



ANALYSIS OF THE HEDGING PERFORMANCE OF SELECT TERM STRUCTURE MODELS IN A REAL WORLD EVOLUTION FRAMEWORK

MARCEL ANDRE CRESNIK

A thesis submitted in partial fulfillment of the

Master of Science

in

Mathematics

Supervisor:

PD Dr. Volkert Paulsen

Münster, August 2015

Eidesstattliche Erklärung

Hiermit versichere ich, Marcel Andre Cresnik, dass ich die vorliegende Arbeit selbstständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel verwendet habe. Sinngemäß oder wörtlich Übernommenes habe ich durch Angabe von Herkunft oder Anmerkung kenntlich gemacht.

Marcel Andre Cresnik, Münster, 31. August 2015

Acknowledgments

I am extremely grateful to Dr. Riccardo Rebonato for his original proposal of considering this particular subject of research, and even more so for providing continued guidance and highly valuable advice throughout the writing of this thesis.

I would like to thank my supervisor PD Dr. Volkert Paulsen for giving me the opportunity to write this thesis and for his trust and encouragement with regard to this work and beyond.

Furthermore, my gratitude goes to zeb.information.technology GmbH & Co. KG, and to Christoph Moll in particular for his continued assistance and helpful advice throughout.

Contents

1	Introduction	1
2	Preliminaries	5
2.1	Interest Rates and Bonds	5
2.2	Interest Rate Options	10
2.3	Hedging Fundamentals	14
3	The Term Structure Models	17
3.1	Motivation	17
3.2	The Hull-White Model	17
3.2.1	Calibration to the Yield Curve and Cap Prices	18
3.2.2	Calibration to Swaption Prices	27
3.3	The SABR Model	29
3.3.1	Valuation of European Options	30
3.3.2	Calibration	31
3.4	The SABR-LIBOR Market Model	31
3.4.1	Parameterization of the Volatility Functions g and h	33
3.4.2	Calibration of the LMM-SABR to Caplets	35
3.4.3	Calibration to Swaptions	38
4	A Model for the Evolution of the Yield Curve and Smile Surface Under the Physical Measure	41
4.1	The Need for a Real-World Evolution Framework	41
4.2	Evolution of the Yield Curve	42
4.2.1	Motivation	42
4.2.2	Window Selection	46
4.2.3	The σ -Style Sampling Model	47
4.2.4	Calibration of the Functions σ and Σ	50
4.2.5	Modeling the Long Rate	51
4.2.6	Modeling the Remaining Tenors	53
4.2.7	Estimating the Volatility Adjustment Parameters	53
4.3	Evolving the Volatility Surface	55
4.3.1	General Idea	55
4.3.2	Parameterization of Historical Data	57
4.3.3	Evolution of σ_0^τ and ν^τ	58
4.3.4	Treatment of k_0^τ and ξ^τ	62
4.3.5	Evolution of SABR ρ	62

4.3.6	Obtaining Synthetic Implied Volatilities	64
5	Hedging	67
5.1	Methodology	67
5.1.1	Data Selection and Processing	68
5.1.2	Construction of a Delta Hedged Portfolio	71
5.1.3	Measurement of the Hedging Performance	76
5.2	Hedging Performance in the Historical Framework	79
5.3	Hedging Performance in the Synthetic Evolution Model	80
6	Conclusion and Outlook	89
	Appendix A	91
A.1	AR(1) Filters	91
A.1.1	Fitting the AR(1) Parameters	92
	References	95

Chapter 1

Introduction

Motivation

The interest rate derivatives market is the most liquid derivatives market in the world. According to the Bank of International Settlements, the total notional amount of outstanding over-the-counter interest rate contracts exceeds 500 trillion USD (as of December 2014).

The academic literature has produced a wide array of models and techniques for the valuation of interest rate derivatives and the management of interest rate risk, some of which are used by financial institutions and other practitioners.

The continuous-time term structure models form a very important class among these frameworks. By assigning risk-neutral dynamics to the yield curve (or a particular part thereof), they imply prices for interest rate derivatives that can be computed numerically or in closed form, depending on the model and the particular derivative.

The prices can be differentiated with regard to any of the model's state variables, parameters or input quantities — numerically or, in certain cases, analytically — so that each model makes an implicit prediction of how the prices of different contingent claims move in relation to one another.

Naturally, the evaluation of such models with regard to their hedging performances is of vital importance. Anyone who plans on conducting an empirical study of this kind faces the major challenge of acquiring a sufficient amount of relevant market data. For one thing, as of August 2015 only a very limited history of data is publicly available. This holds for options data in particular. And even if available, the historical time series might only cover weekly or even monthly observations. For another, due to the evident presence of regime changes, there is little benefit in basing one's choice of a hedging strategy on its empirical performance in a bygone environment reflected by outdated market data.

In this thesis we will extend and apply a method for the simulation of synthetic evolutions of the yield curve and the implied volatility surface under the physical measure. The approach is based on a window sampling mechanism originally devised by Rebonato et al. (2005), which has later been advanced by Deguillaume (2009) and adapted to the evolution of implied volatilities by Andreichenko (2011). The intention of this model is the generation of ample time series of artificial market data sharing the statistical characteristics of actual historical data used as input.

We will test the hedging performance of three separate interest rate models, both within the unmodified historical data set and within the synthetic world simulated by the evolution model. In particular, we will construct delta neutral portfolios around one-period swaptions according to the Black (1976) model, the Hull-White one-factor model put forth by Hull and White (1990), and the SABR model devised by Hagan et al. (2002). For each of these models we will construct single- as well as multi-factor hedges. The hedging strategies will be evaluated according to three metrics, based on the one-day hedging slippages, the maximum drawdown and the terminal portfolio values.

Literature Review

A small number of studies on the performance of hedging models has been conducted and published. Among these are the following:

Gupta and Subrahmanyam (2005) explore the pricing and hedging performance of several term structure models using market data on US dollar caps and floors. They find that the inclusion of a second stochastic factor improves the hedging performance in comparison to a one-factor model. Li and Zhao (2006) deal with the question of whether bonds span interest rate derivatives. In the course of their study, they investigate the hedging of interest rate caps using only zero-coupon bonds. The analysis is conducted in a framework of actual market data. They conclude that stochastic volatility and correlation are important for the successful pricing and hedging of caps, even more so than in the case of swaptions. Fan et al. (2007) ascertain the significance of the number of stochastic drivers and their associated volatility structures with regard to the pricing accuracy and hedging performance in the swaptions market, using real market data for 1998–2000. They conclude that there are notable benefits in using a multi-factor hedging model. An and Suo (2008) examine the dynamic performances of one-factor LIBOR and swap market models in hedging out-of-sample floors and swaptions using real market data for the years 1998–2004. They find that LIBOR market models outperform swap market models, and that adding a humped volatility structure does not significantly improve the models' hedging performances. Rebonato et al. (2008) study the delta and vega hedging of the SABR and SABR-LIBOR market models using real market data, and find that both models are well-specified. Another hedging study that uses actual market prices has been performed by Pietersz and Pelsser (2010) who find that Markov-functional and multi-factor market models are equally capable of hedging Bermudan swaptions. Schröter et al. (2012) study the robustness of the Black-Scholes, the Heston and the SABR model by hedging an Asian option in simulated real-world markets incorporating stochastic volatility, stochastic correlation and jumps in asset price and volatility. They find that the use of a more sophisticated model does not necessarily lead to a better hedging performance or a higher robustness to changes in the market.

Among these publications, only Schröter et al. (2012) make use of synthetic market data. All of the remaining studies rely on market data relating to a more or less limited period of time. Naturally, the generation of realistic synthetic data is a delicate undertaking. However, the yield curve evolution model developed by Rebonato et al. (2005) and extended by Deguillaume (2009) as well as Andreichenko (2011) appears to be a very promising approach. For this reason, we will — after making appropriate modifications to the procedure — employ it in the context of delta hedging interest rate swaptions, and draw a comparison to the results of hedging within a historical framework. The results are promising and encourage further research in this area. Furthermore, we complement the

above-mentioned studies since we are considering the hedging performance of models that include one- and two-factor as well as deterministic- and stochastic volatility models, using historical data that covers “normal” market states as well as the excited market conditions during the financial crisis of 2007–2008 and the European sovereign debt crisis.

Structure

This thesis is organized as follows: In chapter 2 we introduce the interest rates and derivatives used in this study, as well as some fundamental hedging terminology. In chapter 3 we present the relevant interest rate models alongside their calibration mechanisms. In particular, these are the Hull-White and SABR models which are part of the hedging study, as well as the LMM-SABR which will be an elementary part of the generation of synthetic implied volatilities. In chapter 4 the evolution model for the yield curve and implied volatility surface is derived. In chapter 5 the precise methodology of the hedging study is developed. After conducting the experiment, the results are presented and evaluated. Chapter 6 concludes this work and gives an outlook for future studies.

The implementation of all models and the eventual hedging experiment will be performed by use of the functional programming language *F#* which is part of the Microsoft .NET framework. The open-source library *Math.NET Numerics* has been used for cubic spline interpolation, and the trial version of the commercial *Extreme Optimization Numerical Libraries for .NET* has been used for the numerical evaluation of integrals, for the solution of numerical optimization problems and for numerical differentiation. All figures in this thesis have been drawn with the help of the *MathWorks MATLAB* software.

Chapter 2

Preliminaries

2.1 Interest Rates and Bonds

In order to set the ground for the rest of the thesis, we first go over the essential interest rates and derivatives. The definitions and explanations presented here can be found in Brigo and Mercurio (2006) or Andersen and Piterbarg (2010), e.g.

Interest rates reflect the cost of borrowing and lending money and therefore determine the present value of a future cash flow. For any particular currency there is a variety of commonly quoted interest rates — such as mortgage rates, deposit rates or prime borrowing rates. All interest rates reflect the credit risk associated with the borrowing party. This thesis does not deal with default risk. Hence, none of the interest rate derivatives is subject to counterparty or default risk.

One of the most important interest rates is the *London Interbank Offered Rate* — or LIBOR — which is a reference interest rate calculated for 5 currencies and 7 borrowing periods ranging from overnight to 12 months, administered by the Intercontinental Exchange. In order to determine the LIBOR rates, each business day at 11 a.m. a panel of participating major banks is asked to submit the rate at which they could presumably borrow funds on the interbank market. The quoted LIBOR rates then result from taking a trimmed average of the polling submissions.

These LIBOR rates are often used to construct the short end of the yield curve in the USD and GBP markets. In the eurozone, Euribor rates are usually used. These are reference rates computed very similarly to the LIBOR by polling large eurozone banks.

In order to obtain the remainder of the yield curve, usually forward rate agreements or quotes of interest rate swaps are used.

We will use the terms *LIBOR rate* or *forward LIBOR rate* when making a distinction between a simply compounded one-period forward rate (“forward LIBOR rate”) and a general forward swap rate.

As in most of the academic literature, we will make certain simplifying assumptions if necessary or adequate. For one thing, we will always assume that there are no bid-ask spreads.

Definition 2.1.1 (Money-market account and short-rate). Let $B(t)$ be the value of a bank account at time $t \geq 0$. Let $B(0) = 1$ and assume B evolves according to the differential equation

$$dB(t) = r(t)B(t)dt, \quad (2.1)$$

where r is a positive time-dependent function. Equation (2.1) implies

$$B(t) = \exp \left(\int_0^t r(s)ds \right).$$

We call B the *money-market account* and r the *instantaneous spot rate* or *short rate*.

The short rate is the basic state variable of the Hull-White model, which is one of the three models we will examine in our hedging experiment in chapter 5. The model will be introduced in section 3.2.

Note that the short-rate is a theoretical quantity which cannot be directly observed in the market.

Definition 2.1.2 (Zero-coupon bond). A T -maturity *zero-coupon bond* (or *discount bond*) is a contract that promises its holder to be paid the bond's face value at the time of maturity T , absent of any periodic interest payments. The price of a zero-coupon bond at time $t < T$ is denoted by $P(t, T)$. We will always assume that $P(t, T)$ refers to a bond with a face value of one unit of currency.

Furthermore, in this thesis we will only consider riskless zero-coupon bonds. Therefore, the value $P(t, T)$ refers to the time t value of one unit of currency to be received with certainty at time T . Hence, $P(t, T)$ represents a discount factor.

By $P(t, T_1, T_2)$ we will denote the *forward price* of the zero-coupon bond spanning $[T_1, T_2]$, i.e., the price fixed at time t for the purchase of a zero-coupon bond $P(T_1, T_2)$ to be received at a future time $T_1 \geq t$. In the assumed absence of arbitrage, it can be unambiguously computed as

$$P(t, T_1, T_2) = \frac{P(t, T_2)}{P(t, T_1)}.$$

Absent of arbitrage, $P(t, T)$ satisfies the following conditions:

- The mapping $t \mapsto P(t, T)$ is monotonically increasing,
- the mapping $T \mapsto P(t, T)$ is monotonically decreasing and
- $P(T, T) = 1$ for all T .

The graph of the function

$$T \mapsto P(t, T), \quad T \geq t$$

is called the *zero-coupon bond curve at time t* or the *term structure of discount factors*.

Definition 2.1.3 (Simply-compounded spot rate). For a time interval $[t, T]$, where t is the present, the *simply-compounded spot rate* is defined as

$$L(t, T) := \frac{1}{T - t} \left(\frac{1}{P(t, T)} - 1 \right), \quad (2.2)$$

or equivalently via the relation

$$(1 + (T - t)L(t, T))P(t, T) = 1.$$

Thus, $L(t, T)$ is the interest rate corresponding to the zero-coupon bond $P(t, T)$ under the assumption that no compounding takes place, so that the accrual of interest occurs at a constant rate and is hence proportional to the duration of the investment.

In terms of the simply-compounded spot rate, the price of a zero-coupon bond can be expressed as:

$$P(t, T) = \frac{1}{1 + (T - t)L(t, T)}.$$

There is an obvious relation between the simply-compounded spot rate and the short rate introduced in definition 2.1.1. For one unit of currency, one can obtain $1/P(t, T)$ zero-coupon bonds, so that the money market account $B(t)$ evolves according to:

$$\frac{1}{P(t, T)} = 1 + (T - t)L(t, T),$$

and thus earns an instantaneous interest rate of

$$r(t) = \lim_{T \searrow t} L(t, T).$$

Definition 2.1.4. [Forward rate agreement] A *forward rate agreement (FRA)* is a derivative which allows to lock in at the present time t the interest rate over a future time period $[T_1, T_2]$. The buyer of the FRA agrees to pay a fixed rate K on a deposit, the so-called *notional amount*, in exchange for a floating-rate payment based on a *reference rate*, which in this case will be the spot rate $L(T_1, T_2)$ resetting at T_1 .

When the notional amount is equal to one unit of currency, the time T_2 value of the contract will be

$$(T_2 - T_1)(L(T_1, T_2) - K).$$

Using eq. (2.2), this can be rewritten as

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

Since one unit of currency to be received at time T_2 is worth $P(T_1, T_2)$ units of currency at time T_1 , and in turn one unit of currency at time T_1 is worth $P(t, T_1)$ units of currency at time t , multiplying the above expression by $P(t, T_2) = P(t, T_1)P(T_1, T_2)$ yields the present value of the FRA:

$$\begin{aligned} V_{\text{FRA}}(t, T_1, T_2, K) &= P(t, T_1) - P(t, T_2) - P(t, T_2)(T_2 - T_1)K \\ &= (T_2 - T_1)P(t, T_2) \left(\frac{1}{T_2 - T_1} \left(\frac{P(t, T_1)}{P(t, T_2)} - 1 \right) - K \right). \end{aligned}$$

Definition 2.1.5 (Simply-compounded forward rate). The (*simply-compounded*) *forward rate* $f(t, T_1, T_2)$ is defined as the value for K such that the value of the corresponding FRA is rendered zero, i.e.,

$$f(t, T_1, T_2) = \frac{1}{T_2 - T_1} \left(\frac{P(t, T_1)}{P(t, T_2)} - 1 \right). \quad (2.3)$$

Thus, $f(t, T_1, T_2)$ specifies the fixed interest rate for the time period $[T_1, T_2]$ that can be secured at time $t \leq T_1$ at no additional cost.

The length of the period $[T_1, T_2]$, i.e.,

$$\tau = T_2 - T_1,$$

is called the *tenor* of the forward rate $f(t, T_1, T_2)$.

For an arbitrary set of dates

$$T_0 < T_1 < \dots < T_n,$$

note that the price $P(t, T_n)$ of a zero-coupon bond maturing at T_n can be recovered from the corresponding forward rates via the formula

$$P(t, T_n) = P(t, T_0) \prod_{i=0}^{n-1} \frac{1}{1 + (T_{i+1} - T_i) f_i(t)},$$

where we write

$$f_i(t) := f(t, T_i, T_{i+1}).$$

We call

$$f(t, T) := \lim_{\tau \rightarrow 0} f(t, T, T + \tau)$$

the time t *instantaneous forward rate* to time T , which can be thought of as the forward rate observed at time t that spans the interval $[T, T + dT]$.

There is an obvious relation between the price of a zero-coupon bond and the instantaneous forward rate:

$$P(t, T) = \exp \left(- \int_t^T f(t, s) ds \right),$$

so that

$$f(t, T) = - \frac{\partial \ln P(t, T)}{\partial T}. \quad (2.4)$$

The short rate can also be expressed in terms of the instantaneous forward rate:

$$r(t) = f(t, t) := \lim_{T \searrow t} f(t, T).$$

Definition 2.1.6 (Interest rate swap). An *interest rate swap* (or *swap*) can be considered a multi-period FRA. It is a contract in which two parties agree to exchange one stream of interest rate cash flows, based on a fixed rate, for another stream, based on a floating rate — such as the LIBOR. The streams of cash flows are called the *legs* of the swap.

Payments are made on a pre-defined set of time intervals, called *periods*. The floating rate is *fixed* at the beginning of each period, and both coupons (fixed and floating) are paid out at the end of the period. The dates when the floating rates are observed are called *fixing dates* or *reset dates*, and the dates when the cash flows are exchanged are called *payment dates*.

Various market conventions exist with regard to the day count, maturities and fixing / payment dates, depending on the currency and the underlying index. An overview of common market conventions can be found in the *OpenGamma Interest Rate Instruments*

and *Market Conventions Guide* [Henrard (2013)]. For simplicity we assume that there is no lag between fixing and payment dates.

