mathbb{Q}^{B_d} are given by

$$dL_n(t) = \sigma_n(t)' \left(\sum_{i=\eta(t)}^n \frac{\tau_i L_i(t)}{1 + \tau_i L_i(t)} \sigma_i(t) dt + dW^{B_d}(t) \right),$$

where $W^{B_d}(t)$ is a d-dimensional standard Brownian motion under \mathbb{Q}^{B_d} .

Proof. Recall from Section 2.3 that the numeraire process associated with the spot Libor measure is given by

$$B_d(t) = P(t, T_{\eta(t)}) \prod_{i=0}^{\eta(t)-1} (1 + \tau_i L_i(T_i)).$$

At any time t the random part of this process is the factor $P(t, T_{\eta(t)})$ (all Libor rates with index $i < \eta(t)$ will have been fixed by time t). So essentially just we need to derive the dynamics under the measure $\mathbb{Q}^{T_{\eta(t)}}$. Starting from the dynamics under $\mathbb{Q}^{T_{n+1}}$ and proceeding inductively as in the proof of Proposition 3.1.1 we obtain

$$dW^{n+1}(t) = \sum_{i=\eta(t)}^{n} \frac{\tau_i L_i(t)}{1 + \tau_i L_i(t)} \sigma_i(t) dt + dW^{\eta(t)}(t)$$

as stated.

An immediate consequence of the log-normal property (3.2) is that caplets can be priced within the log-normal LMM by means of Black's formula. The model can therefore be trivially calibrated to a collection of market-implied caplet volatilities $\{\sigma_n^{\text{mkt}}\}_{n=0}^{N-1}$ by choosing the instantaneous Libor volatilities $\{\sigma_n(t)\}_{n=0}^{N-1}$ such that (cp. Equation (2.15) and Remark 2.4.1)

$$(\sigma_n^{\text{mkt}})^2 T_n = \int_0^{T_n} \|\sigma_n(t)\|^2 dt.$$
(3.5)

However, from the above propositions it is also easy to see that a forward Libor rate $L_n(t)$ is a log-normal martingale only under its respective forward measure. Put differently, there exists no measure under which all Libor rates are simultaneously log-normal. Furthermore, the systems of SDEs stated in the above propositions do not allow for closed-form analytical solutions and one therefore has to resort to simulation of discretized versions of the SDEs when pricing derivative securities, that depend on the joint evolution of several Libor rates. For a comparison of various numerical discretization schemes see, e.g. [JS08].

Remark 3.1.2. From a mathematical point of view working under the terminal measure or the spot measure is equivalent. When performing Monte Carlo simulations, however, it is often advantageous to use the spot measure, since the variance of the Monte Carlo error is then typically lower (see [Bra08], p. 42) and possible drift approximation errors are more evenly distributed among the rates (see [BM05], p. 219).

3.1.1 Swap-Rate Dynamics

While pricing caplets in a LMM is, by construction, a fairly easy task, pricing swaptions requires some more work as we will see in the following. Consider, for instance, a payer swaption with strike price K > 0, written on a swap covering the time interval $[T_m, T_n]$, where $0 \le m < n \le N$. In Section 2.4 we have seen that working under the swap measure $\mathbb{Q}^{m,n}$ we may write the time-zero price of such a swaption as (cp. Equation (2.17))

$$V_{\text{swptn}}(0) = A_{m,n}(0)\mathbb{E}^{m,n}\left[(S_{m,n}(T_m) - K)^+ \right].$$

For the dynamics of $S_{m,n}(t)$ under the measure $\mathbb{Q}^{m,n}$ we have (cp. [AP10b]):

Proposition 3.1.3. Let $L_i(t)$, $0 \le t \le T_m$, $i = m, \ldots, n-1$ satisfy (3.1). The dynamics of the swap rate $S_{m,n}(t)$ under the swap measure $\mathbb{Q}^{m,n}$ are given by

$$dS_{m,n}(t) = S_{m,n}(t) \sum_{i=m}^{n-1} w_i(t)\sigma_i(t)' dW^{m,n}(t), \ 0 \le t \le T_m,$$
(3.6)

where the stochastic weights are

$$\begin{split} w_i(t) &= \frac{L_i(t)}{S_{m,n}(t)} \cdot \frac{\partial S_{m,n}(t)}{\partial L_i(t)} \\ &= \frac{L_i(t)}{S_{m,n}(t)} \cdot \frac{\tau_i S_{m,n}(t)}{1 + \tau_i L_i(t)} \left(\frac{P(t,T_n)}{P(t,T_m) - P(t,T_n)} + \frac{A_{i,n}(t)}{A_{m,n}(t)} \right), \end{split}$$

and where $W^{m,n}(t)$ is a d-dimensional standard $\mathbb{Q}^{m,n}$ -Brownian motion.

Proof. Recall that by Lemma 2.3.2 $S_{m,n}(t)$ must be a martingale under the associated swap measure $\mathbb{Q}^{m,n}$. Further, note that $S_{m,n}(t)$ can be written as a function of the Libor rates $L_m(t), \ldots, L_{n-1}(t)$, see Equation (2.10). A straightforward application of Itô's lemma then gives the result.

The SDE (3.6) is rather complicated due to the stochastic weights $w_i(t)$ and does not allow for an analytical solution. Therefore, closed-form pricing of swaptions in a LMM is in general not possible. Nonetheless, having a closer look at the stochastic weights $w_i(t)$, one finds that these are generally slowly varying, such that one obtains as a reasonable approximation for the dynamics

$$dS_{m,n}(t) \approx S_{m,n}(t) \sum_{i=m}^{n-1} w_i(0)\sigma_i(t)' dW^{m,n}(t), \qquad (3.7)$$

in which case the swap rate follows again a log-normal martingale. Accordingly, prices of swaptions can be approximated by using Black's swaption formula (2.18) with volatility parameter (cp. Remark 2.4.1)

$$\sigma_{m,n} \coloneqq \sqrt{\frac{1}{T_m} \sum_{i,j=m}^{n-1} w_i(0) w_j(0) \int_0^{T_m} \sigma_i(t)' \sigma_j(t) dt}$$

This approximation, which was first introduced by Hull & White [HW00], is surprisingly accurate⁶ and allows to efficiently calibrate a LMM to market-observable swaption prices.

Swap Market Models

The standard LMM as introduced above is based on a set of contiguous forward Libor rates. In contrast, Jamshidian [Jam97] introduces so-called *swap market models* (SMM) which take a set of co-terminal⁷ swap rates $\{S_{0,N}(t), S_{1,N}(t), \ldots, S_{N-1,N}(t)\}$

 $^{^6 {\}rm See}$ e.g. [JR03a] and [Reb02] for an analysis of the approximation error. Various other approximations are analyzed in Chapters 6 and 8 of [BM05].

⁷I.e., swap rates that share the same final payment date.

as the fundamental building blocks. The approach was later extended to also cover the case of co-initial swap rates $\{S_{0,1}(t), S_{0,2}(t), \ldots, S_{0,N}(t)\}$, see e.g. [HK00] and [GH04]. Assuming that the swap rates follow log-normal martingales under their respective swap measures, swaptions written on the swap rates underlying a particular model can be priced by means of Black formulas. As a consequence, SMMs can be calibrated almost effortlessly to a set of market-observable swaptions prices. Similar to the LMM case it is possible to derive relations between swap-rate dynamics under different measures, such that the swap rates can be evolved under one common measure.

In principle, all products that can be priced with LMMs can also be priced with SMMs and vice versa. The reason is that swap rates can be written as weighted sums of Libor rates and Libor rates can be expressed as differences of certain swap rates. Nevertheless, the two model classes are not equivalent and they will generally not yield the same prices for a given set of caplets and swaptions. In particular, in a log-normal LMM swaptions cannot be priced (exactly) by means of Black's swaption formula (see above), whereas a log-normal SMM does not produce prices in line with Black's caplet formula⁸. This is due to the fact that Libor rates and swap rates (defined in terms of sums of Libor rates) cannot be simultaneously log-normal under their respective natural measures. The incompatibility of log-normal LMMs and SMMs though, is mostly theoretical. In fact, by means of simulation studies (see Chapter 8 of [BM05]) one can show that swap rates in a log-normal LMM are "very close" to being log-normal, which is also one reason why the approximation (3.7) works so well.

Even though SMMs feature some advantages, especially when it comes to pricing swap-based securities, LMMs are generally easier to handle both mathematically and numerically. Moreover, it is usually more natural to express swap rates in terms of Libor rates rather than doing the opposite, i.e., Libor rates are in some sense the more representative coordinates for describing the yield curve. As a consequence, the LMM has become the benchmark model for pricing both Libor- and swap-rate dependent products, and we will therefore not pursue the swap-rate based modeling approach any further.

3.2 A Stochastic-Volatility Extended LMM

The standard log-normal LMM as introduced in the previous section has become an essential tool for pricing and risk-managing complex interest-rate derivatives. One of its great advantages, and the main motivation for its development, has been the ease with which it can be calibrated to a grid of caplet and swaption volatilities. However, a trivial consequence of the log-normality is, that the model produces flat caplet and swaptions volatility smiles, which are obviously not compatible with the pronounced volatility smiles found in today's interest-rate markets.

Several extensions of the original LMM have therefore been proposed in the literature to incorporate the volatility smile effect by using, for instance, local volatilities

⁸In a log-normal SMM even swap rates that do not belong to the underlying set of swap rates will not be log-normal.

(Andersen & Andreasen [AA00]), jump diffusions or general Lévy processes (Glasserman & Kou [GK03], Eberlein & Özkan [EÖ05]) or stochastic volatilities (Andersen & Andreasen [AA02], Joshi & Rebonato [JR03b], Andersen & Brotherton-Ratcliffe [ABR05], Piterbarg [Pit03], [Pit05a]), to name just a few.

The models that are arguably the most popular examples of these extensions are based on Heston-type dynamics (see e.g. [AA02], [ABR05] or [Pit05a]). Models of this type are fairly tractable due to the existence of closed formulas for pricing European options and efficient simulation schemes (see e.g. [And08] and [VHP10]). Moreover, with the time-averaging techniques introduced in Section 2.5, it is possible to efficiently calibrate such models to the full term-structure of volatility smiles across the swaption grid.

In the next section we present the displaced Heston LMM with time-dependent skew functions as introduced by Piterbarg [Pit03], [Pit05a]. This LMM, which will be referred to as *stochastic-volatility LMM* (SV-LMM) from now on, contains the standard log-normal LMM as well as the *displaced-diffusion* (or *shifted log-normal*) LMM^9 as special cases. It will serve as the basis model for all further theoretical and empirical investigations in this thesis.

3.2.1 Model Description

In order to establish the SV-LMM we simply need to "embed" the displaced Heston model from Section 2.5.3 into the LMM framework. For this, let the stochastic variance process V(t) follow a square-root diffusion

$$dV(t) = \kappa (1 - V(t))dt + \xi \sqrt{V(t)} dZ(t), \ V(0) = 1,$$
(3.8)

with positive constants κ and ξ , and where Z(t) is a one-dimensional standard Brownian motion under a generic pricing measure \mathbb{Q} . The spanning forward Libor rates $L_n(t), \ 0 \le t \le T_n, \ n = 0, \dots, N-1$ are then assumed to have dynamics of the form

$$dL_n(t) = \varphi_n(t, L_n(t)) \sqrt{V(t)} \sigma_n(t)' \left(\sqrt{V(t)} \mu_n(t) dt + dW(t) \right), \qquad (3.9)$$

with

$$\varphi_n(t,L) = \beta_n(t)L + (1 - \beta_n(t))L_n(0),$$

and where W(t) is a d-dimensional standard Brownian motion under \mathbb{Q} , independent of Z(t), $\{\sigma_n(t)\}_{n=0}^{N-1}$ is a collection of d-dimensional (deterministic) instantaneous volatility functions for some fixed $d \in \{1, \ldots, N\}$, and $\{\beta_n(t)\}_{n=0}^{N-1}$ is a collection of one-dimensional skew functions¹⁰ $\beta_n : [0, T_n] \to (0, 1], n = 0, \ldots, N-1$. The numeraire-specific drift terms $\mu_n(t)$, which ensure that the model is arbitrage-free, can be derived similarly as in the standard LMM. For example, if \mathbb{Q} is the spot Libor measure \mathbb{Q}^{B_d} , then the drift terms are given by (see [AP10b])

$$\mu_n(t) = \sum_{i=\eta(t)}^n \frac{\tau_i \varphi_i(t, L_i(t))}{1 + \tau_i L_i(t)} \sigma_i(t), \ 0 \le t \le T_n, \ n = 0, \dots, N-1.$$

⁹Due to the aforementioned similarity between displaced diffusions and CEV processes (see Section 2.5.2), it may also serve as a tractable proxy for the CEV-LMM as introduced in [AA00].

¹⁰The range of the skew functions can be extended to (-1, 1], see [Pit03].

Note carefully that all Libor rates in (3.9) are simultaneously scaled by a single common stochastic-volatility factor $\sqrt{V(t)}$. Moreover, due to the independence of Z(t) and W(t), the dynamics of V(t) are unaffected by measure changes and remain the same under all pricing measures¹¹.

3.2.2 Pricing European Options

Calibrating a model by means of Monte Carlo simulations is usually prohibitively expensive. Therefore, in order for a model to be really useful in practice, fast and accurate approximations for caplet and swaption prices have to be found. In case of the SV-LMM (3.8)-(3.9), this task can be accomplished by using the parameter-averaging techniques introduced in Section 2.5.

Caplets

Deriving caplet pricing formulas in extended LMMs is generally similarly straightforward as in the standard LMM, as the underlyings of caplets – the Libor rates $L_n(t)$ – constitute the fundamental building blocks of every LMM. More concretely, observe that under the T_{n+1} -forward measure, $L_n(t)$ has driftless dynamics

$$dL_n(t) = \varphi_n(t, L_n(t))\sqrt{V(t)}\sigma_n(t)'dW^{n+1}(t), \qquad (3.10)$$

where $W^{n+1}(t)$ is a *d*-dimensional standard Brownian motion under $\mathbb{Q}^{T_{n+1}}$. Drawing on the results of Section 2.5, we find that in the SV-LMM the time-zero price of a caplet written on $L_n(T_n)$ with strike price K > 0, i.e.,

$$V_{\text{caplet}}(0, T_n) = \tau_n P(t, T_{n+1}) \mathbb{E}^{T_{n+1}} [(L_n(T_n) - K)^+],$$

can be well approximated by using the Fourier pricing formula (2.39) in the vanilla model

$$dL_n(t) = \bar{\lambda}_n(\bar{\beta}_n L_n(t) + (1 - \bar{\beta}_n) L_n(0)) \sqrt{V(t)} d\tilde{W}^{n+1}(t), \qquad (3.11)$$

$$dV(t) = \kappa (1 - V(t))dt + \xi \sqrt{V(t)} d\tilde{Z}(t), \ V(0) = 1,$$
(3.12)

with effective skew parameter

$$\bar{\beta}_n = \int_0^{T_n} \beta_n(t) w_{T_n}(t) dt,$$

$$w_{T_n}(t) = \frac{v_n(t)^2 \|\sigma_n(t)\|^2}{\int_0^{T_n} v_n(t)^2 \|\sigma_n(s)\|^2 dt},$$

$$v_n(t)^2 = \int_0^t \|\sigma_n(s)\|^2 ds + \xi^2 e^{-\kappa t} \int_0^t \|\sigma_n(s)\|^2 \frac{e^{\kappa s} - e^{-\kappa s}}{2\kappa} ds$$

with independent one-dimensional $\mathbb{Q}^{T_{n+1}}$ -Brownian motions $\tilde{W}^{n+1}(t)$ and $\tilde{Z}(t)$, and where the effective volatility parameter $\bar{\lambda}_n$ is obtained as the solution to the rootsearch problem (2.46).

¹¹See [AP10b], p. 602 for more details.

Swaptions

Now consider a swaption written on a swap covering the time interval $[T_m, T_n]$, where $0 \le m < n \le N$. The dynamics of the underlying swap rate $S_{m,n}(t)$ under the swap measure $\mathbb{Q}^{m,n}$ can be derived similarly as in Proposition 3.1.3 and are given by (see [AP10b])

$$dS_{m,n}(t) = \sqrt{V(t)} \sum_{i=m}^{n-1} w_i(t)\sigma_i(t)' dW^{m,n}(t), \ 0 \le t \le T_m,$$
(3.13)

where the stochastic weights are

$$w_i(t) = \varphi_i(t, L_i(t)) \frac{\partial S_{m,n}(t)}{\partial L_i(t)},$$

and where $W^{m,n}(t)$ is a *d*-dimensional standard $\mathbb{Q}^{m,n}$ -Brownian motion. By freezing the stochastic weights at their initial values (in the spirit of (3.7)) and assuming that the swap-rate dynamics are again approximately of displaced-Heston type, Piterbarg [Pit03] obtains as an accurate approximation for the dynamics (3.13)

$$dS_{m,n}(t) \approx \left(\beta_{m,n}(t)S_{m,n} + (1 - \beta_{m,n}(t))S_{m,n}(0)\right)\sqrt{V(t)}\sigma_{m,n}(t)'dW^{m,n}(t) \quad (3.14)$$

with

$$\sigma_{m,n}(t) = \sum_{i=m}^{n-1} q_i^{m,n} \sigma_i(t), \qquad (3.15)$$

$$\beta_{m,n}(t) = \sum_{i=m}^{n-1} p_i^{m,n} \beta_i(t), \qquad (3.16)$$

$$q_i^{m,n} = \frac{L_i(0)}{S_{m,n}(0)} \cdot \frac{\partial S_{m,n}(0)}{\partial L_i(0)},$$
(3.17)

$$p_i^{m,n} = \frac{\sigma_i(t)'\sigma_{m,n}(t)}{\|\sigma_{m,n}(t)\|^2}.$$
(3.18)

The derivation of the above quantities is straightforward but somewhat lengthy and is therefore omitted at this point, see [Pit03] for more details.

Note that the swap-rate dynamics (3.14) are of the same form as the exact SDE (3.10) for the Libor rate $L_n(t)$. Hence, the same formulas as in the caplet case can be used to derive effective (time-independent) parameters $\bar{\beta}_{m,n}$ and $\bar{\lambda}_{m,n}$, which in conjunction with Proposition 2.5.7 yield a semi-closed formula for the time-zero price of swaptions. For an analysis of the accuracy of this formula the reader is referred to [Pit03].

3.3 Parameterization of the LMM

So far our discussion of the LMM framework has been relatively generic in the sense that we did not make any further assumptions about the instantaneous Libor volatility functions $\sigma_n(t)$ and (in case of the SV-LMM) the skew functions $\beta_n(t)$. However, in order to make a model operational (and to reduce the degrees of freedom to a reasonable level), we need to impose some additional structure on these parameter functions.

3.3.1 Volatility Structure

For convenience we define the *absolute* or *scalar instantaneous volatility* $\lambda_n(t)$ of the forward Libor rate $L_n(t)$ by

$$\lambda_n(t) \coloneqq \|\sigma_n(t)\|, \ n = 0, \dots, N-1,$$

and the factor loadings $R_n(t)$ by

$$R_n(t) \coloneqq \frac{\sigma_n(t)}{\|\sigma_n(t)\|}, \ 0 \le t \le T_n, \ n = 0, \dots, N-1,$$

such that

$$\sigma_n(t) = \lambda_n(t) R_n(t), \ 0 \le t \le T_n, \ n = 0, \dots, N-1.$$

The instantaneous Libor correlations¹² $\rho(t)$ are then defined by

$$\rho_{ij}(t) \coloneqq R_i(t)' R_j(t), \ 0 \le t \le \min\{T_i, T_j\}, \ i, j = 0, \dots, N-1.$$
(3.19)

Let us assume for a moment that we want to calibrate a log-normal LMM to a collection of ATM caplet prices. As stated earlier, caplet prices are only affected by the (scalar) Libor volatilities (and not by Libor correlations) and we can perfectly recover the market prices by specifying the Libor volatilities $\lambda_n(t)$ such that (cp. Equation (3.5))

$$(\sigma_n^{\text{mkt}})^2 T_n = \int_0^{T_n} \lambda_n(t)^2 dt,$$

where σ_n^{mkt} denotes the implied ATM caplet volatility for maturity T_n . Obviously, there exists an infinite number of solutions to this problem with the trivial solution being $\lambda_n(t) \equiv \sigma_n^{\text{mkt}}$.

While for caplets, which are simple European options, only the total (integrated) variance of the Libor rates is of importance, prices of exotic products may indeed be affected by the time-dependent evolution of the Libor volatilities. The price of a complex derivative security is essentially given by the sum of hedging costs incurred throughout its life time. This includes not only the initial cost of setting up the hedging portfolio, but also all future re-hedging costs. So if caplets are used for dynamically hedging a certain product, then a model should not only recover today's caplet prices but also the future prices, i.e., volatilities.

Generally the overall shape of the term structure of caplet volatilities does not change too much over time. In particular, most of the time it can be observed that the volatility term structure is humped shaped¹³, as is demonstrated in Figure 3.1. Therefore, if we do not have any further information (or view) on the future evolution of caplet volatilities, it is generally most reasonable to require the future caplet volatilities term structure (as implied by the model) to look similar to today's. This means that the Libor volatility functions $\lambda_n(t)$ should be time-homogeneous¹⁴

 $^{^{12}\}mbox{I.e.},$ the correlation between forward-rate increments.

¹³For a financial explanation for the existence of a volatility hump see, e.g. [Reb04], p. 672.

¹⁴No matter whether it be a standard or an extended LMM.



Figure 3.1: Bootstrapped ATM euro caplet implied volatilities as of 1/14/2008.

(or time-stationary) in the sense that they are only functions of (residual) time to maturity $T_n - t$, rather than of calender time t and maturity T_n , i.e.,

$$\lambda_n(t) = g(T_n - t), \ 0 \le t \le T_n, \ n = 0, \dots, N-1,$$

for some function $g(\cdot)$. A parametric form for the function $g(\cdot)$ that has been commonly used in the past is due to [Reb99b]:

$$g(T-t) = (a+b(T-t))\exp(-c(T-t)) + d, \quad a,b,c,d \in \mathbb{R}_+.$$
(3.20)

This functional form, commonly known as *abcd*-parameterization, is capable of producing a wide range of empirically observed volatility term structures. Nevertheless, it is generally impossible to perfectly fit the initially observed caplet volatilities with a purely time-homogeneous parameterization¹⁵. To address this drawback several "separable" extensions of the type

$$\lambda_n(t) = f(T_n)h(t)g(T_n - t)$$

have been proposed in the literature¹⁶, where the functions f(T) and h(t) are usually required to be close to unity, such that the corresponding Libor volatilities are at the very least "almost" time-homogeneous.

However, relying on parametric or separable forms is often too restrictive and inflexible for practical applications, particularly if not only caplet but also swaption volatilities are to be used as calibration targets. Therefore, we shall take a more general approach and assume the Libor volatility functions $\lambda_n(t)$ to be piecewise constant functions of the form

$$\lambda_n(t) = \sum_{i=0}^n \mathbb{1}_{[T_{i-1}, T_i]}(t) \lambda_{ni}, \ \lambda_{ni} > 0,$$
(3.21)

¹⁵Note that in order to match the market-observed volatilities with a purely time-homogeneous $g(\cdot)$, we need to have $(\sigma_n^{\text{mkt}})^2 T_n - (\sigma_{n-1}^{\text{mkt}})^2 T_{n-1} = \int_{T_{n-1}}^{T_n} g(u)^2 du > 0, n = 1, \ldots, N-1$, which is not always satisfied in practice. So even with an arbitrarily flexible parameterization for $g(\cdot)$, a perfect fit can not always be achieved.

¹⁶See [BM05], Section 6.3 for an overview.

with the convention that $T_{-1} := 0$. In order to reduce the number of free parameters, we further parameterize the volatility functions via a "grid-based" approach (cp. [AP10b]). For this, let us define

$$\lambda(t, T_n - t) \coloneqq \lambda_n(t),$$

and introduce a grid of times and times to maturity

$$\{(t_i, \mathcal{T}_j) | i = 1, \dots, N_t, \ j = 1, \dots, N_{\mathcal{T}}\},\tag{3.22}$$

for some integers N_t and $N_{\mathcal{T}}$. Furthermore, let Λ be a $N_t \times N_{\mathcal{T}}$ -dimensional matrix, whose elements will be interpreted as

$$\tilde{\lambda}(t_i, \mathcal{T}_j) = \Lambda_{ij}.$$

With Λ we specify $\lambda_i(t) = \tilde{\lambda}(t, T_i - t)$ on the grid (t_i, \mathcal{T}_j) . The remaining function values for all *i* and *t* are then obtained by bilinear inter-/extrapolation¹⁷.

Observe that if we require all rows of Λ to be identical, then the corresponding Libor volatility functions will be fully time-homogeneous. In particular, the functional form (3.20) can be closely approximated by setting

$$\Lambda_{ij} = (a + b \mathcal{T}_j) \exp(-c \mathcal{T}_j) + d.$$

Even after "parameterizing" $\lambda_n(t)$ via Λ , the number of parameters might still be quite large (depending on N_t and N_T). In order to avoid overfitting and to obtain stable calibration results, we will therefore use suitable penalty functions (see next section), by means of which we can control how smooth and constant the rows and columns of the matrix Λ should be.

Parameterizing and calibrating the instantaneous Libor correlations $\rho(t)$ will be discussed extensively in Chapters 5–7. We shall therefore assume throughout the rest of this chapter, that these correlations (and hence the factor loadings $R_n(t)$) have been exogenously specified.

3.3.2 Skew Structure

After having made some structural assumptions about the Libor volatilities $\lambda_n(t)$, by means of which we can control the model implied ATM volatilities of caplets and swaptions, we are left with specifying the skew functions $\beta_n(t)$, which determine the skewness of the model-implied caplet and swaption volatility smiles.

As with the volatility functions, we will take a grid-based approach. More specifically, we assume that the skew functions are of the form

$$\beta_n(t) = \sum_{i=0}^n \mathbb{1}_{[T_{i-1}, T_i]}(t) \beta_{ni}, \ \beta_{ni} \in (-1, 1]$$
(3.23)

and define

$$\tilde{\beta}(t, T_n - t) \coloneqq \beta_n(t).$$

¹⁷Note, that the entries in the bottom right part of the matrix with $t_i + T_j > T_{N-1}$ are of course redundant.

For our applications we will always use the same grid (3.22) for both the skew and the volatility functions (although using different grids would clearly be possible). We proceed by defining a $N_t \times N_T$ -dimensional matrix B, whose elements will be interpreted as

$$\beta(t_i, \mathcal{T}_j) = B_{ij}.$$

The remaining function values of $\beta_n(t) = \tilde{\beta}(t, T_n - t)$ for all n and t are then again obtained by bilinear inter-/extrapolation. The regularity and time-homogeneity of these functions are controlled by means of certain penalty functions, which will be described in the next section.

3.4 Calibration

In this section we give a brief review of the general calibration procedure if only caplets and swaptions are used as calibration targets (cp. [AP10b]).

Suppose that, based on the exotic products that are later to be priced and riskmanaged with the model, we have fixed the tenor structure, have decided upon the number of factors d to be used and have exogenously specified the correlation structure¹⁸.

Depending on the characteristics of the exotic products, we also assume that a set of (market-observable) caplet and swaption smiles – which are considered to contain "information" about the price of the exotic product and are potentially used as hedging instruments – has been selected.

3.4.1 Pre-Calibration: Effective Swaption and Caplet Parameters

The first step of the calibration procedure consists of the so-called *pre-calibration*, where for each caplet smile an effective volatility and an effective skew parameter is extracted. More precisely, if $X(\cdot)$ denotes the underlying process of one of the caplets or swaptions from the calibration set, then we assume that its dynamics (under the appropriate forward/swap measure) are of displaced Heston type:

$$dX(t) = (\bar{\beta}_x X(t) + (1 - \bar{\beta}_x) X(0)) \bar{\sigma}_x \sqrt{V(t)} dU(t),$$

$$dV(t) = \kappa (1 - V(t)) dt + \xi \sqrt{V(t)} dZ(t), \quad V(0) = 1,$$

where U(t) and Z(t) are independent one-dimensional standard Brownian motions. These (separate) vanilla models are then calibrated¹⁹ to the implied volatility smiles. When solving for the market implied smile parameters $\bar{\sigma}_x^*$ and $\bar{\beta}_x^*$, we simultaneously optimize for a "global" speed of mean reversion κ^* and volatility of variance ξ^* , such that all market smiles are matched as well as possible.

 $^{^{18}\}mathrm{See}$ Chapter 5 and 7 for more information on choosing the number of factors and the correlation structure.

¹⁹By least-square fitting implied volatilities or weighted prices.

3.4.2 Main Calibration: Time-Dependent Parameters

The purpose of the main calibration is to obtain the time-dependent model parameters $\lambda_n(t)$ resp. $\beta_n(t)$ in (3.21) resp. (3.23) from the effective parameters determined in the pre-calibration.

Suppose that the grid of times and times to maturities (3.22) has been fixed²⁰. For given matrices Λ and B, the forward-rate parameters $\beta_n(\cdot)$ and $\sigma_n(\cdot)$ and the (model) swap-rate parameters²¹ $\beta_{m,n}(\cdot)$ and $\sigma_{m,n}(\cdot)$ can be related to effective (constant) model parameters $\bar{\sigma}_x$ and $\bar{\beta}_x$ by using the parameter-averaging formulas from Section 3.2. As such, $\bar{\sigma}_x$ and $\bar{\beta}_x$ can be considered as functions of Λ and B and calibrating the model ultimately comes down to performing a least-squares-style minimization

$$(\Lambda^*, B^*) = \underset{\Lambda, B}{\operatorname{argmin}} \left\{ \sum_x w_x^{\sigma} \left(\bar{\sigma}_x(\Lambda, B) - \bar{\sigma}_x^* \right)^2 + \sum_x w_x^{\beta} \left(\bar{\beta}_x(\Lambda, B) - \bar{\beta}_x^* \right)^2 \right\}$$
(3.24)

$$+\sum_{i,j} \left[w_t^{\lambda} \left(\frac{\partial \Lambda_{ij}}{\partial t_i} \right)^2 + w_{\mathcal{T}}^{\lambda} \left(\frac{\partial \Lambda_{ij}}{\mathcal{T}_j} \right)^2 + w_{t^2}^{\lambda} \left(\frac{\partial^2 \Lambda_{ij}}{\partial t_i^2} \right)^2 + w_{\mathcal{T}^2}^{\lambda} \left(\frac{\partial^2 \Lambda_{ij}}{\partial \mathcal{T}_j^2} \right)^2 \right]$$
(3.25)

$$+\sum_{i,j}\left[w_t^{\beta}\left(\frac{\partial B_{ij}}{\partial t_i}\right)^2 + w_{\mathcal{T}}^{\beta}\left(\frac{\partial B_{ij}}{\mathcal{T}_j}\right)^2 + w_{t^2}^{\beta}\left(\frac{\partial^2 B_{ij}}{\partial t_i^2}\right)^2 + w_{\mathcal{T}^2}^{\beta}\left(\frac{\partial^2 B_{ij}}{\partial \mathcal{T}_j^2}\right)^2\right]\right\}, \quad (3.26)$$

where the index x references the different caplet/swaption-smiles from the calibration set, and $w_x^{\sigma}, w_x^{\beta} \in \mathbb{R}_+$ are exogenously specified weights, chosen according to the importance/reliability of the respective market parameters. The penalty functions in (3.25)-(3.26) are included to regularize the problem and to control the smoothness and time-homogeneity of the parameter functions. The partial derivative terms are to be interpreted as first- and second-order finite differences along the rows/columns of the matrices Λ and B. With the weights $w_t^{\cdot}, w_{\mathcal{T}}^{\cdot}, w_{t^2}^{\cdot}, w_{\mathcal{T}^2}^{\cdot} \in \mathbb{R}_+$, we can control the constancy and smoothness of the rows/columns²² of the matrices.