Formally, one defines a tenor structure

$$0 \leq T_m < \dots < T_n, \quad \tau_i = T_{i+1} - T_i, \quad i = m, m+1, \dots, n-1. \quad (2.5)$$

At the end of each period $[T_i, T_{i+1}]$, the fixed rate payer pays simple interest based on a pre-defined fixed rate K , corresponding to the length of the agreed-upon time interval. The floating rate payer pays simple interest based on the LIBOR rate for the given period, as observed at time T_i . From the perspective of the fixed rate payer, at each payment date T_{i+1} the following net cash flow is exchanged:

$$\tau_i(f_i(T_i) - K), \quad i = m, \dots, n-1.$$

By viewing the swap as a portfolio of FRAs, its value at time $t \leq T_0$ can be calculated as

$$V_{\text{Swap}}(t) = \sum_{i=m}^{n-1} V_{\text{FRA}}(t, T_i, T_{i+1}, K) \quad (2.6)$$

$$= \sum_{i=m}^{n-1} \tau_i P(t, T_{i+1})(f_i(t) - K). \quad (2.7)$$

In particular, this shows that the swap's value depends only on the current level $f_i(t)$ of each forward rate, and is independent of the interest rate dynamics.

We note that eq. (2.7) can alternatively be expressed as

$$V_{\text{Swap}}(t) = P(t, T_m) - P(t, T_n) - K A^{m,n}(t)$$

where $A^{m,n}(t)$ is the so-called *annuity of the swap*

$$A^{m,n}(t) = \sum_{i=m}^{n-1} \tau_i P(t, T_{i+1}).$$

Definition 2.1.7 (Forward swap rate). The forward swap rate $S^{m,n}(t)$ is the value for the fixed rate K such that the value of the corresponding interest rate swap is rendered zero, i.e.,

$$\begin{aligned} S^{m,n}(t) &= \frac{\sum_{i=m}^{n-1} \tau_i P(t, T_{i+1}) f_i(t)}{\sum_{i=m}^{n-1} \tau_i P(t, T_{i+1})} \\ &= \frac{P(t, T_m) - P(t, T_n)}{\sum_{i=m}^{n-1} \tau_i P(t, T_{i+1})} \\ &= \frac{P(t, T_m) - P(t, T_n)}{A^{m,n}(t)}. \end{aligned} \quad (2.8)$$

2.2 Interest Rate Options

We will now introduce the two main interest rate options used in this thesis and describe their valuation.

Definition 2.2.1 (Interest rate cap/floor). A *caplet* is a European call option on a simply compounded forward rate for a given time period with the characteristic property that the rate is fixed at the beginning of the period but the option payoff will be paid out at the end of the period. The underlying interest rate could be the LIBOR rate $f_i(t)$, e.g. Given the strike rate K , a unit notional caplet of this kind would pay

$$\tau_i (f_i(T_i) - K)^+$$

at time T_{i+1} , where again $\tau_i = T_{i+1} - T_i$. Since the LIBOR rate f_i fixes at time T_i , the caplet payoff is deterministic from this point in time onward. We will call T_i the *expiry time / date* of the caplet.

Similarly, a *floorlet* is a European put option on an interest rate. In the above setting, the holder would receive a payment of

$$\tau_i (K - f_i(T_i))^+$$

at time T_{i+1} .

A *cap* is a collection of caplets on successive forward rates, all having the same strike rate K . E.g., if we are still using the tenor structure given in eq. (2.5), the corresponding cap would comprise $(n - m - 1)$ individual caplets on LIBOR rates f_i and pay

$$\tau_i (f_i(T_i) - K)^+$$

at times T_{i+1} , for all $i \in \{m, \dots, n - 1\}$.

Analogously, a *floor* is a collection of floorlets.

The holder of a cap is guaranteed an upper limit on the corresponding interest rate (hence the terminology). E.g., assume that an institution is indebted and is obliged to pay the LIBOR rates $\tau_i f_i(T_i)$ at the end of each period $[T_i, T_{i+1}]$, where $i = m, \dots, n - 1$. If the institution holds a cap with corresponding notional, at each payment date T_{i+1} the overall interest rate paid equals

$$\tau_i (f_i(T_i) - (f_i(T_i) - K)^+) = \tau_i \min\{f_i(T_i), K\}.$$

The market price of a cap / caplet is usually quoted as a Black implied volatility $\hat{\sigma}$. At time 0 the implied volatility of the (unit notional) cap considered above is the value for $\hat{\sigma}$ so that the cap's market price equals the following sum of Black's formulas:

$$V_{\text{Cap}}(0) = \sum_{i=m}^{n-1} P(0, T_{i+1}) \tau_i \text{Bl}^+(K, f_i(0), T_i, \hat{\sigma}), \quad (2.9)$$

where

$$\text{Bl}^+(K, F, T, \sigma) = F\Phi(d_1(K, F, T, \sigma)) - K\Phi(d_2(K, F, T, \sigma)) \quad (2.10)$$

$$d_1(K, F, T, \sigma) = \frac{\ln(\frac{F}{K}) + \frac{\sigma^2 T}{2}}{\sigma\sqrt{T}} \quad (2.11)$$

$$d_2(K, F, T, \sigma) = \frac{\ln(\frac{F}{K}) - \frac{\sigma^2 T}{2}}{\sigma\sqrt{T}}, \quad (2.12)$$

and where Φ denotes the cumulative distribution function of the standard normal distribution. Note that in eq. (2.9) it is assumed that all individual caplets contained in the cap share the same Black volatility.

Similarly, a floor would be priced by replacing the term Bl^+ in eq. (2.9) with

$$\text{Bl}^- = -\text{Bl}^+.$$

Definition 2.2.2. A cap (floor) with payment dates as given above is said to be *at-the-money (ATM)* at time $t < T_m$ if its strike rate satisfies

$$K = K_{\text{ATM}} = S^{m,n}(t).$$

The cap is said to be *in-the-money (ITM)* if $K < K_{\text{ATM}}$, and *out-of-the-money (OTM)* if $K > K_{\text{ATM}}$. The opposite holds for a floor.

Note that for a single caplet on a LIBOR rate f_m , the ATM strike K_{ATM} equals

$$S^{m,m+1}(t) = \frac{P(t, T_m) - P(t, T_{m+1})}{\tau_m P(t, T_{m+1})} = f_m(t).$$

Following Brigo and Mercurio (2006) we will now see that a cap (floor) can actually be regarded as a portfolio of put (call) options on zero-coupon bonds. This representation will come in useful in the context of calibrating the Hull-White model in section 3.2.

Consider once again the cap based on the tenor structure given in eq. (2.5), i.e., a cap consisting of $(n - m - 1)$ unit notional caplets on LIBOR rates f_i paying

$$\tau_i(f_i(T_i) - K)^+ \quad (2.13)$$

at times T_{i+1} where $i = m, \dots, n - 1$. Denote by

$$V_{\text{Cpl}}(t, i, K)$$

the time t value of the caplet resetting at time $T_i > t$. The caplet payoff, eq. (2.13), to be received at T_{i+1} is already fixed at T_i and is thus equivalent to receiving

$$P(T_i, T_{i+1})\tau_i(f_i(T_i) - K)^+$$

at T_i . This can be rewritten as

$$\begin{aligned} \left(1 - \frac{1 + \tau_i K}{1 + \tau_i f_i(T_i)}\right)^+ &= (1 - (1 + \tau_i K)P(T_i, T_{i+1}))^+ \\ &= (1 + \tau_i K) \left(\frac{1}{1 + \tau_i K} - P(T_i, T_{i+1})\right)^+, \end{aligned}$$

which is clearly the payoff of $1 + \tau_i K$ unit notional put options with strike $1/(1 + \tau_i K)$ on the forward price $P(t, T_i, T_{i+1})$ of the zero-coupon bond spanning $[T_i, T_{i+1}]$. Therefore, the caplet is equivalent to a bond put option:

$$V_{\text{Cpl}}(t, i, K) = (1 + \tau_i K) V_{\text{Put}}(t, K', P(t, T_i, T_{i+1}), T_i),$$

where

$$K' = \frac{1}{1 + \tau_i K}.$$

The cap price can be obtained by summation of the underlying caplet prices:

$$V_{\text{Cap}}(t) = \sum_{i=m}^{n-1} (1 + \tau_i K) V_{\text{Put}}(t, K', P(t, T_i, T_{i+1}), T_i). \quad (2.14)$$

Analogously, the price of the floorlet resetting at T_i corresponds to

$$V_{\text{Fl}}(t, i, K) = (1 + \tau_i K) V_{\text{Call}}(t, K', P(t, T_i, T_{i+1}), T_i)$$

so that the price of a floor with identical underlying tenor structure equals

$$V_{\text{Floor}}(t) = \sum_{i=m}^{n-1} (1 + \tau_i K) V_{\text{Call}}(t, K', P(t, T_i, T_{i+1}), T_i). \quad (2.15)$$

Proposition 2.2.3 (Put-call parity for caps and floors). *The following parity holds:*

$$V_{\text{Cpl}}(t, i, K) - V_{\text{Fl}}(t, i, K) = V_{\text{FRA}}(t, T_i, T_{i+1}, K).$$

Thus, the value of a cap minus the value of a floor is equal to the forward price of a swap:

$$V_{\text{Cap}}(t) - V_{\text{Floor}}(t) = V_{\text{Swap}}(t).$$

Proof. A simple no-arbitrage argument with regard to the payoff at time T_i gives the regular put-call parity for bond options

$$\begin{aligned} &V_{\text{Put}}(t, K', P(t, T_i, T_{i+1}), T_i) \\ &= V_{\text{Call}}(t, K', P(t, T_i, T_{i+1}), T_i) - P(t, T_{i+1}) + K' P(t, T_i), \end{aligned}$$

which in turn shows that

$$\begin{aligned}
V_{\text{Cpl}}(t, i, K) &= (1 + \tau_i K) V_{\text{Put}}(t, K', P(t, T_i, T_{i+1}), T_i) \\
&= (1 + \tau_i K) [V_{\text{Call}}(t, K', P(t, T_i, T_{i+1}), T_i) - P(t, T_{i+1}) + K' P(t, T_i)] \\
&= V_{\text{Fl}}(t, i, K) - (1 + \tau_i K) P(t, T_{i+1}) + P(t, T_i) \\
&= V_{\text{Fl}}(t, i, K) + V_{\text{FRA}}(t, T_i, T_{i+1}, K).
\end{aligned}$$

The second part follows from eq. (2.6). \square

We will now introduce the most important option in this thesis.

Definition 2.2.4 (Interest rate swaption). A *swap option* or *swaption* is an option that gives the owner the right to enter into an interest rate swap at a given future date at a pre-determined rate, called the *strike rate* or *fixed rate* of the swaption. A *payer swaption* is the option to pay the fixed leg, and a *receiver swaption* is the option to receive the fixed leg on the swap.

In most cases, a swaption expires at the first reset date of the underlying swap. The set of reset and payment dates is called the *tenor structure* of the swaption.

Consider the payer swaption with strike K corresponding to the tenor structure given in eq. (2.5), where T_m is both the option maturity and the first reset date. At time T_m the swaption will pay

$$\begin{aligned}
V_{\text{PSwaption}}(T_m) &= (V_{\text{Swap}}(T_m))^+ \\
&= \left(\sum_{i=m}^{n-1} \tau_i P(T_m, T_{i+1}) (f_i(T_m) - K) \right)^+. \tag{2.16}
\end{aligned}$$

In contrast to the case of an interest rate cap, this value cannot be decomposed into a sum of terms where each addend depends only on a single LIBOR rate. Instead, a swaption's payoff — and consequently its value — depends on the joint evolution of the LIBOR rates contained in the underlying swap.

The payoff of a swaption can also be written in terms of the swap rate and the swap annuity. In particular, eq. (2.16) can be re-expressed as

$$\begin{aligned}
V_{\text{PSwaption}}(T_m) &= \sum_{i=m}^{n-1} \tau_i P(T_m, T_{i+1}) (S^{m,n}(T_m) - K)^+ \\
&= A^{m,n}(T_m) (S^{m,n}(T_m) - K)^+. \tag{2.17}
\end{aligned}$$

Definition 2.2.5. A payer or receiver swaption with maturity T_m written on a swap with reset dates T_m, \dots, T_{n-1} and payment dates T_{m+1}, \dots, T_n is said to be *at-the-money (ATM)* at time $t < T_m$ if its strike price satisfies

$$K = K_{\text{ATM}} = S^{m,n}(t).$$

The payer swaption is said to be *in-the-money (ITM)* if $K < K_{\text{ATM}}$, and *out-of-the-money (OTM)* if $K > K_{\text{ATM}}$. The reverse holds for a receiver swaption.

Similar to caps and floors, the price of a swaption is usually given in terms of a Black implied volatility $\hat{\sigma}_{m,n}$. At time 0, the implied volatility of the (unit notional) payer swaption considered above is the value for $\hat{\sigma}_{m,n}$ so that the payer swaption's market price is equal to

$$V_{\text{PSwaption}}(0) = A^{m,n}(0) \text{Bl}^+(K, S^{m,n}(0), T_m, \hat{\sigma}_{m,n}). \quad (2.18)$$

For a receiver swaption, the corresponding pricing equation reads

$$V_{\text{RSwaption}}(0) = A^{m,n}(0) \text{Bl}^-(K, S^{m,n}(0), T_m, \hat{\sigma}_{m,n}). \quad (2.19)$$

We note that no immediate comparison between Black implied volatilities of caps and swaptions can be made, since cap implied volatilities refer to forward rates while swaption implied volatilities apply to (forward) swap rates. Instead, in this context Black implied volatilities serve only as a practical way of communicating prices.

We will now consider a different interpretation of a swaption which will simplify its valuation under the Hull-White model. Using the second part of eq. (2.8), the swaption payoff given by eq. (2.17) can be re-written as

$$\begin{aligned} V_{\text{PSwaption}}(T_m) &= \sum_{i=m}^{n-1} \tau_i P(T_m, T_{i+1}) \left(\frac{1 - P(T_m, T_n)}{\sum_{i=m}^{n-1} \tau_i P(T_m, T_{i+1})} - K \right)^+ \\ &= \left(1 - P(T_m, T_n) - \sum_{i=m}^{n-1} K \tau_i P(T_m, T_{i+1}) \right)^+ \\ &= \left(1 - \sum_{i=m}^{n-1} c_{i+1} P(T_m, T_{i+1}) \right)^+, \end{aligned} \quad (2.20)$$

where

$$c_i = \begin{cases} K \tau_i & \text{for } i = m+1, \dots, n-1 \\ 1 + K \tau_i & \text{for } i = n. \end{cases}$$

The term $\sum_{i=m}^{n-1} c_i P(T_m, T_{i+1})$ is simply the time T_m price of a coupon-bearing bond paying c_i at times T_i , $i = m+1, \dots, n$. A payer swaption can thus be viewed as a put option on the forward price of a coupon-bearing bond. In the same way, one can see that a receiver swaption is equivalent to a call option on a coupon bearing bond. This relationship will later be used to calibrate the Hull-White model to market prices of swaptions.

2.3 Hedging Fundamentals

A financial institution that engages in over-the-counter options trading encounters the problem of managing risk. If the institution sells to their customer a complex derivative that is not actively traded on an exchange, they cannot simply eliminate the associated risk by taking the opposite position in the market.

Instead, the exact factors of risk related to an option will have to be identified so they can be neutralized by taking positions in more liquidly traded products, such as plain vanilla derivatives. For this purpose, the so-called *Greeks* are introduced. Further details can be found in Hull (2012).

Definition 2.3.1 (The Greeks). Let Π be a portfolio of derivatives depending on a single asset whose price is S . The *Greeks* (or *risk sensitivities* or *hedge parameters*) are the quantities representing the sensitivity of the portfolio value with regard to a change in the parameters used to arrive at this value. Some of the most commonly quoted Greeks are defined as follows:

The *delta* describes the rate by which the portfolio's value changes as the value of the underlying changes:

$$\Delta = \frac{\partial \Pi}{\partial S}.$$

Similarly, the *gamma* of the portfolio,

$$\Gamma = \frac{\partial^2 \Pi}{\partial S^2},$$

represents the rate of change in the value with regard to the delta, the *theta*

$$\Theta = \frac{\partial \Pi}{\partial t}$$

defines the rate of change with regard to the passage of time, and the *vega*

$$\mathcal{V} = \frac{\partial \Pi}{\partial \sigma}$$

is the rate of change with regard to a change in the volatility σ of the portfolio.

Note that this is not a complete enumeration. One can also consider further second (or even third) order sensitivities.

A portfolio is called *delta-neutral* [or *gamma- / theta- / vega-neutral*] if its delta [or gamma / theta / vega] is zero. In this case the portfolio value will be *locally* immune to a change in the respective underlying parameter.

Depending on the particular derivative under consideration, some of the Greeks might be given in a model-independent way. Usually, especially in the context of options trading, the above sensitivities need to be computed by use of a model ascribing dynamics to the underlying.

Furthermore, the above are only *local* quantities. Even if one knew the exact process for the underlying as well as the exact dependencies of the portfolio value on all parameters, one could not construct a totally riskless portfolio without the ability to re-hedge continuously. Since in reality trading is occurring in discrete time, a certain “slippage” cannot be avoided.

In this thesis we will focus on the delta, which is the most important first order dependency. In particular, we will — according to several models — construct a delta neutral portfolio around various one-period swaptions, and analyze the respective slippage under each of the strategies.

In-model and out-of-model hedging

The practice of hedging against uncertainty captured by the model at consideration is called *in-model hedging*. The practice of hedging against risks that are non-existent according to the model is called *out-of-model hedging*.

A typical example of in-model hedging is a delta hedge within the Black-Scholes framework. In this setting, the price of the underlying is a “moving part” in the model and the Black-Scholes formula predicts by how much the value of a European call or put option will change relative to a change in the value of the underlying. A delta hedge simply addresses and neutralizes this model-predicted source of uncertainty. A vega-hedge on the other hand is an example of an out-of-model hedge, since the volatility of the underlying is assumed to be a static parameter in the Black-Scholes framework.

The distinction between in-model and out-of-model hedging is essential when analyzing the plausibility and predictive qualities of a particular model. As Gupta and Subrahmanyam (2005) point out, an out-of-model-hedge is, naturally, inconsistent with the predictions made by the model and is therefore not a useful indicator when assessing the accuracy of said predictions.

In our study we will only hedge the delta exposure, and hence we will stay in the realm of in-model hedging. However, the approach can easily be extended to the hedging against risk in any of the other parameters or state variables of a model, and beyond¹.

¹Rebonato et al. (2009) describe the concept of functional-dependence hedging which aims at capturing uncertainty with regard to sources of risk that a model does not immediately know about. In this case hedge ratios are computed by numerically bumping input quantities to the model, such as prices of vanilla options used in the process of calibrating the model, rather than model parameters / state variables.

Chapter 3

The Term Structure Models

3.1 Motivation

Among practitioners, derivatives models are used for various purposes. Plain-vanilla-options traders, e.g., can easily access market prices for the options they are working with, but need to rely on the predictions made by a model when determining how option prices move in relation to the underlying or to one another. Hence, they rely on the *prescriptive* capabilities of the model. Complex-derivatives traders on the other hand depend on a model's *descriptive* abilities to arrive even at the prices of the products they are considering [Rebonato (2004)].

The models capable of fulfilling these tasks come in various complexities. The Black (1976) model is a very simple (but robust) representative of this class. The relevant formulas have been introduced in section 2.2. In this chapter we introduce an additional three interest rate models that are relevant to this thesis. Among them are the Hull-White model and the SABR model, which will be part of our hedging experiment in chapter 5, as well as the LMM-SABR model which will be part of the evolution model outlined in chapter 4.

The Hull-White model is a popular one-factor short rate model and has been chosen as a classical representative for this type of model. The SABR model is a two-factor stochastic-volatility forward rate model developed by Hagan et al. (2002), which has shown very promising hedging results in empirical studies (see, e.g., Rebonato et al. (2008)). The LMM-SABR model, which has been put forth by Rebonato (2007), is a stochastic-volatility extension of the popular deterministic-volatility multi-factor LIBOR Market Model. Its ability to model multiple forward rates within a single framework while being able to recover the prices of plain-vanilla-options, including smile effects, to a very high degree will be paramount to the simulation of artificial market scenarios in chapter 4.

3.2 The Hull-White Model

The *one-factor Hull-White model*, also known as the *extended Vasicek model*, was originally proposed by Hull and White (1990), and was intended as an extension of the term structure

model put forward by Vasicek (1977).

It is a short-rate model under which the short-rate r evolves according to the stochastic differential equation

$$dr(t) = (\vartheta(t) - a(t)r(t)) dt + \sigma(t)dW(t) \quad (3.1)$$

under the risk-neutral measure \mathbb{Q} . Here, W is a standard \mathbb{Q} -Brownian motion, and a , σ and ϑ are time-dependent functions. We will focus on the case where a and σ are positive constants, so that eq. (3.1) simplifies to

$$dr(t) = (\vartheta(t) - ar(t)) dt + \sigma dW(t). \quad (3.2)$$

Under these dynamics, the short rate $r(t)$ reverts with speed a to the time-dependent level $\vartheta(t)/a$.

3.2.1 Calibration to the Yield Curve and Cap Prices

In order to specify the model, we need to set the parameters a and σ as well as the time-dependent function $\vartheta(t)$ for all $t \geq 0$.

This process is called the *calibration* of the model. In this section we follow the framework laid out in Brigo and Mercurio (2006) and add a few of the omitted arguments and derivations.

The general idea will be to calibrate a and σ to market prices of interest rate caps or swaptions, and subsequently choose $\vartheta(t)$ so as to recover the current term structure of interest rates.

Proposition 3.2.1. *For any $s < t$, a solution to the stochastic differential equation (3.2) is given by*

$$r(t) = e^{-a(t-s)}r(s) + \int_s^t e^{-a(t-u)}\vartheta(u)du + \sigma \int_s^t e^{-a(t-u)}dW(u).$$

In particular, conditional on \mathcal{F}_s , $r(t)$ is normally distributed with expectation

$$\mathbb{E}^{\mathbb{Q}}[r(t)|\mathcal{F}_s] = e^{-a(t-s)}r(s) + \int_s^t e^{-a(t-u)}\vartheta(u)du$$

and variance

$$\mathbb{V}^{\mathbb{Q}}[r(t)|\mathcal{F}_s] = \frac{\sigma^2}{2a} \left(1 - e^{-2a(t-s)}\right).$$

Proof. We will use variation of constants. Consider the homogeneous partial differential equation

$$\frac{dy(t)}{dt} = -ay(t), \quad y(0) = 1,$$

which solves for

$$y(t) = e^{-at}.$$

Let $r(t)$ be a solution of eq. (3.2). Then

$$\begin{aligned}
d\frac{r(t)}{y(t)} &= r(t)d\frac{1}{y(t)} + \frac{1}{y(t)}dr(t) \\
&= \frac{ar(t)}{y(t)}dt + \frac{1}{y(t)}(\vartheta(t) - ar(t))dt + \frac{\sigma}{y(t)}dW(t) \\
&= \frac{\vartheta(t)}{y(t)}dt + \frac{\sigma}{y(t)}dW(t) \\
&= e^{at}\vartheta(t)dt + e^{at}\sigma dW(t)
\end{aligned}$$

so that

$$r(t) = e^{-a(t-s)}r(s) + \int_s^t e^{-a(t-u)}\vartheta(u)du + \sigma \int_s^t e^{-a(t-u)}dW(u). \quad (3.3)$$

□

The time t price of a zero-coupon bond maturing at time T is given by

$$P(t, T) = \mathbb{E}^{\mathbb{Q}} \left[e^{-\int_t^T r(u)du} \middle| \mathcal{F}_t \right]. \quad (3.4)$$

In order to compute the right-hand side of eq. (3.4), we first determine the distribution of $e^{-\int_t^T r(u)du}$ conditional on \mathcal{F}_t .

Lemma 3.2.2. *For any $t < T$ and conditional on \mathcal{F}_t , the random variable*

$$\int_t^T r(u)du$$

is normally distributed with expectation

$$\mathbb{E}^{\mathbb{Q}} \left[\int_t^T r(u)du \middle| \mathcal{F}_t \right] = r(t)B(t, T) + \int_t^T \vartheta(u)B(u, T)du$$

and variance

$$\mathbb{V}^{\mathbb{Q}} \left[\int_t^T r(u)du \middle| \mathcal{F}_t \right] = \sigma^2 \int_t^T B(u, T)^2 du$$

where

$$B(t, T) := \frac{1}{a} \left(1 - e^{-a(T-t)} \right).$$

Proof. We use eq. (3.3) for $s = 0$. This yields:

$$\begin{aligned}
\int_t^T r(u) du &= r(0) \int_t^T e^{-au} du + \int_t^T \int_0^u e^{-a(u-t')} \vartheta(t') dt' du + \sigma \int_t^T \int_0^u e^{-a(u-t')} dW(t') du \\
&= r(0) \frac{e^{-at} - e^{-aT}}{a} + \int_t^T \int_0^t e^{-a(u-t')} \vartheta(t') dt' du + \int_t^T \int_t^u e^{-a(u-t')} \vartheta(t') dt' du \\
&\quad + \sigma \int_t^T \int_0^t e^{-a(u-t')} dW(t') du + \sigma \int_t^T \int_t^u e^{-a(u-t')} dW(t') du \\
&= r(0) \frac{e^{-at} - e^{-aT}}{a} + \int_0^t \vartheta(t') \int_t^T e^{-a(u-t')} du dt' + \int_t^T \vartheta(t') \int_{t'}^T e^{-a(u-t')} du dt' \\
&\quad + \sigma \int_0^t \int_t^T e^{-a(u-t')} du dW(t') + \sigma \int_t^T \int_{t'}^T e^{-a(u-t')} du dW(t') \\
&= r(0) \frac{e^{-at} - e^{-aT}}{a} + \int_0^t \vartheta(t') \frac{e^{-a(t-t')} - e^{-a(T-t')}}{a} dt' \\
&\quad + \int_t^T \vartheta(t') \frac{1 - e^{-a(T-t')}}{a} dt' + \sigma \int_0^t \frac{e^{-a(t-t')} - e^{-a(T-t')}}{a} dW(t') \\
&\quad + \sigma \int_t^T \frac{1 - e^{-a(T-t')}}{a} dW(t') \\
&= r(0) e^{-at} B(t, T) + e^{-at} B(t, T) \int_0^t \vartheta(t') e^{at'} dt' + \int_t^T \vartheta(t') B(t', T) dt' \\
&\quad + \sigma e^{-at} B(t, T) \int_0^t e^{at'} dW(t') + \sigma \int_t^T B(t', T) dW(t') \\
&= r(t) B(t, T) + \int_t^T \vartheta(t') B(t', T) dt' + \sigma \int_t^T B(t', T) dW(t').
\end{aligned}$$

□

Lemma 3.2.2 implies that, conditional on \mathcal{F}_t , the random variable $e^{-\int_t^T r(u) du}$ is lognormally distributed with expectation

$$\begin{aligned}
P(t, T) &= \mathbb{E}^{\mathbb{Q}} \left[e^{-\int_t^T r(u) du} \middle| \mathcal{F}_t \right] \\
&= e^{\mathbb{E}^{\mathbb{Q}} \left[-\int_t^T r(u) du \middle| \mathcal{F}_t \right] + \frac{1}{2} \mathbb{V}^{\mathbb{Q}} \left[-\int_t^T r(u) du \middle| \mathcal{F}_t \right]} \\
&= e^{-r(t)B(t, T) - \int_t^T \vartheta(u) B(u, T) du + \frac{\sigma^2}{2} \int_t^T B(u, T)^2 du}.
\end{aligned} \tag{3.5}$$

Proposition 3.2.3. *For*

$$\vartheta(t) = \frac{\partial f(0, t)}{\partial T} + af(0, t) + \frac{\sigma^2}{2a} (1 - e^{-2at}) \tag{3.6}$$

the current term structure of interest rates, given by the market instantaneous forward rates

$$f(0, T) = -\frac{\partial \ln P(0, T)}{\partial T},$$

is recovered under the Hull-White model defined by eq. (3.2).

Proof. Consider the instantaneous forward rate

$$\begin{aligned} f(0, T) &= -\frac{\partial \ln P(0, T)}{\partial T} \\ &= \frac{\partial}{\partial T} \left(r(0)B(0, T) + \int_0^T \vartheta(t')B(t', T)dt' - \frac{\sigma^2}{2} \int_0^T B(t', T)^2 dt' \right) \\ &= r(0)e^{-aT} + \vartheta(T)B(T, T) + \int_0^T \vartheta(t')e^{-a(T-t')} dt' \\ &\quad - \frac{\sigma^2}{2} B(T, T)^2 - \sigma^2 \int_0^T B(t', T)e^{-a(T-t')} dt' \\ &= r(0)e^{-aT} + \int_0^T \vartheta(t')e^{-a(T-t')} dt' - \frac{\sigma^2}{a} \int_0^T (1 - e^{-a(T-t')}) e^{-a(T-t')} dt' \\ &= r(0)e^{-aT} + \int_0^T \vartheta(t')e^{-a(T-t')} dt' - \frac{\sigma^2}{a^2} (1 - e^{-aT}) + \frac{\sigma^2}{2a^2} (1 - e^{-2aT}). \end{aligned}$$

Differentiating with regard to the second variable yields

$$\begin{aligned} \frac{\partial f(0, T)}{\partial T} &= -ar(0)e^{-aT} + \vartheta(T) - a \int_0^T \vartheta(t')e^{-a(T-t')} dt' - \frac{\sigma^2}{a} e^{-aT} + \frac{\sigma^2}{a} e^{-2aT} \\ &= \vartheta(T) - af(0, T) - \frac{\sigma^2}{a} (1 - e^{-aT}) + \frac{\sigma^2}{2a} (1 - e^{-2aT}) - \frac{\sigma^2}{a} e^{-aT} + \frac{\sigma^2}{a} e^{-2aT} \\ &= \vartheta(T) - af(0, T) - \frac{\sigma^2}{a} + \frac{\sigma^2}{a} e^{-aT} + \frac{\sigma^2}{2a} - \frac{\sigma^2}{2a} e^{-2aT} - \frac{\sigma^2}{a} e^{-aT} + \frac{\sigma^2}{a} e^{-2aT} \\ &= \vartheta(T) - af(0, T) - \frac{\sigma^2}{2a} (1 - e^{-2aT}) \end{aligned}$$

and thus

$$\vartheta(T) = \frac{\partial f(0, T)}{\partial T} + af(0, T) + \frac{\sigma^2}{2a} (1 - e^{-2aT}).$$

□

Substituting this result into eq. (3.5) yields the following result:

Proposition 3.2.4. *At time t , the price of a zero-coupon bond maturing at T is given by*

$$P(t, T) = A(t, T)e^{-B(t, T)r(t)} \quad (3.7)$$

where

$$A(t, T) = \frac{P(0, T)}{P(0, t)} e^{B(t, T)f(0, t) - \frac{\sigma^2}{4a} (1 - e^{-2at}) B(t, T)^2}.$$

□

Note that the terms $P(0, t)$ and $P(0, T)$ in proposition 3.2.4 are given by the yield curve observed in the market at time 0.

Using the exact formula for $\vartheta(t)$, we can rewrite proposition 3.2.1 as:

Proposition 3.2.5. *For any $s < t$, a solution to the stochastic differential equation (3.2) is given by*

$$r(t) = e^{-a(t-s)}r(s) + \alpha(t) - e^{-a(t-s)}\alpha(s) + \sigma \int_s^t e^{-a(t-u)} dW(u) \quad (3.8)$$

where

$$\alpha(t) = f(0, t) + \frac{\sigma^2}{2a^2} (1 - e^{-at})^2.$$

In particular, conditional on \mathcal{F}_s , the short rate $r(t)$ is normally distributed with expectation

$$\mathbb{E}^{\mathbb{Q}}[r(t)|\mathcal{F}_s] = e^{-a(t-s)}r(s) + \alpha(t) - e^{-a(t-s)}\alpha(s)$$

and variance

$$\mathbb{V}^{\mathbb{Q}}[r(t)|\mathcal{F}_s] = \frac{\sigma^2}{2a} (1 - e^{-2a(t-s)}).$$

Proof. Let $s < t$. From proposition 3.2.1 we have

$$r(t) = e^{-a(t-s)}r(s) + \int_s^t e^{-a(t-u)}\vartheta(u)du + \sigma \int_s^t e^{-a(t-u)} dW(u).$$

Using eq. (3.6) and integrating by parts, the second term can be expressed as

$$\begin{aligned} \int_s^t e^{-a(t-u)}\vartheta(u)du &= \int_s^t e^{-a(t-u)} \left(\frac{\partial f(0, u)}{\partial u} + af(0, u) + \frac{\sigma^2}{2a} (1 - e^{-2au}) \right) du \\ &= -a \int_s^t e^{-a(t-u)} f(0, u) du + \left[e^{-a(t-u)} f(0, u) \right]_{u=s}^t \\ &\quad + a \int_s^t e^{-a(t-u)} f(0, u) du + \frac{\sigma^2}{2a} \int_s^t e^{-a(t-u)} (1 - e^{-2au}) du \\ &= f(0, t) - e^{-a(t-s)} f(0, s) + \frac{\sigma^2}{2a} \left(\left[\frac{e^{-a(t-u)}}{a} \right]_{u=s}^t - \left[-\frac{e^{-a(t+u)}}{a} \right]_{u=s}^t \right) \\ &= f(0, t) - e^{-a(t-s)} f(0, s) + \frac{\sigma^2}{2a^2} (1 - e^{-a(t-s)} + e^{-2at} - e^{-a(t+s)}) \\ &= \alpha(t) - e^{-a(t-s)}\alpha(s). \end{aligned}$$

Hence it follows, as claimed, that

$$r(t) = e^{-a(t-s)}r(s) + \alpha(t) - e^{-a(t-s)}\alpha(s) + \sigma \int_s^t e^{-a(t-u)}dW(u).$$

□

This also demonstrates that the model short rate can become negative with nonzero probability, which is usually seen as a shortcoming of the Hull-White model in most applications.¹

We still need to determine the parameters a and σ by recovering market-observed cap prices. For this, we will make use of the fact that a cap can be interpreted as a portfolio of put options written on zero-coupon bonds (see section 2.2) and use the pricing formula eq. (3.7).

The time t price of a unit notional European call option with strike K and maturity T , written on the forward price of a zero-coupon bond spanning $[T, S]$, is given by

$$V_{\text{Call}}(t, K, P(t, T, S), T) = P(t, T)\mathbb{E}^T \left[(P(T, S) - K)^+ \middle| \mathcal{F}_t \right], \quad (3.9)$$

where \mathbb{E}^T denotes the expectation under the T -forward measure \mathbb{Q}_T , i.e., the measure associated with the numéraire $P(t, T)$.

In order to compute $V_{\text{Call}}(t, K, P(t, T, S), T)$, we will first determine the distribution of r under the measure \mathbb{Q}_T .

For this purpose, we will make use of a different and more convenient representation of $r(t)$ which follows immediately from eq. (3.8) using $s = 0$, namely

$$r(t) = \alpha(t) + x(t), \quad (3.10)$$

where

$$x(t) = \sigma \int_0^t e^{-a(t-u)}dW(u). \quad (3.11)$$

Note that $\alpha(t)$ is fully deterministic, and that the stochastic term $x(t)$ is independent of the term structure at time 0.

The process $x(t)$ defined in eq. (3.11) is the solution of the stochastic differential equation

$$dx(t) = -ax(t)dt + \sigma dW(t), \quad x(0) = 0, \quad (3.12)$$

where $W(t)$ is a standard Brownian motion under the risk-neutral measure \mathbb{Q} . This can easily be confirmed by variation of constants.

We would like to determine the dynamics of $x(t)$ under the forward measure \mathbb{Q}_T . To accomplish this we make use of the change-of-numéraire technique.