Instead of performing one "big" optimization, where both Λ and B are calculated simultaneously, Piterbarg [Pit03] suggests to split the problem into a sequence of skew and volatility calibrations, and in this way make the calibration procedure faster and more stable. One can justify this by the fact that the effective model volatilities $\bar{\sigma}_x$ depend only very mildly on the skew functions $\beta_i(t)$, and the volatility and skew calibrations can therefore be considered as being two "nearly-orthogonal" problems.

 $^{^{20}}$ Most often the points in the grid will be chosen to include (or coincide with) the maturities/tenors of the caplets and swaptions from the calibration set.

²¹Obtained by using the approximation (3.14).

 $^{^{22}\}text{Recall}$ that forcing the columns of Λ and B to be constant will result in a perfectly time-homogeneous model.

To summarize, the main calibration procedure can be carried out in three steps:

	C 111	
Main	(alibration	•
Iviaiii		

- **Step 1:** Set the skew parameters B_{ij} all to the same value B, chosen for example to be the average of all effective market skews $\bar{\beta}_x^*$. Calibrate the model volatilities $\lambda_n(t)$ to the $\bar{\sigma}_x^*$;
- **Step 2:** Using model volatilities calculated in the previous step, the skews $\beta_n(t)$ are now calibrated to the $\bar{\beta}_x^*$;
- **Step 3:** Finally, the model volatilities $\lambda_n(t)$ are re-calibrated to the $\bar{\sigma}_x^*$, with the updated skews $\beta_n(t)$ from the previous step.

Steps 2 and 3 can be repeated several times, although often one cycle (even without Step 3) is already enough to obtain a good fit. Concrete calibration examples will be provided in Chapter 7.

Remark 3.4.1. The advantage of calibrating (in the main calibration) to implied vanilla model parameters (determined in the pre-calibration step) instead of calibrating outright to option prices is twofold: First, working with implied parameters usually gives a more homogeneous error norm, as option prices may vary considerably with respect to strike and option maturity (although this could be resolved by choosing appropriate weighting functions). Second, and more importantly, the performance of the calibration procedure is improved, since we avoid applying timeconsuming Fourier option-pricing formulas.

Remark 3.4.2. Once the model has been calibrated to market-observable vanilla option prices, prices and hedge ratios for exotic interest-rate products can be obtained by performing Monte Carlo simulations. For details on numerical aspects see, e.g. [Gla04], [And08], [JS08] and [VHP10].

Chapter 4

Efficient Pricing of CMS Spread Options¹

Being able to efficiently compute CMS spread option (CMSSO) prices in a LMM is important for several reasons. For example, when using a LMM for pricing exotic interest-rate products with CMS spread-related payoff structures, we would like to be able to easily check whether the model prices CMSSOs inline with the market. More importantly, with an efficient pricing method at our disposal we can include CMSSOs in the general calibration procedure and in this way extract information about the Libor correlations from the market.

While existing literature, e.g. [BM05] and [BKS10], deals mainly with approximation methods for pricing CMSSOs in standard (log-normal) LMMs, we present in this chapter a new rapid and accurate method for approximating prices of CMSSOs in the SV-LMM (3.8)–(3.9). In Section 4.5 we also compare the performance of our approach to that of an approximation introduced in a recent paper by Antonov & Arneguy [AA09].

4.1 Approximating the Swap-Rate Dynamics

Using again the general set-up from the last chapter, assume that we have fixed integers n, m and m' such that $0 \le n < m < m' \le N$. Recall from Section 2.6.2 that pricing CMS spread caps essentially comes down to calculating prices of CMS spread caplets², with (undiscounted) prices of the form (see Equation (2.68))

$$\mathbb{E}^{T_n+\delta}\left[\left(S_{n,m'}(T_n)-S_{n,m}(T_n)-K\right)^+\right],\tag{4.1}$$

where K > 0 and $\delta \ge 0$. In the following we will often just mean (4.1) when we speak of CMS spread options. Furthermore, we will use again the short-hand notation $S_i(t)$, i = 1, 2 to denote $S_{n,m'}(t)$ and $S_{n,m}(t)$. In order to avoid confusion with the Libor skew and volatility functions $\beta_i(t)$ and $\sigma_i(t)$, $i = 0, \ldots, N-1$, we denote the swap rate skew and volatility functions by $\beta_i^{SR}(t)$ and $\sigma_i^{SR}(t)$, i = 1, 2, respectively.

Notice that under the forward measure $\mathbb{Q}^{T_n+\delta}$ both swap rates will have non-zero drift terms. Using the approximation (3.14) we have that the swap-rate dynamics

¹The contents of this chapter will appear in [KL11].

 $^{^{2}}$ Here and in the sequel, we will often just consider the call case for clarity of exposition. All statements easily carry over to the floor case.

under $\mathbb{Q}^{T_n+\delta}$ are (approximately) of the form

$$dS_i(t) = \varphi_i(t, S_i(t)) \left(\mu_i(t) dt + \sigma_i^{SR}(t)' \sqrt{V(t)} dW^{T_n + \delta}(t) \right), \ i = 1, 2,$$
(4.2)

$$dV(t) = \kappa (1 - V(t))dt + \xi \sqrt{V(t)} d\tilde{Z}(t), \ V(0) = 1, \ 0 \le t \le T_n,$$
(4.3)

with

$$\varphi_i(t,s) \coloneqq \left(\beta_i^{SR}(t)s + (1 - \beta_i^{SR}(t))S_i(0)\right), \ i = 1, 2,$$

independent one- resp. *d*-dimensional $\mathbb{Q}^{T_n+\delta}$ -Brownian motions Z(t) and $W^{T_n+\delta}(t)$, and with $\beta_i^{SR}(t)$ and $\sigma_i^{SR}(t)$ as in (3.15) and (3.16), respectively.

Drift Terms

In order to calculate the drift terms of the swap rates we first use the product rule to calculate the dynamics of $S_i(t)A_i(t)/P(t, T_n + \delta)$:

$$d\left(\frac{S_i(t)A_i(t)}{P(t,T_n+\delta)}\right) = S_i(t) d\left(\frac{A_i(t)}{P(t,T_n+\delta)}\right) + \frac{A_i(t)}{P(t,T_n+\delta)} dS_i(t) + d\langle S_i(t), \frac{A_i(t)}{P(t,T_n+\delta)} \rangle.$$
(4.4)

Clearly, $S_i(t)A_i(t)/P(t, T_n + \delta)$ and $A_i(t)/P(t, T_n + \delta)$ are martingales under $\mathbb{Q}^{T_n+\delta}$ and thus their drift terms must be zero. Collecting only *dt*-terms from (4.4), we obtain for the drift term

$$\varphi_i(t, S_i(t))\mu_i(t)dt = -\frac{P(t, T_n + \delta)}{A_i(t)} d\langle S_i(t), \frac{A_i(t)}{P(t, T_n + \delta)} \rangle.$$

Next, we can write the swap rates $S_i(t)$ as well as the annuity-bond ratios $A_i(t)/P(t, T_n + \delta)$ as functions of the underlying Libor rates $L(t) := (L_0(t), \ldots, L_{N-1}(t))$ at time t:

$$S_i(t) \coloneqq F_i(L(t))$$
 and $\frac{A_i(t)}{P(t, T_n + \delta)} \coloneqq G_i(L(t)), i = 1, 2.$

Applying Itô's Lemma, we finally obtain for the drift terms

$$\varphi_i(t, S_i(t))\mu_i(t) = -\frac{P(t, T_n + \delta)}{A_i(t)}\nabla F_i(L(t))D(t)\Sigma(t)D(t)\nabla G_i(L(t))'V(t),$$

with $\mathbb{R}^{N-1 \times N-1}$ -valued functions

$$D(t) = \operatorname{diag} \left(\beta_k(t)L_k(t) + (1 - \beta_k(t))L_k(0), \ k = 0, \dots, N-1\right),$$

$$\Sigma(t) = \left(\sigma_k(t)'\sigma_l(t)\right)_{k,l=0}^{N-1}.$$

With these fully path-dependent drifts there is of course no hope to obtain closedform solutions for the system of SDEs (4.2)-(4.3). Usually the so-called "freezing-thedrift" technique is applied, which means fixing the drift terms as an approximation

³Here, $\langle X(t), Y(t) \rangle$ denotes the quadratic covariation of the two processes X(t) and Y(t).

at their initial values. This, however, works reasonably well only if the processes are considered over relatively small periods of time or if the occurring volatilities are quite small.

Instead of just fixing the drift terms at their initial values, we propose to use the following approximation

$$\mu_{i}(t) \approx -\frac{P(0, T_{n} + \delta)}{A_{i}(0)\varphi_{i}(0, S_{i}(0))} \nabla F_{i}(L(0)) D \overline{\Sigma} D \nabla G_{i}(L(0))' V(t)$$

=: $\bar{\mu}_{i}V(t), \ i = 1, 2,$ (4.5)

with

$$D \coloneqq D(0) = \operatorname{diag}(L_k(0), \ k = 0, \dots, N-1),$$

and the "average covariance matrix"

$$\overline{\Sigma} = \left(\frac{1}{T_n} \int_0^{T_n} \sigma_k(t)' \sigma_l(t) dt\right)_{k,l=0}^{N-1}$$

By taking average covariances we account for generally time-dependent, i.e., nonconstant Libor volatility functions $\sigma_n(t)$. More importantly, we retain the stochastic variance factor V(t), which will prove to be quite effective in capturing the nature of the convexity adjustments, which are implicitly given by the drift terms. Intuitively, this can be explained by the fact that the volatility smiles/skews of the underlying rates generally have a strong influence on the size of the convexity adjustments (cp. Remark 2.6.1), and the stochastic-variance factor V(t) is the main determinant of the curvature of the volatility smile.

Swap-Rate Correlation

Next, with the definition

$$\varphi_i(t,s) \coloneqq \left(\beta_i^{SR}(t)s + (1 - \beta_i^{SR}(t))S_i(0)\right), \ i = 1, 2,$$

we obtain for the quadratic variation and covariation of $S_1(t)$ and $S_2(t)$

$$d\langle S_{i}(t)\rangle = \varphi_{i}(t, S_{i}(t))^{2} \|\sigma_{i}^{SR}(t)\|^{2} V(t) dt, \ i = 1, 2,$$

$$d\langle S_{1}(t), S_{2}(t)\rangle = \varphi_{1}(t, S_{1}(t))\varphi_{2}(t, S_{2}(t))\sigma_{1}^{SR}(t)\sigma_{2}^{SR}(t)V(t) dt,$$

such that the instantaneous correlation⁴ among the swap rates at time t is given by

$$\operatorname{Corr}\left(dS_{1}(t), dS_{2}(t)\right) = \frac{\sigma_{1}^{SR}(t)'\sigma_{2}^{SR}(t)}{\|\sigma_{1}^{SR}(t)\|\|\sigma_{2}^{SR}(t)\|}.$$
(4.6)

Note carefully that the instantaneous correlation in the SV-LMM is indeed deterministic (and the same as in a non-stochastic-volatility displaced/log-normal LMM) – a consequence of the fact that all Libor/swap rates are being multiplied by the same stochastic-volatility factor $\sqrt{V(t)}$.

⁴I.e., the correlation of the increments of the swap-rate processes.

In the subsequent calculations we will replace the time-dependent instantaneous correlation given in (4.6) by the following *effective* or *term correlation*⁵ (cp. [BM05], p. 286)

$$\bar{\rho} \coloneqq \frac{\int_0^{T_n} \sigma_1^{SR}(t)' \sigma_2^{SR}(t) dt}{\sqrt{\int_0^{T_n} \|\sigma_1^{SR}(t)\|^2 dt} \sqrt{\int_0^{T_n} \|\sigma_2^{SR}(t)\|^2 dt}}.$$
(4.7)

Swap-Rate Skews and Volatilities

In order to obtain the effective swap rate skew and volatility parameters $\bar{\beta}_i$ and $\bar{\sigma}_i$, we make again use of the parameter averaging formulas from Section 3.2.

Finally, with the various approximations from above, we obtain for the dynamics of the swap rates under $\mathbb{Q}^{T_n+\delta}$

$$dS_i(t) \approx \left(\bar{\beta}_i S_i(t) + (1 - \bar{\beta}_i) S_i(0)\right) \left(\bar{\mu}_i V(t) dt + \bar{\sigma}_i \sqrt{V(t)} dU_i(t)\right), \ i = 1, 2, \quad (4.8)$$

with $U_1(t)$ and $U_2(t)$ being two one-dimensional $\mathbb{Q}^{T_n+\delta}$ -Brownian motions (independent of Z(t)), such that $dU_1(t)dU_2(t) = \bar{\rho}dt$.

4.2 Iterated Expectations

By applying the standard displaced-diffusion transform (see p. 35) we can always reduce displaced Heston type SDEs like (4.8) to the standard (log-normal) Heston case. So for clarity of exposition, let us assume for the moment that $\bar{\beta}_i = 1$, i = 1, 2. Then our pricing problem essentially reduces to calculating expectations of the form⁶

$$\pi(0) = \mathbb{E}\left[(S_1(T_n) - S_2(T_n) - K)^+ \right], \tag{4.9}$$

with K > 0 and two stochastic processes $S_i(t)$, i = 1, 2, having dynamics

$$dS_{1}(t) = S_{1}(t) (\bar{\mu}_{1}V(t)dt + \bar{\sigma}_{1}\sqrt{V(t)}dU_{1}(t)), \ S_{1}(0) > 0,$$

$$dS_{2}(t) = S_{2}(t) (\bar{\mu}_{2}V(t)dt + \bar{\sigma}_{2}\sqrt{V(t)}dU_{2}(t)), \ S_{2}(0) > 0,$$

$$dV(t) = \kappa(1 - V(t))dt + \xi\sqrt{V(t)}dZ(t), \ V(0) = 1,$$

where

$$dU_1(t)dU_2(t) = \bar{\rho}dt,$$

$$dU_i(t)dZ(t) = 0, \ i = 1, 2.$$

Conditioning on the path of the stochastic-variance process and using iterated expectations, we may write the pricing formula (4.9) as

$$\pi(0) = \mathbb{E} \left[\mathbb{E} \left[\left(S_1(T_n) - S_2(T_n) - K \right)^+ \right| \{ V(t), 0 \le t \le T_n \} \right] \right],$$

⁵The term rather than the instantaneous correlation is important for spread options. Note that there can be a significant decorrelation effects due to time-dependent volatilities. We shall return to this topic later.

⁶All expectations in the following are to be taken with respect to $\mathbb{Q}^{T_n+\delta}$.

4.2. Iterated Expectations

Due to the independence of the Brownian motions $U_i(t)$ and Z(t), we have that $S_i(t)$, i = 1, 2, given the path of V(t) are (jointly) geometric Brownian motions with (conditionally) deterministic time-dependent drift and volatility functions, as pointed out by Hull & White [HW87]. As we have seen earlier, for the terminal distribution of a geometric Brownian motion only the total integrated squared volatility (and the integrated drift) is important and not the particular "path" of the time-dependent functions. Hence, defining the integrated variance (up to time T_n) by

$$\overline{V}(T_n) \coloneqq \int_0^{T_n} V(t) dt,$$

we have that $S_1(T_n)$ and $S_2(T_n)$ given $\overline{V}(T_n)$ are jointly log-normal, and using again iterated expectations we may thus write

$$\pi(0) = \mathbb{E}\Big[\mathbb{E}\Big[\left(S_1(0)e^{Y_1} - S_2(0)e^{Y_2} - K\right)^+ | \overline{V}(T_n)]\Big], \qquad (4.10)$$

with the random vector $Y := (Y_1, Y_2)'$ being conditionally a bivariate normal

$$Y | \overline{V}(T_n) \sim \mathcal{N}\left(\begin{pmatrix} (\bar{\mu}_1 - \bar{\sigma}_1^2/2)\overline{V}(T_n) \\ (\bar{\mu}_2 - \bar{\sigma}_2^2/2)\overline{V}(T_n) \end{pmatrix}, \, \overline{V}(T_n) \begin{pmatrix} \bar{\sigma}_1^2 & \bar{\sigma}_1 \bar{\sigma}_2 \bar{\rho} \\ \bar{\sigma}_1 \bar{\sigma}_2 \bar{\rho} & \bar{\sigma}_2^2 \end{pmatrix} \right).$$

Now recall that $Y_1|Y_2$, for two jointly normal random variables Y_1 and Y_2 , is again normally distributed. Exploiting this fact and applying once more the law of iterated expectations, one can derive (cp. [BM05], p. 604) a pseudo-analytical formula for the inner expectation⁷ in Equation (4.10)

$$\mathbb{E}\Big[\big(S_1(0)e^{Y_1} - S_2(0)e^{Y_2} - K\big)^+ \big| \,\overline{V}(T_n) = v\,\Big] = \int_{-\infty}^{+\infty} g(u,v)\frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}u^2}du,$$

where

$$g(u,v) = S_1(0) \exp\left(m(u,v) + s(v)^2/2\right) \Phi\left(\frac{\log\left(\frac{S_1(0)}{h(u,v)}\right) + m(u,v) + s(v)^2}{s(v)}\right) - h(u,v) \Phi\left(\frac{\log\left(\frac{S_1(0)}{h(u,v)}\right) + m(u,v)}{s(v)}\right)$$
(4.11)

and

$$h(u,v) = S_2(0) \exp \left\{ \bar{\mu}_2 T - \bar{\sigma}_2^2 v/2 + \bar{\sigma}_2 \sqrt{v} u \right\} + K$$

$$m(u,v) = \bar{\mu}_1 T - \bar{\sigma}_1^2 v/2 + \bar{\rho} \bar{\sigma}_1 \sqrt{v} u,$$

$$s(v)^2 = \bar{\sigma}_1^2 (1 - \bar{\rho}^2) v.$$

⁷This formula holds for the case where $S_1(0), S_2(0)$ and K are all greater than zero. We will consider a more general case later.

If we were able to efficiently evaluate the density f of the integrated variance \overline{V} at time T_n , then the calculation of the expectation in (4.9) would just be a matter of calculating the two-dimensional integral

$$\int_{0}^{+\infty} \int_{-\infty}^{\infty} g(u,v) \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}u^2} du f(v) dv, \qquad (4.12)$$

with $g(\cdot, \cdot)$ given in (4.11). Unfortunately, the density f is not known in closed form and in order to evaluate it one needs to numerically invert its Laplace transform. This makes the numerical computation of (4.12) generally too slow to be usable for calibration purposes. In the next section we therefore derive a new method for speeding up the numerical evaluation of the density of $\overline{V}(T_n)$. The approach is based on carefully choosing the integration contour for the Laplace inverse transform and allows the fast and accurate evaluation of the density in its entire domain, a result which is also applicable outside the context of spread-option pricing.

4.3 The Density of the Integrated Variance

4.3.1 The Branch-Cut Corrected Laplace Transform

The Laplace transform of the integrated variance $\overline{V}(T)$ for T > 0 is given by (see e.g. [LL08])

$$\begin{aligned} \hat{f}(z) &= \mathbb{E}\Big[e^{-z\overline{V}(T)}\Big] \\ &= \exp\left\{\frac{2\kappa}{\xi^2}\log\left(\frac{2\gamma e^{-\gamma T/2}}{\gamma + \kappa + (\gamma - \kappa)e^{-\gamma T}}\right) + \frac{\kappa^2 T}{\xi^2} + \left(\frac{2\gamma e^{-\gamma T}}{\gamma + \kappa + (\gamma - \kappa)e^{-\gamma T}} - 1\right)\frac{V(0)(\gamma - \kappa)}{\xi^2}\right\}, \end{aligned}$$
(4.13)

with

$$\gamma = \gamma(z) = \sqrt{\kappa^2 + 2z\xi^2}$$

and $z \in \mathbb{C}$ such that $\hat{f}(z)$ exists. As we show in Appendix A.1, all singularities of \hat{f} lie on the negative real axis. Before we can proceed with the numerical inversion of the Laplace transform we need to take care of the complex logarithm appearing in (4.13). Recall that the complex logarithm is a multi-valued function. If we restrict the logarithm to its principal branch (as is done in most software packages), then it is discontinuous along the negative real axis and evaluating the density via numerically inverting the Laplace transform will therefore possibly result in wrong values. In Appendix A.2 we use similar ideas as in Kahl & Jäckel [KJ05] to derive a "continuified" version of the Laplace transform (4.13), which is given by

$$\hat{f}(z) = \exp\left\{\frac{2\kappa}{\xi^2}A(z) + B(z)\right\},$$

where

$$\begin{split} A(z) &= \operatorname{Log} \left| 2\gamma e^{-\gamma T/2} \right| + i \left(\operatorname{Arg}(2\gamma e^{-\gamma T/2}) + 2\pi n \right) - \operatorname{Log}\left(\gamma + \kappa + (\gamma - \kappa)e^{-\gamma T}\right), \\ B(z) &= \frac{\kappa^2 T}{\xi^2} + \left(\frac{2\gamma e^{-\gamma T}}{\gamma + \kappa + (\gamma - \kappa)e^{-\gamma T}} - 1 \right) \frac{V(0)(\gamma - \kappa)}{\xi^2}, \\ n &= \left\lfloor \frac{\operatorname{Arg}(\gamma) - \frac{T}{2} |\gamma| \sin(\operatorname{Arg}(\gamma)) + \pi}{2\pi} \right\rfloor, \\ \gamma &= \sqrt{\kappa^2 + 2z\xi^2}, \end{split}$$

where $\lfloor . \rfloor$ denotes rounding to the nearest smaller integer and Log(.) denotes the principal value of the logarithm. In what follows, when we write \hat{f} we are always referring to this this representation of the Laplace transform.

4.3.2 Calculating the Bromwich Integral

With the Laplace transform from above, the density f of $\overline{V}(T)$ is given in terms of the inverse Laplace transform

$$f(x) = \frac{1}{2\pi i} \int_{a-i\infty}^{a+i\infty} e^{xs} \hat{f}(s) ds, \ x \ge 0,$$
(4.14)

where the integration in the complex plane is done along the so-called *Bromwich line* $\operatorname{Re}(s) = a$, with a being greater than the real part of all singularities of \hat{f} . In our case, as mentioned above, all singularities of the Laplace transform lie on the negative real line and a can therefore be greater than or equal to zero. Note that with a = 0, the above inverse integral becomes identical to the standard inverse Fourier transform. Parameterizing the Bromwich line by s = a + iu and taking symmetry into account, we can write Equation (4.14) as

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{x(a+iu)} \hat{f}(a+iu) du$$
$$= \frac{1}{\pi} \int_{0}^{\infty} \operatorname{Re}\left(e^{x(a+iu)} \hat{f}(a+iu)\right) du.$$
(4.15)

The integrand in (4.15) is often rapidly oscillating as depicted in Figure 4.1. Additionally, the transform $\hat{f}(z)$ may decay slowly as $\text{Im}(z) \to \infty$, making the evaluation of the density via direct numerical integration of the Bromwich integral generally too costly (or inaccurate) for our applications⁸.

The problem of the occurrence of highly oscillatory and/or slowly decaying integrands is of course not specific to our case, but is often faced when dealing with the numerical inversion of Laplace transforms. A standard strategy for avoiding these

⁸Note that we possibly need to evaluate the density f several hundred times when numerically integrating (4.12).



Figure 4.1: Integrand $g(u) \coloneqq \operatorname{Re}\left(e^{x(a+iu)}\hat{f}(a+iu)\right)$ in (4.15) for a=0 and x=15. Parameters: $V(0)=1, \ \kappa=0.15, \ \xi=1.3, \ T=5$.

problems is to make use of Cauchy's theorem and to deform the Bromwich (half-)line into a contour s, such that $\operatorname{Re}(s) \to -\infty$ at the upper end. This has the effect that the exponential factor in (4.14) forces the integrand to decay rapidly. The approach of deforming the Bromwich line was originally pioneered by Talbot [Tal79], who used a contour of the form

$$s(\varphi) = a + b\varphi(\cot\varphi + ic), \ 0 \le \varphi < \pi, \tag{4.16}$$

where the parameters a, b and c must be chosen such that all singularities of the Laplace transform are to the left of the contour. Besides this technical requirement (which is due to Cauchy's theorem) there is, however, a priori no general rule for choosing a, b and c, and finding parameters such that one obtains a nice and smooth integrand is often a challenge in its own right.



Figure 4.2: Left: Upper half of Talbot's contour (4.16). Right: Linear contour (4.17).

Note also that in our case one needs to find an optimal contour for each and every x, for which the density is to be evaluated, as the shape of the integrand does not only depend heavily on the parameters of the stochastic-variance process, but also on the value of x. Furthermore, observe that the real part of Talbot's contour is bounded. This can lead to an integrand that falls off rapidly along the first part of the contour, but then becomes irregular due to coming too close to the sometimes slowly decaying singularities along the negative real axis. In order to circumvent this

problem we therefore propose to use a "linear contour" of the form

$$s(u) = a + u(bi - a), \ 0 \le u < \infty,$$
(4.17)

with a, b > 0. The consequence is that $\text{Im}(s(u)) \to \infty$ as $u \to \infty$, i.e., we "stay away" from the singularities at the left end (see Figure 4.2). However, one problem remains to be solved: The optimal choice of a and b.

4.3.3 The Optimal Linear Contour

Using the contour (4.17), the integral (4.15) becomes

$$f(x) = \frac{1}{\pi} \int_0^\infty \operatorname{Im} \left(e^{xs(u)} \hat{f}(s(u)) s'(u) \right) du$$

= $\frac{1}{\pi} \int_0^\infty \operatorname{Im} \left(e^{x(a+u(bi-a))} \hat{f}(a+u(bi-a))(bi-a) \right) du.$ (4.18)

First note, that for large values of x the integrand will be unnecessarily blown up by the exponential factor (at least in the region of the contour where $\operatorname{Re}(s(u)) > 0$). In order to avoid this, it is clear that a should be chosen approximately inverse proportional to the value of x, and we will therefore simply set a = 1/x. Having fixed a, we are now only left with choosing the second parameter b. Observe that the main behavior of the integrand in (4.18) is determined by $\operatorname{Im}\left(e^{x(a+u(bi-a))}\hat{f}(a+u(bi-a))\right)$ and $\operatorname{Re}\left(e^{x(a+u(bi-a))}\hat{f}(a+u(bi-a))\right)$, since s'(u) = bi - a is just a constant scaling factor. Having a closer look at these two functions in Figure 4.3, we can see that choosing b too large (red lines) generally results in picking up too many oscillations along the path of integration. On the other hand, when choosing b too small (orange lines) we might get too close to the singularities along the negative real line, and the integrand can also become unsmooth.

In order to find an "optimal" value for b, we need to have some kind of measure for the oscillatory behavior of the integrand. If we denote by $\eta(z)$ the exponent of the Laplace transform, i.e., $\hat{f}(z) = e^{\eta(z)}$, then we need to consider (if we are just concerned with the behavior along the positive imaginary axis)

$$e^{xui+\eta(ui)} = \exp\left\{\operatorname{Re}(\eta(ui)) + i\operatorname{Im}(xui+\eta(ui))\right\}$$
$$= e^{\operatorname{Re}(\eta(ui))}\left[\cos(xu+\operatorname{Im}(\eta(ui))) + i\sin(xu+\operatorname{Im}(\eta(ui)))\right].$$

Now, if we assume that the oscillatory behavior of the integrand is mainly due to the sine and cosine terms, then we can take as a measure for the frequency, for example, the smallest value b^* such that

$$\cos\left(b^*x + \operatorname{Im}(\eta(b^*i))\right) = 0,$$

that is

$$b^*x + \operatorname{Im}(\eta(b^*i)) = \frac{\pi}{2}.$$
 (4.19)

This non-linear equation, however, does not possess an analytical solution and one needs to use numerical root-finding algorithms in order to solve for b^* . Furthermore,



Figure 4.3: $Re(e^{xz}\hat{f}(z))$ and $Im(e^{xz}\hat{f}(z))$ in the upper complex plane for x = 3, V(0) = 1, $\kappa = 0.15$, $\xi = 1.3$, T = 1. The lines depict the path of integration for b = 15 (red lines), b = 0.09 (orange lines) and $b^* = 0.768$ (cyan lines). In all three cases a was set to 1/3. The irregular behavior of the integrand in the vicinity of (-2, 0) is due to one of the singularities along the negative real axis.

Equation (4.19) has to be solved for each and every x, for which the density is to be evaluated. As if this was not enough, depending on the parameter combination (and x), b^* can be close to zero or greater than 500, making it hard to even give a good initial guess for the root finding algorithm. At first glance, it therefore seems like this route would be too costly for being usable. However, having a closer look at

$$h(u) \coloneqq xu + \operatorname{Im}(\eta(ui)) - \frac{\pi}{2}$$

one finds that this function is generally convex and for large values of u almost linear⁹

⁹In fact, one can show that asymptotically it behaves like $xu - \frac{\kappa T + V(0)}{\xi} \sqrt{u}$.

(see also Figure 4.4).

Thus, Newton's method applied to h(u) converges pretty fast and reliably towards b^* , even if the initial guess b_0 is set, for instance, to 1,000. As h is close to linear for large values of u, four Newton iterations are usually more than enough to obtain an accurate approximation¹⁰ of b^* . Hence, if we approximate the first derivative, needed



Figure 4.4: h(u) for parameters V(0) = 1, $\kappa = 0.15$, $\xi = 1.3$, T = 5 and for values of x ranging between 0.18 (bottommost line) and 10 (uppermost line).

when applying Newton's method, via finite differences, then for four iterations we just need eight evaluations of h in order to find b^* . The initial choice $\pi/2$ in (4.19) was of course arbitrary and numerical experiments have shown that the integrands will be generally even smoother if one chooses a value around $\pi/8$.

Having determined the optimal contour parameters a and b, we could in principle proceed by numerically integrating (4.18). However, direct numerical integration of the improper integral would result in a truncation error. In order to avoid this, one can change variables and transform the original integration range $[0, \infty)$ to the finite interval [0, 1]. When analyzing the asymptotic behavior of the integrand one finds that it decays at least exponentially (see Appendix A.3). Thus, we may use a logarithmic transform of the integration range. Actually, we obtained excellent results for practically all relevant parameter scenarios using the transform

$$u \to -3\log(\tilde{u}).$$

With this transformation (4.18) becomes

$$f(x) = -\frac{1}{\pi} \int_0^1 \text{Im}\left(e^{x\tilde{s}(u)}\hat{f}(\tilde{s}(u))\tilde{s}'(u)\right) du$$
(4.20)

with

$$\tilde{s}(u) = a - 3\log(u)(bi - a),$$

$$\tilde{s}'(u) = -\frac{3}{u}(bi - a),$$

¹⁰Recall that we are not really interested in a highly accurate approximation of b^* , since we only need it as a measure for the order of magnitude of the oscillatory behavior of the integrand.

and where a and b are chosen as described above.

In Figure 4.5 we show the integrand in (4.20). Observe that it is perfectly smooth over the entire range of considered x-values. In particular, this smooth integrand (on a finite interval) should be compared with the oscillatory integrand (on an unbounded interval) shown in Figure 4.1. Using an advanced integration algorithm like, for instance, the adaptive Gauss-Lobatto scheme [GG00], 50–80 evaluations of the integrand are mostly enough to obtain a precision in the order of magnitude of 10^{-6} for the value of the density, whereas one needs up to 4,000 evaluations of the inverse Fourier integrand (4.15), in order to obtain a similar accuracy¹¹. The efficiency of our method is also demonstrated by the fact, that sampling the density (with parameters as given in Figure 4.5) at 400 equally spaced points $x \in \{0.05, \ldots, 20\}$ takes only 0.069 seconds, opposed to 1.521 seconds when using Fourier inversion.



Figure 4.5: Integrand $g(u) := \operatorname{Im}\left(e^{x\tilde{s}(u)}\hat{f}(\tilde{s}(u))\tilde{s}'(u)\right)$ in (4.20) for parameters V(0) = 1, $\kappa = 0.15$, $\xi = 1.3$, T = 5 and for various values of x.

Figure 4.6 shows the density of the integrated variance for various parameter values. One can clearly see, that the proposed method for evaluating the density yields perfectly smooth plots, even in the tails, where direct Fourier inversion is usually prone to numerical instabilities due to the highly oscillatory nature of the integrand.

Finally, as a simple sanity check, we consider the normalization integral $\int_0^\infty f(v)dv$ to see how much it deviates from 1. Numerical integration¹² of the various densities presented in Figure 4.6 yields a maximum absolute error of 1.1×10^{-6} , which should be negligible for most practical applications.

4.4 Calculating CMS Spread Option Prices

In this section we return to our original goal: Pricing CMSSOs, which comes down to calculating the expectation given in (4.9). In order to solve the system of SDEs (4.8),

¹¹Admittedly, there might exist more suitable integration schemes tailored towards the integration of highly oscillating functions.

¹²Integration was performed on [0,1.5] in the case T = 0.25 and on [0,150] in the case T = 5.



Figure 4.6: Density f(x) of the integrated variance $\overline{V}(T)$ for V(0) = 1, $\kappa = 0.15$, T = 0.25 (top), T = 5 (bottom) and various values of ξ .

we use the standard displacement transform and define the stochastic processes

$$X_i(t) \coloneqq \bar{\beta}_i S_i(t) + (1 - \bar{\beta}_i) S_i(0), \ i = 1, 2.$$
(4.21)

Applying Itô's Lemma yields for the dynamics of these processes

$$dX_{i}(t) = \bar{\beta}_{i}dS_{i}(t)$$

$$= \bar{\beta}_{i}\left(\bar{\beta}_{i}S_{i}(t) + (1 - \bar{\beta}_{i})S_{i}(0)\right)\left(\bar{\mu}_{i}V(t)dt + \bar{\sigma}_{i}\sqrt{V(t)}dU_{i}(t)\right)$$

$$= X_{i}(t)\left(\tilde{\mu}_{i}V(t)dt + \tilde{\sigma}_{i}\sqrt{V(t)}dU_{i}(t)\right), \qquad (4.22)$$

where $\tilde{\mu}_i \coloneqq \bar{\beta}_i \bar{\mu}_i$ and $\tilde{\sigma}_i \coloneqq \bar{\beta}_i \bar{\sigma}_i$, i = 1, 2. With the results from Section 4.2, we then have that conditionally on the integrated variance

$$X_i(T_n) \stackrel{d}{=} X_i(0) \exp\left\{\left(\tilde{\mu}_i - \frac{\tilde{\sigma}_i^2}{2}\right) \overline{V}(T_n) + \tilde{\sigma}_i \sqrt{\overline{V}(T_n)} Z_i\right\} \rightleftharpoons \tilde{X}_i(T_n), \ i = 1, 2, \ (4.23)$$

where

$$\begin{pmatrix} Z_1 \\ Z_2 \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \bar{\rho} \\ \bar{\rho} & 1 \end{pmatrix}\right).$$

Reversing the transformation (4.21) finally gives us for the swap rates (conditional on $\overline{V}(T_n)$)

$$S_i(T_n) \stackrel{d}{=} \frac{1}{\overline{\beta}_i} \left(\tilde{X}_i(T_n) - (1 - \overline{\beta}_i) S_i(0) \right), \ i = 1, 2.$$

Using the law of iterated expectations we obtain for $K \in \mathbb{R}$ and $w = \pm 1$

$$\mathbb{E}\left[\left(wS_{1}(T_{n})-wS_{2}(T_{n})-wK\right)^{+}\right] =$$

$$=\mathbb{E}\left[\mathbb{E}\left[\left(wS_{1}(T_{n})-wS_{2}(T_{n})-wK\right)^{+}|\overline{V}(T)\right]\right]$$

$$\approx\mathbb{E}\left[\mathbb{E}\left[\left(\frac{w}{\overline{\beta}_{1}}\left(\tilde{X}_{1}(T_{n})-(1-\overline{\beta}_{1})S_{1}(0)\right)\right.\right.\right.\right.$$

$$\left.-\frac{w}{\overline{\beta}_{2}}\left(\tilde{X}_{2}(T_{n})-(1-\overline{\beta}_{2})S_{2}(0)\right)-wK\right)^{+}\left|\overline{V}(T)\right]\right]$$

$$=\mathbb{E}\left[\mathbb{E}\left[\left(\frac{w}{\overline{\beta}_{1}}\tilde{X}_{1}(T_{n})-\frac{w}{\overline{\beta}_{2}}\tilde{X}_{2}(T_{n})-w\tilde{K}\right)^{+}\left|\overline{V}(T)\right.\right]\right],$$
here

W

$$\tilde{K} = \frac{1 - \bar{\beta}_1}{\bar{\beta}_1} S_1(0) - \frac{1 - \bar{\beta}_2}{\bar{\beta}_2} S_2(0) + K$$

and with $\tilde{X}_i(T_n)$ as given in (4.23). Applying the generalized spread-option formula for two jointly log-normal random variables given in Appendix A.4, we finally obtain for the time-zero price of a CMSSO

$$P(0, T_n + \delta) \mathbb{E} \left[(wS_1(T_n) - wS_2(T_n) - wK)^+ \right] \approx \\ \approx P(0, T_n + \delta) \int_0^{+\infty} \int_{-\infty}^{+\infty} g(u, v) \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}u^2} du f(v) dv, \quad (4.24)$$

where

$$g(u,v) = \begin{cases} \left(\frac{S_{1}(0)}{\bar{\beta}_{1}}e^{h(u,v)} - \hat{K}(u,v)\right) \mathbb{1}_{\{w=1\}} &, \hat{K}(u,v) \leq 0, \\ w\frac{S_{1}(0)}{\bar{\beta}_{1}}e^{h(u,v)}\Phi\left(w\frac{\log\left(\frac{S_{1}(0)}{\bar{\beta}_{1}\bar{K}(u,v)}\right) + h(u,v) + \frac{1}{2}\tilde{\sigma}_{1}^{2}(1-\bar{\rho}^{2})v}{\bar{\sigma}_{1}\sqrt{v(1-\bar{\rho}^{2})}}\right) &, \hat{K}(u,v) > 0, \\ -w\hat{K}(u,v)\Phi\left(w\frac{\log\left(\frac{S_{1}(0)}{\bar{\beta}_{1}\bar{K}(u,v)}\right) + h(u,v) - \frac{1}{2}\tilde{\sigma}_{1}^{2}(1-\bar{\rho}^{2})v}{\bar{\sigma}_{1}\sqrt{v(1-\bar{\rho}^{2})}}\right) &, \hat{K}(u,v) > 0, \\ \hat{K}(u,v) = \frac{S_{2}(0)}{\bar{\beta}_{2}}\exp\left\{\left(\tilde{\mu}_{2} - \frac{1}{2}\tilde{\sigma}_{2}^{2}\right)v + \tilde{\sigma}_{2}\sqrt{v}u\right\} + \tilde{K} \\ h(u,v) = \left(\tilde{\mu}_{1} - \frac{1}{2}\bar{\rho}^{2}\tilde{\sigma}_{1}^{2}\right)v + \bar{\rho}\,\tilde{\sigma}_{1}\sqrt{v}u, \end{cases}$$

and where the density f(v) can be evaluated by using the method developed in Section 4.3.

Remark 4.4.1. Note that Formula (4.24) remains valid even for other stochasticvolatility extensions of the LMM, provided that the variance process and the forward rates are uncorrelated and that an appropriate method for evaluating or approximating the integrated variance density f(v) is used.

Moment Explosions

As we have remarked in Section 2.6.1, CMS linked payoffs implicitly or explicitly depend on higher-order moments of the underlying swap rate under the corresponding swap measure. In order to make sure that MC-prices of CMSSOs generated by the SV-LMM are meaningful, it suffices to check (by using Proposition 2.6.2) that $\mathbb{E}^{A_i}[S_i(T_n)^2]$, i = 1, 2 are finite.

However, it is not so obvious whether this is also a sufficient condition for the spread-option pricing formula (4.24) (which is only an approximation) to be well-defined. For this we need the first-order moments of $X_i(T_n)$, i = 1, 2 to be finite¹³.

Proposition 4.4.1. Consider the processes

$$dX(t) = X(t) \left(\mu V(t) dt + \sigma \sqrt{V(t)} dU(t) \right), X(0) > 0,$$

$$dV(t) = \kappa (1 - V(t)) dt + \xi \sqrt{V(t)} dZ(t), V(0) = 1,$$

where μ, σ, κ, ξ are positive constants and where U(t) and Z(t) are independent onedimensional Brownian motions. Define

$$D \coloneqq \kappa^2 - \xi^2 \nu (2\mu + \sigma^2 (\nu - 1))$$

and fix $\nu \geq 1$. The moment $\mathbb{E}[X(T)^{\nu}]$ will be finite for all T > 0 if $D \geq 0$. If D < 0, then $\mathbb{E}[X(T)^{\nu}]$ will be finite for $T < T^*$ and infinite for $T^* \geq T$, where T^* is given by

$$T^* = \frac{2}{\sqrt{-D}} \left(\pi + \arctan\left(-\frac{\sqrt{-D}}{\kappa}\right) \right).$$

Proof. The explosion time may be derived by using similar arguments as in the proof of Prop. 2.6.2, which is sketched in [AP07]. First we make the change of variable $Y(t) = \log(X(t))$. An application of the Feynman-Kac formula then shows that $u(0, Y(0), V(0)) = \mathbb{E}[X(T)^{\nu}]$, where u(t, y, v) satisfies the PDE

$$u_t + (\mu - \sigma^2/2) v u_y + \kappa (1 - v) u_v + \frac{1}{2} \sigma^2 v u_{yy} + \frac{1}{2} \xi^2 v u_{vv} = 0$$

subject to the final condition $f(T, y, v) = e^{\nu y}$. Making the Ansatz $u(t, y, v) = e^{\nu y} e^{A(T-t)+\nu B(T-t)}$, where A(0) = B(0) = 0, results in the following system of ODEs $(\tau \equiv T - t)$

$$A'(\tau) = \kappa B(\tau),$$

$$B'(\tau) = \frac{\xi^2}{2} B^2(\tau) - \kappa B(\tau) + \nu(\mu + (\nu - 1)\sigma^2/2)$$

¹³For $\tilde{\mu}_i = 0$ this is always true. However, the drift terms $\tilde{\mu}_i V(t) dt$ in (4.22) capture the necessary convexity adjustments, so in general $\tilde{\mu}_i > 0$, which may again lead to moment explosions.

The ODE for $B(\tau)$ can be solved in closed form and is given by

$$B(\tau) = \frac{1}{\xi^2} \left(\frac{2R(R-\kappa)}{(R+\kappa)e^{R\tau} + R - \kappa} - R + \kappa \right)$$
(4.25)

with $R = \sqrt{D} = \sqrt{\kappa^2 - \xi^2 \nu (2\mu + \sigma^2 (\nu - 1))}$. The explosion of $\mathbb{E}[X(T)^{\nu}]$ is directly related to the explosion of $B(\tau)$. If $D \ge 0$ then $B(\tau)$ is finite for all $\tau > 0$. However, if D < 0 then the denominator in (4.25) may become zero. We have

$$(R+\kappa)e^{R\tau} + R - \kappa = 0 \iff \tau = \frac{1}{R}\left(\log\left(\frac{\kappa - R}{\kappa + R}\right) + 2in\pi\right),\tag{4.26}$$

where $n \in \mathbb{Z}$ and $\text{Log}(\cdot)$ denotes the principal branch of the complex logarithm. The right hand side of (4.26) can be rewritten as

$$\tau = \frac{2}{\sqrt{-D}} \left(\arctan\left(-\frac{\sqrt{-D}}{\kappa}\right) + n\pi \right).$$

Choosing n = 1 yields the first positive root of the above denominator, which completes the proof.

By means of Proposition 4.4.1 we can easily check that the prices produced by Formula (4.24) are reliable¹⁴. However, observe that for typical market parameters we virtually always have $D = \kappa^2 - 2\xi^2 \tilde{\mu}_i > 0$, such that the explosion times (for $\mathbb{E}[X_i(T_n)], i = 1, 2)$ will almost always be infinite.

4.5 Numerical Results

In this section we provide evidence that the approximation developed above is indeed fast and highly accurate. In the first part we want to investigate how the approximation performs in a setting with fully time-dependent Libor parameters, whereas in the second part we compare our method with the approach introduced in a recent paper by Antonov & Arneguy (AA) [AA09].

4.5.1 Time-Dependent Parameters

In order to test our approximation we use a realistic time-dependent (but timehomogeneous) parameter setting. The underlying model is based on 6-month Libor rates ($\tau_i \equiv 0.5$, N = 39) with a typical upward sloping initial yield curve, given in Table B.4, Appendix B. Instantaneous (scalar) Libor volatilities are generated via the *abcd*-parameterization (3.20) with a = 0.04, b = 0.32, c = 1.1 and d =0.17. Instantaneous and effective Libor volatilities are shown in Figure 4.7. Skew parameters decrease linearly from 90% to 40% with decreasing time to maturity

$$\beta_n(t) = \left(1 - \frac{T_n - t}{T_{N-1}}\right) 0.4 + \frac{T_n - t}{T_{N-1}} 0.9, \ 0 \le n \le N - 1.$$

¹⁴ As we have noted earlier, numerical pricing routines usually have a built-in dampening effect. In particular, when evaluating Formula (4.24) by using numerical integration, the domain of integration is typically truncated at some large finite value. Therefore, the produced numbers might look reasonable even if, in fact, the option prices do not exist.

4.5. Numerical Results

For the correlation among the Libor rates we use the time-homogeneous specification (cp. [Reb04], p. 691)

$$\rho_{i,j}(t) = \exp\{-|T_i - T_j|\nu \exp\{-\eta \min(T_i - t, T_j - t)\}\}, \ t \le \min(T_i, T_j),$$

with $\nu = 0.11$, $\eta = 0.22$ and perform a rank reduction¹⁵ to 5 factors (see Figure 4.8). For the stochastic-variance process we use again a speed of mean-reversion of $\kappa = 0.15$ and a volatility of variance of $\xi = 1.3$.



Figure 4.7: Time-homogeneous instantaneous and effective Libor volatilities.



Figure 4.8: Instantaneous Libor correlations $\rho_{ij}(0)$.

Table 4.1 gives the test results for 10Y–2Y CMSSOs with 5Y and 10Y maturity and a payment lag of half a year ($\delta = 0.5$). In both cases our approximation yields excellent results over the entire strike ranges, with absolute errors generally remaining well below 1 bp. Note that the absolute error of 0.2 bp for the 5Y option with strike 0.442% corresponds to an error in "implied correlation" of only 0.17%, i.e., bumping the input correlation for the approximation by 0.17 percentage points yields the MCprice.

5Y Maturity									
Strike	-0.558%	-0.308%	-0.058%	0.192%	0.442%	0.692%	0.942%	1.192%	1.442%
MC	88.2	69.6	52.1	36.9	25.2	17.5	12.5	9.2	6.9
	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)
Approx	87.7	69.2	52.1	37.1	25.4	17.3	12.0	8.5	6.2
Abs. Err.	0.5	0.4	0.1	0.2	0.2	0.2	0.5	0.7	0.7

10Y Maturity									
Strike	-0.509%	-0.259%	-0.009%	0.241%	0.491%	0.741%	0.991%	1.241%	1.491%
MC	72.5	58.0	44.9	33.9	25.9	20.4	16.5	13.6	11.5
	(0.2)	(0.2)	(0.2)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)
Approx	71.7	57.7	44.9	34.2	26.0	20.0	15.7	12.7	10.4
Abs. Err.	0.8	0.4	0.1	0.3	0.1	0.4	0.7	0.9	1.0

Table 4.1: Test results, in basis points (bp), for 10Y-2Y CMS spread options. Numbers in parentheses denote standard deviations $(2^{18}-1 = 262,143 \text{ paths}, 16 \text{ steps per year})$. Convexity-unadjusted forward spreads: $S_1(0) - S_2(0) = 0.335\%$ (5Y), 0.218% (10Y). Convexity-adjusted forward spreads (MC-values): $\mathbb{E}^T[S_1(T) - S_2(T)] = 0.442\%$ (5Y), 0.491% (10Y).



Figure 4.9: Implied normal spread volatilities in basis points for 10Y–2Y CMS spread options with 5Y maturity (top) and 10Y maturity (bottom).

In order to provide some further intuition for the quality of the approximations, we present in Figure 4.9 implied normal spread volatilities¹⁶ for the prices given in Table 4.1. Also with respect to this "error metric" the approximation performs excellently.

To conclude, let us note that there are basically two types of errors made when calculating CMSSO prices via formula (4.24). The first one is due to the numerical calculation of the integrals (including the evaluation of the density of $\overline{V}(T)$) and plays

 $^{^{15}\}mathrm{More}$ details on rank reducing correlation matrices will be provided in Chapter 6. $^{16}\mathrm{See}$ p. 48.

in our case only a minor role¹⁷. The second one, being the "actual" approximation error, comes from the assumptions being made about the dynamics of the swap rates under the measure $\mathbb{Q}^{T_n+\delta}$ (implicitly given by Equation (4.8)). In order to further reduce this error one could principally use more sophisticated drift-approximations¹⁸ and let the percentage drift terms, for instance, also depend on the stochastic drivers of the swap rates. Nevertheless, using the stochastic variance as the only stochastic driver of the percentage drift terms seems to capture already a great part of the variability of the drift terms, and the numerical results suggest, that the quite simple approximation proposed above is already good enough to obtain excellent results.

4.5.2 Constant Parameters

In [AA09] two approximations for CMSSOs are presented, the better of which is based on first changing the measure from the $T_n + \delta$ -forward measure to a so-called "spreadmeasure", under which the swap-rate spread $S(t) \coloneqq S_1(t) - S_2(t)$ is a martingale

$$\mathbb{E}^{T_n+\delta}\big[(S(T_n)-K)^+\big] = \mathbb{E}^S\big[(S(T_n)-K)^+M(T_n)\big]$$

They then use a (non-linear) regression approximation for the stochastic factor $M(T_n)$, which is due to the change of measure, and finally arrive at

$$\mathbb{E}^{T_n+\delta} [(S(T_n) - K)^+] \approx \mathbb{E}^S [(S(T_n) - K)^+ (A(T_n) + B(T_n)\Delta S(T_n) + C(T_n)(\Delta S(T_n))^+)] = A(T_n)\mathbb{E}^S [(S(T_n) - K)^+] + B(T_n)\mathbb{E}^S [(S(T_n) - K)^+ \Delta S(T_n)] + C(T_n)\mathbb{E}^S [(S(T_n) - K)^+ (\Delta S(T_n))^+], \quad (4.27)$$

where $\Delta S(T_n) = S(T_n) - S(0)$ and the coefficients $A(T_n)$, $B(T_n)$ and $C(T_n)$ are obtained by solving a linear system of equations containing the first two moments of $\Delta S(T_n)$ and $(\Delta S(T_n))^+$. In order to calculate expectations involving the spreadoption payoff function, they provide a two-dimensional Fourier representation of the payoff, which is based on the complex Gamma function, and then use two-dimensional Fourier inversion. The Fourier representation of the spread payoff was developed simultaneously by Hurd & Zhou [HZ10].

For numerical comparison we set up a 3-factor LMM with a 20Y annual model tenor and run Monte Carlo simulations with the parameter scenario¹⁹ given in [AA09]. Although we were not able to reproduce their prices exactly, the obtained prices generally differ only by a few basis points and thus absolute differences between Monte Carlo values and the respective approximations should be comparable.

In Table 4.2 and Figure 4.10 we compare the approximation errors of our method with the corresponding values given in [AA09]. For 10Y–2Y CMS spread options with 5Y maturity both approximation methods yield excellent results over the entire strike

¹⁷Provided one chooses reasonable cutoff and accuracy levels.

¹⁸Observe, though, that not only the drift terms but also the diffusion terms are just approximate. ¹⁹Note, that they use as input a somewhat unrealistic yield curve, which increases linearly at the beginning, has a jump of more than 100 basis points at 14Y maturity and then decreases linearly.

5Y Maturity									
Strikes	-0.672%	-0.422%	-0.172%	0.078%	0.328%	0.578%	0.828%	1.078%	1.328%
MC	116.3	98.0	80.9	65.4	52.2	41.4	33.2	26.9	22.2
	(0.2)	(0.2)	(0.2)	(0.2)	(0.2)	(0.2)	(0.2)	(0.2)	(0.2)
Approx	116.9	98.7	81.6	66.0	52.6	41.7	33.2	26.8	21.9
Abs. Err.	0.6	0.7	0.7	0.6	0.4	0.2	0.0	0.1	0.2
Abs. Err. (AA)	1.0	0.5	0.2	0.1	0.2	0.1	0.0	0.0	0.1
			1()Y Matur	ity				
Strikes	-1.321%	-1.071%	-0.821%	-0.571%	-0.321%	-0.071%	0.179%	0.429%	0.679%
MC	109.6	94.9	81.0	68.2	56.9	47.1	39.1	32.6	27.4
	(0.3)	(0.3)	(0.2)	(0.2)	(0.2)	(0.2)	(0.2)	(0.2)	(0.2)
Approx	109.7	94.9	81.0	68.1	56.6	46.6	38.3	31.7	26.4
bs. Err.	0.1	0.1	0.0	0.1	0.3	0.5	0.7	0.9	1.0
Abs. Err. (AA)	1.0	1.2	1.3	1.1	0.5	0.7	2.0	3.0	3.7

Table 4.2: Test results, in basis points (bp), for 10Y-2Y CMS spread options. Numbers in parentheses denote standard deviations $(2^{19} - 1 = 524,287 \text{ paths}, 16 \text{ steps per year})$. Parameter scenario as in [AA09]. Convexity-unadjusted forward spreads: $S_1(0) - S_2(0) = 0.327\%$ (5Y), -0.319% (10Y). Convexity-adjusted forward spreads (MC-values): $\mathbb{E}^T[S_1(T) - S_2(T)] = 0.578\%$ (5Y), 0.055% (10Y).

range and absolute errors generally remain below 1 bp. For spread options with 10Y maturity our approximation method works again remarkably well and errors remain in the 1 bp range. In particular, our method yields more accurate results for all strike values than the AA-approximation, with a maximum absolute error of 1.0 bp opposed to 3.7 bp.

Based on the above results, we find that the accuracy of formula (4.24) is excellent, even for the quite challenging parameter setting given in [AA09], with constant but rather high Libor volatilities between 20% and 39 % and a volatility of variance ξ of 130%. Note also that the approximation is pretty fast: The computation of the spread-option prices in Table 4.2 took only 48 ms on average and thus the approximation should be suitable for being used within calibration routines.

Remark 4.5.1. Instead of integrating against the densities of the normal distribution and of the integrated variance there is a another way of calculating prices: Use the drift approximation derived in Section 4.1, calculate the joint characteristic function of $(\tilde{X}_1(T), \tilde{X}_2(T))'$, combine it with the Fourier transform of the (transformed) spread payoff (see [HZ10] or [AA09]) and then use a two-dimensional Fourier inversion to finally calculate the price. While this is principally possible, it is, as far as our experience goes, generally slower. Using the same adaptive integration scheme and the same precision level as before it took 755 ms on average to calculate the spread option prices given in Table 4.2 via the two-dimensional Fourier inversion. Furthermore, note that when using the AA-approximation given in (4.27), one generally needs to calculate not one but several of these two-dimensional Fourier integrals.



Figure 4.10: Absolute approximation errors in basis points with respect to Monte Carlo for 10Y–2Y CMS spread options with 5Y maturity (top) and 10Y maturity (bottom).

4.6 Conclusion

In this chapter we considered the efficient pricing of CMSSOs within the SV-LMM. In Section 4.3 we introduced a new method for calculating the density of an integrated CIR process, which is based on carefully choosing the integration contour for the necessary Laplace inverse transform. This method allows the fast and accurate evaluation of the density in its entire domain, a result which is also applicable outside the context of spread-option pricing. In Section 4.1 we proposed an approach for approximating the drift terms of the swap rates under the forward measure and derived a semi-analytical formula for CMS spread option prices. The accuracy and speed of our approximation was finally demonstrated in Section 4.5, where it was also compared with the approximation introduced in [AA09]. We have shown that our method is, at least for the used parameter scenario, more accurate (in terms of maximum absolute error) and faster.
Chapter 5

New Correlation Parameterizations

The correct specification of the volatility and correlation structures underlying a LMM is essential for obtaining accurate prices and hedge ratios from the model. While information about the forward-rate volatilities is usually quite easy to extract from market prices of caps and swaptions (cp. Section 3.4), the situation is more delicate with regard to the Libor correlations. Statistically estimating the forward-rate correlations from time series of forward-rate curves is, apart from some practical difficulties, straightforward, and using such historically estimated correlation matrices for valuing less correlation-sensitive products may be acceptable. For products which exhibit a more pronounced dependence on forward-rate correlations, however, this approach is not ideal due to its "backward-looking" nature.

It is therefore tempting to imply the sought correlations directly from market instruments and in this way incorporate the prevailing market views. But also this approach leads to some difficulties, since swaptions, which are usually the main calibration instruments besides caps, depend only relatively mildly on the underlying forward-rate correlations. In fact, as pointed out by Rebonato [Reb02], swaption prices show an "almost total lack of dependence on the shape of the correlation function". Intuitively, this can be explained by the fact that swap rates are just weighted averages of forward Libor rates (see Equation (2.10)). So rather different correlation functions (the volatility functions being fixed) will produce almost identical swaption prices/volatilities, as long as the average correlations among the forward rates are roughly the same. It can therefore be a rather challenging task to obtain stable calibration results when trying to simultaneously back out volatilities and correlations from caps and swaptions alone¹. To overcome this problem, one can consider augmenting the set of calibration instruments with CMSSOs, which are far more correlation sensitive and have become relatively liquid in recent years.

No matter whether the instantaneous forward-rate correlations of a LMM are exogenously specified, by using some historically estimated correlation matrix, or obtained implicitly, by calibrating the model to market prices of interest-rate options, in both cases one needs some parsimonious yet flexible parameterization for the Libor correlations $\rho_{ij}(t)$. Although there exists quite a number of different parameterizations (see e.g. [Reb04] or [Sch05]), most of these are relatively inflexible or are not guaranteed to always yield positive definite matrices. In this chapter we therefore

 $^{^{1}}$ In [Sch05] a regularization method is introduced which may help to overcome some of the stability problems.

present a new generic method for constructing correlation parameterizations that are always positive definite and derive new flexible low-parametric forms.

5.1 General Considerations

The instantaneous Libor correlations, as we have defined them in Equation (3.19), may generally exhibit a functional dependence on calendar time t and on the maturities T_n of the respective forward rates:

$$\rho_{ij}(t) = \rho(t, T_i, T_j), \ i, j = 0, \dots, N-1.$$

However, as with the Libor volatilities (see Section 3.3.1), it is in general a convenient and financially desirable feature for the correlation function $\rho(t, T_i, T_j)$ to display a time-homogeneous behavior. Therefore, the correlation structures that we shall work with in the following will most often be of the form

$$\rho_{ij}(t) = \rho(T_i - t, T_j - t),$$

i.e., they will depend solely on the residual time to maturity of the respective forward rates. For ease of notation we will furthermore make the common assumption that the Libor correlations are piecewise constant on the tenor structure. In this case², we just need to consider the initial matrix

$$\rho_{ij} \coloneqq \rho_{ij}(0), \ i, j = 0, \dots, N-1$$
(5.1)

and can set

$$\rho_{ij}(t) = \rho_{i-\eta(t),j-\eta(t)} \tag{5.2}$$

with $\eta(t)$ as in (2.14). In the next paragraph we will shortly review the general procedure for estimating historical correlations (see also [Sch05] and [AP10b]).

5.1.1 Historical Correlations

For some fixed year fraction τ (e.g. 0.25 or 0.5) and some integer N, let us define so-called *sliding forward rates*³ $L_k^{\rm sl}(t)$ by

$$L_k^{\rm sl}(t) = L(t, t + k\tau, t + (k+1)\tau), \ k = 1, \dots, N.$$

Suppose that for each t_i from a given set of calendar times⁴ t_0, \ldots, t_{N_t} , we have constructed a forward-rate curve $L_k^{\rm sl}(t_i)$, $k = 1, \ldots, N$, bootstrapped from observed prices of deposits, FRAs, Futures and swaps, as described in Section 2.2.5. In Figure 5.1 we show a time-series of EUR forward-rate curves spanning January 2004 to April 2010. Assuming time-homogeneity⁵, the covariance matrix of the normalized

²And assuming roughly equidistant tenor dates T_i .

³With a fixed time to maturity rather than a fixed time of maturity.

⁴The number of calendar times N_t has no relation with the grid dimension N_t from Section 3.3.