¹unless of course one uses the Hull-White model to deliberately model possibly negative rates

Proposition 3.2.6. Let \mathbb{Q}^M and \mathbb{Q}^N be equivalent martingale measures with associated numéraires M and N , respectively. Assume that under a common measure \mathbb{P} equivalent to both numéraire measures, the numéraires evolve according to

$$\begin{aligned} dM(t) &= (\dots)dt + \sigma_M(t)dW_t \\ dN(t) &= (\dots)dt + \sigma_N(t)dW_t. \end{aligned}$$

Let W^M and W^N be standard Brownian motions corresponding to the measures \mathbb{Q}^M and \mathbb{Q}^N , respectively. Then W^M and W^N are related by

$$dW^M(t) = dW^N(t) - \left(\frac{\sigma_M(t)}{M(t)} - \frac{\sigma_N(t)}{N(t)} \right). \quad (3.13)$$

Proof. See Brigo and Mercurio (2006). □

Lemma 3.2.7. Under the T -forward measure \mathbb{Q}_T , the process $x(t)$ defined by eq. (3.11) evolves according to

$$dx(t) = (-B(t, T)\sigma^2 - ax(t))dt + \sigma dW^T(t), \quad x(0) = 0,$$

where W^T is a \mathbb{Q}^T -Brownian motion defined by

$$dW^T(t) = dW(t) + \sigma B(t, T)dt.$$

In particular, for $s \leq t \leq T$ it holds that

$$x(t) = e^{-a(t-s)}x(s) - M^T(s, t) + \sigma \int_s^t e^{-a(t-u)}dW^T(u), \quad (3.14)$$

where

$$M^T(s, t) = \frac{\sigma^2}{a^2} \left(1 - e^{-a(t-s)} \right) - \frac{\sigma^2}{2a^2} \left(e^{-a(T-t)} - e^{-a(T+t-2s)} \right).$$

Proof. We apply proposition 3.2.6 using $\mathbb{Q}^N = \mathbb{Q}$ and $\mathbb{Q}^M = \mathbb{Q}_T$.

Under the risk-neutral measure, we know that the bond price $P(t, T)$, which is the numéraire corresponding to \mathbb{Q}_T , follows a stochastic differential equation of the type

$$dP(t, T) = \mu_T(t)dt + \sigma_T(t)dW(t),$$

where $W(t)$ is a Wiener process under \mathbb{Q} .

Applying Itô's lemma to eq. (3.7), we see that

$$\begin{aligned} dP(t, T) &= \frac{\partial P(t, T)}{\partial t}dt + \frac{\partial P(t, T)}{\partial r}dr + \frac{1}{2} \frac{\partial^2 P(t, T)}{\partial r^2} d\langle r \rangle(t) \\ &= \frac{\partial P(t, T)}{\partial t}dt - B(t, T)P(t, T) [(\vartheta(t) - ar(t))dt + \sigma dW(t)] + \frac{1}{2} B(t, T)^2 P(t, T) \sigma^2 dt \\ &= \left[\frac{\partial P(t, T)}{\partial t} - B(t, T)P(t, T) (\vartheta(t) - ar(t)) + \frac{1}{2} B(t, T)^2 P(t, T) \sigma^2 \right] dt \\ &\quad - \sigma B(t, T)P(t, T)dW(t). \end{aligned}$$

It follows that

$$\sigma_T(t) = -\sigma B(t, T)P(t, T).$$

Under the risk neutral measure, the numéraire is given by the money market account

$$B(t) = \exp \left(\int_0^t r(s) ds \right),$$

which is clearly driftless under \mathbb{Q} .

Equation (3.13) states that a standard Brownian motion W^T under the forward measure \mathbb{Q}_T follows the dynamics

$$dW^T = dW(t) + \sigma B(t, T)dt,$$

so that eq. (3.12) implies

$$\begin{aligned} dx(t) &= -ax(t)dt + \sigma dW(t) \\ &= (-\sigma^2 B(t, T) - ax(t)) dt + \sigma dW^T(t). \end{aligned}$$

Finally, eq. (3.14) follows from variation of constants. □

An immediate result is:

Corollary 3.2.8. *For any $s < t$, a solution to the stochastic differential equation (3.2) is given by*

$$r(t) = \alpha(t) + e^{-a(t-s)}x(s) - M^T(s, t) + \sigma \int_s^t e^{-a(t-u)}dW^T(u) \quad (3.15)$$

where W^T is a standard Brownian motion under the T -forward measure \mathbb{Q}_T .

In particular, conditional on \mathcal{F}_s , under \mathbb{Q}_T the random variable $r(t)$ is normally distributed with expectation

$$\mathbb{E}^T [r(t)|\mathcal{F}_s] = \alpha(t) + e^{-a(t-s)}x(s) - M^T(s, t)$$

and variance

$$\mathbb{V}^T [r(t)|\mathcal{F}_s] = \frac{\sigma^2}{2a} \left(1 - e^{-2a(t-s)} \right).$$

□

This result can now be used to determine the prices of European call and put options on zero-coupon bonds.

Proposition 3.2.9. *The time t price of a unit notional European call option with strike K and maturity T , written on the forward price of a zero-coupon bond spanning $[T, S]$, is given by*

$$V_{\text{Call}}(t, K, P(t, T, S), T) = P(t, S)\Phi(h) - KP(t, T)\Phi(h - \sigma_p)$$

where Φ denotes the cumulative distribution function of the standard normal distribution and

$$\sigma_p = \sigma \left(\frac{1 - e^{-2a(T-t)}}{2a} \right)^{1/2} B(T, S)$$

$$h = \frac{1}{\sigma_p} \ln \left(\frac{P(t, S)}{KP(t, T)} \right) + \frac{\sigma_p}{2}.$$

The price of a unit notional European put option with analogous characteristics is given by

$$V_{\text{Put}}(t, K, P(t, T, S), T) = KP(t, T)\Phi(-h + \sigma_p) - P(t, S)\Phi(-h).$$

Proof. See Brigo and Mercurio (2006). □

Using the fact that a cap (floor) can be interpreted as a portfolio of put (call) options on zero-coupon bonds (cf. eqs. (2.14) and (2.15)), this yields:

Corollary 3.2.10. *Define the tenor structure*

$$0 \leq T_m < \dots < T_n, \quad \tau_i = T_{i+1} - T_i, \quad i = m, m+1, \dots, n-1, \quad (3.16)$$

and consider the corresponding unit notional cap with strike K , which consists of $(n - m - 1)$ unit notional caplets on LIBOR rates f_i paying

$$\tau_i(f_i(T_i) - K)^+$$

at times T_{i+1} .

The time t value of this cap is given by

$$V_{\text{Cap}} = \sum_{i=m}^{n-1} P(t, T_i)\Phi(-h_i + \sigma_p^i) - (1 + \tau_i K) P(t, T_{i+1})\Phi(-h_i), \quad (3.17)$$

where

$$\sigma_p^i = \sigma \left(\frac{1 - e^{-2a(T_i-t)}}{2a} \right)^{1/2} B(T_i, T_{i+1})$$

$$h_i = \frac{1}{\sigma_p^i} \ln \left(\frac{P(t, T_{i+1})(1 + \tau_i K)}{P(t, T_i)} \right) + \frac{\sigma_p^i}{2}.$$

The price of the corresponding floor with analogous characteristics is given by

$$V_{\text{Floor}} = \sum_{i=m}^{n-1} (1 + \tau_i K) P(t, T_{i+1})\Phi(h_i) - P(t, T_i)\Phi(h_i - \sigma_p^i). \quad (3.18)$$

□

This result can be used to calibrate the parameters a and σ . Let

$$\text{Cap}_{\text{market}}(i), \quad i = 1, \dots, N,$$

be a given set of market cap prices, and denote by $V_{\text{Cap}}^{a,\sigma}(i)$ the model price of each cap according to eq. (3.17) given a and σ .

Parameters a and σ can be obtained by numerically minimizing the squared absolute or relative distances between market and model prices, i.e.,

$$\arg \min_{a,\sigma} \sum_{i=1}^N \left(\text{Cap}_{\text{market}}(i) - V_{\text{Cap}}^{a,\sigma}(i) \right)^2 \quad (3.19)$$

or

$$\arg \min_{a,\sigma} \sum_{i=1}^N \left(\frac{\text{Cap}_{\text{market}}(i) - V_{\text{Cap}}^{a,\sigma}(i)}{\text{Cap}_{\text{market}}(i)} \right)^2. \quad (3.20)$$

3.2.2 Calibration to Swaption Prices

Besides caps, European swaptions can also be priced analytically in the Hull-White model. As seen in section 2.2, a swaption is equivalent to a put or call option on a coupon bearing bond. In several term structure models, including the Hull-White model, such an option can be valued as a weighted sum of options written on zero-coupon bonds.

Consider a European option with strike K and maturity T , written on the forward price of a bond paying n coupons after the option maturity. Denote by T_1, \dots, T_n the payment times and by c_1, \dots, c_n the payment values associated with the bond.

Let

$$\Pi(t, r(t)) = \sum_{i=1}^n c_i P(t, T, T_i)$$

be the time t forward price of the coupon bearing bond and let r^* be the value for the spot rate at time T for which

$$\Pi(T, r^*) = K.$$

Further, let

$$K_i = P(T, T_i | r(T) = r^*), \quad i = 1, \dots, n,$$

be the time T price of a zero-coupon bond maturing at T_i conditional on $r(T) = r^*$.

Then the price of the European option is given by

$$V_{\text{Call}}(t, K, \Pi(t), T) = \sum_{i=1}^n c_i V_{\text{Call}}(t, K_i, P(t, T, T_i), T) \quad (3.21)$$

or

$$V_{\text{Put}}(t, K, \Pi(t), T) = \sum_{i=1}^n c_i V_{\text{Put}}(t, K_i, P(t, T, T_i), T), \quad (3.22)$$

depending on whether it is a call or a put option. This relationship is based on a technique introduced by Jamshidian (1989). A formal derivation of eqs. (3.21) and (3.22) can also be found in Brigo and Mercurio (2006).

Corollary 3.2.11. *Consider the unit notional payer swaption with maturity T_m and strike rate K which gives the owner the right to enter at time T_m into a swap with reset dates T_m, \dots, T_{n-1} and payment dates T_{m+1}, \dots, T_n , so that the holder pays the fixed leg and receives the floating leg.*

Let

$$c_i = \begin{cases} K\tau_i & \text{for } i = m+1, \dots, n-1 \\ 1 + K\tau_i & \text{for } i = n \end{cases}$$

and let r^* be the value of the short rate at time T_m for which

$$\sum_{i=m+1}^n c_i A(T_m, T_i) e^{-B(T_m, T_i)r^*} = 1.$$

Further let

$$K_i = A(T_m, T_i) e^{-B(T_m, T_i)r^*}, \quad i = m+1, \dots, n.$$

Then the time t value of the payer swaption is given by

$$V_{\text{PSwaption}} = \sum_{i=m+1}^n c_i V_{\text{Put}}(t, K_i, P(t, T_m, T_i), T_m). \quad (3.23)$$

The price of the analogous receiver swaption is given by

$$V_{\text{RSwaption}} = \sum_{i=m+1}^n c_i V_{\text{Call}}(t, K_i, P(t, T_m, T_i), T_m).$$

Proof. Using the identity

$$P(T_m, T_i | r(T_m) = r^*) = A(T_m, T_i) e^{-B(T_m, T_i)r^*},$$

which has been shown in proposition 3.2.4, this follows immediately from Jamshidian's (1989) technique since the payer (receiver) swaption can be viewed as a put (call) option with strike 1 on the forward price of a coupon-bearing bond (see section 2.2). \square

As in the caplet calibration mechanism (see eqs. (3.19) and (3.20)), we can use this result to calibrate the parameters a and σ by minimizing the squared absolute or relative distances between market and model prices of swaptions.

3.3 The SABR Model

The SABR model has been put forth by Hagan et al. (2002). The acronym SABR stands for *Stochastic Alpha, Beta, Rho*, where α, β and ρ represent variables of the model in the original notation (the state variable α is the initial volatility, and the parameters β and ρ denote a CEV-type exponent and a correlation term, respectively). To avoid confusion and since we are dealing with a variety of models and their parameters simultaneously, we will follow Rebonato et al. (2009) and use the symbol σ to denote volatilities.

Prior to the emergence of the SABR model, local volatility models developed by Dupire (1994) as well as Derman and Kani (1994) were the most common method of managing smile / skew risk. However, Hagan et al. (2002) point out that local volatility models make wrong predictions about the dynamics of the volatility smile / skew. Empirical evidence shows that a change in the forward price is usually accommodated by a shift of the implied volatility curve in the same direction, which is the opposite of the predictions made by the above-mentioned local volatility models.

As Hagan et al. show, this means that hedges calculated from these models are incorrect. Consider for example a call option on the forward rate f struck at K and expiring at time T . There exists a function $\sigma_{\text{loc}}(f, K)$ such that the price of a European call (and similarly a put) option under the local volatility model can be computed using Black's formula [eq. (2.10)]

$$\text{Bl}^+(K, f, T, \sigma_{\text{loc}}(f, K)).$$

Differentiation with regard to f yields the option Delta

$$\Delta = \frac{\partial \text{Bl}^+}{\partial f} + \frac{\partial \text{Bl}^+}{\partial \sigma} \frac{\partial \sigma_{\text{loc}}(K, f)}{\partial f}.$$

The first term is the delta as predicted by the Black model, and the second term is a correction term stating by how much the local volatility model's delta deviates from the one predicted by the Black model. Under the local volatility model, the correction term has the wrong sign, and the resulting hedges are worse than those suggested by the Black model.

To overcome this issue, Hagan et al. (2002) developed the SABR model, which is a stochastic volatility model capable of evolving a single forward rate f_t of maturity T in isolation. Let τ denote the tenor of f_t . Under its associated terminal measure $\mathbb{Q}_{T+\tau}$, i.e., the measure under which f_t and its volatility are martingales, the rate f_t is assumed to follow the dynamics

$$df_t = \sigma_t (f_t)^\beta dW_t \quad (3.24)$$

$$d\sigma_t = \sigma_t \nu dZ_t \quad (3.25)$$

$$d\langle W, Z \rangle_t = \rho dt \quad (3.26)$$

for $t \leq T$, where W and Z are correlated Wiener processes. For a given set of expiry-dependent parameters $\nu \in \mathbb{R}$, $\beta \in [0, 1]$ and $\rho \in [-1, 1]$, the SABR dynamics are thus fully determined by the initial values f_0 and σ_0 . We note that among these two state variables only the initial forward rate f_0 is directly observable from market values.

In the SABR model, f_t could represent any forward rate, such as a forward LIBOR rate or a forward swap rate. The former allows to model interest rate caplets / floorlets, and the

latter allows to model swaptions. If f_t is a forward LIBOR rate, the numéraire is given by $P(t, T + \tau)$. In the case where a swap rate is modeled, the numéraire is given by the swap annuity.

3.3.1 Valuation of European Options

While the initial value f_0 is given by the market, the parameters σ_0 , ν and ρ will be obtained by calibrating the model to the prices of European vanilla options (i.e., caplets if f_t denotes a forward LIBOR rate, or swaptions in case f_t represents a forward swap rate).

Hagan et al. (2002) derive an approximation formula for the Black implied volatility of a call / put option on an underlying that obeys the SABR model. This allows for the analytical pricing of options under the SABR model by use of the Black formula, and also provides a tractable calibration algorithm. In particular, at time $t = 0$ the (approximate) Black implied volatility of a European call option on the underlying f with strike K and expiry t_{ex} is given by

$$\sigma^{\text{SABR}}(f_0, K, t_{\text{ex}}, \sigma_0, \rho, \nu, \beta) = A \cdot \left(\frac{z}{\chi(z)} \right) \cdot B \quad (3.27)$$

with

$$\begin{aligned} A &= \frac{\sigma_0}{(f_0 K)^{\frac{1-\beta}{2}} \left[1 + \frac{(1-\beta)^2}{24} \ln^2 \frac{f_0}{K} + \frac{(1-\beta)^4}{1920} \ln^4 \frac{f_0}{K} + \dots \right]} \\ B &= \left[1 + \left(\frac{(1-\beta)^2}{24} \frac{(\sigma_0)^2}{(f_0 K)^{1-\beta}} + \frac{\rho \beta \nu \sigma_0}{4(f_0 K)^{\frac{1-\beta}{2}}} + \frac{2-3\rho^2}{24} \nu^2 \right) \cdot t_{\text{ex}} + \dots \right] \\ z &= \frac{\nu}{\sigma_0} (f_0 K)^{\frac{1-\beta}{2}} \ln \frac{f_0}{K} \\ \chi(z) &= \ln \left(\frac{\sqrt{1-2\rho z + z^2} + z - \rho}{1-\rho} \right). \end{aligned}$$

In the case of an at-the-money option, eq. (3.27) reduces to:

$$\sigma^{\text{SABR}}(f_0, f_0, t_{\text{ex}}, \sigma_0, \rho, \nu, \beta) = \frac{\sigma_0}{f_0^{1-\beta}} \left[1 + \left(\frac{(1-\beta)^2}{24} \frac{(\sigma_0)^2}{f_0^{2-2\beta}} + \frac{\rho \beta \nu \sigma_0}{4(f_0)^{\frac{1-\beta}{2}}} + \frac{2-3\rho^2}{24} \nu^2 \right) \cdot t_{\text{ex}} \right].$$

The formula has later been improved by various authors, such as Oblój (2008) and Paulot (2009), but is still widely used in its original form.

If the SABR state variable f represents a forward swap rate, the above formula can be used to value a payer swaption. At the time of maturity T , the swaption will pay

$$A(T) (f(T) - K)^+.$$

The swap annuity $A(t)$ is the associated numéraire under the swap measure so that the time 0 value of the swaption is

$$\mathbb{E}^{T+\tau} \left[\frac{A(0)}{A(T)} A(T) (f(T) - K)^+ \right] = A(0) \mathbb{E}^{T+\tau} \left[(f(T) - K)^+ \right]$$

where $\mathbb{E}^{T+\tau}$ denotes the expectation under the terminal measure.

All we need to compute this expression is the price of a European call option under the SABR model which is easily obtained by evaluating the Black formula using the implied volatility given by Hagan's formula, i.e.,

$$V_{\text{PSwaption}}(0) = A(0) \text{BI}^+(K, f(0), T, \hat{\sigma}^{\text{SABR}}). \quad (3.28)$$

Note that a caplet is simply a (degenerate) one-period swaption (where $A(t) = P(t, T + \tau)$) and can also be valued according to eq. (3.28).

3.3.2 Calibration

Consider a set of caplets or swaptions with expiry T , strikes K_1, \dots, K_N and Black implied volatilities $\hat{\sigma}(K_i)$. The SABR parameters can be calibrated via eq. (3.27) by solving

$$(\sigma_0, \rho, \nu) = \arg \min_{\substack{\sigma_0, \nu \in \mathbb{R}_+, \\ |\rho| \leq 1}} \sum_{i=1}^N \left[\hat{\sigma}(K_i) - \sigma^{\text{SABR}}(f_0, K_i, T, \sigma_0, \rho, \nu, \beta) \right]^2. \quad (3.29)$$

For any given single exercise date, the SABR model can be used to fit the implied volatility curve. The predicted dynamics of the smile are consistent with empirical observations. Assets with a single exercise date — including caplets and swaptions — can thus be hedged using the SABR model.

Furthermore, its capability to parameterize the volatility smile as a tuple (σ_0, ρ, ν) allows for the SABR model to be used as a tool for the interpolation of the full smile from a discrete set of market quotes.

The hedging performance of the SABR model is tested in Rebonato et al. (2009, chapter IV) as well as Rebonato et al. (2008), where the authors find that the model is very well specified with regard to delta and vega hedging.

3.4 The SABR-LIBOR Market Model

One major drawback of the SABR model is its inability to model rates of different expiries simultaneously. In particular, complex derivatives depending on multiple forward rates cannot be captured by the SABR model.

The classic log-normal LIBOR Market Model (LMM), primarily developed by Brace et al. (1997), Jamshidian (1997) and Miltersen et al. (1997) provides this feature. However, the

LMM is based on the assumption of deterministic volatilities, and due to the assumed log-normality it is inconsistent with a volatility smile as can be observed in the interest rate cap / floor market (see e.g., Jarrow et al. (2007)).

While certain stochastic-volatility extensions of the LMM existed and some were capable of capturing the smile to a certain degree (see again Jarrow et al. (2007)), none of them were compatible with the SABR model. Rebonato (2007) set out to provide a stochastic-volatility extension of the LMM that was capable of recovering SABR caplet prices. The resulting SABR-LIBOR Market Model (LMM-SABR) was extended further and studied by Rebonato et al. (2009).

For $i = 1, 2, \dots, N$, consider the forward rates f_t^i expiring at T_i and paying at $T_i + \tau$. Denote by s_t^i the instantaneous volatility of the i th forward rate. Under its own terminal measure $\mathbb{Q}_{T_i+\tau}$, which is the measure corresponding to the numéraire $P(t, T_i + \tau)$, for $t \leq T_i$ the LMM-SABR dynamics of the forward rate f_t^i are given by

$$df_t^i = s_t^i (f_t^i)^{\beta^i} dW_t^i \quad (3.30)$$

$$s_t^i = g_t^i k_t^i \quad (3.31)$$

$$\frac{dk_t^i}{k_t^i} = \mu^i dt + h_t^i dZ_t^i, \quad (3.32)$$

where g_t^i and h_t^i are deterministic functions, $\beta^i \in [0, 1]$, and W^i, Z^i are correlated standard Brownian motions.

We follow the second route described in Rebonato (2007) and only consider the case $\mu^i = 0$. Note that under $\mathbb{Q}_{T_i+\tau}$ the driftlessness in eqs. (3.30) and (3.32) only holds for the single forward rate $f_t^i = f(t, T_i, T_i + \tau)$. If all forward rates are modeled simultaneously under the same measure, the resulting equations of motion contain no-arbitrage drift terms.

Denote the correlation structure by

$$d\langle W^i, W^j \rangle_t = \rho_{ij} dt \quad (3.33)$$

$$d\langle Z^i, Z^j \rangle_t = r_{ij} dt \quad (3.34)$$

$$d\langle W^i, Z^j \rangle_t = R_{ij} dt. \quad (3.35)$$

We can write this in terms of a matrix as

$$P = \begin{pmatrix} \rho & R \\ R^T & r \end{pmatrix}. \quad (3.36)$$

Note that ρ, r and P need to be valid correlation matrices (i.e., symmetric and positive semi-definite matrices with unit diagonal). We call P the *super-correlation matrix* of the LMM-SABR.

The instantaneous volatility s_t^i has been decomposed into a purely deterministic component g_t^i and a stochastic part k_t^i . Note that eqs. (3.31) and (3.32) imply

$$ds_t^i = s_t^i \mu^i dt + h_t^i s_t^i dZ_t^i,$$

so that h_t^i represents the instantaneous volatility of volatility.

The model is fully specified once the functions g_t^i and h_t^i as well as the super-correlation matrix are defined and the initial values f_0^i , k_0^i and s_0^i are set.

3.4.1 Parameterization of the Volatility Functions g and h

As for the functions g_t^i and h_t^i , we will follow Rebonato et al. (2009) and use the *abcd*-parameterization originally proposed by Rebonato (2004):

$$\begin{aligned} g_t^i &= g(T_i - t) \\ h_t^i &= \xi_i h(T_i - t), \end{aligned}$$

where

$$g(\tau) = (a + b\tau) \exp(-c\tau) + d \quad (3.37)$$

$$h(\tau) = (\alpha + \beta\tau) \exp(-\gamma\tau) + \delta, \quad (3.38)$$

and where ξ_i are correction factors. While a fully time-homogeneous volatility structure might be a desirable feature, a perfect fit to a set of plain-vanilla derivatives is even more important in many cases. In the presence of the correction factors ξ_i , the resulting volatility of volatility will not be strictly time-homogeneous but will allow for a better fit to market quotes. If the model is well-specified, the required correction terms will be close to unity. Rebonato et al. (2009) show that this is indeed the case when calibrating the volatility structure to market prices of caplets.

Both g and h should be well-behaved instantaneous volatility functions. Therefore, we will impose certain constraints on the parameters of g and h . As explained by Rebonato et al. (2009), the *abcd*-parameterization allows to give g and h monotonically decreasing or humped shapes. In most cases, the function h will be monotonically decreasing and g will either exhibit a humped shape (under “normal” market conditions) or decrease monotonically (under “excited” market conditions).

Rebonato (2006) explains the humped shape typically observed in instantaneous volatilities as follows. Forward rates close to expiry are primarily affected by the actions of monetary authorities. In general these institutions tend to communicate their intentions well before significant rate decisions are made. Therefore, forward rates usually show a declining level of volatility as they reach expiry. Expectations about very long-dated rates on the other hand are primarily influenced by expectations about long-term inflation. Central banks often follow an inflation target which reduces the volatility of long-term nominal rates (under the assumption that real interest rates are little volatile). Thus, the largest uncertainty can be found in the medium term, usually between six and twenty-four months.

In “excited” periods on the other hand, there is a lack of consensus about the imminent decisions of the monetary institutions so that short-term rates become highly volatile and the hump disappears.

A selection of parameter choices and the associated shapes of the *abcd*-function are shown in fig. 3.1.




Figure 3.1: Possibles shapes the $abcd$ -parameterization defined in eqs. (3.37) and (3.38) can recover. Source: (Rebonato et al., 2009, Fig. 2.1).

We would like g to have a limit as τ tends to infinity and thus we require $c > 0$. As τ goes to infinity, $g(\tau)$ will converge to d which can be interpreted as the approximate instantaneous volatility of a very long-dated rate, which of course shall be strictly positive as well. As t tends toward T_i , g_t^i approaches the instantaneous volatility of a rate arbitrarily close to expiry, which shall again be positive:

$$g(0) = a + d > 0.$$

Consider the derivative of the $abcd$ -function:

$$g'(\tau) = (b - ac - bc\tau) \exp(-c\tau).$$

The function g will be strictly decreasing if and only if $g'(\tau) < 0$ for all $\tau > 0$. This is the case if and only if — in addition to the constraints described above — one of the following two conditions holds true

- $a > 0$ and $b = 0$ or
- $a, b > 0$ and $c \geq \frac{b}{a}$.

On the other hand, g will describe a humped shape if there exists some $\tau^* > 0$ where

$$\frac{\partial}{\partial \tau} g(\tau^*; a, b, c, d) = 0,$$

i.e.,

$$\tau^* = \frac{1}{c} - \frac{a}{b} \tag{3.39}$$

for some $\tau^* > 0$. This requires $b > 0$. Note that in this case

$$g''(\tau^*) = -bc \exp(-c\tau^*) < 0$$

so that the extremum is indeed a maximum.

If necessary we could also use eq. (3.39) to impose conditions on the location of the maximum.

Equivalent constraints can be defined for the parameters of the function h .

3.4.2 Calibration of the LMM-SABR to Caplets

We will now give a very concise overview of the calibration approach we will perform in section 4.3. For a detailed description see Rebonato et al. (2009).

Obtaining SABR Parameters

We assume that we have market quotes f_0^i , $i = 1, 2, \dots, N$, for a set of forward rates corresponding to a spanning tenor structure of

$$0 = T_0 < T_1 < \dots < T_{N+1}.$$

Furthermore, we assume that for $i = 1, 2, \dots, N$, we know the Black implied volatilities of a set of caplets expiring at T_i and paying at T_{i+1} . This allows us to perform a (separate) SABR calibration for each of the T_i which yields parameters²

$$\left(\sigma_0^{T_i}, \rho_{\text{SABR}}^{T_i}, \nu^{T_i}, \beta_{\text{SABR}}^{T_i} \right), \quad i = 1, 2, \dots, N.$$

Calibration of g and h

The calibration approach for the parameters of g and h will be based on the objective to recover the input caplet prices as closely as possible.

If the stochastic component $k_t^i = k_0^i$ were constant, so that the volatility term $s_t^i = g_t^i k_t^i$ were deterministic, and if $\beta = 1$, the LMM-SABR implied Black volatility would be exactly equal to

$$\hat{\sigma}_i^2 = \frac{1}{T_i} \int_0^{T_i} s_u^2 du = \frac{(k_0^i)^2}{T_i} \int_0^{T_i} (g_u^i)^2 du. \quad (3.40)$$

In the deterministic-volatility lognormal setting, the initial values $k_0^{T_i}$ would thus be fully determined by eq. (3.40), namely

$$k_0^T = \frac{\hat{\sigma}_i}{\hat{g}_i}$$

where

$$\hat{g}_i = \sqrt{\frac{1}{T_i} \int_0^{T_i} (g_u^i)^2 du} \quad (3.41)$$

²Strictly speaking, σ_0 is a *state variable* in the SABR model, not a *parameter*.

denotes the root-mean-squared volatility. Of course, the above calculation is invalid in the case of a stochastic k^i , so that the initial value k_0^i will have to be chosen in a different way. First note that the SABR volatility process $\sigma_t^{T_i}$ satisfies

$$\mathbb{E} \left[\sigma_t^{T_i} \right] = \sigma_0^{T_i}.$$

As for the calibration of g and k_0^i , Rebonato et al. (2009) propose the following heuristic approach: First, the parameters of g will be chosen so that \hat{g}_i matches $\sigma_0^{T_i}$ as closely as possible, i.e.,

$$(a, b, c, d) = \arg \min \sum_{i=1}^N \left(\sigma_0^{T_i} - \hat{g}_i \right)^2. \quad (3.42)$$

The minimization has to be performed numerically, under the necessary (or desired) constraints.

The initial loadings k_0^i will subsequently be chosen so as to compensate for the remaining discrepancies and bring about exact recovery of $\sigma_0^{T_i}$:

$$k_0^i = \frac{\sigma_0^{T_i}}{\hat{g}_i}, \quad i = 1, 2, \dots, N. \quad (3.43)$$

Note that the state variables k_0^i will be close to 1 if the calibration of \hat{g}_i in eq. (3.42) already provides a close fit to the initial SABR volatilities $\sigma_0^{T_i}$.

The idea behind this procedure is to have the deterministic factor g_t^i represent the (average) physical behavior of the volatility function so that the stochasticity merely induces deviation from this deterministic behavior.

Next we wish to calibrate h by attempting to make LMM-SABR caplet prices match SABR caplet prices:

$$\text{Call}_{\text{LMM}} \approx \text{Call}_{\text{SABR}}. \quad (3.44)$$

Unfortunately, no immediate closed-form solution or sufficiently precise approximation formula exists for the pricing of caplets (or floorlets) under the LMM-SABR model.

In his original publication of the LMM-SABR, Rebonato (2007) suggested choosing the parameters of h by performing a fit of the root-mean-squared volatility

$$\hat{h}_i = \sqrt{\frac{1}{T_i} \int_0^{T_i} h(T_i - s)^2 ds}$$

to the SABR volatility of volatility parameters ν^{T_i} , $i = 1, \dots, N$, i.e., by solving

$$(\alpha, \beta, \gamma, \delta) = \arg \min \sum_{i=1}^N \left(\nu^{T_i} - \hat{h}_i \right)^2. \quad (3.45)$$

Subsequently, Rebonato and White (2009) presented an improvement of this method based on the following reasoning: In a deterministic-volatility setting (e.g., the deterministic-volatility Black (1976) framework or the classic LMM), the price of a caplet only depends on the terminal distribution of the underlying forward rate. In this case all that is need to

ensure that a caplet is priced correctly is to make sure that its average volatility over the lifetime of the forward rate is correct. This is the rationale behind eq. (3.40). In eq. (3.45) the same concept is simply applied to the volatility of volatility, i.e., h is calibrated by ensuring that the average volatility of volatility comes close to a benchmark value — the SABR ν in this case.

Rebonato and White (2009) note that this reasoning is not strictly valid in the stochastic volatility case. Since the volatility of a forward rate and its volatility of volatility interact, it is not only relevant that each one of them has the correct value *on average* but it is also important *when* certain values occur. This can be seen intuitively: At a time when the volatility s_t^i of the forward rate is extremely low (because g_t^i is very low, e.g.), a high volatility of volatility h_t^i will only have a small effect on the evolution of the forward rate. If instead such a high volatility of volatility occurred at a time when s_t^i was at a high level, the effect would be much stronger.

Rebonato and White (2009) devise a different calibration mechanism where this aspect is considered. In order to perform a calibration based on eq. (3.44), Rebonato and White note that for the minimization of the difference $\text{Call}_{\text{LMM}} - \text{Call}_{\text{SABR}}$, it is sufficient to have a representation formula for this very difference; there is no need to derive an individual formula for each of the two terms.

The basic idea is to set equal the exact (non-closed) formulas for Call_{LMM} and $\text{Call}_{\text{SABR}}$, and then make similar approximations with regard to each of the expressions. Even though neither of the two approximations is accurate in itself, eq. (3.44) still holds approximately since the applied transformations produce roughly the same errors on both sides of the equation. Without presenting the algebraic manipulations in full detail, the resulting calibration mechanism is

$$(\alpha, \beta, \gamma, \delta) = \arg \min \sum_{i=1}^N \left[\nu^{T_i} - \frac{k_0^i}{\sigma_0^{T_i} T_i} \left(2 \int_0^{T_i} (g_u^i)^2 (\hat{h}_u^i)^2 u du \right)^{1/2} \right]^2 \quad (3.46)$$

where

$$\begin{aligned} \hat{h}_u^i &= \sqrt{\frac{1}{u} \int_0^u h(T_i - s)^2 ds} \\ &= \sqrt{\frac{1}{u} \int_0^u [(\alpha + \beta(T_i - s)) \exp(-\gamma(T_i - s)) + \delta]^2 ds}. \end{aligned}$$

Similar to eq. (3.43), correction terms ξ_i are introduced to compensate for each of the residuals in eq. (3.46):

$$\xi_i = \nu^{T_i} / \left[\frac{k_0^i}{\sigma_0^{T_i} T_i} \left(2 \int_0^{T_i} (g_u^i)^2 (\hat{h}_u^i)^2 u du \right)^{1/2} \right]. \quad (3.47)$$

For our purposes we will use the improved calibration method. Nevertheless, we also considered the simple method (eq. (3.45)) since it allows us to obtain a rough prior estimate of what shape we can expect for the h function. We will make use of this in section 4.3 where we need to calibrate the LMM-SABR for a large number of consecutive dates.

Calibration of the Correlation Structure

The only elements of the super correlation matrix P that influence the prices of caplets are the diagonal elements of the submatrix R , i.e., the correlations between each forward rate and its own volatility.

In order to recover the SABR caplet prices, one has to ensure that

$$R_{ii} = \rho_{\text{SABR}}^{T_i}, \quad i = 1, 2, \dots, N.$$

The full calibration matrix can be obtained by econometric estimation, cf. Rebonato et al. (2009), or by performing a best fit using market prices of correlation-dependent derivatives, such as swaptions.

In the context of our study we have no need for the estimation of the full correlation matrix.

3.4.3 Calibration to Swaptions

As an alternative to the calibration to caplets, the LMM-SABR can also be calibrated to market prices of swaptions. In our experiment, we are only considering (degenerate) one-period swaptions. These are equivalent to caplets, and the pricing formulas coincide. For the sake of completeness, we list the LMM-SABR swaption pricing formulas. In conjunction with the full correlation matrix P , these formulas can be used to calibrate the LMM-SABR to market prices of (multi-period) swaptions. All derivations as well as the estimation process for the correlation matrix P can be found in Rebonato et al. (2009) or Rebonato and White (2009).

To arrive at a tractable calibration method, Rebonato and White (2009) develop accurate closed-form approximations for the price of a swaption given a set of LMM-SABR parameters.

Rebonato and White (2009) proceed as follows: Assume that the evolution of a swap rate under the forward-rate based LMM-SABR can be represented by a SABR system,

$$dS_t^{m,n} = (S_t^{m,n})^{B^{m,n}} \Sigma_t^{m,n} dZ_t^{m,n} \quad (3.48)$$

$$\frac{d\Sigma_t^{m,n}}{\Sigma_t^{m,n}} = V^{m,n} dW_t^{m,n} \quad (3.49)$$

$$d\langle Z_t^{m,n}, W_t^{m,n} \rangle = R_{\text{SABR}}^{m,n} dt, \quad (3.50)$$

under the swap measure.

The system is fully specified once the parameters $V^{m,n}$, $B^{m,n}$, $R_{\text{SABR}}^{m,n}$ as well as the initial conditions $S_0^{m,n}$ and $\Sigma_0^{m,n}$ have been set.