Figure 5.1: Evolution of the EUR forward-rate curve (1/14/2008 - 4/26/2010). Semi-annual forward rates ($\tau = 0.5$) were bootstrapped from weekly observations of 6M and 12M deposits and swaps with tenors of 2, 3, 4, 5, 7, 10, 15 and 20 years, using log-cubic interpolation in discount-price coordinates.

forward-rate increments⁶

$$I_k(t_i) \coloneqq \frac{L_k^{\rm sl}(t_i) - L_k^{\rm sl}(t_{i-1})}{\sqrt{t_i - t_{i-1}}}, \ k = 1, \dots, N, \ i = 1, \dots, N_t.$$

can then be estimated via the standard sample covariance estimator

$$\hat{C} = \frac{1}{N_t - 1} \sum_{i=1}^{N_t} (I(t_i) - \bar{I}) (I(t_i) - \bar{I})' \in \mathbb{R}^{N \times N}$$

with column vectors $I(t_i) = (I_1(t_i), \ldots, I_N(t_i))'$, $i = 1, \ldots, N_t$ and standard sample mean $\overline{I} = \frac{1}{N_t} \sum_{i=1}^{N_t} I(t_i)$. If we assume a LMM with time-homogeneous (scalar) Libor volatilities and time-homogeneous correlations of the form (5.1) - (5.2), then the sample correlations

$$\hat{\rho}_{ij} = \frac{\hat{C}_{i+1,j+1}}{\sqrt{\hat{C}_{i+1,i+1}\hat{C}_{j+1,j+1}}}, \ i,j = 0,\dots,N-1$$

may be used as a rough approximation (neglecting any drift terms) for the Libor correlations ρ_{ij} of the LMM.

In Figures 5.4 and 5.6 (which we will comment on later) we present two concrete examples of historically estimated correlation matrices.

 $^{{}^{5}}$ Time-homogeneous forward Libor rate volatilities/correlations imply constant sliding forward-rate volatilities/correlations.

⁶Here we use absolute increments, as forward rates generally tend to be more normal than lognormal. Using e.g. log-returns is of course also possible. However, for small time increments $t_i - t_{i+1}$ the particular choice is not overly important.

Remark 5.1.1. The historical correlation matrix $(\hat{\rho}_{ij})_{ij}^{N-1}$ will generally be of full rank. So we implicitly assumed that we are working with a full-factor model, i.e., d = N. The problem of reducing the rank of Libor correlation matrices will be dealt with in Chapter 6.

In practice, obtaining meaningful correlations from historical data is not as simple as it might look at first glance. First, correlations will generally not be constant over time, such that the estimation horizon will have a significant influence on the results. On the one hand, the used estimation period should be long enough to obtain statistically meaningful estimates, but on the other hand, it should not be too long if we are to capture current market conditions. Second, non-synchronous prices of the instruments, from which the forward-rate curves are constructed, can introduce nonnegligible errors. To minimize effects due to non-synchronicity (and other noise in the data) it is therefore advisable to not use too short a sampling frequency (weekly often works reasonably well). Third, the interpolation method used in the bootstrapping procedure often has a considerable influence on the estimation results, see e.g. [AL03]. In particular, so-called "ringing effects" in forward-rate curves may lead to highly peaked correlation matrices with counterintuitive (e.g. negative) entries. To avoid ringing effects it is best to choose a "relatively linear" (in forward-rate coordinates) interpolation method and to use only forward rates with maturities that correspond to tenors of the underlying instruments from which the curves are constructed. Last but not least, the statistical estimation of correlations by itself may suffer from certain robustness problems, see e.g. [JKB95] or [Wil05].

In view of these problems and because of the backward looking nature of historical parameters in general, it is most often preferable to extract the Libor correlation via an implied calibration to market prices of correlation-sensitive products, as we will demonstrate in Chapter 7. Nevertheless, historically estimated correlation matrices can be used to get a feel for the general shape of Libor correlation matrices and to come up with reasonable correlation parameterizations.

5.1.2 Stylized Facts

No matter whether historical or implied correlations are to be used within a model, in both cases one needs a parsimonious and smooth parameterization of the correlation matrix. In case of historically estimated correlation matrices, this is due to the fact that these are often quite noisy and may contain counterintuitive entries (see above). By fitting a low-parametric functional form to the historical correlations, one tries to obtain a reasonably smooth matrix, which captures only the main features of the historical data. In case of the implied calibration approach the reason is that one cannot expect to obtain reliable results for all the N(N-1)/2 entries of the correlation matrix.

Before we will have a closer look at some candidate parametric forms for the matrix $\rho = (\rho_{ij})_{i,j=0}^{N-1}$, we need to fix some general requirements that ρ must comply with in order to be regarded as a valid correlation matrix:

- (A1) ρ must be real and symmetric,
- (A2) $\rho_{ii} = 1, \ i = 0, \dots, N-1,$

(A3) ρ must be positive semi-definite.

Besides these minimum requirements, which any valid correlation matrix must satisfy, there is conventional wisdom that a forward-rate correlation matrix should also exhibit the following features:

- (B1) $i \mapsto \rho_{ij}, i \ge j$ is decreasing,
- (B2) $i \mapsto \rho_{i+p,i}$ is increasing for fixed $p \in \{1, \dots, N-2\}$.

The first requirement comes from the observation that the farther apart two forward rates are, the less correlated they typically are. Furthermore, one usually expects same distance rates to move more "in line" at the long end than at the short end of the forward-rate curve, i.e., the 2Y and the 5Y rate should be more decorrelated than, say, the 12Y and the 15Y rate. This results in requirement (B2), that the entries along the sub-diagonals should be increasing. The above stylized facts can be readily observed from the historically estimated correlation matrices depicted in Figures 5.4 and 5.6.

5.2 Classical Parameterizations

The simplest parametric form for a forward-rate correlation matrix is arguably the

Classical exponential form:

$$\rho_{ij} = \exp\left(-\beta|i-j|\right), \ \beta \ge 0. \tag{5.3}$$

The matrix clearly satisfies (B1) for any $\beta > 0$ and one can also show that it is always an admissible correlation matrix, i.e., it satisfies (A1)–(A3). Unfortunately, expression (5.3) cannot reproduce feature (B2) since the sub-diagonals are simply constant. In order to have more degrees of freedom and also to take condition (B2) into account, Rebonato [Reb99a] proposes the following

Rebonato 3-parametric form I:

$$\rho_{ij} = \rho_{\infty} + (1 - \rho_{\infty}) \exp(-|i - j|(\beta - \alpha \max(i, j)))), \quad (5.4)$$
$$-1 < \rho_{\infty} < 1, \ \beta > 0, \ 0 \le \alpha \le \frac{\beta}{N - 1}.$$

However, while (5.4) may produce realistic correlations, Schoenmakers & Coffey [SC03] point out that the matrices are not guaranteed to be positive semi-definite for all parameter combinations. One therefore always has to check, whether the matrix has negative eigenvalues and, if necessary, "repair" it by using for example one of

the methods introduced in [RJ99] or [Hig02]. This, of course, adds an undesirable computational burden, at least if we want to use full-factor matrices⁷.

In order to fix the above problem Rebonato [Reb04] provides a further functional form

Rebonato 3-parametric form II (Reb3):

$$\rho_{ij} = \rho_{\infty} + (1 - \rho_{\infty}) \exp\left(-\beta |i - j| \exp\{-\alpha \min(i, j)\}\right), \qquad (5.5)$$
$$\beta > 0, \ \alpha \in \mathbb{R}, \ -1 < \rho_{\infty} < 1,$$

and claims that all eigenvalues of ρ_{ij} can be shown to be always positive. However, it is quite easy to find otherwise realistic parameters that also let (5.5) have negative eigenvalues⁸. One would therefore be again forced to check during a calibration, whether the resulting matrix is positive semi-definite.

Parametric forms that satisfy conditions (A1)-(A3) and (B1)-(B2) by their very construction can be derived from the semi-parametric family introduced by Schoenmakers & Coffey (SC) [SC03].

Semi-parametric Schoenmakers-Coffey family:

Let b_0, \ldots, b_{N-1} be a strictly increasing sequence of coefficients such that

$$1 = b_0 < b_1 < \dots < b_{N-1} \quad \text{and} \quad \frac{b_0}{b_1} < \frac{b_1}{b_2} < \dots < \frac{b_{N-2}}{b_{N-1}}.$$
 (5.6)

Then set

$$\rho_{ij} = \frac{b_j}{b_i}, \quad 0 \le j < i \le N-1.$$
(5.7)

Conditions (B1) and (B2) are directly enforced by the monotonicity constraints in (5.6), and it can also be shown that (5.7) is always positive definite. The above family is called semi-parametric as the number of parameters needed is N-1, opposed to the N(N-1)/2 entries of the entire correlation matrix.

As the conditions in (5.6) might be difficult to handle in practice, Schoenmakers & Coffey [SC03] show that a sequence of the above type can always be characterized in terms of another, not necessarily increasing sequence of N-1 coefficients $\Delta_i \geq 0$, $1 \leq i \leq N-1$ via

$$b_i = \exp\left(\sum_{k=1}^{N-1} \min(k, i)\Delta_k\right), \quad 0 \le i \le N-1.$$
(5.8)

More low-parametric correlation structures can now be easily derived from the above parameterization by choosing simple functional forms for either the b's or the $\Delta's$. Imposing for example a linear behavior to the $\Delta's$ results, after some rearrangements, in the following popular two-parametric form

 $^{^{7}}$ If a PCA-style factor reduction is applied, then the factors corresponding to smaller (and negative) eigenvalues are discarded anyway.

⁸See the calibration examples below.

Schoenmakers-Coffey 2-parametric form I (SC2):

$$\rho_{ij} = \exp\left(-\frac{|i-j|}{N-1}\left(-\log \rho_{\infty} + \eta h(i,j)\right)\right), \quad 0 \le i, j \le N-1, \quad (5.9)$$

$$h(i,j) = \frac{i^2 + j^2 + ij - 3Ni - 3Nj + 6i + 6j + 2N^2 - 7N + 5}{(N-2)(N-3)}$$

$$\rho_{\infty} \in (0,1), \ \eta \in [0, -\log \rho_{\infty}].$$

Notice, that $\rho_{\infty} = \rho_{1,N-1}$ is the correlation between the farthest apart forward rates, whereas η can be used to control the rate of correlation decay.

Another two-parametric form can be constructed by using

$$b_i = \exp(i^{\alpha}\beta), \quad i = 0, \dots, N-1, \ 0 < \alpha < 1, \ \beta > 0.$$

Introducing $\rho_{\infty} \coloneqq 1/b_{N-1}$ we obtain

Schoenmakers-Coffey 2-parametric form II :

$$\rho_{ij} = \exp\left(\log \rho_{\infty} \left| \left(\frac{i}{N-1}\right)^{\alpha} - \left(\frac{j}{N-1}\right)^{\alpha} \right| \right), \quad 0 \le i, j \le N-1.$$
(5.10)

Even though the two above parameterizations are generally not quite as flexible as for example (5.4), one can rest assured that they always produce admissible correlation matrices.

5.3 New Flexible Correlation Parameterizations

When calibrating a model to caps and swaptions or when fitting historical correlation matrices, it may be reasonable to use only one- or two-parametric forms, in order to obtain a sufficient smoothing effect and to avoid overfitting. In contrast, if we augment the set of calibration instruments with CMSSOs, we should be able to back out more detailed information about the underlying correlation structure. Put differently, if we assume that the volatilities are mainly determined by prices of caps and swaptions, then we cannot expect to obtain a reasonable fit to, say, 5-10 prices of CMS spread options with different maturities with just two correlation parameters.

In order to obtain more flexible parameterizations that are positive definite, a natural attempt would be to parameterize the Δ -coefficients of the SC-family in a very general way. Unfortunately, controlling the shape of the correlation matrix via the Δ - or *b*-coefficients is not very intuitive and it is not evident how a more general parameterization of these coefficients should look like. Apart from that, there is a more fundamental problem with the SC-family when it comes to designing more flexible correlation structures. Considering expression (5.7) we find that the entries of the first column of the correlation matrix are given by $\rho_{i1} = 1/b_i$. Consequently, fixing the first column of the matrix determines all N-1 parameters b_i and therefore also all the remaining entries of the matrix. In particular, the last row of the matrix $\rho_{N-1,i} = b_i/b_{N-1}$ is always approximately inverse-proportional to the first column of the matrix, which makes it difficult to control the "front-end" and "back-end" correlations separately and introduces constraints on possible shapes of the correlation matrix.

Below we present an alternative way of characterizing matrices from the SCfamily, which gives rise to a new construction principle for correlation matrices.

5.3.1 Alternative Characterization of the SC-Family

Let us fix some coefficients b_1, \ldots, b_{N-1} such that (5.6) holds and denote the corresponding correlation matrix from the SC-family by ρ . If we consider the Cholesky decomposition L of this matrix

$$\rho = LL',$$

we find that

$$L_{ij} = \begin{cases} b_i^{-1}, & j = 0, \\ b_i^{-1} \sqrt{b_j^2 - b_{j-1}^2}, & 0 < j < i, \\ 0, & \text{otherwise.} \end{cases}$$
(5.11)

Likewise, the Cholesky decomposition can be written as

$$L = AC$$

where

$$A = \text{diag}(b_0^{-1}, \dots, b_{N-1}^{-1})$$

and C has entries

$$C_{ij} = \begin{cases} 1, & j = 0, \\ \sqrt{b_j^2 - b_{j-1}^2}, & 0 < j < i, \\ 0, & \text{otherwise} \end{cases}$$

i.e., C is a lower triangular matrix with constant columns (see Figure 5.2). The column coefficients

$$c_0 \coloneqq 1, \ c_j \coloneqq \sqrt{b_j^2 - b_{j-1}^2}, \quad j = 1, \dots, N-1,$$

can alternatively be defined recursively via

$$c_j \coloneqq \sqrt{b_j^2 - \sum_{k=0}^{j-1} c_k^2}, \quad j = 1, \dots, N-1,$$
 (5.12)

which reveals that for the *i*-th row of C, denoted by r_i , we have

$$||r_i|| = b_i, \quad i = 1, \dots, N-1.$$



The above observations suggest that one should be able to obtain more flexible semiparametric forms by suitably reparameterizing the column coefficients c_i of the matrix C or even allowing non-constant columns, while keeping the row norms fixed at b_i to preserve the general structure.

5.3.2 Reformulation of the Cholesky Decomposition

Inspired by the above findings, we present the following general construction principle for correlation matrices:

Lemma 5.3.1. Let N > 2, $I := \{1, \ldots, N-1\}$, $I_0 := I \cup \{0\}$. Assume $f : I_0 \to [-1, 1]$ and $h : I \times I \to \mathbb{R}$ are functions such that f(0) = 1 and $h(i, 1) \neq 0$, $i \in I$. Define Lto be a lower triangular matrix with entries

$$L_{ij} = \begin{cases} f(i), & j = 0, \\ h(i,j)\sqrt{\frac{1-f(i)^2}{a_i}}, & 0 < j \le i, \\ 0, & otherwise, \end{cases}$$
(5.13)

where

$$a_0 = 1, \ a_i = \sum_{k=1}^i h(i,k)^2, \ i \in I.$$

Then the matrix

$$\rho_{ij} = (LL')_{ij} = f(i)f(j) + \sqrt{\frac{1 - f(i)^2}{a_i} \cdot \frac{1 - f(j)^2}{a_j}} \sum_{k=1}^{\min(i,j)} h(i,k)h(j,k), \quad i,j \in I_0$$
(5.14)

is a proper correlation matrix with first column/row

$$\rho_{i0} = \rho_{0i} = f(i), \ i \in I_0.$$

If, furthermore, |f(i)| < 1 and $h(i,i) \neq 0 \ \forall i \in I$, then ρ is of full rank and $L = (L_{ij})_{i,j\in I_0}$ is its unique Cholesky decomposition.

Proof. It is easy to see that ρ is real and symmetric. Furthermore, it is clearly positive semi-definite since for arbitrary $x \in \mathbb{R}^N$ we have

$$x'\rho x = x'LL'x = (x'L)(x'L)' = ||xL||^2 \ge 0.$$

Finally, we have

$$\rho_{ii} = \sum_{k=0}^{i} L_{ik}^2 = f(i)^2 + \sum_{k=1}^{i} h(i,k)^2 \frac{1 - f(i)^2}{a_i} = 1.$$

The last claim follows from linear algebra.

Parametric correlation structures can now be easily derived by specifying simple parametric forms for the functions f and h.

Remark 5.3.2. In principle, it is also possible to directly parameterize the Cholesky decomposition of a forward correlation matrix. However, it is generally more difficult to come up with simple parametric forms, as the Cholesky decompositions of otherwise smooth forward correlation matrices may look quite erratic, and one has to simultaneously take care of the normalization. In contrast, splitting up the problem into finding parametric forms for f and h is usually easier. Since f represents the "front-end" correlations in the first column of ρ , it is natural to require f to be some positive function with f(0)=1, which decreases monotonically to some asymptotic correlation level. The behavior of the remaining entries of the correlation matrix can be controlled by means of the function h, where using a smooth functional form for h will generally result in a smooth transition from the entries in the first column/row to the 1's on the diagonal.

Remark 5.3.3. Comparing (5.11) and (5.13) it can be easily seen that for a given sequence $b_i, i \in I_0$, which satisfies conditions (5.6), one can recover the corresponding SC correlation matrix from (5.14) by choosing

$$f(i) = 1/b_i$$
, and $h(i, j) = \sqrt{b_j^2 - b_{j-1}^2}, i, j \in I.$

Observe, that while matrices from the SC-family naturally satisfy conditions (B1) and (B2), this is not necessarily the case for correlation matrices constructed via "arbitrary" generating functions f and h. However, for our applications this is not a problem since our primary goal is to provide more flexible correlation structures that can be used when calibrating to prices of CMSSOs. Provided conditions (B1) and (B2) are reasonable, they should naturally come out of the market data, rather than be imposed. We will return to this point in Chapter 7.

Finally, we note that the result of Lemma 5.3.1 also holds in the opposite direction:

Lemma 5.3.4. Let ρ be an arbitrary $N \times N$ correlation matrix of full rank. Then there exist functions $f: I_0 \to [-1, 1]$ and $h: I \times I \to \mathbb{R}$ such that

$$\rho_{ij} = f(i)f(j) + \sqrt{\frac{1 - f(i)^2}{a_i} \cdot \frac{1 - f(j)^2}{a_j}} \sum_{k=1}^{\min(i,j)} h(i,k)h(j,k), \quad i, j \in I_0$$
(5.15)

and the Cholesky decomposition L of ρ can be written as

$$L_{ij} = \begin{cases} f(i), & j = 0, \\ h(i,j)\sqrt{\frac{1-f(i)^2}{a_i}}, & 0 < j \le i, \\ 0, & otherwise, \end{cases}$$
(5.16)

with

$$a_i = \sum_{k=1}^i h(i,k)^2, \ i \in I$$

Proof. If ρ is real symmetric positive definite, then its Cholesky decomposition $L = (L_{ij})_{i,j \in I_0}$ exists and is unique. Now simply set $f(i) \coloneqq L_{i0} = \rho_{i0}, i \in I_0$ and $h(i,j) \coloneqq L_{ij}, i, j \in I$, and verify that (5.15) and (5.16) hold.

5.3.3 New Parametric Forms

Below we present two concrete examples of choices for the "generating functions" f and h. The resulting correlation structures are quite flexible and can be written in a nice and compact form.

For the first parameterization let us take

$$f(i) = \exp\left(-\beta i^{\alpha}\right), \quad \alpha, \beta > 0, \ i \in I_0,$$
$$h(i,j) = \exp\left(-\frac{\gamma}{N-2} j\right), \quad \gamma \in \mathbb{R}, \ i,j \in I.$$
(5.17)

Then, after adding a further parameter $\rho_{\infty} \in [0, 1)$ for controlling the asymptotic correlation level (see Appendix C) and after some simplifications, we obtain the following

New 4-parametric form (4P):

$$\rho_{ij} = \rho_{\infty} + (1 - \rho_{\infty}) \left[\exp\left(-\beta(i^{\alpha} + j^{\alpha})\right) + \psi(i, j)\sqrt{\left(1 - \exp\{-2\beta i^{\alpha}\}\right)\left(1 - \exp\{-2\beta j^{\alpha}\}\right)} \right], \quad i, j \in I_{0}, \quad (5.18)$$

$$\psi(i, j) = \begin{cases} 1, & \min(i, j) = 0, \\ \sqrt{\frac{\min(i, j)}{\max(i, j)}}, & \min(i, j) > 0, \quad \gamma = 0, \\ \sqrt{\frac{1 - \exp\left(-\frac{2\gamma}{N - 2}\min(i, j)\right)}{1 - \exp\left(-\frac{2\gamma}{N - 2}\max(i, j)\right)}}, & \min(i, j) > 0, \quad \gamma \neq 0, \\ \alpha, \beta > 0, \quad \gamma \in \mathbb{R}, \quad \rho_{\infty} \in [0, 1). \end{cases}$$

This parametric form has two parameters α and β for controlling the front-end correlations, one parameter γ for controlling the correlations at the back end, and one parameter ρ_{∞} for the asymptotic correlation level. It is considerably more flexible than the 2-parametric forms (5.9) and (5.10). It is generally also more flexible than the 3-parametric form (5.4) and additionally even always positive definite.

5. New Correlation Parameterizations

Note that h in (5.17) does not depend on the row index i and thus the matrix C from Section 5.3.1 has again constant columns. In order to gain additional flexibility this restriction can be relaxed. As a refinement of (5.17) let us take

$$h(i,j) = \exp\left(-\left(\frac{i-1}{N-2}\gamma + \frac{N-1-i}{N-2}\delta\right)\left(\frac{j}{i}-1\right)\right), \ \gamma \in \mathbb{R}, \ i,j \in I,$$
(5.19)

which yields, again after applying the ρ_{∞} -extension, the following

New 5-parametric form (5P):

$$\rho_{ij} = \rho_{\infty} + (1 - \rho_{\infty}) \left[\exp\left(-\beta(i^{\alpha} + j^{\alpha})\right) + \frac{\vartheta_{ij}}{\sqrt{\vartheta_{ii}\vartheta_{jj}}} \sqrt{\left(1 - \exp\{-2\beta i^{\alpha}\}\right) \left(1 - \exp\{-2\beta j^{\alpha}\}\right)} \right], \quad i, j \in I_0,$$
(5.20)

$$\vartheta_{ij} = \begin{cases} 1, & \min(i, j) = 0, \\ \min(i, j), & \min(i, j) > 0, \ \xi_i \xi_j = 1, \\ \frac{(\xi_i \xi_j)^{\min(i, j)} - 1}{1 - 1/(\xi_i \xi_j)}, & \min(i, j) > 0, \ \xi_i \xi_j \neq 1, \end{cases}$$

$$\xi_i = \exp\left(-\frac{1}{i} \left(\frac{i - 1}{N - 2}\gamma + \frac{N - 1 - i}{N - 2}\delta\right)\right),$$

$$\alpha, \beta > 0, \ \gamma, \delta \in \mathbb{R}, \ \rho_{\infty} \in [0, 1).$$

The interpretation of the parameters is essentially the same as above, only that we now have two parameters for controlling the transition from the diagonal to the first column of the correlation matrix: δ acts more on the rows in the upper part whereas γ acts mostly on the rows in the lower part of the correlation matrix. Note, that if we fix $\delta = 0$ we obtain again a 4-parametric version, which hardly differs from (5.18). Some examples for possible shapes of correlation matrices, that can be generated with (5.20), are shown in Figure 5.3.



Figure 5.3: A variety of shapes that can be generated by the 5Pparameterization. The corresponding parameter values are given in the table (numbering from top, left to right). In all cases N = 40. The asymptotic correlation level ρ_{∞} was kept fixed at 0.2.

5.4 Fitting Historical Correlations

In Figures 5.4-5.7 we present two concrete examples, where the parametric forms SC2, Reb3 and 5P were fitted to historical correlations. For estimating the historical matrices we used in both cases 4 years of weekly data from the EUR market, spanning Jan. 2004 to Jan. 2008 and Apr. 2006 to Apr. 2010, respectively. For each date in the respective data set a 6M forward-rate curve was constructed from market quotes for deposits and swaps. The correlations were then calculated from the obtained



Figure 5.4: Historically estimated forward-rate correlations (Jan. 2004 – Jan. 2008) and fitted parameterizations. From top, left to right: Historical, SC2, Reb3, 5P.

Corrform	2004 - 2008	2006 - 2010
SC2	0.132	0.187
Reb3	0.096	0.133
$5\mathrm{P}$	0.061	0.076

Table 5.1: Root mean square errors with respect to the historical correlation matrices.

time series of 6M (sliding) forward rates with maturities of 0.5, 1, 2, 3, 4, 5, 7, 10, 20 and 30 years. Finally, the parametric forms were fitted to these correlations by using simple least-squares. The optimal parameters are given in Table D.5 and D.6 in Appendix D. As can be seen from the plots in Figures 5.4-5.7 (and the root mean square errors (RMSE) in Table 5.1), the new 5-parametric form is capable of fitting the historical matrices far better than the other parameterizations, while

110

still providing a sufficient smoothing effect. We also note that for both data sets the best-fit parameters for the Reb3 parameterization result in non-positive definite matrices. Hence, before being able to use these matrices one is forced to somehow "repair" them as mentioned above.



Figure 5.5: Historically estimated forward-rate correlations (Jan. 2004 – Jan. 2008) and fitted parameterizations. First columns (top) and last rows (bottom) of the correlation matrices.

According to our experience the 5P-parameterization is generally flexible enough for most applications, not only for fitting historical matrices but also when used for calibrating LMMs to market prices of caps, swaptions and CMSSOs, as we will demonstrate in Chapter 7. If, however, this correlation structure is still considered as not having enough degrees of freedom, one can pursue a more "semi-parametric" approach. For this, one can introduce some grid of row and column indices $\{i_k\} \times$ $\{j_l\}, \ k = 1, \ldots, n, \ l = 1, \ldots, m,$ a vector $\theta \in [0, 1]^n$ and a matrix $H \in \mathbb{R}^{n \times m}_+$. The "parameters" θ_k and H_{kl} are then interpreted as

$$f(i_k) = \theta_k,$$

$$h(i_k, j_l) = H_{kl},$$

where the remaining values of the generating functions f and h are obtained by using suitable interpolation schemes⁹. With such a grid-based approach one can have "arbitrary" many parameters, but one also runs the risk of overfitting if no further regularity conditions are imposed on θ and H.

⁹For example, Kruger-Splines for f and some bicubic interpolation in log-coordinates for h would be possible.



Figure 5.6: Historically estimated forward-rate correlations (Apr. 2006 – Apr. 2010) and fitted parameterizations. From top, left to right: Historical, SC2, Reb3, 5P.



Figure 5.7: Historically estimated forward-rate correlations (Apr. 2006 – Apr. 2010) and fitted parameterizations. First columns (top) and last rows (bottom) of the correlation matrices.

5.5 Conclusion

In this chapter we have presented a new generic method for constructing parameterizations of forward-rate correlation matrices, that are always positive definite. New flexible parametric forms were derived, where different portions of the correlation matrix can intuitively be controlled by means of the parameters. Furthermore, we have provided evidence that our new 5-parametric form can fit historically estimated matrices better than the standard parameterizations.

Ultimately, the main field of application of the new flexible parameterizations, however, will be the implied calibration to CMSSOs, where they will allow for an easy handling and excellent market fits. Another great advantage of parametric forms derived via our generic construction principle is that they can be readily combined with the DCT rank-reduction method. This new method, which we will present in the next chapter, allows to significantly reduce the computational cost of the calibration procedure.

Chapter 6

DCT Rank-Reduced Parameterizations

In practice, the number of factors d used for a LMM is typically much lower than the number of modeled forward rates N (e.g. 3–10 factors for a model with 60 or 80 forward rates). The reasons for using low-factor models are manifold. For example, when empirically analyzing historical yield curve movements, one often finds (using principle component analysis (PCA)) that a low number of factors¹ suffices to explain a large amount² of the variability of forward rates. In particular, we do not have 60 or even more (meaningful) factors. Apart from that, the shapes of interest-rate curves generated by low-factor LMMs will generally look more realistic, see e.g. [Fri07], p. 368. On the numerical side, using a low number factor will save computing time when performing MC simulations, since the state-dependent Libor rate drift terms (which constitute a large part of the total computational cost) can be calculated more efficiently (see [Jos03]) and the MC variance is generally lower³.

Ultimately, the choice for a particular number of factors d will also be influenced by the products that are to be priced with the model. Some interest-rate products, such as Bermudan swaptions for example, are not very sensitive to the number of factors and using two or even just a single⁴ factor may be sufficient, see e.g. [AA01]. However, in case of strongly correlation-sensitive products such as CMS spread-linked structures⁵, for instance, a larger number of factors (up to say 10) is necessary to capture the market implied correlations and to achieve a sufficient decorrelation among the rates.

The obvious problem now is that almost all⁶ practically relevant correlation pa-

¹Due to the shape of the first three eigenvectors, the corresponding factors can often be interpreted as "level", "slope" and "curvature".

 $^{^{2}}$ In case of the two historical data sets from the last chapter, 5 factors are sufficient to explain more than 97% of the variance of the forward-rate increments.

³Note also that Quasi-MC methods typically work better in lower dimensions.

⁴In this case all forward rates are driven by a single Brownian motion and hence all rates are instantaneously perfectly correlated. Note carefully, however, that this does not mean that the term correlations of Libor and swap rates in the model are also all equal to one. In fact, term correlations are not only determined by the instantaneous Libor correlations ρ_{ij} but also by the shape of the Libor volatilities $\lambda_n(t)$. Time-dependent volatilities can lead to a significant decorrelation effect, see [Fri07], Section 21.2 for illustrative examples.

⁵See Chapter 4 of [Sch05] for other examples.

⁶The only two low-rank low-parametric forms that we are aware of are the ones given in [RBM07], Section 10.1 and [AA01], Appendix C. However, both are overly restrictive with regard to flexibility and number of factors, and therefore generally cannot provide reasonable fits to market data.

rameterizations, including those introduced in the previous chapter, yield full-rank correlation matrices and hence some kind of rank-reduction method has to be applied before the matrices can actually be used within the model. This adds of course an undesirable computational burden and may significantly slow down the calibration process. In the following we will shortly review the general problem of rank-reducing correlation matrices and then present a new simple method for reducing the rank of given positive-definite parametric forms. Due to its simplicity and the low computationally cost, the method can be considered as a way of generating new low-rank low-parametric forms from given full-rank parameterizations.

6.1 Existing Methods for Rank-Reducing Correlation Matrices

The problem of finding a low-rank correlation matrix nearest to a given correlation matrix appears in many areas in finance and quite a number of different approaches (see e.g. [PG04], Section 2 for a good literature review) exist for solving this problem. From a mathematical point of view, the problem is usually formulated as follows: Find

$$X^* = \underset{X \in \mathbb{R}^{N \times d}}{\operatorname{argmin}} \{ \| XX' - C \|_F : \operatorname{diag}(XX') = \mathbf{1} \},$$
(6.1)

where $d \leq N$, $\|.\|_F$ denotes the Frobenius norm, $\|A\|_F^2 \coloneqq \operatorname{tr}(AA')$, and where C is a given correlation matrix, i.e., $C \in \mathbb{R}^{N \times N}$ satisfies (A1)–(A3). The matrix $C^* \coloneqq X^*X^{*'}$ is then the low-rank approximation to C with $\operatorname{rank}(C^*) \leq d$.

Different methods that aim^7 at finding X^* in (6.1) are presented for example in [ZW03], [PG04], [GP07] or [RBM07]. All of these methods involve some kind of numerical optimization or iterative algorithm and are therefore (especially for large correlation matrices) computationally expensive.

In practice, speed is often more important than a high degree of accuracy⁸ and one therefore often resorts to the so-called *modified PCA* or *eigenvalue zeroing method* (see e.g. [Bri02]), which is based on the spectral or eigenvalue decomposition of the correlation matrix. Since C is real and symmetric it can be written as

$$C = Q\Lambda Q',$$

with an orthonormal matrix $Q \in \mathbb{R}^{N \times N}$ and a diagonal matrix Λ , containing the corresponding eigenvalues in descending order. A rank-*d* approximation⁹ to the original correlation matrix is then given by

$$\bar{C} = \bar{Y}\bar{Y}',\tag{6.2}$$

$$\bar{Y}_{i.} = \frac{Y_{i.}}{\|Y_{i.}\|_2}, \ i = 0, \dots, N-1,$$
(6.3)

$$Y = Q\Lambda_d^{1/2},\tag{6.4}$$

⁷Most of these methods do not in general guarantee to converge to the (globally) optimal solution. However, it is quite easy to check the optimality of a potential solution by a criterion based on Lagrangian multipliers; see [GP07].

⁸Recall that the statistical estimation of correlations is often difficult and the corresponding confidence intervals may therefore be quite large anyway, see Section 5.1.1.

⁹Provided the d - th eigenvalue is not equal to zero.

where $\Lambda_d^{1/2} \in \mathbb{R}^{N \times d}$ consists only of the first *d* columns of $\Lambda^{1/2}$ and where A_i denotes the *i*-th row of a matrix *A*. Without the normalization step in (6.3), the approach would correspond to ordinary PCA, which gives the optimal approximation to *C* in the standard least-square sense, i.e.,

$$Y = \operatorname*{argmin}_{X \in \mathbb{R}^{N \times d}} \{ \| XX' - C \|_F \}.$$

$$(6.5)$$

It is clear that the subsequent rescaling of Y, which is necessary in order for \overline{C} to have unit diagonal, affects the optimality of the PCA-based solution and we will in general have

$$\|\bar{C} - C\|_F > \|C^* - C\|_F.$$

For regular forward-rate correlation matrices and a modest number of factors d, however, the ordinary PCA-based solution YY' is often not too far from having a unit diagonal. The difference between $\bar{Y}\bar{Y}'$ and the optimal solution $X^*X^{*'}$ will therefore generally be rather small, as the normalization step will then have little effect. For comparisons between the optimal and the modified PCA-based solution see for example [Bri02] or Section 6.3 below. Since the modified PCA method is easy to implement, reasonably fast and accurate it is quite popular among practitioners.

Application to the LMM

As we have discussed in the previous chapter, there exist two possible approaches when it comes to specifying the Libor correlation matrix (5.1). First, one can use historically estimated correlations. In this case the rank-reduction problem is exactly (6.1), where C is the empirical correlation matrix, often smoothed via some parametric form, and the objective is to find the nearest rank-d approximation to the given matrix. Here, the relative slowness of the methods mentioned in the previous section is generally not a serious problem, as the historically estimated correlation matrices are usually not updated too often, and the rank-reduction task must be only performed once outside of the general LMM calibration routine.

In the second approach, which we are mainly interested in, the correlations are implied from market prices of traded options. In order to reduce the effective number of parameters, and since the N(N-1)/2 entries of the forward-rate correlation matrix cannot all be inferred from market instruments, one usually uses one of the parameterizations introduced in Chapter 5. The correlation parameters together with the other model parameters are then obtained by minimizing some calibration objective function.

As we have noted above, almost all of the existing correlation parameterizations yield full-rank matrices, so that these matrices cannot be used directly. Instead, for a given parameter vector θ , first a rank reduction

$$\rho(\theta) \longrightarrow \hat{\rho}^d(\theta)$$

must be performed. This rank reduction step must be performed over and over again inside the calibration loop and thus using the modified PCA method for this task is usually the only computationally viable approach. Note carefully, that for the above application it is not of great importance whether $\|\rho(\theta) - \hat{\rho}^d(\theta)\|_F$ is small or not, as long as $\hat{\rho}^d(\theta)$ is able to generate similar shapes of correlation matrices as the original parametric form $\rho(\theta)$. In effect, rather than solving a problem like (6.1), one simply tries to obtain with the rank-reduction step a low-rank parametric form. For such a "simple task", even the modified PCA method can be considered an overkill. Although being generally much faster than other rank-reduction methods, it still requires a full numerical eigenvalue decomposition and is therefore one of the computationally more expensive calculation steps performed during the calibration. Especially if one tries to calibrate time-dependent correlations, the PCA rank reductions can account for a non-negligible part of the required total calibration time. In the next section we introduce a new simple approach for rank-reducing positive definite parametric forms, which is much faster than the modified PCA method.

6.2 The DCT Rank Reduction Method

The approach we present below is based on applying the discrete cosine transform (DCT) to the rows of the Cholesky decomposition of a correlation matrix. Similar to the discrete Fourier transform (DFT), the DCT transforms a discrete function or signal from the spatial domain to the frequency domain and expresses it in terms of a sum of oscillating functions with different frequencies and amplitudes. While the DFT uses cosine and sine terms (in the form of complex exponentials), the DCT only uses cosine terms, and the output remains real (for real input). There exist different types of DCTs (usually denoted by type I, II, III and IV), which differ by implying different boundary conditions, see [BYR06]. In the following we will only consider the DCT of type III, which we have found to work best for the applications that we have in mind. In this case N real numbers x_k , $k = 0, \ldots, N-1$ are transformed into N real numbers \hat{x}_k , $k = 0, \ldots, N-1$ according to the formula

$$\hat{x}_k = \frac{1}{\sqrt{N}} \left(x_0 + \sqrt{2} \sum_{n=1}^{N-1} x_n \cos\left(\frac{\pi}{N} n(k+1/2)\right) \right), \ k = 0, \dots, N-1.$$
(6.6)

Written in this form¹⁰, the transform is its own inverse, i.e.,

$$x_k = \frac{1}{\sqrt{N}} \left(\hat{x}_0 + \sqrt{2} \sum_{n=1}^{N-1} \hat{x}_n \cos\left(\frac{\pi}{N} n(k+1/2)\right) \right), \ k = 0, \dots, N-1$$
(6.7)

and x is expanded into a sum of sinusoids of increasing frequency. For "smooth" sequences x, the absolute values of the coefficients \hat{x}_k usually decrease quite rapidly, such that the original x_k 's can be approximated fairly well by using only the first few low frequency components in (6.7), i.e.,

$$x_k \approx \frac{1}{\sqrt{N}} \left(\hat{x}_0 + \sqrt{2} \sum_{n=1}^{d-1} \hat{x}_n \cos\left(\frac{\pi}{N} n(k+1/2)\right) \right), \ k = 0, \dots, N-1$$
(6.8)

 $^{^{10}{\}rm Note}$ that there exist different normalization conventions, not all of which directly yield orthogonal transforms.

for some $d \ll N$. For truncation error estimates and convergence properties of Fourier cosine series see [Boy01].

The DCT transform can also be written in matrix form

$$(x_0,\ldots,x_{N-1})\Psi = (\hat{x}_0,\ldots,\hat{x}_{N-1}),$$

with an orthonormal $N \times N$ matrix Ψ having entries

$$\Psi_{ij} = \begin{cases} \frac{1}{\sqrt{N}} & , i = 0, \\ \sqrt{\frac{2}{N}} \cos\left(\frac{\pi}{N}i(j+1/2)\right) & , i > 0, \end{cases}$$

for i, j = 0, ..., N-1. From this it is easy to see that for $x, y \in \mathbb{R}^N$, the DCT transform satisfies a version of Plancherel's theorem

$$\sum_{n=0}^{N-1} x_n y_n = \sum_{n=0}^{N-1} \hat{x}_n \hat{y}_n,$$

since

$$x \, y' = x \, \Psi \, \Psi' \, y' = \hat{x} \, \hat{y}'. \tag{6.9}$$

Now let $A \in \mathbb{R}^{N \times N}$ be a symmetric positive definite matrix with square-root decomposition

$$A = BB'$$

for some matrix¹¹ $B \in \mathbb{R}^{N \times N}$. If we combine (6.8) and (6.9) from above and assume that the rows of B are sufficiently smooth, then we should have as a good rank-d approximation to A

$$A \approx (B\Psi_d)(B\Psi_d)',$$

where Ψ_d denotes, as before, the $N \times d$ -matrix consisting only of the first d columns of the DCT matrix Ψ . The above result can be adapted to forward-rate correlation matrices and their corresponding Cholesky decompositions. To this end let $\rho \in \mathbb{R}^{N \times N}$ be a forward-rate correlation matrix with Cholesky decomposition

$$\rho = LL'.$$

If we ignore for a moment the first column entries of L, then, provided ρ is sufficiently smooth, also the rows of L will be smooth, at least up to the main diagonal beyond which all entries of L are clearly zero. The first column of L requires some special treatment, as it is always equal to the first column of ρ (see Section 5.3). In particular, its first entry L_{00} is always equal to one. To account for these properties we do not multiply L by Ψ_d directly but rather by the first d columns of the following matrix

$$\tilde{\Psi} = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & & & \\ \vdots & \Psi^{(N-1)} \\ 0 & & & \end{pmatrix},$$
(6.10)

¹¹Note that B is only unique up to a unitary transformation.

where $\Psi^{(N-1)}$ is an $N-1 \times N-1$ DCT matrix. After a normalization, which is again necessary in order for the final matrix to have unit diagonal, we obtain the following rank-d approximation to ρ

$$\hat{\rho}^d = \bar{Z}\bar{Z}',\tag{6.11}$$

$$\bar{Z}_{i.} = \frac{Z_{i.}}{\|Z_{i.}\|_2}, \ i = 0, \dots, N-1,$$
(6.12)

$$Z = L\tilde{\Psi}_d. \tag{6.13}$$

Observe that if $\rho = \rho(\theta)$ is a parametric form from the SC-family or is derived via the construction principle presented in Section 5.3, then the corresponding Cholesky decompositions are either directly given in closed form or can be calculated very efficiently, see Appendix C.1. We may therefore write $L = L(\theta)$, and (6.11) - (6.13)(requiring, in effect, only some matrix multiplications) can be regarded as a simple method for deriving low-rank parameterizations $\hat{\rho}^d = \hat{\rho}^d(\theta)$ from given (positivedefinite) full-rank parameterizations.

Remark 6.2.1. Letting $d \to N$ we clearly have $\hat{\rho}^d \to \rho$.

Remark 6.2.2. Instead of interpreting the above approach in terms of the DCT and its interpolating properties, it can also be interpreted as expressing the Cholesky factor in a different basis. To see this, let the singular value decomposition of the Cholesky factor L be given by $L = U\Sigma V'$ for some orthonormal matrices $U, V \in \mathbb{R}^{N \times N}$ and a diagonal matrix $\Sigma \in \mathbb{R}^{N \times N}$. Further, let $\rho = LL' = Q\Lambda Q'$ be the eigenvalue decomposition of ρ with an orthonormal matrix $Q \in \mathbb{R}^{N \times N}$ and a diagonal matrix $\Lambda \in \mathbb{R}^{N \times N}$. Then, as is known, we have¹² U = Q and $\Sigma = \Lambda^{1/2}$. It is therefore easy to see that $Q\Lambda_d^{1/2} = LV_d$, which is the matrix square root used in the PCAbased approach. Comparing LV_d with (6.13), we can interpret the above approach as replacing the right-singular vectors of L by some "general purpose" basis vectors. Although the right-singular vectors are in some sense optimal, they are tailored towards the particular L and do not necessarily "work well" for other Cholesky factors. In contrast, as we will see in the next section, the usage of $\tilde{\Psi}_d$ will generally yield smooth and realistically shaped forward-rate correlation matrices for quite a wide range of Cholesky factors L.

Remark 6.2.3. Once the dimension N and the number of factors d has been fixed, the DCT matrix $\tilde{\Psi}_d$ remains always the same and could in principle be cached. In practice, the potential speed advantage, however, will most often be negligible.

6.3 Numerical Results

In Figures 6.1 and 6.2 we consider two scenarios for possible shapes of forward-rate correlation matrices, where in both cases, the full-rank matrices were generated by the 5P form. The matrix in Figure 6.1 corresponds to the historically fitted matrix from Figure 5.5, while the matrix in Figure 6.2 is implied from a calibration to market data,

¹²Provided we use the same ordering for the eigenvalues and the singular values.

6.3. Numerical Results

see Chapter 7 below. As can be seen from the figures, the DCT method produces quite realistic correlation matrices, relatively close to the initial full-rank matrices. We can therefore expect this implicit low-rank parameterization to exhibit the same dependence on the different parameters as the original full-rank parameterization.

We also show in Figures 6.1 and 6.2 the matrices obtained by applying the modified PCA method and the majorization method from [PG04], the latter of which yields the respective optimal¹³ low-rank approximations to the given full-rank matrices. We would like to emphasize again, however, that closely matching these optimal solutions is not our primary concern, and the plots are merely included for the reader's interest.



Figure 6.1: Rank reduction to 4 factors. From top, left to right: Full-Rank, DCT, Majorization, PCA.

In Table 6.1 we present the required computing times for rank-reducing a 60×60 and a 120×120 correlation matrix to 4 factors by using different methods. The parameters for the underlying 5P form are the same as those used in Figure 6.1. As we have already noted above, the Cholesky decomposition of the 5P form is available at practically no extra cost, and together with the DCT method we obtain virtually instantly the low-rank correlation matrix. For a 60×60 correlation matrix the DCT method is over 100 times faster than the modified PCA¹⁴ method, as can be seen

¹⁴For the eigenvalue decomposition we used a combination of Househoulder reduction to tridiagonal form followed by the QL algorithm with implicit shifts, see [PTVF01]. This is still one of the

 $^{^{13}}$ See Footnote 7 in Section 6.1.



Figure 6.2: Rank reduction to 7 factors. From top, left to right: Full-Rank, DCT, Majorization, PCA.

from Table 6.1. For N = 120, which would correspond to, say, a 30Y LMM with quarterly fixings, the difference is even more significant.

For parameterizations other than the Schoenmakers-Coffey forms or our 4P and 5P form, the Cholesky decomposition might not always be directly available. For this reason we also report in Table 6.1 the required computing times, if prior to calculating (6.11)–(6.13), the standard Cholesky algorithm must be performed. While both the Cholesky decomposition and the eigenvalue decomposition have computational complexity $O(N^3)$, the Cholesky decomposition is generally much faster. The combination of the Cholesky decomposition and the DCT method is therefore still a factor of almost 30 times faster than the modified PCA method.

To give the reader an idea of how using DCT rank-reduced parametric forms might affect calibration times: In Chapter 7 we perform calibrations of a 10 factor 30Y LMM with semiannual fixings and time-dependent correlations to real market data¹⁵. Although both the PCA-based and the DCT based calibration yield virtually identical calibration results (in terms of market fit), the first approach required 71.4s

most efficient ways of computing the eigenvalues and eigenvectors of real symmetric matrices of the considered size.

¹⁵As correlation structure we choose our 4P-form with time-dependent ρ_{∞} parameter. $\rho_{\infty}(t)$ is "parameterized" by using a grid of 5 knot points t = 0, 5, 10, 15, 20 years and interpolating linearly between them.

6.4. Conclusion

while the latter took only $31.1s^{16}$.

N	6	60	120		
	time (s)	speedup	time (s)	speedup	
PCA	10.56	_	64.89	-	
DCT+Chol	0.38	28	2.27	29	
DCT	0.09	112	0.38	173	
Majorization	0.39	_	1.97	_	

Table 6.1: First three rows: Computing times for 5,000 runs of the different rank reduction methods (reduction to 4 factors). The columns denoted 'speedup' give the speedup factors compared to the PCA method. For comparison we also give in the last row the computing times for one run of the majorization method.

6.4 Conclusion

In this chapter we have presented a new method for rank reducing forward-rate correlation parameterizations. The proposed approach is easy to implement, computationally inexpensive and together with the existing positive definite parameterizations implicitly generates a new family of flexible low-parametric forms of arbitrary rank. These low-parametric forms allow for a much more efficient calibration of LMMs than the currently common approach of using PCA rank-reduced parameterizations, which is computationally much more demanding.

¹⁶In both cases roughly 20,000 (low-rank) correlation matrices had to be calculated.

Chapter 7

Extracting Correlations from the Market

In the past, the volatility structures of LMMs were most often calibrated to market prices of caps and/or swaptions, while the Libor correlations ρ_{ij} were specified exogenously and would typically have originated from an empirical analysis in the style of that presented in Section 5.1.

However, as we have discussed earlier, it would generally be preferable to use correlations obtained from an implied calibration to market prices. In this way one would capture the prevailing market conditions and obtained a model, that prices exotic securities inline with the prices of products used as hedging instruments.

In previous chapters we have established a new framework for parameterizing and rank-reducing forward-rate correlations and presented an efficient method for pricing CMSSOs within the SV-LMM. In the following we will apply these methods to actually extract the correlation information implicitly contained in market prices of interest-rate options. Moreover, we will compare the performance of various correlation parameterizations and also provide some pricing applications.

7.1 Including CMS Spread Options

Before deciding on where and how to incorporate market prices of CMS spread options in the general calibration procedure, it is worth checking how the different model parameters affect the prices of such options. Suppose, that we have fixed two swap rates denoted¹ by S_1 and S_2 with fixing date T and strike price K. Then, the "parameters" that enter the CMSSO pricing formula (4.24) are (besides the initial swap rates $S(0) = (S_1(0), S_2(0))'$) the effective swap-rate volatilities $\bar{\sigma} = (\bar{\sigma}_1, \bar{\sigma}_2)'$, the swap-rate skews $\bar{\beta} = (\bar{\beta}_1, \bar{\beta}_2)'$, the effective swap-rate correlation $\bar{\rho}$ and the stochasticvolatility parameters κ and ξ^2 .

Figure 7.1 depicts the sensitivities of CMS spread option prices (expressed in terms of normal implied spread volatilities) with respect to the various parameters. Clearly, the main determinants of the implied spread volatility are the swap-rate volatilities $\bar{\sigma}$ and the swap rate correlation $\bar{\rho}$. The swap rate volatilities affect the level and skew of the spread smile, while the correlation almost solely affects its level. As with caplets and swaptions, the stochastic-volatility parameters κ and ξ control

¹Here we use the notation of Chapter 4.

²For the moment we neglect the convexity adjustments, i.e., the drift terms $\tilde{\mu}_i$.

the convexity of the volatility smile, which can be seen from the sign-changes of the respective sensitivities. Furthermore, it can be seen from Figure 7.1 that CMS spread options exhibit only a very modest sensitivity with respect to the swap rate skews, especially if compared to the sensitivities with respect to the swap rate volatilities and the correlation.

We therefore propose to include prices (or implied volatilities³) of CMS spread options only in Steps 1 and 3 of the main calibration procedure⁴ and simultaneously calibrate the volatility parameters Λ_{ij} and the correlation parameters of the chosen correlation structure. Step 2 remains untouched, and the model skew parameters are, as before, just calibrated to the market implied effective caplet and swaption skews.



Figure 7.1: Sensitivities of implied normal spread volatilities in basis points with respect to model parameters. Parameter values: $S(0) = (0.04, 0.04)', \ \bar{\sigma} = (0.15, 0.15)' \ (top), \ \bar{\sigma} = (0.15, 0.18)' \ (bottom), \ \tilde{\mu} = (0, 0)', \ \kappa = 0.06, \ \xi = 0.95, \ \bar{\rho} = 0.87, \ T = 5 \ (years).$

Even though Formula (4.24) allows a reasonably rapid computation of spread option prices, it is still computationally more expensive than the computation of the various effective swaption and caplet parameters/prices during the calibration. If the calibration procedure as introduced above is considered too slow, it is generally

 $^{^{3}}$ As with caps or swaptions it is usually preferable to calibrate to implied volatilities instead of prices for a more balanced error norm.

⁴See page 74.

possible to speed things up (without sacrificing too much accuracy) by calibrating to implied effective correlations instead of CMS spread option prices/volatilities. In the following we therefore outline the necessary steps for extracting market-implied effective correlations from CMS spread option prices.

If we denote the spread-option pricing formula (4.24) by CMSSO, then the problem is to solve the following equation for $\bar{\rho}_{imp}$

$$CMSSO(K, T; S(0), \tilde{\mu}, \bar{\sigma}, \beta, \xi, \kappa, \bar{\rho}_{imp}) = MarketPrice(K, T).$$
(7.1)

For this we need to fix all other parameters. For ξ and κ we may just plug in ξ^* and κ^* from the pre-calibration. Further, if we assume that the calibrated model will closely match the market-implied swap-rate volatilities $\bar{\sigma}$ and skews $\bar{\beta}$, we may just use these parameter values in (7.1), leaving us with the drift-related parameters $\tilde{\mu}_i = \bar{\beta}_i \bar{\mu}_i$, i = 1, 2 (cp. Equation (4.22)). The drift coefficients $\bar{\mu}_i$ can be written as (cp. Equation (4.5))

$$\bar{\mu}_i = \tilde{F}_i(L(0)) \,\overline{\Sigma} \, \tilde{G}_i(L(0))', \ i = 1, 2,$$

with

$$\overline{\Sigma} = \left(\frac{1}{T} \int_0^T \sigma_k(t) \sigma_l(t)' dt\right)_{k,l=0}^{N-1}$$
(7.2)

and where \tilde{F}_i and \tilde{G}_i are some \mathbb{R}^N -valued functions of the initial forward-rate curve $L(0) = (L_0(0), \ldots, L_{N-1}(0))'$. For the covariance matrix $\overline{\Sigma}$ we would need the instantaneous model forward-rate volatility functions $\sigma_k(t)$, which we are just about to calibrate. As a proxy we therefore propose to use

$$\overline{\Sigma} = \left(\bar{\sigma}_k^c \,\bar{\sigma}_l^c \,\rho_{kl}\right)_{k,l=0}^{N-1},\tag{7.3}$$

where $\bar{\sigma}_k^c$ are the market-implied effective caplet volatilities and $(\rho_{kl})_{k,l=0}^{N-1}$ is some simple "average" correlation matrix, e.g., $\rho_{kl} = 0.5(1 + \exp(-0.15|T_k - T_l|))$.

Having fixed all the parameters we can now solve (7.1) for $\bar{\rho}_{imp}$. After this has been done for all relevant CMS spread option maturities, we can perform the main calibration as described above, except that we now calibrate to implied correlations instead of CMS spread option prices/volatilities. The model implied correlations, which are compared in Steps 1 and 3 with the respective market implied ones, are calculated via Equation (4.7). After the first run of Steps 1 and 2 we suggest to update the drift terms $\tilde{\mu}$ by using the proper model covariance matrix (7.2) instead of the proxy (7.3) and then recalculate the market-implied correlations for the following runs.

Remark 7.1.1. The above method for calculating "market-implied" correlations is based on the assumption that the model implied swap-rate parameters (in particular the swap-rate volatilities) of the calibrated model will closely match the market implied ones. If this is not the case, then it may happen that the calibrated model matches the "market implied" correlations perfectly, even though the market and model prices for CMS spread options diverge. In such a case the calibration to prices/volatilities instead of correlations is to be preferred. Our experience shows, however, that the model is generally capable of reproducing the market swap-rate parameters quite well, and the calibration strategy via implied correlations therefore produces usually perfectly usable results (see the calibration examples below) with shorter computing times.

Remark 7.1.2. When using spread-option prices/volatilities, it is possible to calibrate several strikes per maturity. In contrast, we can only have one implied correlation per maturity, i.e., for each maturity we can calibrate the model only to one point of the respective "spread smile" if we use the route via implied correlations. We return to this point in Section 7.2.2.

To summarize, the calibration procedure can be carried out in the following steps:

Calibration to Caplets, Swaptions and CMSSOs

- **Step 0**: Perform the pre-calibration as described in Section 3.4.1. If necessary calculate market-implied correlations $\bar{\rho}_{imp}$;
- **Step 1:** Set the skew parameters B_{ij} all to the same value \overline{B} , chosen for example to be the average of all effective market skews $\overline{\beta}_x^*$. Calibrate the model volatilities $\lambda_i(t)$ and the parameters of the chosen model correlation structure to the $\overline{\sigma}_x^*$ and the CMSSO prices/volatilities/correlations;
- **Step 2:** Using model volatilities and correlations calculated in the previous step, the skews $\beta_i(t)$ are now calibrated to the $\bar{\beta}_x^*$;
- **Step 3**: (Only if implied correlations are used and if this is the first time that Step 3 is performed: Update the market-implied correlations $\bar{\rho}_{imp}$ as described above). Finally, the model volatilities $\lambda_i(t)$ and the correlation parameters are re-calibrated to the $\bar{\sigma}_x^*$ and the CMSSO prices/volatilities/ correlations, with the updated skews $\beta_i(t)$ from the previous step.

Steps 2 and 3 can be repeated several times, although often one cycle (Steps 1 to 3) is already enough to obtain a good fit (see examples below).

7.2 Calibration Examples

7.2.1 Data Description

Below we calibrate 30Y SV-LMMs based on 6M Libor rates, i.e., N = 60 and $\tau_k \approx 0.5$ (up to day-count-conventions), to market data as of 01/14/2008 and 04/26/2010. The following selection of market instruments is used:

• Implied Black volatilities of 6M EUR caplets with maturities 1, 2, 5, 7, 10, 15, 20, 25 and 30 years, stripped from market volatilities of caps/floors. For each maturity 9 relative strikes⁵ are used (0.00% being ATM): -2.00%, -1.00%, -0.5%, -0.25%, 0.00%, 0.25%, 0.5%, 1.00%, 2.00%.

⁵Provided the strike price is greater than zero.

- A grid of implied Black volatilities of EUR swaptions with maturities 1, 2, 5, 7, 10, 15, 20 years and tenors 2, 5, 7, 10, 15, 20, 25 years provided maturity + tenor ≤ 30 years (see also Tables D.7 and D.9 in Appendix D). For each maturity/tenor combination again nine relative strikes are used (see above).
- Implied normal volatilities of EUR 10Y–2Y CMS spread caplets with maturities 1, 2, 5, 7, 10, 15 and 20 years stripped from market prices of 10Y–2Y CMS spread caps/floors⁶. The quoted price matrix on 01/14/2008 contains prices of CMS spread caps/floors with strike prices -0.50%, -0.25%, 0.00%, 0.25%, 0.50%, 0.75%, 1.00%. On 04/26/2010 the strike prices of the quoted options are -0.25%, -0.10%, 0.00%, 0.25%, 0.50%, 0.75%, 1.00%. In both cases only the center strikes are used for the calibration, i.e., 0.25% for 01/14/2008 and 0.50% in case of 04/26/2010.



Figure 7.2: Initial 6M forward-rate curves.

The 2008 data set corresponds to a more normal market environment with a relatively flat forward-rate curve (see Figure 7.2) and with Black-implied (ATM) caplet and swaption volatilities in the 10% - 18% range. In contrast, the 2010 data set corresponds to a more excited market environment with caplet and swaption volatilities ranging from roughly 20% to 50% and with a very steep forward-rate curve.

7.2.2 Calibration Results

For the calibrations presented below we choose the following weights for the different instrument classes: 20% for caplets, 75% for swaptions and 5% for CMS spread options⁷. Within each group all instruments receive the same weighting. The penalty functions by means of which we can control the behavior of Λ and B, are chosen, such that we obtain reasonably smooth and time-homogeneous volatility and skew functions. For each data set the same penalty function settings are used for all the calibration tasks. For the parameterization of Λ and B we use a grid with $N_t = N_T =$ 8 and $t_i = 0, 3, 9, 13, 19, 29, 39, 49$ and $\mathcal{T}_j = 0, 3, 9, 19, 29, 39, 49, 59$. Concerning initial values, we set all entries of Λ and B to the averages of the effective market volatilities and effective market skews, respectively. For the correlation parameterizations we

⁶Reuters page ICAPSPREADS1.

⁷With these weightings we account for the fact that caps and swaptions are more liquid than CMS spread options and hence their prices are more reliable. The other reason is that we use 36 swaption smiles but only 9 caplet smiles and 7 CMS spread options.

use as initial values the respective best-fit parameters obtained from the historical fitting procedure⁸, see Tables D.5 and D.6 in Appendix D. We use 10 factors, which suffices in most cases to satisfactorily recover the full-rank correlation matrices and provides enough flexibility for spread-linked products. For the rank reductions we use in all cases the modified PCA method (6.2) - (6.4) to make the computing times comparable⁹. We always perform only one cycle of Steps 1 to 3 from the main calibration.

Market Data as of 01/14/2008

In Table 7.1 we present calibration results in terms of root mean square errors (RMSE) for the 2008 data set. Since the results for the two calibration approaches are relatively similar, no matter whether implied CMS spread volatilities or implied correlations are used, we will consider for the moment only the results for the latter approach and give some comments on the first one at the end of this section.

Correlation-	Caplets (Black-vol)		Swaptions (Black-vol)		CMSSOs (bp vol)				
Structure	ATM	Smile	ATM	Smile	K = 0.25%	Smile	Time (min)		
Calibration to CMSSO implied correlations									
SC2	0.24%	0.51%	0.26%	0.33%	3.0	3.8	0:30		
Reb3	0.26%	0.49%	0.27%	0.34%	2.2	3.2	0:25		
5P	0.24%	0.48%	0.25%	0.31%	0.7	2.4	0:23		
Calibration to CMSSO implied volatilities									
SC2	0.38%	0.44%	0.24%	0.28%	2.5	3.6	3:13		
Reb3	0.34%	0.50%	0.30%	0.36%	1.2	2.9	1:53		
5P	0.25%	0.46%	0.25%	0.32%	0.3	2.6	2:17		
Calibration to caplets and swaptions only									
5P (hist.)	0.26%	0.47%	0.25%	0.33%	2.8	4.2	0:02		

Table 7.1: Calibration errors (RMSE) for the 2008 data set. Errors for caplets and swaptions are quoted in terms of Black-implied volatilities. Errors for CMSSOs are quoted in terms of normal implied bp-volatilities. The last column gives the computing time (in minutes) required for the main calibration. The pre-calibration step requires 1:31 minutes and is identical for all scenarios.

The pre-calibration step, in which the effective market caplet/swaption parameters and the stochastic-volatility parameters are calibrated, requires 1:31 minutes. The corresponding parameter values can be found in Table D.7 in Appendix D. In the upper part of Table 7.1 we report the results for the correlation structures SC2, Reb3 and 5P, obtained from the main calibration if market-implied correlations are used, as described in Section 7.1. The columns denoted "ATM" give the RMSEs for at-the-money caplets and swaptions¹⁰. The columns denoted "Smile" give the overall RMSEs if all strike prices are considered¹¹. All three correlation parameterizations

130

 $^{^{8}}$ We note, that the specific choice of the initial values has no influence on the final result, only on computing times. We use the historical best-fit parameters only to make computing times comparable in some way.

⁹Note that in case of the Reb3 form the faster DCT rank reduction method cannot be applied since this parametric form is not guaranteed to always yield positive semi-definite matrices. ¹⁰9 caplets and 36 swaptions.

¹¹81 caplets, 324 swaptions and 49 CMS spread options.

are capable of fitting the caplet and swaption volatilities almost equally well and the computing times remain well under one minute.