If one had knowledge of these values, one could use the Hagan formula (eq. (3.27)) to price a payer or receiver swaption on the swap rate $S^{m,n}$.

Rebonato and White (2009) succeed in deriving approximation formulas for the above parameters and initial values that are implied by a given set of parameters and initial

values of the forward-rate based LMM-SABR. The accuracy of the predicted swaption prices is shown by comparing them to the results of a full Monte-Carlo simulation of the LMM-SABR system.

First we note that a forward swap rate $S_t^{m,n}$ (cf. eq. (2.8)) can be expressed as a weighted sum of the associated forward rates:

$$S_t^{m,n} = \sum_{i=m}^{n-1} \omega_i^{m,n}(t) f_t^i \quad (3.51)$$

where

$$\omega_i^{m,n}(t) = \frac{\tau_i P(t, T_{i+1})}{\sum_{j=m}^{n-1} \tau_j P(t, T_{j+1})}.$$

To increase readability, we will omit the superscripts m, n in the following formulas. The approximations derived by Rebonato and White (2009) for the unknown parameters in eqs. (3.48) to (3.50) are as follows:

$$\Sigma_0 \approx \left[\frac{1}{T_m} \sum_{i,j=m}^{n-1} \left(\rho_{ij} W_i^0 W_j^0 k_0^i k_0^j \int_0^{T_m} g_t^i g_t^j dt \right) \right]^{1/2} \quad (3.52)$$

$$V \approx \frac{1}{\Sigma_0 T_m} \left[2 \sum_{i,j=m}^{n-1} \left(\rho_{ij} r_{ij} W_i^0 W_j^0 k_0^i k_0^j \int_0^{T_m} g_t^i g_t^j \hat{h}_{ij}(t)^2 t dt \right) \right]^{1/2} \quad (3.53)$$

where

$$\hat{h}_{ij}(t) = \left(\frac{1}{t} \int_0^t h^i(s) h^j(s) ds \right)^{1/2}$$

$$W_i^0 = \omega_i(0) \frac{(f_0^i)^\beta}{S_0^B}.$$

The correlation parameter R_{SABR} is approximated as

$$R_{\text{SABR}} \approx \sum_{i,j=m}^{n-1} \Omega_{ij} R_{ij}$$

where

$$\Omega_{ij} = \frac{2 \rho_{ij} r_{ij} W_i^0 W_j^0 k_0^i k_0^j \int_0^{T_m} g_t^i g_t^j \hat{h}_{ij}(t)^2 t dt}{(V \Sigma_0 T_m)^2}.$$

The SABR exponent B is approximated as

$$B \approx \sum_{i=m}^{n-1} \omega_i(0) \beta_i, \quad (3.54)$$

where β_i are the forward rate exponents in the LMM-SABR dynamics. Since we are using a single β ($= 1/2$) for all forward rates, eq. (3.54) simplifies to

$$B \approx \beta.$$

The initial swap rate S_0 is given by eq. (3.51) in conjunction with the initial forward rates.

In the case of a one-period swaption, minimization of the mean-squared distance between the approximate Σ_0 (eq. (3.52)) and the forward-rate based SABR σ_0 coincides with the calibration procedure presented in eq. (3.42). Furthermore, a best fit of the approximate V (eq. (3.53)) to the forward-rate based SABR ν is identical to the approach given in eq. (3.46).

Hence, the two calibration procedures are perfectly consistent.

Chapter 4

A Model for the Evolution of the Yield Curve and Smile Surface Under the Physical Measure

4.1 The Need for a Real-World Evolution Framework

One might be inclined to evolve the term structure of interest rates by simply employing a pricing model. Rebonato et al. (2005) distinguish between two subclasses of such models. “Fundamental” models describe the evolution of the yield curve under the real-world measure and reach conclusions about the prices of derivative products by the explicit specification of risk premiums. “Reduced-form” models describe the evolution of interest rates directly under the risk-neutral measure. In principle, fundamental pricing models could be used to simulate the real-world evolution of the yield curve. Even though they were never intended to describe the dynamics of the term structure in the most realistic way, they provide a parsimonious description of the most fundamental features of the yield curve dynamics. However, as Rebonato et al. (2005) point out, their descriptions lack too many statistical subtleties to be able to provide satisfactory results, especially since realism has often been waived in favor of tractability. Reduced-form models are even less suitable for this purpose. Since they work directly under the pricing measure, most of their parameters implicitly contain components linked to risk aversion so that the interest rate dynamics contain drift terms that are different from what can be observed in real-world time series. Especially over a long horizon, these drift terms will completely dominate the simulation and the results will look very different from reality. Moreover, the choice of a pricing measure is often ambiguous.

Hence, it is clear that we require a model that captures the statistical properties of historical movements in the term structure of interest rates, and produces synthetic realizations under the physical measure.

We will use a variation of a particular sampling approach originally suggested by Rebonato et al. (2005), which has subsequently been adapted by Deguillaume (2009) and Andreichenko (2011) to the modeling of interest rates and implied volatilities, respectively. However, we will make a few changes to the method.

We note that an important ingredient in the original approach by Rebonato et al. (2005) is the presence of pseudo-arbitrageurs modeled in the form of spring constants, intended to straighten out “kinks” and avoid excessive slopes in the synthetic yield curves. We follow Deguillaume (2009) as well as Andreichenko (2011) and instead make use of an Ornstein-Uhlenbeck type mean reversion force for this purpose. This decision is motivated by the desire to have a single approach that can be consistently applied to all involved parameters, and includes no judgment on the viability of the original approach which was intended at the evolution of yield curves alone.

4.2 Evolution of the Yield Curve

4.2.1 Motivation

In this section we wish to specify an approach designed by Deguillaume (2009) which will allow us to generate stochastic realizations of the *yield curve* under the physical measure. The evolution of the implied volatility surface will be treated in section 4.3.

As an introduction, the design of the approach will be illustrated in a simplified manner.

The general idea is to start with today’s yield curve and iteratively generate future rates by re-applying randomly selected historical movements of the curve. In order to generate economically sensible results, we will have to translate the historical rate movements to today’s market situation in a way consistent with our view on the behavior of interest rates.

Denote by $y(0)$ the current level of the (synthetic) rate at consideration, and by

$$Y(\omega_j + 1) - Y(\omega_j)$$

the amount by which the rate changed between two randomly chosen consecutive days ω_j and $\omega_j + 1$ out of our historical data set. As a motivating example, assume for a moment that changes in the level of rates took place in a simple additive way, such that a historical 1-day change could be characterized by

$$Y(\omega_j + 1) = Y(\omega_j) + \varepsilon_{\omega_j},$$

where ε_{ω_j} is a random variate drawn from some distribution ϕ .

Under this assumption, in our simulation we could generate tomorrow’s (synthetic) rate as

$$y(1) = y(0) + [Y(\omega_j + 1) - Y(\omega_j)] = y(0) + \varepsilon_{\omega_j}. \quad (4.1)$$

If instead we assumed that changes could be characterized by multiplicative movements, e.g.,

$$Y(\omega_j + 1) = Y(\omega_j)\varepsilon_{\omega_j},$$

we could define $y(1)$ by the relation

$$y(1) = y(0)\varepsilon_{\omega_j},$$

which is equivalent to

$$\frac{y(1) - y(0)}{y(0)} = \frac{Y(\omega_j + 1) - Y(\omega_j)}{Y(\omega_j)}, \quad (4.2)$$

or

$$\ln(y(1)) = \ln(y(0)) + [\ln(Y(\omega_j + 1) - Y(\omega_j))].$$

Using the bijective function $\ln(\cdot)$, the proportional world has thus been translated into an additive world similar to eq. (4.1). Since unfortunately changes in real world interest rates appear to neither manifest in a strictly additive nor a proportional manner, we need to find a different bijection if we wish to make a similar transformation in the context of real world rates.

We can interpret eq. (4.2) in a different way. Assume the evolution of interest rates could be characterized by

$$\frac{Y(\omega_j + 1) - Y(\omega_j)}{Y(\omega_j)} = \nu\varepsilon_{\omega_j},$$

where ε_{ω_j} is a random variate drawn from a distribution ϕ with *unit variance*. If we evolved

$$y(1) - y(0) = y(0)\nu\varepsilon_{\omega_j},$$

where ε_{ω_j} is drawn from ϕ , we could informally interpret $y(0)\nu$ as the volatility of y at time 0. In this sense, we could re-write eq. (4.2) as

$$\frac{y(1) - y(0)}{\nu y(0)} = \frac{Y(\omega_j + 1) - Y(\omega_j)}{\nu Y(\omega_j)} \quad (4.3)$$

or

$$y(1) - y(0) = (Y(\omega_j + 1) - Y(\omega_j)) \frac{\nu y(0)}{\nu Y(\omega_j)}$$

which means we are evolving y by adding to today's rate $y(0)$ historical changes

$$Y(\omega_j + 1) - Y(\omega_j)$$

scaled by the “volatility” quotient

$$\frac{\nu y(0)}{\nu Y(\omega_j)}.$$

We will go with this interpretation and make two adaptations to the world of eq. (4.3). We will model the “volatility” by a deterministic function

$$\sigma : \mathbb{R}_+ \rightarrow \mathbb{R}_+$$

depending on the level of the rate, and carry over eq. (4.3) to a continuous setting. The latter will be done by assuming that in the course of the movement between two levels, say $y(0)$ and $y(1)$, the scaling factor $\sigma(s)$ has been continuously applied to any $s \in [y(0), y(1)]$. Formally, the updated version of eq. (4.3) reads

$$\int_{y(0)}^{y(1)} \frac{ds}{\sigma(s)} = \int_{Y(\omega_j)}^{Y(\omega_j+1)} \frac{ds}{\sigma(s)}. \quad (4.4)$$

In order to translate this equation into an additive setting, note that eq. (4.4) is equivalent to

$$\Sigma(y(1)) - \Sigma(y(0)) = \Sigma(Y(\omega_j + 1)) - \Sigma(Y(\omega_j)), \quad (4.5)$$

if we write

$$\Sigma(y) := c_1 + c_2 \int_{c_3}^y \frac{ds}{\sigma(s)} \quad (4.6)$$

for arbitrary $c_1, c_2, c_3 \in \mathbb{R}$, $c_2 > 0$. Of course, we will have to make sure that all integrals in eqs. (4.4) and (4.6) actually exist and are finite. Also note that since σ is strictly positive, any Σ defined in this way will be strictly increasing and thus have an inverse Σ^{-1} which will allow us to recover the rate $y(1)$ once we have obtained $\Sigma(y(1))$.

Note that we have not yet given a full evolution model. At this point all we have done is make the step from understanding historical changes in a simple additive manner to the more sophisticated interpretation of eq. (4.4). Furthermore, by making use of the function Σ we have found our way back to the additive interpretation of interest rate movements.

If the daily increments of interest rates were independent and identically distributed, one could produce a statistically satisfactory synthetic evolution of the yield curve

$$y^\tau(0), \dots, y^\tau(j_{\max}),$$

where τ denotes the tenor of each rate, by performing the following method:

1. Obtain a large set of interest rate data, $Y^\tau(0), Y^\tau(1), \dots, Y^\tau(N)$, for various tenors τ .
2. Select initial values $y^\tau(0)$ for each tenor τ .
3. Generate a set of dates $\omega_1, \dots, \omega_{j_{\max}}$, randomly drawn from a discrete uniform distribution on the set $\{0, 1, \dots, N-1\}$.
4. For each tenor τ , iteratively evolve

$$\begin{aligned} \Sigma(y^\tau(j+1)) &= \Sigma(y^\tau(j)) + [\Sigma(Y^\tau(\omega_j + 1)) - \Sigma(Y^\tau(\omega_j))], \quad j = 1, \dots, j_{\max}, \\ y^\tau(j+1) &= \Sigma^{-1}(\Sigma(y^\tau(j+1))). \end{aligned}$$

Note that the approach samples a full cross-section of the yield curve simultaneously. E.g., the historical increments for the 1 year rate and the 20 year rate will be sampled from the same date ω_j . Therefore, the evolution is capable of recovering the correlation structure between one-day movements in different parts of the yield curve. Obviously, this is an indispensable feature of a yield curve evolution model.

However, Nyholm and Rebonato (2007) point out that after a few “months” of evolution, synthetic yield curves generated by a simple sampling mechanism as the one described above will generally bear little resemblance with yield curves observed historically. In particular, the simulated yield curves will often have excessive slopes and “kinks”, where the term kink refers to the occurrence of the phenomenon

$$y^{\tau_1} < y^{\tau_2} \text{ and } y^{\tau_2} > y^{\tau_3}$$

or

$$y^{\tau_1} > y^{\tau_2} \text{ and } y^{\tau_2} < y^{\tau_3},$$

for three successive tenors $\tau_1 < \tau_2 < \tau_3$. In order to resolve these effects, Rebonato et al. (2005) introduce a spring mechanism which exerts a deterministic force straightening out the shape of the yield curve. Furthermore, a mean-reversion component is added to the short and long ends of the curve. As we have noted, we will go with Deguillaume (2009) and Andreichenko (2011) and replace the spring mechanism with a mean-reverting stochastic process.

Another drawback of the one-day sampling mechanism described above is its inability of reproducing the serial autocorrelation seen in historical yield changes. In addition to the correlations of one-day changes in different parts of the yield curve, Rebonato et al. (2005) show that historical time series of interest rates exhibit a strong positive serial autocorrelation for n -day returns. In the above approach, the stochastic process $(\omega_j)_{j=1,2,\dots}$ randomly jumps to individual historical days so that any serial correlation present in the historical time series of rates is clearly lost.

Therefore, the process $(\omega_j)_{j=1,2,\dots}$ will be modified. Instead of drawing each ω_j independently from the set of available dates, like Rebonato et al. (2005) [as well as Deguillaume (2009) and Andreichenko (2011)] we will consider a full historical window of length W . In the case of $W = 40$, e.g., we would start by randomly drawing ω_1 from the set of available dates and subsequently define $\omega_{j+1} = \omega_j + 1$ for $j = 2, \dots, 40$. The subsequent value, ω_{41} , would again be drawn randomly.

The updated method can be described as follows:

1. Obtain a large set of interest rate data, $Y^\tau(0), Y^\tau(1), \dots, Y^\tau(N)$, for various tenors τ .
2. Define a fixed window length W .
3. Select initial values $y^\tau(0)$ for each tenor τ .
4. The future yield curve is simulated by the following process:
 - (a) Start with the initial yield curve $y^\tau(0)$ and repeat steps 4b to 4d until the desired number of data points, j_{\max} , have been obtained.
 - (b) Select from the historical data set a window of W days.
 - (c) Extract the movement the yield curve underwent within this window.
 - (d) Generate W days in the simulated future by re-applying the extracted movement to the current yield curve.

We will now develop the exact mathematical concept of how steps 4b to 4d are to be realized so as to ensure that the simulation will produce realistic outcomes.

Notation 4.2.1. For the remainder of the section, we will stay with the notation introduced above, i.e.:

- W : window length in the sampling mechanism
- $N + 1$: number of data points (business days) contained in the available data
- τ : tenor of the rate under consideration, e.g., $\tau = 1\text{Yr}$

- $Y^\tau(0), \dots, Y^\tau(N)$: historical time series for the rate corresponding to the tenor τ
- $y^\tau(0), y^\tau(1), \dots$: simulated time series for the rate corresponding to the tenor τ

Until now we have only considered the generic term “interest rate”. The approach can be adapted to any particular kind of interest rates. In our experiment, y^τ will denote forward LIBOR rates (more specifically, n Yr into 1Yr forward swap rates).

4.2.2 Window Selection

Since the number of days spanned by our historical data is of course finite, a decision has to be made on how to proceed if the end of the time series is reached before the current window has been concluded, or whether to avoid the issue altogether by ensuring that the starting point of each window is always sufficiently far away from the end of the series.

We will once again go with Deguillaume (2009) and simply jump to the beginning of the time series whenever the end is reached before the current window has been completed. This method has the benefit that each of the N historical realizations of 1-day changes will have an equal (unconditional) probability of being selected. This simplifies several formulas. A downside to this choice is the fact that regime changes are not captured with full consistency.

Additionally, one could introduce random jumps so that at each step a random variable is drawn which might induce an immediate jump to a different window. This concept is used by Andreichenko (2011), e.g., but will not be applied here.

Formally:

Definition 4.2.2. By λ_j we denote the starting point of the j th window. Let

$$(\lambda_j)_{j=1,2,\dots}$$

be a stochastic process where each realization is independently drawn from the discrete uniform distribution on $\{0, 1, \dots, N-1\}$, i.e.,

$$\mathbb{P}(\lambda_j = k) = \begin{cases} 1/N & \text{if } k \in \{0, 1, \dots, N-1\} \\ 0 & \text{otherwise.} \end{cases}$$

Further, let

$$(\omega_j)_{j=1,2,\dots}$$

be a stochastic process such that for each future day j in the simulation, ω_j will be its historical counterpart. It will depend on the starting point of the window attributed to the j th simulated day and on the position of the j th day in its associated window. Formally,

$$\omega_j = \begin{cases} \lambda_{j/W} & \text{if } j \equiv 0 \pmod{W} \\ \omega_{j-1} + 1 & \text{if } j \not\equiv 0 \pmod{W} \text{ and } \omega_{j-1} \neq N-1 \\ 0 & \text{if } j \not\equiv 0 \pmod{W} \text{ and } \omega_{j-1} = N-1. \end{cases}$$

4.2.3 The σ -Style Sampling Model

We will now determine how exactly the changes found in a historical window are to be applied to today's yield curve.

Following the basic idea laid out in section 4.2.1, we will attempt to find suitable functions σ and Σ .

First we wish to gain an intuition of how the variance scales with the level of rates. From the website of the Federal Reserve Bank of St. Louis we obtain time series of US dollar constant maturity treasury rates for a total of eight tenors between 1 and 30 years. We note that the time series span periods of different lengths. E.g., the 1 year rate ranges from 1962 until 2014, and the 30 year rate covers the period 2006–2014.