Concerning the fit to CMSSOs the results are as was to be expected when regarding the number of correlation parameters: The 2-parametric SC-form provides the worst fit, whereas our 5-parametric form gives the best fit. This can also be seen from Figure 7.3 where we show implied CMS spread option volatilities for the options with strike K = 0.25%, which were used as calibration targets. The volatility term-structures implied by SC2 and Reb3 are generally too "flat". As an aside, we note that the obtained best-fit parameters for the Reb3-form result again in a non-positive definite matrix.



Figure 7.3: 2008 data set: 10Y-2Y CMS spread option implied volatilities for K = 0.25%.



Figure 7.4: 2008 data set: 10Y–2Y CMS spread option volatility surface, market (blue) and model (orange). Correlation parameterization 5P.

For comparison we also show in Figure 7.3 the model-implied spread volatilities

if historical correlations are used¹² and the model is just calibrated to caplets and swaptions (see also the last row of Table 7.1). In this case CMS spread options are generally undervalued, especially for longer maturities. A partial explanation for this can be given if we look at the calibrated correlation matrices shown in Figure 7.6. Comparing the historical matrix with the calibrated parametric forms we find that the historical correlations are generally higher, which in turn yields lower CMS spread option prices (the corresponding swap-rate volatilities being fixed)¹³. As a side note, we point out that (neglecting drift terms) 10Y–2Y CMS spread options carry only information about the correlation matrix entries highlighted in Figure 7.5, while the remaining entries must still be implicitly inferred from prices of swaptions.



Figure 7.5: Entries of the forward-rate correlation matrix that can theoretically be inferred from prices of 10Y–2Y CMS spread options.

Figure 7.4 shows the market-implied spread volatility surface together with the model implied one, if the correlation structure 5P is used. The model implied volatility smiles are especially at the long end too right-skewed. As we have noted above, the considered SV-LMM has no parameters that solely affect the spread smiles¹⁴. The main drivers of the skewness and the curvature of the spread volatility smiles are the respective swap-rate volatilities and the stochastic-volatility parameters, which are, however, already "fixed" by the market swaption smiles. The shapes of the spread smiles are therefore rather a consequence of the swaption-volatilities than being calibrated and we can only use the forward-rate correlations to calibrate the levels of the spread smiles.

For the examples above we have only used CMSSOs with strike prices 0.25% for the calibration. This is of course arbitrary and one could also use other strike prices that might be more important for a specific exotic product that is to be priced. If we calibrate to CMSSO prices/volatilities instead of implied correlations, then, in principle, we can also use several options per maturity and in this way try to achieve a better fit between the market and model spread smiles. This, however, is likely to work only if we sacrifice a little bit of the swaptions fit.

 $^{^{12}}$ More precisely, we use 5P fitted to the historical correlation matrix, see Figure 5.4.

¹³Recall, however, that the instantaneous forward-rate correlations ρ_{ij} are not the only factors that determine the terminal correlation of the swap rates underlying a spread option, see Footnote 4 on page 115.

¹⁴This would require some sort of multi-stochastic-volatility extended LMM, which would be far more difficult to handle.


Figure 7.6: Calibrated correlation matrices for the 2008 data set. From top, left to right: 5P fitted to Historical, SC2, Reb3, 5P.

Observe that for the examples above, we have used a relatively large number of calibration instruments. In order to obtain better fits (not only to CMSSOs), one can of course reduce the number of calibration instruments and judiciously choose only a small subset of caplets and swaptions, which might be of particular importance for pricing and hedging a certain exotic security.

Calibration errors can also be reduced by putting less weight on the various penalty functions and by allowing less regular parameter functions. For the above calibration tasks, however, we aimed to obtain rather smooth and time-homogeneous parameter functions, as can be seen from Figure 7.7, where we plot the calibrated time-dependent forward-rate volatilities and skews for the 5P form. The corresponding functions for the other two correlation parameterizations look very similar, but are mostly a little bit less "constant", partly probably as a compensation for the less flexible correlation structures.

Lastly, we note that all of the above findings also hold if we calibrate to CMS spread volatilities instead of implied correlations, and the results are pretty much the same as can be seen from Table 7.1. The fitting errors for CMS spread options are obviously a little bit smaller. In case of the SC2 and Reb3 structures, however, the better fit to CMS spread options comes with slightly larger caplet errors and with less regular volatility functions. Forcing the volatility functions to be similarly smooth as before, we obtain practically the same fitting errors as those given in the first two rows of Table 7.1. In case of the 5P form, the caplet and swaptions errors are basically the same (and the volatility functions equally smooth) as before, only



Figure 7.7: 2008 data set: Calibrated time-dependent forward-rate volatilities $\sigma_i(t)$ (top) and skews $\beta_i(t)$ (bottom) for i = 4, 9, 14, 19, 29, 39, 49, 59. Correlation parameterization 5P.

that the CMS spread errors are marginally smaller. We note that the RMSE of 0.3 bp for the 0.25% strike prices is already within the accuracy of our approximation formula for CMSSO prices. For the calibrated parameter values a fine-stepped Monte Carlo simulation yields that the maximum approximation error (in terms of implied spread volatility) is 0.6 bp with a RMSE of 0.3 bp.

As we have noted in Section 7.1, the calibration approach via implied correlations relies on the fact that the market effective swaption parameters will be closely matched by the calibrated model. If this is the case, as can be seen above, then the two approaches yield almost identical results and the calibration via implied correlations is generally to be preferred, due to much shorter computing times.

Market Data as of 04/26/2010

The calibration results for the 2010 data set, which corresponds to a rather excited market environment, are given in Table 7.2. Fitting simultaneously all the calibration instruments is a more challenging task as with the 2008 data set, and overall the root mean square errors are consequently somewhat larger.

The SC2 and Reb3 forms yield again worse results than the 5P form. Observe that in case of a calibration to implied correlations the swaption errors are more or less identical for all three correlation parameterizations, while the caplet fitting errors are slightly larger for SC2 and Reb3. Reducing the weights for CMSSOs, such that we obtained similar caplet and swaption fitting results as with the 5P form would result in CMS spread errors of 5-6 bp for K = 0.5% and 7-8 bp for the smiles.

Similar results are obtained when calibrating to implied spread volatilities instead of implied correlations. Fitting errors are generally a little bit smaller, although in

Correlation-	Caplets ((Black-vol)	Swaption	s (Black-vol)	CMSSOs (l	op vol)			
Structure	ATM	Smile	ATM	Smile	K = 0.50%	Smile	Time (min)		
Calibration to CMSSO implied correlations									
SC2	0.78%	2.76%	0.44%	0.63%	3.8	6.8	0:20		
Reb3	0.79%	2.80%	0.44%	0.63%	4.0	6.9	0:43		
5P	0.65%	3.05%	0.42%	0.61%	1.9	5.5	1:02		
Calibration to CMSSO implied volatilities									
SC2	0.66%	2.65%	0.53%	0.72%	3.4	6.5	4:53		
Reb3	0.74%	2.59%	0.49%	0.67%	2.5	6.0	6:24		
5P	0.67%	2.92%	0.39%	0.58%	1.7	5.3	5:54		
Calibration to caplets and swaptions only									
5P (hist.)	0.41%	2.37%	0.44%	0.60%	5.7	7.8	0:04		

Table 7.2: Calibration errors (RMSE) for the 2010 data set. Errors for caplets and swaptions are quoted in terms of Black-implied volatilities. Errors for CMS spread options are quoted in terms of normal implied bp-volatilities. The last column gives the computing time (in minutes) required for the main calibration. The pre-calibration step requires 2:09 minutes and is identical for all scenarios.

case of SC2 and Reb3 this comes again with slightly less regular volatility functions.

As a side note, observe that the computing times in Table 7.2 are for calibrations from a "cold start", with not particularly well chosen initial values. We would expect to see significantly reduced computing times if we took as initial values, for instance, the parameters from a previous calibration.

As can be seen from Figure 7.8, the overall implied correlation level ist substantially lower in comparison to the 2008 data set, partly, probably, as a reflection of the increased market uncertainty. Also the historical correlations are somewhat lower now. Nevertheless, the corresponding implied spread volatilities are, as before, too low for options with long maturities, as depicted in Figure 7.9. Notice, that this does not change too much if we use only 2 years worth of data (instead of 4 years) for the historical estimation, in order to base the correlations only on more recent data.

Quite interesting is the fact that the requirement (B2) from Section 5.1.2, i.e., "same-distance rates are more decorrelated at the front-end than at the back-end", can hardly be observed from the implied correlation matrices in Figure 7.8. The subdiagonals in case of the $SC2^{15}$ and the Reb3 form are only very marginally upward sloping, and in case of the 5P form are even slightly decreasing at the back end. Although intuitively this requirement makes sense and can also be observed from the historical correlation matrices, it seems that it is not supported by the market data. The market premia for long-term CMS spread options are just too high for being compatible with pronounced upward sloping sub-diagonals and the resulting high correlations at the back end of the correlation matrix. It is therefore questionable whether one should really impose (B2) on correlation structures, if they are to be used for implied calibrations.

In Figure 7.10 we show the spread volatility surface corresponding to using the 5P form. The smiles for long maturities are again too right skewed, although the

¹⁵Recall that requirement (B2) is directly built into correlation matrices from the Schoenmakers-Coffey family.



Figure 7.8: Calibrated correlation matrices for the 2010 data set. From top, left to right: 5P fitted to Historical, SC2, Reb3, 5P.

general volatility levels and the short-maturity smiles are matched reasonably well.

Finally, in Figure 7.11 we show the calibrated time-dependent forward-rate volatilities and skews for the 5P-parameterization (the corresponding functions for the other two correlation parameterization look more or less the same). This figure shows that the market obviously anticipates a decrease in volatility, with instantaneous forwardrate volatilities coming down to more normal levels in the 10% range.

7.3 Pricing Applications

In the following we will briefly investigate how different correlation calibrations may affect prices of CMS spread-linked products. More precisely, we use SV-LMMs calibrated to the 2008 market data set and compute prices of the following products¹⁶:

• A 10Y-2Y CMS spread ratchet cap with a tenor of 20 years, first fixing in 6 months and initial strike set to ATM, i.e., the *i*-th cash flow (fixed at T_i and paid at T_{i+1}) is given by

$$C_i = \tau_i (S_{i,i+20}(T_i) - S_{i,i+4} - K_i)^+, \ K_i = S_{i-1,i-1+20}(T_{i-1}) - S_{i-1,i-1+4},$$

with $K_0 = 0.0044$.

¹⁶All products are based on a semi-annually tenor structure.



Figure 7.9: 2010 data set: 10Y–2Y CMS spread option implied volatilities for K = 0.50%.



Figure 7.10: 2010 data set: 10Y–2Y CMS spread option volatility surface, market (blue) and model (orange). Correlation parameterization 5P.

• A 10Y-2Y CMS spread sticky cap with a tenor of 20 years, first fixing in 6 months, initial strike set to ATM and a spread of 15bp, i.e., the *i*-th cash flow (fixed at T_i and paid at T_{i+1}) is given by

$$C_{i} = \tau_{i}(S_{i,i+20}(T_{i}) - S_{i,i+4} - K_{i})^{+},$$

$$K_{i} = \min\{S_{i-1,i-1+20}(T_{i-1}) - S_{i-1,i-1+4}(T_{i-1}), K_{i-1}\} + s,$$

with $K_0 = 0.0044$ and s = 0.0015.

• A callable 10Y-2Y CMS spread note (or steepener note) with maturity in 20 years and principal value N = 100. During the first 3 years, the owner of the note receives fixed coupons of 8% (payed semi-annually in-arrears). Thereafter, the note pays the 10Y-2Y CMS spread floored at 0% multiplied by a gearing factor of 14. The



Figure 7.11: 2010 data set: Calibrated time-dependent forward-rate volatilities $\sigma_i(t)$ (top) and skews $\beta_i(t)$ (bottom) for i = 4, 9, 14, 19, 29, 39, 49, 59. Correlation parameterization 5P.

coupons are hence given by

$$C_{i} = \begin{cases} N\tau_{i}K, & i = 0, \dots, 5, \\ N\tau_{i}g \cdot (S_{i,i+20}(T_{i}) - S_{i,i+4}(T_{i}))^{+}, & i = 6, \dots, 39, \end{cases}$$

with K = 0.08 and g = 14. After an initial lock-out period of 3 years, the bond can be called by the issuer. If the bond is called at time T_i , then at time T_{i+1} the last coupon C_i as well as the principal is paid and no further cash flows occur.

In Table 7.3 we present Monte Carlo prices of the above products computed by using differently calibrated LMMs. In all cases, the exercise strategy for the callable note was calculated by using the Longstaff-Schwartz [LS01] algorithm. Besides LMMs based on the implied and the historical (time-homogeneous) 5P correlation structures from the last section, we also used a SV-LMM with a time-dependent 4P correlation structure¹⁷, where the overall correlation level parameter ρ_{∞} was allowed to depend on calendar time (all other parameters were kept constant). The level function $\rho_{\infty}(t)$ was "parameterized" by using a grid of 5 knot points t = 0, 5, 10, 15, 20 years and interpolating linearly between them. The calibrated parameter values are given in Table D.11 in Appendix D. As we have already anticipated in Section 6.3, using the DCT method for rank-reducing the correlation matrices significantly reduces the computational burden of calibrating a model with a time-dependent correlation structure. The PCA-based calibration requires 71.4s while the DCT-based approach takes only 31.1s.

As can be seen from Table 7.3, the LMM with the historical correlation matrix yields prices that are significantly lower than the corresponding values pro-

¹⁷Implicitly calibrated to the same market data as described in the last section.

duced by the LMMs with the implied correlation structures. Compared to the timehomogeneous implied 5P version, prices are generally off by 5–20%. Only the prices of the callable note seem to be comparable. The reason is that the callability feature partly cancels the differently priced CMS spread caplets implicitly contained in the note. Considering a non-callable version of the steepener note, we find again that the "historical" model produces significantly lower prices. The values of the callability options¹⁸ even differ by more than 40%.

For the given market data set, the model with the time-homogeneous 5P version and the model with the time-dependent 4P form both yield virtually the same market fit in terms of calibration errors with respect to caplets, swaptions and CMSSOs. Hence, these models will generally produce almost the same prices for Europeantype options, which can be readily observed by comparing the values of the CMSS coupons¹⁹ contained in the steepener notes. However, from the other prices in Table 7.3 we can see that the time-dependent correlations may have an impact on prices of more exotic products. While the price of the callable note seems not to be overly sensitive with respect to the time-dependent behavior of correlations, the situation is different with the exotic CMSS caps. The model based on the time-dependent correlation structure implies significantly higher prices than the model with the timehomogeneous correlation parameterization. Observe that the cash flows of sticky and ratchet CMSS caps are not only sensitive with respect to the decorrelation among the swap rates defining the spread, but they also depend on the decorrelation among contiguous spread fixings, as their strikes at a given time depend on the previous spread realization. The time-dependent correlation structure, with its monotonically decreasing correlation level parameter $\rho_{\infty}(t)$, obviously leads to a stronger decorrelation effect and hence to higher prices.

Correlation-	CMS	S caps	CMSS note						
structure	sticky	ratchet	callable	${\rm non-callable}$	coupons	call. option			
5P hist.	193.7	122.1	98.17	113.54	52.26	15.37			
	(0.7)	(0.2)	(0.05)	(0.14)					
5P	237.4	127.8	99.70	126.22	64.94	26.52			
	(0.9)	(0.2)	(0.04)	(0.16)					
4P time-dep.	292.2	146.0	100.20	125.74	64.46	25.54			
_	(1.1)	(0.3)	(0.04)	(0.16)					

Table 7.3: Monte Carlo prices of different exotic products. 130,000 paths, 16,000 training paths for the Longstaff-Schwartz exercise strategy. Prices of caps in basis points. Values in parentheses denote Monte Carlo errors. The last four columns give the prices of the callable CMS spread note, the non-callable CMS spread note, the plain CMS spread coupons contained in the non-callable note and the value of the callability option.

From the market fits to the (European-type) calibration instruments alone, we cannot judge which model specification yields the "correct" exotics prices. For the considered market data set it seems not necessary to allow for time-dependent correlation parameters, provided that we use a sufficiently flexible correlation parameterization (such as the 5P form). Note that even with a SC2 form we can obtain similar

 $^{^{18}\}mbox{I.e.}$ the price difference between the callable and the non-callable note.

¹⁹Which are in effect just multiples of strike zero European CMSS caplets.

market fits as with the time-homogeneous 5P form, if we allow the two parameters of the SC2 form to depend on time. In this case, however, introducing time-dependent parameters is obviously only necessary to increase the "degrees of freedom" of the less flexible SC2 form and not because the market data really implies a time-dependent behavior.

As we have noted earlier, if the market fit is reasonably accurate and unless we have a particular view on the future evolution of the market, then in most cases it is advisable to use a time-homogeneous parameterization in order to obtain robust calibration results and to avoid overfitting. Nevertheless, in some market environments it might become necessary to use time-dependent correlation parameterizations. In this case, as we have demonstrated, models with time-homogeneous and time-dependent correlation structures may yield significantly different prices for CMS spread-linked exotics, even if they coincide on the prices of the calibration instruments.

7.4 Conclusion

In this chapter we have discussed two different methods for calibrating the SV-LMM to caplets, swaptions and CMSSOs and calibrated models to real market data. It has been demonstrated that with our new 5-parametric form better market fits can be achieved than with the other analyzed correlation structures. It also turned out, that pronounced upward sloping sub-diagonals of correlation matrices, as they can often be observed from empirically estimated matrices, and as they are sometimes directly built into correlation parameterizations, seem not to be compatible with market prices of long-term CMSSOs. Implied long-term correlations are usually substantially lower than the corresponding historically estimated correlations. If this was a persistent feature and if suitable correlation-sensitive products were liquidly traded, then it should be possible to run some kind of statistical arbitrage strategies. Lastly, we have demonstrated that time-dependent correlation structures may have a perceptible impact on prices of certain exotic products.

Appendices

A Laplace Transform of $\overline{V}(T)$

A.1 Singularities

 $\text{Lemma A.1. } \gamma + \kappa + (\gamma - \kappa)e^{-\gamma T} \not\in (-\infty, 0) \text{ for } z \in \{x + iy \in \mathbb{C} | x \in \mathbb{R}, y \ge 0\} \setminus (-\infty, 0).$

Proof. First recall that for $z \in H := \{x + iy \in \mathbb{C} | x \in \mathbb{R}, y \ge 0\} \setminus (-\infty, 0)$ we have that

$$\gamma = \sqrt{\kappa^2 + 2z\xi^2} \in \{x + iy \in \mathbb{C} | x, y \ge 0\} \setminus [0, i\infty)$$

(provided we take the principal square root). Now assume that $\gamma + \kappa + (\gamma - \kappa)e^{-\gamma T} \in (-\infty, 0)$ for some $u \in H$.

$$\Rightarrow \qquad \operatorname{Im}(\gamma + \kappa) = -\operatorname{Im}\left((\gamma - \kappa)e^{-\gamma T}\right)$$
$$\Rightarrow \qquad (\operatorname{Im}(\gamma))^2 = \left(\operatorname{Im}((\gamma - \kappa)e^{-\gamma T})\right)^2$$

We also must have that

$$\operatorname{Re}\left((\gamma-\kappa)e^{-\gamma T}\right) < -\operatorname{Re}(\gamma+\kappa) < 0$$

$$\Rightarrow \left(\operatorname{Re}((\gamma-\kappa)e^{-\gamma T})\right)^2 > \left(\operatorname{Re}(\gamma+\kappa)\right)^2.$$

But this implies

$$\begin{aligned} |\gamma + \kappa|^2 &< |(\gamma - \kappa)e^{-\gamma T}|^2 \\ &= |\gamma - \kappa|^2 e^{-2\operatorname{Re}(\gamma)T} \\ &= ((\operatorname{Re}(\gamma) - \kappa)^2 + \operatorname{Im}(\gamma)^2) e^{-2\operatorname{Re}(\gamma)T} \\ &\stackrel{\operatorname{Re}(\gamma)T > 0}{<} ((\operatorname{Re}(\gamma) - \kappa)^2 + \operatorname{Im}(\gamma)^2) \\ &\stackrel{\operatorname{Re}(\gamma), \kappa > 0}{<} ((\operatorname{Re}(\gamma) + \kappa)^2 + \operatorname{Im}(\gamma)^2) \\ &= |\gamma + \kappa|^2, \end{aligned}$$

which leaves us with a contradiction. Thus, $\gamma + \kappa + (\gamma - \kappa)e^{-\gamma T} \notin (-\infty, 0)$.

Lemma A.2. All singularities of \hat{f} lie on the negative real line.

Proof. First observe that singularities may occur either if

$$2\gamma e^{-\gamma T/2} = 0$$

in the logarithm or if

$$\gamma + \kappa + (\gamma - \kappa)e^{-\gamma T} = 0 \tag{A.4}$$

in the denominator of the fractions. However,

$$2\gamma e^{-\gamma T/2} = 0 \Leftrightarrow z = -\frac{\kappa^2}{2\xi^2},$$

which is in $(-\infty, 0)$, since κ and ξ are positive real constants.

Now consider the second case. Because of symmetry reasons, we just need to consider \hat{f} in the upper complex half. Hence, assume that (A.4) holds for some $z \in H := \{x + iy | a \in \mathbb{R}, b \ge 0\} \setminus (-\infty, 0]$. Since we always take the principal value of the square root, γ lies in the upper right quadrant of the complex plane (without the positive imaginary axis), i.e., $\operatorname{Re}(\gamma) > 0$ and $\operatorname{Im}(\gamma) \ge 0$. But then (A.4) implies

$$\begin{split} \gamma + \kappa &= -(\gamma - \kappa)e^{-\gamma T} \\ \Rightarrow & |\gamma + \kappa| = |\gamma - \kappa|e^{-\operatorname{Re}(\gamma)T} \\ \Rightarrow & \sqrt{\operatorname{Im}(\gamma)^2 + (\operatorname{Re}(\gamma) + \kappa)^2} = \sqrt{\operatorname{Im}(\gamma)^2 + (\operatorname{Re}(\gamma) - \kappa)^2}e^{-\operatorname{Re}(\gamma)T} \\ \Rightarrow & \operatorname{Im}(\gamma)^2 + (\operatorname{Re}(\gamma) + \kappa)^2 = \left(\operatorname{Im}(\gamma)^2 + (\operatorname{Re}(\gamma) - \kappa)^2\right)e^{-2\operatorname{Re}(\gamma)T} \\ \Rightarrow & 0 \leq \underbrace{\operatorname{Im}(\gamma)^2}_{\geq 0} \underbrace{\left(1 - e^{-2\operatorname{Re}(\gamma)T}\right)}_{>0} = \underbrace{\left(\operatorname{Re}(\gamma) - \kappa\right)^2}_{\operatorname{since } \operatorname{Re}(\gamma), \kappa > 0} \underbrace{e^{-2\operatorname{Re}(\gamma)T}}_{<1} - (\operatorname{Re}(\gamma) + \kappa)^2 < 0, \end{split}$$

which leaves us with a contradiction. Thus,

$$\gamma + \kappa + (\gamma - \kappa)e^{-\gamma T} \neq 0 \text{ for } z \in H$$

and all singularities of \hat{f} lie on the negative real line.

A.2 Derivation of the "Continuified" Laplace Transform

The Laplace transform of the integrated variance is given by

$$\begin{split} \hat{f}(z) &= \exp\left\{\frac{2\kappa}{\xi^2}\log\left(\frac{2\gamma e^{-\gamma T/2}}{\gamma + \kappa + (\gamma - \kappa)e^{-\gamma T}}\right) + \right.\\ &\left. + \frac{\kappa^2 T}{\xi^2} + \left(\frac{2\gamma e^{-\gamma T}}{\gamma + \kappa + (\gamma - \kappa)e^{-\gamma T}} - 1\right)\frac{V(0)(\gamma - \kappa)}{\xi^2}\right\}, \end{split}$$

where $\gamma = \sqrt{\kappa^2 + 2z\xi^2}$. Now recall that the complex logarithm and the square root are multi-valued functions. Although one can choose a unique principal value for such functions (as is usually done in software packages), these functions are not continuous in the entire complex plane.

A. Laplace Transform of $\overline{V}(T)$

In case of the square root γ we may readily restrict it to its principal value (i.e., the square root with non-negative real part) as \hat{f} is an even function in γ . In case of the logarithm, however, we need to do some adjustment in order to account for the branch cut of the principal value along the negative real axis. First, we may write

$$\log\left(\frac{2\gamma e^{-\gamma T/2}}{\gamma+\kappa+(\gamma-\kappa)e^{-\gamma T}}\right) = \log\left(2\gamma e^{-\gamma T/2}\right) - \log\left(\gamma+\kappa+(\gamma-\kappa)e^{-\gamma T}\right),$$

where we must keep in mind that this identity holds in the complex case only modulo $2\pi i$ (we will deal with determining the right branch later). The second term on the right hand side is continuous in the upper complex plane²⁰, as the argument of the logarithm never crosses the negative real line (see Lemma A.1). However, the argument of the first logarithm on the right hand side may cross the negative real axis several times along continuous integration paths in the upper complex plane. So in order to obtain a continuous function we have to make some branch cut corrections.

Consider the principal value of the complex square root $\gamma = \sqrt{\kappa^2 + 2z\xi^2}$. In our case we have that

$$z \in H := \{x + iy \in \mathbb{C} | x \in \mathbb{R}, y \ge 0\} \setminus (-\infty, 0),$$

which can be reinterpreted in terms of polar coordinates as

$$z = r_z e^{i\varphi_z}, \ r_z \ge 0, \ \varphi_z \in [0,\pi).$$

Even if we multiply z by a positive real number $(2\xi^2)$ and add a positive real number (κ^2) , the argument φ_z will still be in $[0, \pi)$. If we use the principal value of a complex square root, we have for a complex number $y = re^{i\varphi}$, $r \ge 0$, $\varphi \in (-\pi, \pi]$

$$\sqrt{y} = \sqrt{r}e^{i\varphi/2}.$$

In our case, we thus have

$$\gamma = \sqrt{\kappa^2 + 2z\xi^2} = r_\gamma e^{i\varphi_\gamma}$$
, with $r_\gamma \ge 0$ and $\varphi_\gamma \in [0, \pi/2)$.

In particular, γ is continuous for $z\in H.$ In the following we shall use the representations

$$\gamma = r_{\gamma} e^{i\varphi}, \ r_{\gamma} \ge 0, \ \varphi_{\gamma} \in [0, \pi/2),$$

or
$$\gamma = a_{\gamma} + ib_{\gamma}, \ (a, b) \in \{(a, b) \in \mathbb{R}^2 | a, b \ge 0\} \setminus (i0, i\infty)$$

as is convenient. We therefore have for the argument of the logarithm

$$2\gamma e^{-\gamma T/2} = 2r_{\gamma} e^{i\varphi_{\gamma}} e^{-(a_{\gamma}+ib_{\gamma})T/2}$$
$$= 2r_{\gamma} e^{-a_{\gamma}T/2} e^{i(\varphi_{\gamma}-b_{\gamma}T/2)}.$$

²⁰In the inverse Laplace integration step we only need to evaluate the Laplace transform in the upper complex (half-)plane.

Since the logarithm is discontinuous along the negative real line, we have to pay special attention to points, where

$$\varphi_{\gamma} - b_{\gamma} \frac{T}{2} = k\pi, \ k \in \{m \in \mathbb{Z} \mid m = 2n+1, n \in \mathbb{Z}\}.$$
(A.5)

Now we have

$$b_{\gamma} = \operatorname{Im}(\gamma) = r_{\gamma} \sin(\varphi_{\gamma}).$$

Hence, (A.5) can be rewritten as

$$\varphi_{\gamma} - \frac{T}{2}r_{\gamma}\sin(\varphi_{\gamma}) = k\pi, \ k \in \{m \in \mathbb{Z} \mid m = 2n+1, n \in \mathbb{Z}\}.$$

In order to make the logarithm continuous we first rewrite its argument as

$$2\gamma e^{-\gamma T/2} = r^* e^{i(\varphi^* + 2\pi n^*)}$$

with

$$\begin{aligned} r^* &= \left| 2\gamma e^{-\gamma T/2} \right|, \\ \varphi^* &= \operatorname{Arg} \left(2\gamma e^{-\gamma T/2} \right), \\ n^* &= \left\lfloor \frac{\varphi_{\gamma} - \frac{T}{2} r_{\gamma} \sin(\varphi_{\gamma}) + \pi}{2\pi} \right\rfloor \\ &= \left\lfloor \frac{\operatorname{Arg}(\gamma) - \frac{T}{2} |\gamma| \sin(\operatorname{Arg}(\gamma)) + \pi}{2\pi} \right\rfloor, \end{aligned}$$

where $\lfloor . \rfloor$ denotes rounding to the nearest smaller integer and Arg(.) denotes the principal argument. Now, applying the logarithm yields

$$\log (2\gamma e^{-\gamma T/2}) = \log(r^*) + i(\varphi^* + 2\pi n^*).$$

This was the crucial step for "continuifying" the Laplace transform: note that the first part on the right hand side $(\log(r^*) + i\varphi^*)$ corresponds to the principal value of the logarithm, which would be returned by most software packages. The last term $2\pi n^*$ is the branch cut correction.

With this formulation of the logarithm, the Laplace transform is clearly continuous. Finally, we have to make sure that everything is "anchored" on the right branch, since

$$\log\left(\frac{2\gamma e^{-\gamma T/2}}{\gamma+\kappa+(\gamma-\kappa)e^{-\gamma T}}\right) = \log(r^*) + i(\varphi^*+2\pi n^*) - \log\left(\gamma+\kappa+(\gamma-\kappa)e^{-\gamma T}\right)$$

only holds modulo $2\pi i$. However, it is easy to verify (recall that the Laplace transform evaluated at a positive real number must be real valued) that

$$\operatorname{Log}(r^*) + i(\varphi^* + 2\pi n^*) - \operatorname{Log}\left(\gamma + \kappa + (\gamma - \kappa)e^{-\gamma T}\right)$$

with Log(.) denoting the principal branch already yields the right values (without adding a multiple of $\pm 2\pi i$).