We proceed as follows: For each tenor τ we compute the time series of 1-day changes

$$\Delta^\tau(j) = y^\tau(j+1) - y^\tau(j).$$

Now order the sequences

$$(\Delta^\tau(j))_{j=0,1,\dots}$$

and

$$(y^\tau(j))_{j=0,1,\dots}$$

so that the latter sequence is in ascending order. Split the (re-ordered) sequence

$$(\Delta^\tau(j))_{j=0,1,\dots}$$

into sections of length 200,

$$X_k^\tau := \{X_{200k}, X_{200k+1}, \dots, X_{200(k+1)-1}\}, \quad k = 0, 1, \dots$$

For each of these sections X_k^τ , compute its mean μ_k^τ and standard deviation σ_k^τ . Moreover, compute the standard deviation σ^τ of the entire sequence $(\Delta^\tau(i))_{i=0,1,\dots}$. Now draw the scatter plot

$$\left\{ \left(\mu_k^\tau, \frac{\sigma_k^\tau}{\sigma^\tau} \right) \middle| \tau \in \{1\text{Yr}, \dots, 30\text{Yr}\}, k = 0, 1, \dots \right\}.$$

The resulting plot, fig. 4.1, shows the standard deviation of each bucket of one-day movements scaled by the encompassing time series' standard deviation plotted against the average rate level within the bucket. The scaling by $1/\sigma^\tau$ is done to make multiple time series comparable. The resulting plot indicates which proportion of a sequence's inherent standard deviation can be found at each respective rate level.

Figure 4.1 encourages us to go with Deguillaume (2009) and model σ as a piecewise linear function.

Definition 4.2.3 (σ -function, Deguillaume (2009)). A function $\sigma : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ that can be expressed in the form

$$\sigma(y) = \sigma_G \begin{cases} y/y_L & \text{if } 0 < y \leq y_L \\ 1 & \text{if } y_L \leq y \leq y_R \\ (1 + K(y - y_R)) & \text{if } y_R \leq y \end{cases}$$




Figure 4.1: Illustration of how the standard deviation of one-day movements in USD constant maturity treasury rates, ranging from 1964 until 2014, scales with the rate level.

where

$$\begin{aligned} y_L, y_R, \sigma_G &> 0, \\ K &\geq 0, \quad \text{and} \\ y_L &\leq y_R \end{aligned}$$

will be called a σ -function.

As outlined in section 4.2.1, we wish to define

$$\Sigma(y) := c_1 + c_2 \int_{c_3}^y \frac{ds}{\sigma(s)}, \quad (4.7)$$

for some $c_1, c_2, c_3 \in \mathbb{R}$, $c_2 > 0$. Once again going with Deguillaume (2009) we use

$$c_1 = c_3 = y_L$$

and

$$c_2 = \sigma_G.$$

Section 4.2.4 shows one example of a σ -function.

Definition 4.2.4 (Σ -function, Deguillaume (2009)). Given the above σ -function σ , define the corresponding Σ -function as

$$\Sigma : \mathbb{R}_+ \rightarrow \mathbb{R}, \Sigma(y) := y_L + \sigma_G \int_{y_L}^y \frac{ds}{\sigma(s)}.$$

Since σ is strictly positive, Σ is a bijection. Both Σ and Σ^{-1} are strictly increasing twice continuously differentiable bijections between \mathbb{R}_+ and \mathbb{R} .

Note that the choice $c_2 = \sigma_G$ has the desirable consequence that Σ is independent of σ_G .

Proposition 4.2.5 (Deguillaume (2009)). *Explicitly, Σ can be computed as*

$$\Sigma(y) = \begin{cases} y_L (1 + \ln(y/y_L)) & \text{if } 0 < y \leq y_L \\ y & \text{if } y_L \leq y \leq y_R \\ y_R + \ln(1 + K(y - y_R)) / K & \text{if } y_R \leq y, \end{cases}$$

and its inverse as

$$\Sigma^{-1}(z) = \begin{cases} \Sigma(y_L) \exp(z/\Sigma(y_L) - 1) & \text{if } z \leq \Sigma(y_L) \\ z & \text{if } \Sigma(y_L) \leq z \leq \Sigma(y_R) \\ \Sigma(y_R) + [\exp(K(z - \Sigma(y_R))) - 1] / K & \text{if } \Sigma(y_R) \leq z. \end{cases}$$

Proof. See (Deguillaume, 2009, chapter 6.2). □

Conjecture 4.2.6 (Deguillaume et al. (2013)). *There exists a unique Σ -function, Σ , such that for any given interest rate time series $(y_j)_{j \in \mathbb{N}}$, there exists a unique $\sigma_G \in \mathbb{R}$, a unique distribution ϕ with mean 0 and variance 1, and a unique function $\mu : \mathbb{R}_+ \rightarrow \mathbb{R}$ such that*

$$\Sigma(y_{j+1}) - \Sigma(y_j) = \sigma_G \varepsilon_j + \mu(\Sigma(y_j)), \quad (4.8)$$

where ε_j are random variates chosen from ϕ .

Note that it is not assumed that $\varepsilon_1, \varepsilon_2, \dots$ are independent or that σ_G , ϕ and μ are identical for each time series.

Conjecture 4.2.6 has a fundamental consequence: Knowledge of Σ and the respective parameters σ_G and μ for each time series allows us to translate any assortment of interest rate time series into a single additive Σ -world.

Note that since our model is based on a sampling mechanism, historical realizations will serve as random variates so that we have no need to know or even approximate the distribution ϕ .

In order to serve its purpose in the context of our simulation model, it is of course of little relevance whether conjecture 4.2.6 actually holds true in the strict mathematical sense (whatever this would mean in a real-world application). Deguillaume et al. (2013) and Deguillaume (2009) show the validity of the implications of conjecture 4.2.6 numerically using historical interest rate data. The results are very encouraging and support the particular design of Deguillaume's yield curve evolution model.

Further details on conjecture 4.2.6 can be found in Deguillaume et al. (2013).

4.2.4 Calibration of the Functions σ and Σ

Since according to conjecture 4.2.6 the parameters y_L , y_R and K should be universal across all interest rate time series, we will calibrate Σ to multiple time series simultaneously. The following method has been developed by Deguillaume (2009).

Let

$$Y^1, \dots, Y^n$$

denote historical time series of interest rates, and let $N_i \in \mathbb{N}$ be the length of the i th time series.

In order to find optimal parameters y_L , y_R , K , we will employ a calibration metric.

For a given set of parameters

$$\pi = (y_L, y_R, \Sigma),$$

let Σ_π be the corresponding Σ -function. Furthermore, for each time series Y^i let

$$\Delta_\pi^i(j) = \Sigma_\pi(Y^i(j+1)) - \Sigma_\pi(Y^i(j)), \quad 1 \leq j \leq N_i - 1,$$

and define μ_π^i to be the mean, $\sigma_{\pi, Gi}$ to be the standard deviation and $\sigma'_{\pi, Gi}$ to be the mean absolute deviation of each sequence

$$\Delta_\pi^i(1), \Delta_\pi^i(2), \dots, \Delta_\pi^i(N_i - 1), \quad i = 1, \dots, n.$$

The calibration metric will be defined as

$$\sum_{i=1}^n \frac{1}{\sigma'_{\pi, Gi}} \sum_{j=1}^{N_i-1} |\Delta_\pi^i(j) - \mu_\pi^i| - \sigma'_{\pi, Gi}. \quad (4.9)$$

Optimal parameters $\pi^* = (y_L, y_R, K)$ are found by minimizing this metric over the set

$$\{0 < y_L < y_R, K > 0\}.$$

Furthermore, we define

$$\sigma_{Gi} := \sigma_{\pi^*, Gi}.$$

The idea behind this approach is as follows: If Σ is working as intended, for any converted series

$$\Sigma(Y^i(1)), \Sigma(Y^i(2)), \dots, \Sigma(Y^i(N_i - 1)), \quad i = 1, \dots, n,$$

the magnitude of 1-day movements $\Delta_\pi^i(j)$ should be similar across all dates j . If the start and end points $Y^i(1)$ and $Y^i(N)$ differ, the series has an inherent drift of

$$\mu_\pi^i = \frac{Y^i(N_i) - Y^i(1)}{N_i - 1},$$

which should be excluded from our consideration. It is corrected for by the subtraction

$$\Delta_\pi^i(j) - \mu_\pi^i,$$

which appears in the formula for σ'_G and in eq. (4.9).




Figure 4.2: Illustration of a σ -function for $y_L = 1.07\%$, $y_R = 6.83\%$, $K = 19.83$. The function has been calibrated to time series of nominal UK government liability curve data and USD swap rates. The UK data spans the time period of 1979 through 2015. It has been obtained from the website of the Bank of England^a. The USD data has been obtained from the website of the Federal Reserve Bank of St. Louis^b. It pertains to the time period 2000 through 2014. Both sets of data cover eight tenors ranging from 1 through 30 years.

^a<http://www.bankofengland.co.uk/statistics/Pages/yieldcurve/archive.aspx>

^b<https://research.stlouisfed.org/fred2/>

The terms $|\Delta_\pi^i(j) - \mu_\pi^i| - \sigma'_{\pi, Gi}$ in eq. (4.9) represent by how far the drift-corrected absolute movements at each position j in each series Y^i deviate from the average drift-corrected movement $\sigma'_{\pi, Gi}$, which is exactly what we wish to minimize. Note that in eq. (4.9) no division by $N_i - 1$ takes place, which means that for each time series the deviations are summed (not averaged), and thus each series is weighted proportionally to the number of points it contains.

The division by σ'_G ensures that a time series is not weighted more strongly the higher its average level of (drift-corrected) deviations is.

Section 4.2.4 shows the result of applying the calibration mechanism described above to USD and UK data of swap rates and government yields, respectively.

Notation 4.2.7. In the following we will write $\bar{y} = \Sigma(y)$ as well as $\bar{Y} = \Sigma(Y)$.

4.2.5 Modeling the Long Rate

The Σ -function will be used to transfer all historical time series into a single additive world, and most of the modeling will take place there.

We will define separate — but similar — models for the (additive world) long rate and for the remaining tenors. In both cases we will employ a discretized version of the Ornstein-Uhlenbeck process using daily time-steps, and take as random variates historical observations. While the process for the long rate will be evolved independently, we will not model the remaining rates directly but instead simulate their spreads with regard to the long rate.

The long rate model described here has been presented in Deguillaume (2009).

Let τ_l denote the long rate tenor. Then the stochastic process of simulated one-day changes will be evolved by a discrete variation of an Ornstein-Uhlenbeck model. In particular,

$$\bar{y}^{\tau_l}(j+1) - \bar{y}^{\tau_l}(j) = a^{\tau_l}(\bar{b}^{\tau_l} - \bar{y}^{\tau_l}(j)) + v^{\tau_l}\Delta^{\tau_l}(j), \quad j = 0, 1, \dots \quad (4.10)$$

We have used the following notation:

- $\Delta^{\tau_l}(j) = \bar{Y}^{\tau_l}(\omega_j+1) - \bar{Y}^{\tau_l}(\omega_j) - \delta^{\tau_l}$ is the drift-corrected, randomly selected historical one-day increment assigned to day j , where δ^{τ_l} is defined by

$$\delta^{\tau_l} := \frac{1}{N} [\bar{Y}^{\tau_l}(N) - \bar{Y}^{\tau_l}(0)] \quad (4.11)$$

and serves the purpose of ensuring that the sampled historical residuals have no inherent drift.

- $a^{\tau_l} \in (0, 1)$ is the mean-reversion speed in the additive world.
- \bar{b}^{τ_l} is the deterministic mean which the additive world rate \bar{y}^{τ_l} reverts to. This could be a constant such as the spot rate $\bar{y}^{\tau_l}(0)$ or any other pre-defined value, or we could even give \bar{b}^{τ_l} a more complex shape and for example let it depend on the current level $\bar{y}^{\tau_l}(j)$. For the purpose of our hedging experiment, we will content ourselves with a constant.
- v^{τ_l} is a volatility adjustment parameter. It has the purpose of compensating for the variance-dampening effect which is caused by the mean reversion force. Without an adjustment, the simulated series would have a lower variance than the historical increments.

Remark 4.2.8. *One might be tempted to sample random variates with an approach such as*

$$\Delta^{\tau_l}(j) = \frac{1}{v^{\tau_l}} [\bar{Y}^{\tau_l}(\omega_j+1) - \bar{Y}^{\tau_l}(\omega_j) - a^{\tau_l}(\bar{b}^{\tau_l} - \bar{Y}^{\tau_l}(\omega_j))].$$

In this case, $\Delta^{\tau_l}(j)$ would be defined as the random variate that would be necessary to explain the historical movement under the assumption that the original time series did indeed follow the process given by eq. (4.10). Unfortunately, this would produce very unrealistic simulations since actual interest time series can in general not be explained by this (or any other similarly tractable) process with high enough accuracy (otherwise the evolution model would also be somewhat redundant). In section 4.3.3, a similar argument will be made with regard to the volatility surface, and the problem will be illustrated on the basis of actual data and simulations.

4.2.6 Modeling the Remaining Tenors

In order to obtain plausible synthetic yield curves, it is not only necessary that each individual rate evolves in a realistic manner. We also need to take care that all rates evolve reasonably in relation to one another.

Therefore, we will apply the above model to the (additive world) spreads between each of the remaining rates and the long rate. In particular, we evolve

$$[\bar{y}^\tau(j+1) - \bar{y}^{\tau_l}(j+1)] - [\bar{y}^\tau(j) - \bar{y}^{\tau_l}(j)] = a^\tau [\bar{b}^\tau - (\bar{y}^\tau(j) - \bar{y}^{\tau_l}(j))] + v^\tau \Delta^\tau(j). \quad (4.12)$$

The parameters in eq. (4.12) are defined as follows:

- $\Delta^\tau(j) = [\bar{Y}^\tau(\omega_j + 1) - \bar{Y}^{\tau_l}(\omega_j + 1)] - [\bar{Y}^\tau(\omega_j) - \bar{Y}^{\tau_l}(\omega_j)] - \delta^\tau$ is the drift-corrected, randomly selected historical one-day increment assigned to day j , where δ^τ is defined by

$$\delta^\tau = \frac{1}{N} [(\bar{Y}^\tau(N) - \bar{Y}^{\tau_l}(N)) - (\bar{Y}^\tau(0) - \bar{Y}^{\tau_l}(0))].$$

- $a^\tau \in (0, 1)$ is the mean-reversion speed in the additive world.
- \bar{b}^τ is the deterministic mean which the additive world spread $\bar{y}^\tau - \bar{y}^{\tau_l}$ is reverting to. Just as in eq. (4.10), we will model \bar{b}^τ using a constant.
- v^τ is a volatility adjustment parameter which serves the same purpose as v^{τ_l} does in eq. (4.10).

We note that Deguillaume (2009) chooses a slightly different approach for the remaining rates. In particular, he samples the historical residuals $\Delta^\tau(j)$ directly from the series \bar{Y}^τ while letting the mean reversion force apply to the spreads $\bar{y}^\tau - \bar{y}^{\tau_l}$ (as in eq. (4.12)).

4.2.7 Estimating the Volatility Adjustment Parameters

The mean reversion forces in the models described above have a variance-dampening effect so that the simulated series would have a lower variance than their historical counterparts. In order to offset this, the historical samples Δ^{τ_l} and Δ^τ are multiplied with correction factors v^{τ_l} and v^τ , respectively.

To estimate these terms, we follow Deguillaume (2009) and consider a regular continuous-time Ornstein-Uhlenbeck process

$$dX(t) = a^*(b - X(t))dt + \nu dW(t), \quad (4.13)$$

where W is a standard Brownian motion. It is a standard result that eq. (4.13) has the exact solution

$$X(t) = X(0)e^{-a^*t} + b(1 - e^{-a^*t}) + \nu \int_0^t e^{a^*(s-t)} dW(s). \quad (4.14)$$

Hence, $X(t)$ is normally distributed with expectation

$$\mathbb{E}[X(t)] = X(0)e^{-a^*t} + b(1 - e^{-a^*t}) \quad (4.15)$$

and variance

$$\mathbb{V}[X(t)] = \nu^2 \int_0^t e^{2a^*(s-t)} ds = \frac{\nu^2}{2a^*} (1 - e^{-2a^*t}). \quad (4.16)$$

We will use eq. (4.15) to determine the mean reversion levels for the spreads of the remaining tenors when we apply the evolution model described in this section to the particular case where the variables y^τ represent forward LIBOR rates. The specific method for the choice of b will be explained in section 5.3.

Equation (4.16) will be used to approximately give the model the desired variance.

In order to accomplish this, we write the solution eq. (4.14), conditional on $X(t)$, in terms of a unit (one-day, e.g.) increment (as we did in eqs. (4.10) and (4.12)):

$$\begin{aligned} X(t+1) &= X(t)e^{-a^*t} + b(1 - e^{-a^*}) + \nu \int_t^{t+1} e^{a^*(s-(t+1))} dW(s) \\ \Leftrightarrow X(t+1) - X(t) &= (1 - e^{-a^*})(b - X(t)) + \nu \int_t^{t+1} e^{a^*(s-(t+1))} dW(s). \end{aligned}$$

This shows that the continuous-time mean reversion speed a^* corresponds to the mean reversion rate $(1 - e^{-a^*})$ in the discretized version of the process. Differently put, we need to choose

$$a^* = -\log(1 - a),$$

where a is the reversion rate of the discrete-time process, so that the continuous-time model has roughly the same variance as its discretized counterpart.

Consider the evolution of the long rate process (the argument is identical for the remaining rates). Define

$$\sigma_{G\tau_l}^2 := \mathbb{V}[\Delta_{\tau_l}]$$

as the empirical variance of the historical residuals

$$\Delta_{\tau_l}(j), j = 0, 1, \dots$$

For a given window

$$\{jW + k \mid k = 0, 1, \dots, W - 1\},$$

we wish to estimate the variance of the W consecutive increments given by our sampling mechanism.

The one-day increments $\Delta^{\tau_l}(k)$ are not independent. Still, an application of the central limit theorem shows that for large n the sum of increments

$$\sum_{k=0}^{n-1} \Delta^{\tau_l}(k)$$

is approximately normally distributed with expectation zero. This encourages us to approximate

$$\sigma_{G\tau_l}^2 \approx \frac{1}{W} \mathbb{V} \left[\sum_{k=jW}^{(j+1)W-1} \Delta^{\tau_l}(k) \right],$$

so that the variance of W consecutive sampled historical increments is approximately given by $\sigma_{G\tau_l}^2 W$.