A. Laplace Transform of $\overline{V}(T)$

To summarize, a "continuified" version of the Laplace transform is given by

$$\begin{split} \hat{f}(z) &= \exp\left\{\frac{2\kappa}{\xi^2}A(z) + B(z)\right\},\\ A(z) &= \operatorname{Log}\left|2\gamma e^{-\gamma T/2}\right| + i\left(\operatorname{Arg}(2\gamma e^{-\gamma T/2}) + 2\pi n\right) - \operatorname{Log}\left(\gamma + \kappa + (\gamma - \kappa)e^{-\gamma T}\right),\\ B(z) &= \frac{\kappa^2 T}{\xi^2} + \left(\frac{2\gamma e^{-\gamma T}}{\gamma + \kappa + (\gamma - \kappa)e^{-\gamma T}} - 1\right)\frac{V(0)(\gamma - \kappa)}{\xi^2},\\ n &= \left\lfloor\frac{\operatorname{Arg}(\gamma) - \frac{T}{2}|\gamma|\sin(\operatorname{Arg}(\gamma)) + \pi}{2\pi}\right\rfloor,\\ \gamma &= \sqrt{\kappa^2 + 2z\xi^2}. \end{split}$$

A.3 Asymptotic Behavior

In this section we consider the asymptotic behavior of

$$e^{s(u)x}\hat{f}(s(u)) = \exp\left\{s(u)x + \frac{2\kappa}{\xi^2}A(s(u)) + B(s(u))\right\}$$

with s(u) = 1/x + u(bi - 1/x), x, b > 0. Actually, due to Euler's formula it suffices to just consider the real part of the exponent on the right hand side. For the first part of the exponent we have

$$\operatorname{Re}(s(u)x) = 1 - u.$$

Next, we have

$$\operatorname{Re}(A(s(u))) = \log(2|\gamma|e^{-\operatorname{Re}(\gamma)T/2}) - \operatorname{Re}(\log(\gamma + \kappa + (\gamma - \kappa)e^{-\gamma T}))$$
$$= \log 2 + \log(|\gamma|) - \frac{T}{2}\operatorname{Re}(\gamma) - \operatorname{Re}(\log(\gamma + \kappa + (\gamma - \kappa)e^{-\gamma T}))$$
$$\simeq -c\sqrt{u}, \ (u \to \infty)$$

for some c > 0. Similarly,

$$\operatorname{Re}(A(s(u))) \simeq -d\sqrt{u}, \ (u \to \infty)$$

for some d > 0. In total we therefore have that

$$\operatorname{Re}(s(u)x + \frac{2\kappa}{\xi^2}A(s(u)) + B(s(u))) \simeq -u - M\sqrt{u}, \ (u \to \infty)$$

for some M > 0, which implies that the decay of $e^{s(u)x} \hat{f}(s(u))$ is asymptotically at least of exponential order.

A.4 Generalized Spread-Option Formula

Let X and Y be two jointly log-normal random variables

$$X = X_0 \exp\{\mu_x + v_x Z_1\}$$
$$Y = Y_0 \exp\{\mu_y + v_y Z_2\}$$

where

$$\begin{pmatrix} Z_1 \\ Z_2 \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \right),$$

with $\mu_x, \mu_y \in \mathbb{R}, v_x, v_y, X_0, Y_0 > 0$ and $\rho \in (-1, 1)$. Consider a "generalized spread payoff" of the form

$$H = (waX - wbY - wK)^+$$

with a, b > 0, $w = \pm 1$ and $K \in \mathbb{R}$. Using iterated expectations we have for the expectation of the payoff

$$\mathbb{E}[H] = \mathbb{E}\left[\mathbb{E}[(waX - wbY - wK)^+ | Z_2]\right].$$
(A.6)

Next, observe that the density of $Z_1|Z_2$ is given by

$$f_{Z_1|Z_2}(z_1, z_2) = \frac{1}{\sqrt{2\pi(1-\rho)^2}} \exp\left\{-\frac{(z_1-\rho z_2)^2}{2(1-\rho^2)}\right\},\,$$

i.e., $Z_1|Z_2$ is again normally distributed with mean ρZ_2 and variance $1 - \rho^2$. Using this fact and defining $\tilde{K}(z_2) := Y_0 \exp\{\mu_y + v_y z_2\} + K$, we obtain for the inner expectation in (A.6)

$$\mathbb{E}[(waX - wbY - wK)^+ | Z_2 = z_2] =$$

= $\int_{-\infty}^{\infty} \left(waX_0 e^{\mu_x + v_x z_1} - w\tilde{K}(z_2) \right)^+ f_{Z_1|Z_2}(z_1, z_2) dz_1$
=: $I(z_2).$

After tedious but straightforward calculations, one obtains for the value of the integral $I(z_2), z_2 \in \mathbb{R},$

$$I(z_{2}) = \begin{cases} \left(aX_{0}\exp\{\mu_{x} + \rho v_{x}z_{2} + \frac{1}{2}(1-\rho^{2})v_{x}^{2}\} - \tilde{K}(z_{2})\right)\mathbb{1}_{\{w=1\}} &, \ \tilde{K}(z_{2}) \leq 0, \\ waX_{0}\exp\{\mu_{x} + \rho v_{x}z_{2} + \frac{1}{2}(1-\rho^{2})v_{x}^{2}\} \\ &\times \Phi\left(w\frac{\log\left(\frac{aX_{0}}{\bar{K}(z_{2})}\right) + \mu_{x} + \rho v_{x}z_{2} + (1-\rho^{2})v_{x}^{2}}{v_{x}\sqrt{1-\rho^{2}}}\right) \\ &\quad -w\tilde{K}(z_{2})\Phi\left(w\frac{\log\left(\frac{aX_{0}}{\bar{K}(z_{2})}\right) + \mu_{x} + \rho v_{x}z_{2}}{v_{x}\sqrt{1-\rho^{2}}}\right) \\ &\quad -w\tilde{K}(z_{2})\Phi\left(w\frac{\log\left(\frac{aX_{0}}{\bar{K}(z_{2})}\right) + \mu_{x} + \rho v_{x}z_{2}}{v_{x}\sqrt{1-\rho^{2}}}\right) \end{cases}$$

where $\Phi(\cdot)$ denotes the standard normal CDF. Finally, the expectation of the payoff is given by

$$\mathbb{E}[H] = \int_{-\infty}^{+\infty} I(z_2) \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z_2^2} dz_2.$$

146

Index n	Fixing times T_n	$L_n(0)$	Eff. Libor vols $\bar{\lambda}_n$	Eff. Libor skews $\bar{\beta}_n$
0	0.5	3.011%	21.00%	40.00%
1	1	3.211%	25.35%	40.69%
2	1.5	3.374%	26.74%	41.18%
3	2	3.510%	26.73%	41.55%
4	2.5	3.629%	26.10%	41.86%
5	3	3.733%	25.26%	42.12%
6	3.5	3.826%	24.42%	42.36%
7	4	3.910%	23.64%	42.58%
8	4.5	3.986%	22.97%	42.79%
9	5	4.056%	22.39%	43.00%
10	5.5	4.121%	21.91%	43.22%
11	6	4.182%	21.50%	43.44%
12	6.5	4.239%	21.16%	43.66%
13	7	4.292%	20.88%	43.89%
14	7.5	4.342%	20.63%	44.13%
15	8	4.389%	20.43%	44.38%
16	8.5	4.434%	20.25%	44.64%
17	9	4.477%	20.09%	44.91%
18	9.5	4.517%	19.95%	45.18%
19	10	4.556%	19.83%	45.46%
20	10.5	4.593%	19.72%	45.75%
21	11	4.629%	19.62%	46.05%
22	11.5	4.663%	19.53%	46.36%
23	12	4.696%	19.45%	46.67%
24	12.5	4.728%	19.37%	46.98%
25	13	4.759%	19.30%	47.31%
26	13.5	4.788%	19.24%	47.64%
27	14	4.817%	19.18%	47.97%
28	14.5	4.845%	19.12%	48.31%
29	15	4.871%	19.07%	48.65%
30	15.5	4.897%	19.02%	49.00%
31	16	4.923%	18.97%	49.35%
32	16.5	4.947%	18.92%	49.71%
33	17	4.971%	18.88%	50.07%
34	17.5	4.994%	18.84%	50.43%
35	18	5.017%	18.80%	50.80%
36	18.5	5.039%	18.76%	51.17%
37	19	5.061%	18.73%	51.54%
38	19.5	5.082%	18.69%	51.92%

B Time-Dependent Parameter Scenario

Table B.4: Initial yield curve and effective Libor parameters for the time-
homogeneous parameter setting.

C Standard ρ_{∞} -extension

Lemma C.1. Let $\rho \in \mathbb{R}^N$ be a proper forward-rate correlation matrix, *i.e.*,

(i) ρ is symmetric, positive-definite,

- (*ii*) $\rho_{ii} = 1, \ 0 \le i \le N 1,$
- (*iii*) $\rho_{ij} \in [0, 1), \ 0 \le j < i \le N-1.$

Furthermore, let $\rho_{\infty} \in [0,1)$ and define a new matrix $C = (c_{ij})_{i,j=0}^{N-1}$ via

$$c_{ij} = \rho_{\infty} + (1 - \rho_{\infty})\rho_{ij}, \ i, j = 0, \dots, N-1.$$

Then C is again a proper forward-rate correlation matrix.

Proof. With $R_{\infty} \coloneqq (\rho_{\infty})_{i,j=0}^{N-1}$ we have

$$C = R_{\infty} + (1 - \rho_{\infty})\rho.$$

For arbitrary $x \in \mathbb{R}^N \setminus \{0\}$ it holds

$$x'Cx = \underbrace{x'\mathbb{R}_{\infty}x}_{\geq 0} + \underbrace{(1-\rho_{\infty})}_{>0} \underbrace{x'\rho x}_{>0} > 0,$$

i.e., C is positive definite. Moreover, C is clearly symmetric and

$$c_{ii} = 1 \cdot \rho_{\infty} + (1 - \rho_{\infty})c_{ii} = 1, \quad 0 \le i \le N - 1,$$

$$c_{ij} = 1 \cdot \rho_{\infty} + (1 - \rho_{\infty})c_{ij} \in [0, 1), \quad 0 \le j < i \le N - 1,$$

as a convex combination.

C.1 The Cholesky Decomposition of ho_∞ -extended Parametric Forms

Let $C \in \mathbb{R}^{N \times N}$ be a correlation matrix with given Cholesky decomposition C = LL. Suppose we need to calculate the Cholesky decomposition of

$$\tilde{C} = R + (1 - \rho_{\infty})C = R + (1 - \rho_{\infty})LL'$$
 (C.7)

where $\rho_{\infty} \in [0, 1)$ and $R = (\rho_{\infty})_{i,j=1}^{N}$. Instead of performing the standard Cholesky algorithm, we can calculate the Cholesky decomposition of \tilde{C} more efficiently by applying so-called rank-one updating methods to the already given L, see [GGMS74]. One possible approach is based on Givens rotations and works as follows. First observe, that we can write

$$\tilde{C} = aa' + (1 - \rho_{\infty})LL'$$
$$= (a \mid \bar{L})(a \mid \bar{L})',$$

where $a = (\sqrt{\rho_{\infty}}, \dots, \sqrt{\rho_{\infty}})'$ and $\bar{L} = \sqrt{1 - \rho_{\infty}} L$. The idea is then to multiply the matrix $(a \ \bar{L})$ by a sequence of Givens matrices in such a way, that we obtain a matrix $(0 \ \tilde{L})$ with a lower triangular matrix \tilde{L} (the required Cholesky factor of \tilde{C}).

The algorithm given below implements this approach and requires only $O(N^2)$ operations, while the standard Cholesky algorithm is of computational complexity $O(N^3)$.

148

Algorithm	1 : Rank one Cholesky updating via Givens rotations
Input:	$a = (\sqrt{\rho_{\infty}}, \dots, \sqrt{\rho_{\infty}})' \in (0, 1)^N$, (if $\rho_{\infty} = 0$ nothing to do)
	$\bar{L} = \sqrt{1-\rho_\infty} L \in \mathbb{R}^{N \times N}$ lower triangular
Output:	a is overwritten with 0
	\overline{L} is overwritten with \widetilde{L}
1:	for j=0 to N-1
2:	sq=sqrt(a[j]^2+L[j][j]^2)
3:	s=a[j]/sq
4:	c=L[j][j]/sq
5:	for i=j to N-1
6:	x=a[i]
7:	y=L[i][j]
8:	a[i]=c*x-s*y
9:	L[i][j]=s*x+c*y
10:	end
11:	end

We note that the 4P- and the 5P form presented in Section 5.3 are exactly of the form (C.7), with L being given in closed form. The above algorithm therefore allows for an almost instant calculation of the Cholesky decompositions of these parametric forms.

D Calibration Results

Corrform	$ ho_{\infty}$	α	β	γ	δ	η
SC2	0.44	-	_	-	-	0.82
Reb3	0.45	0.28	0.21	—	_	_
5P	0.29	0.68	0.27	2.24	-0.34	_

Table D.5: Best-fit parameters for the 2004–2008 historical correlationmatrix.

Corrform	ρ_{∞}	α	β	γ	δ	η
SC2	0.18	-	—	-	—	1.72
Reb3	0.24	0.25	0.31	—	_	_
5P	0.00	0.53	0.50	1.21	-0.33	_

 Table D.6: Best-fit parameters for the 2006–2010 historical correlation matrix.

				Tenor			
Maturity	6M	2Y	5Y	7Y	10Y	20Y	25Y
1Y	18.01%	17.50%	15.84%	14.83%	13.72%	12.19%	11.99%
2Y	17.76%	16.30%	14.82%	13.90%	13.09%	11.95%	11.86%
5Y	15.25%	14.16%	13.28%	12.86%	12.44%	11.72%	11.62%
7Y	14.11%	13.20%	12.65%	12.33%	12.12%	11.61%	—
10Y	13.01%	12.24%	12.02%	11.92%	11.81%	11.40%	—
15Y	11.99%	11.60%	11.58%	11.49%	11.64%	—	—
20Y	11.95%	11.38%	11.35%	11.39%	11.46%	—	-
25Y	11.70%	_		_	—	_	-
30Y	11.08%	_	-	_	—	_	-

				Tenor			
Maturity	6M	2Y	5Y	7Y	10Y	20Y	25Y
1Y	-30.00%	-30.00%	-27.82%	-26.87%	-21.23%	-25.04%	-27.10%
2Y	-20.13%	-30.00%	-25.50%	-25.94%	-21.02%	-23.27%	-24.04%
5Y	5.42%	-6.75%	-8.98%	-10.86%	-11.99%	-17.85%	-18.22%
7Y	5.81%	-5.39%	-9.71%	-12.03%	-13.17%	-17.17%	_
10Y	7.27%	2.29%	-6.99%	-9.35%	-12.58%	-16.05%	_
15Y	9.80%	3.20%	-3.58%	-8.63%	-13.57%	—	_
20Y	-2.72%	9.73%	3.71%	-2.96%	-13.09%	—	-
25Y	10.33%	_	—	_	_	_	-
30Y	16.75%	-	-	_	_	_	_

Table D.7: Effective caplet and swaption volatilities $\bar{\sigma}^*$ (top) and skews $\bar{\beta}^*$ (bottom) for the 2008 data set, obtained from the pre-calibration. The 6M-tenor columns give the effective caplet volatilities and skews, respectively. The best-fit stochastic-volatility parameters are $\kappa^* = 4.50\%$ and $\xi^* = 63.15\%$.

Corrform	$ ho_{\infty}$	α	β	γ	δ	η
SC2	0.35	-	-	-	-	1.05
Reb3	0.10	0.11	0.07	_	—	_
5P	0.45	1.92	0.03	-2.95	1.95	_

 Table D.8: Calibrated correlation parameters for the 2008 data set.

				Tenor				
Maturity	6M	2Y	5Y	7Y	10Y	20Y	25Y	
1Y	53.88%	41.61%	28.22%	24.67%	21.74%	19.97%	21.09%	
2Y	43.04%	33.72%	26.10%	23.70%	21.74%	21.09%	22.16%	
5Y	24.75%	22.91%	20.72%	20.11%	19.73%	20.36%	21.57%	
7Y	20.75%	19.84%	18.96%	18.84%	19.19%	20.18%	—	
10Y	18.23%	18.38%	18.39%	18.65%	19.17%	20.41%	—	
15Y	18.29%	18.69%	20.10%	20.58%	21.30%	_	—	
20Y	20.94%	21.45%	23.33%	23.86%	24.12%	_	—	
25Y	24.55%	-	-	-	-	—	_	
30Y	26.44%	—	—	—	—	—	—	
	Tenor							

				Tenor			
Maturity	6M	2Y	5Y	7Y	10Y	20Y	25Y
1Y	99.90%	60.86%	28.20%	19.11%	10.82%	-29.30%	-40.00%
2Y	62.67%	50.34%	22.14%	13.45%	3.56%	-22.29%	-32.56%
5Y	22.69%	13.93%	7.49%	1.10%	-7.53%	-24.48%	-31.49%
7Y	20.52%	7.38%	3.87%	-2.44%	-9.98%	-25.68%	_
10Y	14.85%	13.04%	10.39%	3.35%	-7.03%	-21.76%	-
15Y	18.27%	3.98%	5.59%	1.73%	-4.74%	_	_
20Y	36.90%	12.53%	14.68%	13.40%	7.64%	_	_
25Y	55.02%	—	—	—	—	—	_
30Y	69.25%	-	-	-	-	-	-

Table D.9: Effective caplet and swaption volatilities $\bar{\sigma}^*$ (top) and skews $\bar{\beta}^*$ (bottom) for the 2010 data set, obtained from the pre-calibration. The 6M-tenor columns give the effective caplet volatilities and skews, respectively. The best-fit stochastic-volatility parameters are $\kappa^* = 6.31\%$ and $\xi^* = 96.22\%$.

Corrparam.	$ ho_{\infty}$	α	β	γ	δ	η
SC2	0.11	—	-	—	—	0.99
Reb3	0.24	0.03	0.10	_	_	-
5P	0.28	3.00	0.04	-4.09	0.99	-

 Table D.10:
 Calibrated correlation parameters for the 2010 data set.

			time t		
Parameter	0	5	10	15	20
α	2.13	2.13	2.13	2.13	2.13
β	0.03	0.03	0.03	0.03	0.03
γ	-2.44	-2.44	-2.44	-2.44	-2.44
ρ_{∞}	0.44	0.41	0.18	0.01	0.00

Table D.11: Implied parameters of the time-dependent 4P form fitted to
the 2008 data set

Abbreviations and Notation

- A': transpose of matrix A.
- $A_{m,n}(t)$: annuity factor, PVBP (present value of a basis point).
- a.s.: almost surely.
- $B_d(t)$: discretely compounded money-market account.
- B(t): continuously compounded money-market account.
- $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot |\mathcal{F}_t]$: conditional expectation w.r.t. \mathcal{F}_t .
- $\eta(t)$: $\eta(t) + 1$ is the index of the first forward rate $L_{\eta(t)+1}(\cdot)$ that has not expired by time t.
- $\mathbb{F} = \{\mathcal{F}_t\}_{0 \le t}$: filtration.
- f(t,T): instantaneous forward rate.
- L(t,T): simply-compounded spot rate, Libor rate.
- $L(t, T_n, T_{n+1})$, $L_n(t)$: simply compounded forward rate, forward Libor rate.
- $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$: filtered probability space.
- P(t,T): time-t price of a zero-coupon bond with maturity T.
- $\Phi(\cdot)$: the Gaussian cumulative distribution function.
- Q: risk neutral measure or generic pricing measure.
- \mathbb{Q}^{B_d} : spot Libor measure.
- \mathbb{Q}^T : *T*-forward measure.
- $\operatorname{Re}(z)$, $\operatorname{Im}(z)$: real and imaginary part of a complex number z.
- R(t,T): continuously-compounded spot-rate.
- r(t): short rate.
- $S_{m,n}(t)$: (forward) swap rate.
- $\tau(T_n, T_{n+1})$: year fraction between time points T_n and T_{n+1} .
- 1W, 1M, 1Y: 1 week, 1 month, 1 year.
- AA: Antonov-Arneguy.
- ATM, ITM, OTM: at-the-money, in-the-money, out-of-the-money.
- BS: Black-Scholes.
- BGM: Brace-Gątarek-Musiela.
- bp: basis point, 1/100 of one percent $(1bp = 10^{-4})$.
- CDF: cumulative distribution function.
- CEV: constant elasticity of variance.
- CIR: Cox-Ingersoll-Ross.
- CMS: constant maturity swap.
- CMSSO: CMS spread option.
- DCT: discrete cosine transform.
- DFT: discrete fourier transform.

- FRA: forward-rate agreement.
- HJM: Heath-Jarrow-Morton.
- LMM: Libor market model.
- MC: Monte Carlo.
- ODE: ordinary differential equation.
- PCA: principal component analysis.
- PDE: partial differential equation.
- PDF: probability density function.
- RMSE: root mean square error.
- SC: Schoenmakers-Coffey.
- SDE: stochastic differential equation.
- SMM: swap market model.
- SV-LMM: stochastic-volatility extended Libor market model.

BIBLIOGRAPHY

Bibliography

- [AA00] L. B. Andersen and J. Andreasen. Volatility skews and extensions of the Libor market model. Applied Mathematical Finance, 7(1), 2000.
- [AA01] L. B. Andersen and J. Andreasen. Factor dependence of Bermudan swaptions: fact or fiction? *Journal of Financial Economics*, 62(1):3–37, October 2001.
- [AA02] L. B. Andersen and J. Andreasen. Volatile Volatilities. Risk, 15(12):163–168, December 2002.
- [AA09] A. V. Antonov and M. Arneguy. Analytical formulas for pricing CMS products in the Libor Market Model with stochastic volatility. Working paper, 2009. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1352606.
- [AB09] F. M. Ametrano and M. Bianchetti. Bootstrapping the Illiquidity: Multiple Yield Curves Construction for Market Coherent Forward Rates Estimation. Working paper, 2009. http://www.bianchetti.org/Finance/ BootstrappingTheIlliquidity-v1.0.pdf.
- [ABR05] L. B. Andersen and R. Brotherton-Ratcliffe. Extended Libor market models with stochastic volatility. *The Journal of Computational Finance*, 9(1):1–40, 2005.
- [AHNW04] E. Ayache, P. Henrotte, S. Nassar, and X. Wang. Can anyone solve the smile problem? Wilmott, pages 78–96, January 2004.
- [AL03] C. Alexander and D. Lvov. Statistical Properties of Forward Libor Rates. ISMA Discussion Papers in Finance, 2003. http://www.icmacentre.ac.uk/pdf/ discussion/DP2003-03.pdf.
- [AMST07] H. Albrecher, P. Mayer, W. Schoutens, and J. Tistaert. The Little Heston Trap. Wilmott, pages 83–92, January 2007.
- [And05] L. B. Andersen. Discount Curve Construction with Tension Splines. Review of Derivaties Research, 10(3):227–267, 2005.
- [And08] L. B. Andersen. Simple and efficient simulation of the Heston stochastic volatility model. *The Journal of Computational Finance*, 11(1):1–42, 2008.
- [AP07] L. B. Andersen and V. V. Piterbarg. Moment explosions in stochastic volatility models. *Finance and Stochastics*, 11(1):29–50, January 2007.
- [AP10a] L. B. Andersen and V. V. Piterbarg. Interest Rate Modeling Volume I: Foundations and Vanilla Models. Atlantic Financial Press, 2010.
- [AP10b] L. B. Andersen and V. V. Piterbarg. Interest Rate Modeling Volume II: Term Structure Models. Atlantic Financial Press, 2010.
- [AP10c] L. B. Andersen and V. V. Piterbarg. Interest Rate Modeling Volume III: Products and Risk Management. Atlantic Financial Press, 2010.

- [BGM97] A. Brace, D. Gątarek, and M. Musiela. The Market Model of Interest Rate Dynamics. *Mathematical Finance*, 7(2):127–147, April 1997.
- [Bjö09] T. Björk. Arbitrage Theory in Continuous Time. Oxford University Press, 3rd edition, 2009.
- [BK91] F. Black and P. Karasinski. Bond and Option Pricing when Short Rates are Lognormal. *Financial Analysts Journal*, 47(4):52–59, 1991.
- [BK04] N. H. Bingham and R. Kiesel. *Risk-Neutral Valuation Pricing and Hedging of Financial Derivatives*. Springer, 2nd edition, 2004.
- [BKS10] D. Belomestny, A. Kolodko, and J. Schoenmakers. Pricing CMS spread options in a Libor market model. *International Journal of Theoretical and Applied Finance*, 13(1):45–62, February 2010.
- [BL78] D. T. Breeden and R. H. Litzenberger. Prices of State-Contingent Claims Implicit in Option Prices. Journal of Business, 51(4):621–651, October 1978.
- [Bla76] F. Black. The pricing of commodity contracts. *Journal of Financial Economics*, 3:167–179, 1976.
- [BM05] D. Brigo and F. Mercurio. Interest Rate Models Theory and Practice. Springer, 2nd edition, 2005.
- [Boy01] J. P. Boyd. *Chebyshev and Fourier Spectral Methods*. Dover Publications, 2nd edition, 2001.
- [Bra08] A. Brace. Engineering BGM. Chapman & Hall/CRC, 2008.
- [Bri02] D. Brigo. A Note on Correlation and Rank Reduction. Draft, 2002. www. damianobrigo.it/correl.pdf.
- [BS73] F. Black and M. Scholes. The Pricing of Options and Corporate Liabilities. The Journal of Political Economy, 81(3):637–654, 1973.
- [BS79] M. J. Brennan and E. S. Schwartz. A continuous time approach to the pricing of bonds. *Journal of Banking and Finance*, 3(1):133–155, 1979.
- [BS03] W. Boenkost and W. M. Schmidt. Notes on convexity and quanto adjustments for interest rates and related options. Working paper, 2003. http://www.econstor. eu/bitstream/10419/27810/1/378773836.PDF.
- [BYR06] V. Britanak, P. C. Yip, and K. R. Rao. Discrete Cosine and Sine Transforms: General Properties, Fast Algorithms and Integer Approximations. Academic Press, 2006.
- [Car94] A. Carverhill. When is the short rate Markovian? *Mathematical Finance*, 4(4):305–312, October 1994.
- [CDG02] P. Collin-Dufresne and R. S. Goldstein. Do Bonds Span the Fixed Income Markets? Theory and Evidence for Unspanned Stochastic Volatility. *The Journal of Finance*, 57(4):1685–1730, August 2002.
- [CIR85] J. C. Cox, J. E. Ingersoll, and S. A. Ross. A Theory of the Term Structure of Interest Rates. *Econometrica*, 53(2):385–407, March 1985.
- [CLV04] U. Cherubini, E. Luciano, and W. Vecchiato. *Copula methods in finance*. Wiley, 2004.
- [CM99] P. Carr and D. B. Madan. Option valuation using the fast Fourier transform.

The Journal of Computational Finance, 2(4):61-73, September 1999.

- [CM01] P. Carr and D. Madan. Optimal positioning in derivative securities. *Quantitative Finance*, 1(1):19–37, 2001.
- [Cox96] J. C. Cox. Notes on Option Pricing I: Constant Elasticity of Variance Diffusions. Working paper, Stanford University, 1975, reprinted in *Journal of Portfolio Management*, 23(1):15–17, 1996.
- [CR76] J. C. Cox and S. A. Ross. The valuation of options for alternative stochastic processes. Journal of Financial Economics, 3:145–166, 1976.
- [CS89] M. Chesney and L. Scott. Pricing European Currency Options: A Comparison of the Modified Black-Scholes Model and a Random Variance Model. Journal of Financial and Quantitative Analysis, 24(3):267–284, September 1989.
- [CS04] R.-R. Chen and L. Scott. Stochastic Volatility and Jumps in Interest Rates: An Empirical Analysis. Working paper, Rutgers University and Morgan Stanley, 2004. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=686985.
- [DK94] E. Derman and I. Kani. Pricing with a Smile. *Risk*, 7(2):32–39, February 1994.
- [DL01] D. Davydov and V. Linetsky. Pricing and hedging path-dependent options under the CEV process. *Management Science*, 47(7):949–965, 2001.
- [Dot78] L. U. Dothan. On the term structure of interest rates. Journal of Financial Economics, 6(1):59–69, 1978.
- [DS02] F. Delbaen and H. Shirakwa. A note on Option Pricing for the Constant Elasticity of Variance Model. *Asia-Pacific Financial Markets*, 9(2):85–99, 2002.
- [Dup94] B. Dupire. Pricing with a Smile. *Risk*, 7(1):18–20, 1994.
- [EÖ05] E. Eberlein and F. Özkan. The Lévy LIBOR model. Finance and Stochastics, 9(3):327–348, 2005.
- [Fri07] C. Fries. Mathematical Finance Theory, Modeling, Implementation. Wiley, 2007.
- [GG00] W. Gander and W. Gautschi. Adaptive Quadrature Revisited. BIT, 40(2):84– 101, March 2000.
- [GGMS74] P. E. Gill, G. H. Golub, W. Murray, and A. Saunders. Methods for Modifying Matrix Factorizations. *Mathematics of Computation*, 28(126):505–535, April 1974.
- [GH04] S. Galluccio and C. Hunter. The Co-initial Swap Market Model. Economic Notes, 33(2):209–232, July 2004.
- [GK03] P. Glasserman and S. G. Kou. The term structure of simple forward rates with jump risk. *Mathematical Finance*, 13(3):383–410, July 2003.
- [Gla04] P. Glasserman. Monte Carlo Methods in Financial Engineering. Springer, 2004.
- [GP07] I. Grubišić and R. Pietersz. Efficient rank reduction of correlation matrices. Linear Algebra and its Applications, 422(2-3), April 2007.
- [Hag03] P. Hagan. Convexity Conundrums: Pricing CMS Swaps, Caps and Floors. Wilmott, pages 38–44, March 2003.
- [Hes93] S. L. Heston. A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *The Review of Financial Studies*,

6(2):327-343, 1993.

- [Hig02] N. J. Higham. Computing the nearest correlation matrix a problem from finance. *IMA Journal of Numerical Analysis*, 22:329–343, 2002.
- [HJM92] D. Heath, R. Jarrow, and A. Morton. Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation. *Econometrica*, 60(1):77–105, 1992.
- [HK00] P. J. Hunt and J. E. Kennedy. *Financial Derivatives in Theory and Practice*. John Wiley & Sons, 2000.
- [HKLW02] P. Hagan, D. Kumar, A. S. Lesniewski, and D. E. Woodward. Managing smile risk. Wilmott, 2(1):84–108, 2002.
- [HL86] T. S. Ho and S.-B. Lee. Term Structure Movements and Pricing Interest Rate Contingent Claims. Journal of Finance, 41(5):1011–1029, December 1986.
- [HW87] J. C. Hull and A. D. White. The Pricing of Options on Assets with Stochastic Volatilities. *The Journal of Finance*, 42(2), June 1987.
- [HW90] J. C. Hull and A. D. White. Pricing Interest-Rate-Derivative Securities. The Review of Financial Studies, 3(4):573–592, 1990.
- [HW00] J. C. Hull and A. D. White. Forward Rate Volatilities, Swap Rate Volatilities and Implementation of the LIBOR Market Model. The Journal of Fixed Income, 10(2):46–62, September 2000.
- [HW03] M. Heidari and L. Wu. Are Interest Rate Derivatives Spanned by the Term Structure of Interest Rates. The Journal of Fixed Income, 13(1):75–86, June 2003.
- [HW06] P. S. Hagan and G. West. Interpolation Methods for Curve Construction. Applied Mathematical Finance, 13(2):89–129, 2006.
- [HW08] P. S. Hagan and G. West. Methods for Constructing a Yield Curve. Wilmott Magazine, 35, 2008.
- [Hym83] J. M. Hyman. Accurate monotonicity preserving cubic interpolation. SIAM Journal on Scientific and Statistical Computing, 4(4):645–654, December 1983.
- [HZ10] T. R. Hurd and Z. Zhou. A Fourier Transform Method for Spread Option Pricing. SIAM Journal on Financial Mathematics, 1:142–157, 2010.
- [ISD06] ISDA. ISDA definitions. 2006.
- [Jam97] F. Jamshidian. LIBOR and swap market models and measures. Finance and Stochastics, 1(4):293–330, September 1997.
- [JK05] P. Jäckel and A. Kawai. The Future is Convex. *Wilmott Magazine*, 15, 2005.
- [JKB95] L. N. Johnson, S. Kotz, and N. Balakrishnan. Continuous Univariate Distributions Volume 2. John Wiley & Sons, 2nd edition, 1995.
- [Jos03] M. S. Joshi. Rapid Computation of Drifts in a Reduced Factor LIBOR Market Model. Wilmott magazine, pages 84–85, May 2003.
- [JR03a] P. Jäckel and R. Rebonato. The link between caplet and swaption volatilities in a Brace-Gatarek-Musiela/Jamshidian framework: approximate solutions and empirical evidence. *The Journal of Computational Finance*, 6(4):41–59, 2003.
- [JR03b] M. S. Joshi and R. Rebonato. A displaced-diffusion stochastic volatility LIBOR

market model: motivation, definition and implementation. *Quantitative Finance*, 3(6):458–469, December 2003.

- [JS08] M. S. Joshi and A. Stacey. New and robust drift approximations for the LIBOR market model. *Quantitative Finance*, 8(4):427–434, June 2008.
- [KG10] M. Kamal and J. Gatheral. Implied volatility surface. In Encyclopedia of Quantitative Finance. John Wiley & Sons, 2010.
- [KJ05] C. Kahl and P. Jäckel. Not-so-complex logarithms in the Heston model. Wilmott, pages 94–103, September 2005.
- [KL11] R. Kiesel and M. Lutz. Efficient Pricing of CMS Spread Options in a Stochastic Volatility LMM. The Journal of Computational Finance, forthcoming, 2011.
- [KN97] G. Kirikos and D. Novak. Convexity Conundrums. Risk, 10(3):60–61, March 1997.
- [KS98] I. Karatzas and S. E. Shreve. Brownian Motion and Stochastic Calculus. Springer, 2nd edition, 1998.
- [KT81] S. Karlin and H. M. Taylor. A Second Course in Stochastic Processes. Academic Press, 1981.
- [Lee04] R. W. Lee. Option Pricing by Transform Methods: Extensions, Unification, and Error Control. *The Journal of Computational Finance*, 7(3):51–86, 2004.
- [Lew01] A. L. Lewis. A simple option formula for general jump-diffusion and other exponential Lévy processes. Working paper, 2001. http://optioncity.net/pubs/ ExpLevy.pdf.
- [Lip02] A. Lipton. The vol smile problem. *Risk*, 15(2):61–65, February 2002.
- [LK07] R. Lord and C. Kahl. Optimal Fourier inversion in semi-analytical option pricing. The Journal of Computational Finance, 10(4):1–30, 2007.
- [LL08] D. Lamberton and B. Lapeyre. Introduction to Stochastic Calculus Applied to Finance. Chapman & Hall/CRC, 2nd edition, 2008.
- [LS01] F. A. Longstaff and E. S. Schwartz. Valuing American Options by Simulation: A Simple Least-Squares Approach. *Review of Financial Studies*, 14(1):113–147, 2001.
- [LZ06] H. Li and F. Zhao. Unspanned Stochastic Volatility: Evidence from Hedging Interest Rate Derivatives. The Journal of Finance, 61(1):341–378, February 2006.
- [McC11] P. McCloud. The CMS triangle arbitrage. Risk, 24(1):126–131, January 2011.
- [Mer08] F. Mercurio. Cash-settled swaptions and no-arbitrage. *Risk*, 21(2):96–98, February 2008.
- [MFE05] A. J. McNeil, R. Frey, and P. Embrechts. *Quantitative Risk Management*. Princeton University Press, 2005.
- [MP06] F. Mercurio and A. Pallavicini. Smiling at convexity. *Risk*, 19(8):64–69, August 2006.
- [MR97] M. Musiela and M. Rutkowski. Continuous-time term structure models: Forward measure approach. *Finance and Stochastics*, 1(4):261–291, September 1997.
- [MR05] M. Musiela and M. Rutkowski. Martingale Methods in Financial Modelling.

Springer, 2nd edition, 2005.

- [MSS97] K. R. Miltersen, K. Sandmann, and D. Sondermann. Closed Form Solutions for Term Structure Derivatives with Log-Normal Interest Rates. The Journal of Finance, 52(1):409–430, March 1997.
- [Nef08] S. N. Neftci. *Principles of Financial Engineering*. Academic Press, 2nd edition, 2008.
- [Nel06] R. B. Nelsen. An Introduction to Copulas. Springer, 2nd edition, 2006.
- [NS87] C. R. Nelson and A. F. Siegel. Parsimonious Modeling of Yield Curves. Journal of Business, 60(4):473–489, October 1987.
- [Øks03] B. Øksendal. Stochastic Differential Equations An Introduction with Applications. Springer, 6th edition, 2003.
- [Pel03] A. Pelsser. Mathematical foundation of convexity correction. Quantitative Finance, 3(1):59–65, 2003.
- [PG04] R. Pietersz and P. J. Groenen. Rank reduction of correlation matrices by majorization. *Quantitative Finance*, 4(6):649–662, December 2004.
- [Pit03] V. V. Piterbarg. A stochastic volatility forward Libor model with a term structure of volatility smiles. Working paper, 2003. http://papers.ssrn.com/sol3/ papers.cfm?abstract_id=472061.
- [Pit05a] V. V. Piterbarg. Stochastic Volatility Model with Time-dependent Skew. Applied Mathematical Finance, 12(2):147–185, June 2005.
- [Pit05b] V. V. Piterbarg. Time to smile. *Risk*, 18(5):71–75, May 2005.
- [PR06] V. V. Piterbarg and M. A. Renedo. Eurodollar futures convexity adjustments in stochastic volatility models. *The Journal of Computational Finance*, 9(3), 2006.
- [PS94] N. D. Pearson and T.-S. Sun. Exploiting the Conditional Density in Estimating the Term Structure: An Application to the Cox, Ingersoll and Ross Model. The Journal of Finance, 49(4):1279–1304, 1994.
- [PSHE09] P. Poulsen, K. R. Schenk-Hoppé, and C.-O. Ewald. Risk minimization in stochastic volatility models: model risk and empirical performance. *Quantitative Fi*nance, 9(6):693–704, September 2009.
- [PTVF01] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. Numerical Recipes – The Art of Scientific Computing. Cambridge University Press, 2001.
- [RBM07] F. Rapisarda, D. Brigo, and F. Mercurio. Parameterizing correlations: a geometric interpretation. IMA Journal of Management Mathematics, 18(1):55–73, January 2007.
- [Reb99a] R. Rebonato. On the simultaneous calibration of multifactor lognormal interest rate models to Black volatilities and to the correlation matrix. *Journal of Computational Finance*, 2(4):5–27, 1999.
- [Reb99b] R. Rebonato. Volatility and Correlation in the pricing of equity, FX and interestrate options. John Wiley & Sons, 1999.
- [Reb02] R. Rebonato. Modern Pricing of Interest-Rate Derivatives The LIBOR Market Model and Beyond. Princeton University Press, 2002.
- [Reb03] R. Rebonato. Which process gives rise to the observed dependence of swaption

implied volatility on the underlying? International Journal of Theoretical and Applied Finance, 6(4):419–442, 2003.

- [Reb04] R. Rebonato. Volatility and Correlation The Perfect Hedger and the Fox. John Wiley & Sons, 2nd edition, 2004.
- [RJ99] R. Rebonato and P. Jäckel. The Most General Methodology to Create a Valid Correlation Matrix for Risk Management and Option Pricing Purposes. The Journal of Risk, 2(2):17–28, 1999.
- [RMW09] R. Rebonato, K. McKay, and R. White. The SABR/Libor Market Model Pricing, Calibration and Hedging for Complex Interest-Rate Derivatives. John Wiley & Sons, 2009.
- [Rog96] L. C. Rogers. Gaussian errors. *Risk*, 9(1):42–45, January 1996.
- [Ron00] U. Ron. A Practical Guide to Swap Curve Construction. Technical Report, Bank of Canada, 2000. http://www.bank-banque-canada.ca/en/res/wp/2000/wp00-17. pdf.
- [RS95] P. Ritchken and L. Sankarasubramanian. Volatility structures of forward rates and the dynamics of the term structure. *Mathematical Finance*, 5(1):55–72, January 1995.
- [Rub83] M. Rubinstein. Displaced Diffusion Option Pricing. The Journal of Finance, 38(1):213–217, March 1983.
- [Rut99] M. Rutkowski. Models of forward Libor and swap rates. Applied Mathematical Finance, 6(1):29–60, March 1999.
- [Rut01] M. Rutkowski. Modelling of Forward Libor and Swap Rates. In E. Jouini, J. Cvitanic, and M. Musiela, editors, Option Pricing, Interest Rates and Risk Management, pages 336–395. Cambridge University Press, 2001.
- [SC03] J. Schoenmakers and B. Coffey. Systematic Generation of Correlation Structures for the Libor Market Model. International Journal of Theoretical and Applied Finance, 6(5):507–519, 2003.
- [Sch89] M. Schroder. Computing the Constant Elasticity of Variance Option Pricing Formula. The Journal of Finance, 44(1):211–219, March 1989.
- [Sch05] J. Schoenmakers. Robust Libor Modelling and Pricing of Derivative Products. Chapman & Hall/CRC, 2005.
- [Sch10] M. Schmelzle. Option Pricing Formulae using Fourier Transform: Theory and Application. Unpublished, 2010. http://pfadintegral.com/docs/ Schmelzle2010%20Fourier%20Pricing.pdf.
- [SG09] S. Svoboda-Greenwood. Displaced Diffusion as an Approximation of the Constant Elasticity of Variance. Applied Mathematical Finance, 16(3):269–286, June 2009.
- [SS91] E. M. Stein and J. C. Stein. Stock Price Distributions with Stochastic Volatility: An Analytic Approach. *The Review of Financial Studies*, 4(4):727–752, 1991.
- [SS97] K. Sandmann and D. Sondermann. A note on the stability of lognormal interest rate models and the pricing of eurodollar futures. *Mathematical Finance*, 7(2):119–125, April 1997.
- [SST04] W. Schoutens, E. Simons, and J. Tistaert. A Perfect Calibration! Now What?

Wilmott, pages 66-78, March 2004.

- [Sve94] L. E. Svensson. Estimating and interpreting forward interest rates: Sweden 1992-1994. NBER Working Paper Series, 1994. http://www.nber.org/papers/w4871. pdf.
- [Tal79] A. Talbot. The accurate numerical inversion of Laplace transforms. Journal of the Institute of Mathematics and Its Applications, 23(1):97–120, 1979.
- [Tan10] C. C. Tan. Demystifying exotic products Interest Rates, Equities and Foreign Exchange. Wiley Finance, 2010.
- [Vas77] O. Vasicek. An equilibrium characterization of the term structure. Journal of Financial Economics, 5(2):177–188, 1977.
- [VHP10] A. Van Haastrecht and A. Pelsser. Efficient, almost exact simulation of the Heston stochastic volatility model. *International Journal of Theoretical and Applied Finance*, 13(1):1–43, February 2010.
- [Wig87] J. B. Wiggins. Option Values under Stochastic Volatility. Journal of Financial Economics, 19(2):351–372, 1987.
- [Wil05] R. R. Wilcox. Introduction to Robust Estimation and Hypothesis Testing. Elsevier Academic Press, 2nd edition, 2005.
- [ZK10] W. Zheng and Y. K. Kwok. Convexity meets replication: hedging of swap derivatives and annuity options. Working paper, 2010. http://papers.ssrn.com/sol3/ papers.cfm?abstract_id=1421734.
- [ZW03] Z. Zhang and L. Wu. Optimal low-rank approximation to a correlation matrix. Linear Algebra and its Applications, 364:161–187, May 2003.

List of Figures

1.1	10Y and 2Y swap-rate time series
2.1	ATM EUR cap volatility curve
2.2	ATM EUR swaption volatilities
2.3	CEV and DD valatility smile
2.4	Displaced Hoston Smiles
2.0 9.6	CMS convertity adjustments
$\frac{2.0}{2.7}$	Displaced Hoston second order moment explosion times
2.1	Displaced fieston second order moment explosion times
3.1	ATM EUR caplet volatility curve
4.1	Fourier inversion integrand
4.2	Integration contours
4.3	Real and imaginary part of the integrand in the complex plane 84
4.4	Graph of $h(u)$
4.5	Smoothed integrand 86
4.6	Density of integrated variance
4.7	Time-homogeneous Libor volatilities
4.8	Time-homogeneous Libor correlations
4.9	Implied spread volatilities for the time-dependent scenario 92
4.10	Absolute approximation errors
5.1	Forward-rate curve time series
5.2	Matrix C
5.3	A variety of shapes that can be generated by the 5P form $\ldots \ldots \ldots 109$
5.4	Parameterizations fitted to historical correlations (2004–2008) 110
5.5	Parameterizations fitted to historical correlations (2004–2008): First
	columns and last rows
5.6	Parameterizations fitted to historical correlations (2006–2010) 112
5.7	Parameterizations fitted to historical correlations (2006–2010): First
	columns and last rows
6.1	Rank reduction to 4 factors
6.2	Rank reduction to 7 factors
7.1	Spread-option greeks
7.2	Initial 6M forward-rate curves
7.3	2008 data set: Calibrated spread volatilities for $K = 0.25\%$

LIST OF FIGURES

7.4	2008 data set: Spread volatility surfaces
7.5	10Y–2Y CMS spread options: Relevant forward-rate correlations $\ . \ . \ . \ 132$
7.6	2008 data set: Calibrated correlation matrices
7.7	2008 data set: Calibrated time-dependent forward-rate volatilities and
	skews
7.8	2010 data set: Calibrated correlation matrices
7.9	2010 data set: Calibrated spread volatilities for $K=0.50\%$ $~$
7.10	2010 data set: Spread volatility surfaces $\hdots \hdots $
7.11	2010 data set: Calibrated time-dependent forward-rate volatilities and
	skews

List of Tables

2.1	Classical short-rate models
$4.1 \\ 4.2$	Test results for the time-dependent parameter scenario
5.1	Root mean square errors with respect to the historical correlation matrices.
6.1	Computing times
7.1	2008 data set: Calibration errors
7.2	2010 data set: Calibration errors $\ldots \ldots 135$
7.3	Prices of exotic products
B.4	Initial yield curve and effective parameters
D.5	Best-fit parameters for the $2004-2008$ historical correlation matrix. $% 149$
D.6	Best-fit parameters for the $2006-2010$ historical correlation matrix. $% 149$
D.7	Effective caplet and swaption volatilities for the 2008 data set 150
D.8	Calibrated correlation parameters for the 2008 data set
D.9	Effective caplet and swaption volatilities for the 2010 data set 151
D.10	Calibrated correlation parameters for the 2010 data set
D.11	Implied parameters of the time-dependent 4P form fitted to the 2008
	data set

LIST OF TABLES

166

Zusammenfassung

Libor-Markt-Modelle zählen zu den Standardmodellen in der Finanzindustrie, wenn es darum geht exotische Zinsderivate zu bewerten und zu hedgen. Aufgrund der Flexibilität ihrer Volatililitäts-Spezifizierung können diese Modelle an eine große Anzahl von Marktinstrumenten kalibriert werden und liefern somit Preise von exotischen Strukturen, die konsistent mit den vorherrschenden Marktbedingungen sind. Im Gegensatz zu anderen Modellen ist die dynamische Struktur von Libor-Markt-Modellen auch reichhaltig genug, um Dekorrelationseffekte zwischen verschiedenen Zinssätzen der Zinskurve realistisch nachzubilden. Dies ist besonders wichtig bei der Bewertung von korrelationssensitiven Produkten, wie zum Beispiel CMS-Spread-Produkten, deren Zahlungsströme jeweils von der Differenz zweier Zinssätze abhängen. Produkte dieser Art haben sich in den letzten Jahren relativ großer Beliebtheit erfreut und exotische Anleihen mit CMS-Spread-abhängigen Coupons wurden in großen Volumina gehandelt. Gleichzeitig hat sich auch ein liquider Markt für "gewöhnliche" CMS-Spread-Optionen entwickelt, und Preise solcher Derivate werden heutzutage von verschiedenen Brokern auf täglicher Basis gequotet. Dennoch ist das korrekte Preisen von CMS-Spread-basierten Produkten immer noch Gegenstand aktiver Forschung. Tatsächlich deckte beispielsweise McCloud [McC11] in einem kürzlich erschienenen Artikel auf, dass es im Jahr 2009 aufgrund von Bewertungsinkonsistenzen über längere Zeiträume hinweg statische Arbitrage-Möglichkeiten zwischen den Märkten für CMS- und CMS-Spread-Optionen gab. Das konsistente Bewerten von CMS-Spread-Optionen relativ zu anderen Marktsektoren ist jedoch nicht die einzige Herausforderung. Auch die verschiedenen CMS-Spread-abhängigen Produkte müssen selbstverständlich konsistent zueinander bewertet werden. Zum einen ist dies notwendig um direkte Arbitrage-Möglichkeiten auszuschließen. Andererseits sollten die Hedging-Kosten beim Preisen von exotischen Produkten korrekt mit einbezogen werden. Schließlich spiegelt der Preis eines exotischen (oder jedes anderen) Derivats letztendlich die während der Laufzeit anfallenden Hedgingkosten wider. Da CMS-Spread-Optionen aufgrund ihrer Liquidität inzwischen als Hedginginstrumente für andere korrelationssensitive exotische Produkte eingesetzt werden können, sollten deren Preise natürlich beim Bewerten der exotischen Instrumente sachgemäß berücksichtigt werden.

Das Ziel der vorliegenden Arbeit besteht darin, effiziente Methoden und Werkzeuge zur Kalibrierung von Libor-Markt-Modellen an Caps, Swaptions und CMS-Spread-Optionen zu entwickeln, um auf diesem Wege marktimplizierte Volatilitäten und Korrelationen zu extrahieren und exotische Produkte marktkonsistent zu bewerten.

Da das standard (log-normale) Libor-Markt-Modell, wie es von Miltersen, Sandmann & Sondermann [MSS97], Brace, Gatarek & Musiela (BGM) [BGM97] und Jamshidian [Jam97] eingeführt wurde, die in den heutigen Märkten beobachtbaren Volatilitäts-Smiles nicht nachbilden kann, sind zahlreiche Erweiterungen des ursprünglichen Libor-Markt-Modells in der Literatur erschienen. Die wohl beliebtesten und am weitesten verbreiteten dieser Erweiterungen basieren auf Dynamiken mit stochastischer Volatilität vom Heston-Typ (siehe z.B. [AA02], [ABR05] oder [Pit05a]). In Kapitel 4 stellen wir eine neue Approximation zum Preisen von CMS-Spread-Optionen innerhalb dieser Modellklasse vor. Grundlage dieser Approximation ist eine neue Methode zum effizienten Auswerten der Dichte eines integrierten Cox-Ingersoll-Ross-Prozesses (CIR) (siehe Abschnitt 4.3), welche auf einer speziellen Wahl des Integrationspfades bei der notwendigen Laplace-Invers-Transformation basiert. Diese Methode kann nicht nur bei der Bewertung von CMS-Spread-Optionen verwendet werden, sondern wann immer eine schnelle und genaue Auswertung der Dichte eines integrierten CIR-Prozesses benötigt wird. Das Preisen von CMS-Spread-Optionen innerhalb der zuvor erwähnten Modellklasse wurde auch in einem kürzlich erschienenen Arbeitspapier von Antonov & Arneguy [AA09] betrachtet. Wir zeigen in Abschnitt 4.5, dass unsere Bewertungsformel hinsichtlich Geschwindigkeit, Genauigkeit und Aufwand der Implementierung im Allgemeinen besser abschneidet als die Bewertungsmethoden von Antonov & Arneguy.

Bisher war es üblich die Volatilitätsparameter von Libor-Markt-Modellen durch Kalibrierung an Caps und Swaptions zu gewinnen, während die Korrelationsparameter meist historisch geschätzt wurden. Auch wenn diese Vorgehensweise bei weniger korrelationssensitiven Produkten akzeptabel ist, so ist sie bei Instrumenten die stark von der Abhängigkeitsstruktur zwischen den einzelnen Zinssätzen abhängen nicht ideal, da die historisch geschätzten Korrelationen in der Regel nicht die aktuellen Marktbedingungen widerspiegeln (vgl. Abschnitt 5.4). Mit Hilfe der von uns vorgestellten Approximationsformel lassen sich Libor-Markt-Modelle mit stochastischer Volatilität gleichzeitig an Caps, Swaptions und CMS-Spread-Optionen kalibrieren. Auf diese Weise können nicht nur marktimplizierte Volatilitäten sondern auch marktimplizierte Korrelationen gewonnen werden und komplexe Produkte somit marktkonsistent bewertet werden.

Unabhängig davon welche Kalibrierungsmethode man für die Libor-Korrelationen verwendet – ob implizit oder historisch – in beiden Fällen benötigt man sparsame und dennoch flexible Parametrisierungen für die Libor-Korrelationsmatrizen. Historisch geschätzte Korrelationsmatrizen sind oft verrauscht und enthalten teilweise unrealistische Einträge. Durch das Fitten von Parametrisierungen an die historisch geschätzten Matrizen versucht man glatte Korrelationsmatrizen zu erhalten, welche nur die wesentlichen Charakteristika der historischen Daten widerspiegeln. Bei der impliziten Kalibrierung hingegen werden Korrelationsparametrisierungen benötigt um Overfitting zu vermeiden und stabile Kalibrierungsergebnisse zu erhalten. In Kapitel 5 präsentieren wir eine neue generische Methode um Korrelationsparametrisierungen zu konstruieren, welche stets positiv definite Matrizen liefern. Des Weiteren stellen wir in Abschnitt 5.3 konkrete Parametrisierungen vor, welche äußert flexibel sind und historisch geschätzte Matrizen wesentlich besser fitten als die bekannten Standard-Parametrisierungen (siehe Abschnitt 5.4).
In der Praxis ist die verwendete Anzahl der Faktoren bei einem Libor-Markt-Modell typischerweise wesentlich kleiner (z.B. 3–10) als die Anzahl der modellierten Libor-Raten (oftmals 40–80). Da praktisch alle der bekannten Korrelationsparametrisierungen Vollfaktor-Matrizen ergeben, muss der Rang dieser Matrizen erst mittels PCA²¹-basierten Methoden reduziert werden, bevor sie tatsächlich innerhalb des Modells verwendet werden können. Im Falle eines historischen Kalibrierungsansatzes stellt diese Vorgehensweise im Allgemeinen keinen sonderlich großen Nachteil dar, da der PCA-Algorithmus nur ein einziges Mal durchgeführt werden muss. Werden jedoch die Korrelationsparameter über eine implizite Kalibrierung geschätzt, so muss der PCA-Algorithmus möglicherweise mehrere tausend Mal durchgeführt werden. In diesem Fall können die dabei notwendigen numerischen Eigenwertzerlegungen einen Großteil des für die gesamte Kalibrierung benötigten Rechenaufwandes ausmachen (vgl. Abschnitt 6.3 und 7.3). In Kapitel 6 entwickeln wir deshalb eine neue Methode um den Rang von parametrischen Libor-Korrelationsmatrizen effizient zu reduzieren. Diese basiert darauf eine Diskrete Kosinustransformation (DCT) auf die Zeilen der Cholesky-Zerlegung der Korrelationsmatrizen anzuwenden (siehe Abschnitt 6.2). Hierbei müssen im Wesentlichen nur einige Matrixmultiplikationen durchgeführt werden. Da außerdem die Cholesky-Zerlegung bei unseren parametrischen Formen in geschlossener Form angegeben werden kann (vgl. Abschnitt 5.3), erhalten wir aus der Kombination von DCT-Methode und unseren parametrischen Formen implizit eine neue Familie von flexiblen Parametrisierungen mit beliebiger Faktoranzahl.

In Kapitel 7 werden schließlich die zuvor entwickelten Methoden und Werkzeuge angewandt und Libor-Markt-Modelle an Marktdaten kalibriert. Wir diskutieren dabei in Abschnitt 7.1 zwei möglich Kalibrierungsansätze und demonstrieren in Abschnitt 7.2, dass mit Hilfe unserer neuen Korrelations-Parametrisierungen die Marktpreise im Allgemeinen besser getroffen werden als mit den üblichen Standard-Parametrisierungen. Ein Hauptresultat unserer empirischen Studien ist dabei, dass implizite Korrelationsmatrizen keine stark ansteigenden Nebendiagonalen aufweisen – ein typisches Merkmal von historisch geschätzten Korrelationsmatrizen, welches manchmal sogar bei Parametrisierungen direkt integriert wird (vgl. Abschnitt 5.2).

Zuletzt verwenden wir in Abschnitt 7.3 die kalibrierten Modelle um einige praxisrelevante Bewertungsbeispiele zu diskutieren. Insbesondere demonstrieren wir dabei, dass Libor-Markt-Modelle mit verschiedenen Korrelationsstrukturen signifikant verschiedene Preise für bestimmte exotische Produkte liefern können, auch wenn die Marktpreise von Caps, Swaptions und CMS-Spread-Optionen von allen Modellen nahezu gleich gut getroffen werden.

²¹Principal component analysis (Hauptkomponentenanalyse).

INDEX

170

Index

Symbols

ρ_{∞} -extension		107 f.,	148
----------------------------	--	---------	-----

Α

abcd-parameterization
accrual factor
American option
annuity
factor
measure
arbitrage opportunity

В

Bachelier model
basis-point volatility
Bermudan option
Bermudan swaption
Black formula
Black model
bond
discount
zero-coupon
bootstrapping14
break-even rate $\ldots \ldots 13$
Breeden-Litzenberger formula 24, 41
Bromwich integral
Bromwich line
Brownian motion7
geometric19

С

callability
cap18, 43
at-the-money
CMS spread ratchet
CMS spread sticky 137
implied volatility 20
caplet
CEV LMM
CEV model 25, 45
change-of-numeraire technique $\dots 9, 61$

characteristic function32
Cholesky decomposition104
CIR process
distribution
CMS cap/floor
CMS caplet/floorlet
CMS convexity adjustment40
CMS options
CMS rate
forward40
CMS spread
call option
caplet
coupon
note
option
sensitivities 125
ratchet cap
sticky cap
volatility
complete market
constant maturity swap (CMS)
contingent claim
continuously-compounded spot rate 10
convexity adjustment
convexity-adjusted forward spread
copula
Gaussian 50
correlation
effective 78
historical 98 110
instantaneous
term 78
correlation matrix 98
general construction principle 105
stylized facts 100
correlation parameterization
3-parametric Rebonato form L 101
3-parametric Rebonato form II 102
4-parametric form
5-parametric form
classical exponential
fit to historical correlations 110

SC 2-parametric form I	103
SC 2-parametric form II	103
coupon	
CMS spread	. 47
exotic	51
Cox-Ingersoll-Ross model	53

D

Ε

effective skew
effective volatility
eigenvalue zeroing116
equivalent martingale measure8, 15
Euribor
European options
exotic coupon
exotic swap
-

F

factor loadings
Feller condition
filtration
fixing date
flatteners 47
floaters 51
floor19, 43
implied volatility
floorlet
forward Libor correlations
stylized facts 100
forward Libor rate 11

log-normally distributed
martingale in forward measure17
forward rate
instantaneous11
simply-compounded11
sliding98
forward-measure approach 59
forward-rate agreement (FRA)10
Fourier option pricing
future

G

Gaussian	$\operatorname{cumulative}$	distribution		•	 . 20
Gaussian	$\mathrm{model}\ldots\ldots$. 56

н

hedging portfolio	$\dots 42$
Heston model	31
historical correlations	, 117
HJM drift condition	57
HJM framework	57

T

implied volatility	20
instantaneous Libor correlations	69
integrated variance	79
inverse floaters	51
Itô integral	. 7
Itô process	. 7
iterated expectation	79

L

INDEX

м

main calibration
market model
market-implied density24
martingale measure
natural
mean-reversion level
measure
annuity
risk-neutral16
spot Libor
swap
T-forward
terminal
moment explosion
explosion time
multi-factor model

Ν

no-call period	. 52
Normal model	, 48
Normal volatility	.48
numeraire	8

0

one-factor model55
option
Bermudan
digital $\dots 51$

Ρ

parameter averaging
path-dependence52
payer swap18
payment date $\dots \dots \dots$
pre-calibration
pricing measure $\dots \dots \dots \dots \dots \dots \dots \dots \dots 9, 15$
principal amount52
principal component analysis $(PCA) \dots 115$
modified 116
probability space7

R

range accrual	52
$rank\ reduction\ problem\ \dots\dots\ 1$	16
rank- d approximation1	16
risk-neutral measure16,	54
risk-neutral valuation formula	16

S

scalar Libor volatility	69
Schoenmakers-Coffey	
semi-parametric family 1	02
shifted log-normal process	27
short rate	12
short-rate model16,	53
multi-factor	56
sigma algebra	. 7
simply-compounded forward rate	11
simply-compounded spot rate	10
skew function	29
sliding forward rates	98
speed of mean reversion	72
spot Libor measure	62
spot rate	
continuously-compounded	10
simply-compounded	10
spot-starting	13
spread option formula	.46
spread volatility	48
spread-option formula	88
square-root diffusion	66
static replication	52
formula	41
steepener note	37
stochastic differential equation (SDE)7.	16
stochastic volatility model	45
stochastic-volatility LMM	66
structured note	52
swap	12
exotic	51
fixed-for-floating 12.	39
forward rate	13
legs	12
option	21
paver	21
rate	13
volatility	22
receiver 12	21
swap market model	64
swap market model	18
swap rate	10
drift term under forward measure	76
forward	13
martingale in swap measure	18
spot	13
swaption 91	42
Bermudan	51
cash-settled	22
implied volatility	22
implica foraulity	

payer	21
receiver	.21
tenor	21

т

T-forward measure
Talbot's contour
tenor 10
structure
term-structure model
$endogenous \dots 54$
exogenous $\dots \dots 55, 57$
one-factor $\dots 55$
terminal measure
trading strategy8
$admissible \dots \dots 8$
replicating8
self-financing
two-factor Gaussian model

U

usual conditions7	,
-------------------	---

v

value process
vanilla model 48
Vasicek model53
volatility
basis-point
cube
Normal
spread
surface24, 48
volatility bootstrapping 20
volatility of variance
volatility skew
volatility smile 23, 48

w

Wiener process.		•	•				•		•	•			•	•										7
-----------------	--	---	---	--	--	--	---	--	---	---	--	--	---	---	--	--	--	--	--	--	--	--	--	---

Υ

year fraction	10
yield curve bootstrap	98
yield enhancement	53

Ζ

zero	bond					•	•									•						•										9)
------	------	--	--	--	--	---	---	--	--	--	--	--	--	--	--	---	--	--	--	--	--	---	--	--	--	--	--	--	--	--	--	---	---

INDEX

175

Ich versichere hiermit, dass ich die Arbeit selbständig angefertigt habe und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt, sowie die wörtlich oder inhaltlich übernommenen Stellen als solche kenntlich gemacht habe.

Ulm, Februar 2011

Matthias Lutz