On the other hand, with the help of eq. (4.16) we estimate the variance of W consecutive increments in our Ornstein-Uhlenbeck type process as

$$(v^{\tau_l})^2 \sigma_{G\tau_l}^2 \frac{1 - e^{-2a^{\tau_l*}W}}{2a^{\tau_l*}},$$

where we have assumed that the mean reversion will scale the variance of the historical one-day increments in the same way as it scaled the variance of the Wiener increments in a regular Ornstein-Uhlenbeck process.

Setting equal the variance of W consecutive sampled increments to the predicted model variance yields

$$\sigma_{G\tau_l}^2 W = (v^{\tau_l})^2 \sigma_{G\tau_l}^2 \frac{1 - e^{-2a^{\tau_l*}W}}{2a^{\tau_l*}}.$$

Therefore, we will define

$$v^{\tau_l} = \sqrt{\frac{2a^{\tau_l*}W}{1 - e^{-2a^{\tau_l*}W}}}$$

for the long rate, and

$$v^{\tau} = \sqrt{\frac{2a^{\tau*}W}{1 - e^{-2a^{\tau*}W}}}.$$

for the remaining rates.

Deguaillume (2009) uses this type of adjustment for the modeling of the long rate, and demonstrates the quality of the approximations numerically.

4.3 Evolving the Volatility Surface

4.3.1 General Idea

In addition to the yield curve simulation presented in the previous section, we wish to obtain synthetic realizations of the implied volatility surface. Similar to the yield curve model, the basis will be a historical sampling mechanism.

Since we are basically using two evolution models to simulate a single “world”, it is of course essential that our approach satisfies the no-arbitrage condition. For this reason we cannot merely take the historical sampling model described above and simultaneously apply it to both interest rates and implied volatilities without making any further adjustments.

Instead, we will develop an approach that extends the semi-parametric historical simulation approach put forward by Rebonato et al. (2005) to the joint evolution of yield curves and implied volatilities. The method described here is based on the work of Andreichenko (2011). Unfortunately, a direct adoption of Andreichenko’s framework did not produce synthetic implied volatilities whose statistical properties were sufficiently similar to historical market data. For this reason, we will make certain changes to the approach.

The LMM-SABR is at the center of the method. A fully calibrated LMM-SABR provides an arbitrage-free parameterization of the implied volatility surface. Andreichenko (2011) proceeds by calibrating the model to caplet implied volatilities for each date in his data set. For each date this yields parameters a, b, c, d for the function g , parameters $\alpha, \beta, \gamma, \delta$ for the function h , SABR correlations ρ^τ as well as initial loadings and correction terms k_0^τ and ξ^τ for each caplet tenor τ . He goes ahead and evolves these parameters with the help of a window sampling mechanism. The resulting set of synthetic parameters can be translated back into implied volatilities via the approximations presented in section 3.4.2.

Note that the remaining entries of the LMM-SABR correlation matrix are not relevant for the pricing of caps.

If one takes sufficient care that the synthetic parameters stay within plausible constraints, Andreichenko's method does indeed yield realistic implied volatility surfaces for any given future date. However, the approach has a significant drawback in the context of our experiment, in the sense that the artificial implied volatilities resulting from the synthetic parameters are far more volatile than their historical counterparts. We illustrate the reason for this by examining the LMM-SABR volatility function

$$g(\tau) = (a + b\tau) \exp(-c\tau) + d.$$

Consider a caplet on the 2Yr \rightarrow 1Yr forward rate. As we have seen in section 3.4.2, the model implied volatility of this caplet depends primarily on the root-mean-squared volatility

$$\hat{g}(2) = \sqrt{\frac{1}{2} \int_0^2 g(\tau)^2 d\tau}. \quad (4.17)$$

If we assume that g takes a humped shape — which in fact it does over the majority of our data set — the result will strongly depend on the location of the maximum, τ^* , and its value, $g(\tau^*)$. As we have seen in section 3.4.1, the maximum is located at

$$\tau^* = \frac{1}{c} - \frac{a}{b},$$

and it holds that

$$g(\tau^*) = \frac{b}{c} \exp\left(\frac{ac}{b} - 1\right) + d.$$

Due to the division by b and c in both expressions, even small movements in b and c will have large impacts on $g(\tau^*)$ and hence on $\hat{g}(2)$. Without exerting strict control over the joint movements of b and c , the resulting implied volatility can show large jumps between two dates. When holding a delta neutral portfolio, even a single unrealistic jump in implied volatilities will cause unrealistically large hedging errors.

One possible remedy for this problem would be to re-parameterize g . Instead of evolving the parameters a, b, c and d , one might try to evolve τ^* or $g(\tau^*)$ directly. Unfortunately, this is no option for us since g does not always attain a humped shape (see section 4.3.6), so that τ^* and $g(\tau^*)$ will not always exist.

However, we will keep with the general idea of evolving parameters that control the actual function values of g and h more directly, rather than simulating the two sets of *abcd*-parameters as individual time series. In particular, we will simulate the set of SABR state variables σ_0^τ and SABR parameters ν^τ for each tenor τ . We will then apply the calibration mechanism described in section 3.4.1 to obtain synthetic functions g and h . This has the additional advantage that σ_0 and ν have more accessible economic interpretations than b and c had in the original approach.

In line with Andreichenko (2011), we will also evolve the SABR correlations ρ^τ , albeit in a slightly different way. We will do without the evolution of k_0^τ and ξ^τ and instead set them to 1, since they are not necessary for the purpose of our experiment. The reasoning behind this will be explained in section 4.3.4.

With the help of this approach we will parameterize the swaption volatility surface for each date in our market history. The simulation of the entire “world” — consisting of the yield curve and the volatility surface — is carried out by simultaneously evolving the above parameters through a historical sampling mechanism and evolving the yield curve using the approach described in section 4.2. The resulting synthetic yield curves and (LMM-)SABR parameters are then combined to produce an arbitrage-free artificial scenario of a future joint evolution of the yield curve and volatility surface.

Note that this approach is completely independent of any predictions the LMM-SABR makes with regard to the forward rate dynamics or the prices of any other derivatives except for the swaptions we use as input to the calibration. The only feature of the LMM-SABR that we are using at this point is the fact that we are able to turn the swaption implied volatility surface (or a subset thereof) into a set of LMM-SABR parameters which can in turn be transformed back into the original volatility surface with sufficient accuracy.

4.3.2 Parameterization of Historical Data

We start by performing a SABR calibration (under $\beta = 0.5$) to the implied volatilities of one-period swaptions for each date in our historical data set (or rather take the pre-calibrated SABR parameters we are given as input data, cf. section 5.1.1). For each date t_j and each tenor τ this yields SABR parameters

$$\sigma_0^\tau(j), \nu^\tau(j), \rho^\tau(j), \quad \tau \in \{0.5, 1, 2, \dots, 10\}.$$

Here we have also included the information on 6m \rightarrow 1Yr swaptions which will help us calibrate the short ends of the volatility functions g and h .

Ultimately we will require for each forward rate f^τ the corresponding SABR correlation $\rho^\tau(j)$, $\tau = 1, 2, \dots, 10$. As for the remaining parameters, we will only work with the volatility functions g and h , rather than with the SABR parameters $\sigma_0^\tau(j)$ or $\nu^\tau(j)$ directly. For this reason and in order to reduce the computational burden, we will only evolve the five original (non-interpolated) time series of $\sigma_0^\tau(j)$ and $\nu^\tau(j)$, where $\tau \in \{0.5, 1, 2, 5, 10\}$.

Overall, we will simulate 20 time series of parameters, namely

$$\begin{aligned} \sigma_0^\tau(j), \quad \tau &\in \{0.5, 1, 2, 5, 10\}, \\ \nu^\tau(j), \quad \tau &\in \{0.5, 1, 2, 5, 10\}, \end{aligned}$$

$$\rho^\tau(j), \quad \tau \in \{1, 2, \dots, 10\}.$$

We will mainly stay with the same notation as in section 4.2. Furthermore, the historical time series of parameters / state variables will be denoted by σ_0^τ , ν^τ and ρ^τ . Their synthetic counterparts will be denoted by $\tilde{\sigma}_0^\tau$, $\tilde{\nu}^\tau$ and $\tilde{\rho}^\tau$, respectively.

4.3.3 Evolution of σ_0^τ and ν^τ

Through statistical analysis, Andreichenko (2011) showed that forward rates, caplet Black implied volatilities, the caplet SABR initial volatilities σ_0 and each of the LMM-SABR *abcd*-parameters exhibit mean reversion.

Similar to the yield curve evolution model presented in section 4.2, we wish to interpret the historical time series of parameters in a way that allows us to generate synthetic paths by re-applying past realizations of a stochastic process. To achieve this, we need a time series model that allows to extract these stochastic realizations.

First, we consider the set of SABR state variables $\sigma_0^\tau(j)$ and parameters $\nu^\tau(j)$.

Andreichenko (2011) uses an AR(1) filter for the modeling of the LMM-SABR *abcd*-parameters corresponding to the volatility functions g and h , i.e., he uses a process of the type

$$X_t = c + \varphi X_{t-1} + \varepsilon_t, \quad (4.18)$$

where c and φ are constants. In a regular AR(1) process the random variates ε_t are *white noise*, i.e., a sequence of independent random variables with zero mean and finite variance. In the evolution model, however, they are replaced by historical residuals.

A short review of the mathematical foundations of AR(1) processes can be found in the appendix (section A.1).

He notes that the above AR(1) process is equivalent to the discretized version of a continuous Ornstein-Uhlenbeck process given that $\varphi \in (0, 1)$ and under the assumption that ε_t is normally distributed (cf. section A.1).

We will go a slightly different route but we will first understand why Andreichenko's method is not adequate for the hedging experiment we are conducting. More importantly, the following analysis shows why we cannot sample historical residuals under the assumption that the historical time series did indeed follow the stochastic process we use in our model — neither when evolving the yield curve nor when evolving the volatility surface (cf. remark 4.2.8).

After calibrating the above process to each of the historical time series of *abcd*-parameters, Andreichenko (2011) samples historical residuals and re-applies them via a window sampling approach. For each parameter π (e.g., $\pi = a$) this means historical residuals are obtained via the formula

$$\varepsilon_j = \pi_{\omega_j+1} - c - \varphi \pi_{\omega_j}, \quad (4.19)$$

where ω_j is a historical date obtained through the sampling mechanism (cf. definition 4.2.2) and where c and φ are the estimated AR(1) parameters pertaining to the time series π_j .




Figure 4.3: Historical evolution of SABR σ_0 obtained by calibrating the SABR model to market prices of 1Yr \rightarrow 1Yr swaptions.

An artificial time series $\tilde{\pi}_j$ is then created by iteratively evolving

$$\tilde{\pi}_{j+1} = c + \varphi \tilde{\pi}_j + \varepsilon_j,$$

starting with some initial value $\tilde{\pi}_0$.

While the approach is indeed capable of producing a realistic term structure of volatility for a given future date, its dynamics would look very different from actual historical data. The reason for this lies in the fact that the historical parameter evolution cannot actually be described by an Ornstein-Uhlenbeck process with sufficient enough accuracy.

Consider the time series of SABR σ_0 pertaining to the 1Yr \rightarrow 1Yr swaption. The historical evolution is shown in fig. 4.3.

Fitting an AR(1) filter, i.e., a discretized Ornstein-Uhlenbeck process, to this time series would imply a certain mean reversion speed and mean reversion level. The historical residuals, which are sampled according to eq. (4.19), represent the random variates necessary to explain the actual movements under the assumption that the process did indeed follow a process of this type. As can be seen in fig. 4.3, for roughly the first two years the state variable σ_0 stayed on a relatively low level after which it increased sharply and subsequently seems to fluctuate around a much higher level.

To illustrate the problem, we fit an AR(1) model to the above time series via a conditional maximum likelihood approach (cf. section A.1). The resulting parameters are

$$c = 2.634 \cdot 10^{-4}$$

$$\varphi = 0.996151.$$




Figure 4.4: 80-day moving average of AR(1) residuals $\varepsilon_j = \sigma_0(j+1) - c - \varphi\sigma_0(j)$.

This corresponds to Ornstein-Uhlenbeck parameters (cf. eq. (A.3))

$$a = -\log(\varphi) = 0.3857\%$$

$$b = \frac{c}{1 - \varphi} = 6.8421\%.$$

We can intuitively understand the issue through an optical examination of fig. 4.3. Even though the process should revert to the long term mean of about 6.8%, none of the reversion force is visible during the first two years. This would imply that the AR(1) residuals were consistently below average during this time, and consistently above average later on.

Figure 4.4 shows the 80-day moving average of historical AR(1) residuals. One can see that in many potential 80-day windows the residuals are strongly biased in a particular direction. Using these residuals in our window sampling approach would produce zigzag-shaped synthetic time series.

The effect has been illustrated in fig. 4.5 which shows a particular 1000-day sample path which has been created by applying the described AR(1) method within our window sampling framework for a window length of 80 days.

Since the SABR σ_0 primarily controls the level of the implied volatility smile, the resulting time series of implied volatilities of the 1Yr into 1Yr ATM swaption would have roughly the same shape. This would have an enormous effect on the hedging performance, especially when only hedging delta risk.

After having understood the need to deviate from Andreichenko's (2011) methodology, we will instead proceed by evolving the parameters σ_0^τ and ν^τ . The approach will be in line with the model we use to simulate the yield curve.

In particular, we evolve the parameters of the longest tenor, $\tilde{\sigma}_0^{10\text{Yr}}$ and $\tilde{\nu}^{10\text{Yr}}$, separately




Figure 4.5: A simulation path obtained by applying the AR(1) method via a window sampling mechanism. This particular path illustrates the zigzag-behavior which can occur if the historical time series does not consistently follow the dynamics postulated by the estimated AR(1) parameters. Compare this with the historical time series depicted in fig. 4.3.

via a discrete variation of an Ornstein-Uhlenbeck process:

$$\tilde{\sigma}_0^{10\text{Yr}}(j+1) - \tilde{\sigma}_0^{10\text{Yr}}(j) = a_{\sigma_0}^{10\text{Yr}}(b_{\sigma_0}^{10\text{Yr}} - \tilde{\sigma}_0^{10\text{Yr}}(j)) + v_{\sigma_0}^{10\text{Yr}} \Delta_{\sigma_0}^{10\text{Yr}}(j), \quad (4.20)$$

$$\tilde{\nu}^{10\text{Yr}}(j+1) - \tilde{\nu}^{10\text{Yr}}(j) = a_{\nu}^{10\text{Yr}}(b_{\nu}^{10\text{Yr}} - \tilde{\nu}^{10\text{Yr}}(j)) + v_{\nu}^{10\text{Yr}} \Delta_{\nu}^{10\text{Yr}}(j). \quad (4.21)$$

As for the remaining parameters, we will evolve the spreads with regard to the longest tenor. For $\tau \in \{0.5\text{Yr}, 1\text{Yr}, 2\text{Yr}, 5\text{Yr}, 10\text{Yr}\}$ we have

$$\begin{aligned} & [\tilde{\sigma}_0^{\tau}(j+1) - \tilde{\sigma}_0^{10\text{Yr}}(j+1)] - [\tilde{\sigma}_0^{\tau}(j) - \tilde{\sigma}_0^{10\text{Yr}}(j)] \\ &= a_{\sigma_0}^{\tau} [b_{\sigma_0}^{\tau} - (\tilde{\sigma}_0^{\tau}(j) - \tilde{\sigma}_0^{10\text{Yr}}(j))] + v_{\sigma_0}^{\tau} \Delta_{\sigma_0}^{\tau}(j), \end{aligned} \quad (4.22)$$

and

$$\begin{aligned} & [\tilde{\nu}^{\tau}(j+1) - \tilde{\nu}^{10\text{Yr}}(j+1)] - [\tilde{\nu}^{\tau}(j) - \tilde{\nu}^{10\text{Yr}}(j)] \\ &= a_{\nu}^{\tau} [b_{\nu}^{\tau} - (\tilde{\nu}^{\tau}(j) - \tilde{\nu}^{10\text{Yr}}(j))] + v_{\nu}^{\tau} \Delta_{\nu}^{\tau}(j). \end{aligned} \quad (4.23)$$

The parameters in eqs. (4.20) to (4.23) are defined as in the yield curve evolution model. In particular, both sets of volatility adjustment parameters $v_{\sigma_0}^{\tau}$ and v_{ν}^{τ} are defined via

$$v^{\tau} = \sqrt{\frac{2a^*W}{1 - e^{-2a^*W}}}, \quad \tau \in \{0.5\text{Yr}, 1\text{Yr}, 2\text{Yr}, 5\text{Yr}, 10\text{Yr}\},$$

where

$$a^* = -\log(1 - a)$$

is the adjusted mean reversion rate pertaining to each particular tenor.

Historical samples are obtained as before, e.g.,

$$\Delta_{\sigma_0}^{10\text{Yr}}(j) = \sigma_0^{10\text{Yr}}(\omega_j + 1) - \sigma_0^{10\text{Yr}}(\omega_j) - \delta_{\sigma_0}^{10\text{Yr}},$$

where $\delta_{\sigma_0}^{10\text{Yr}}$ is the average one-day movement in the historical time series (cf. eq. (4.11)).

4.3.4 Treatment of k_0^τ and ξ^τ

As described in section 3.4.2, the LMM-SABR parameters k_0^τ and ξ^τ are close to unity if the calibration provides a good fit to the empirical volatility surface. For the particular case of our hedging experiment, we will set them to 1 but we note that this approach is not viable in every context.

Assume we are starting the simulation under initial recovery of today's observed market prices. E.g., we might wish to simulate the future evolution of a number of specific swaptions we have acquired today. In this context we might require the initial LMM-SABR parameterization to provide near-perfect recovery of our particular swaptions' implied volatility. In general, this can only be achieved by incorporating non-unity factors $k_0^\tau(0)$ and $\xi^\tau(0)$. If we decided to simply set them to 1 when creating subsequent synthetic realizations, as Andreichenko (2011) points out, this would cause a deterministic jump in the first step of the evolution.

Andreichenko (2011) found a strong mean reversion to 1 for both variables and decided to decay $\log(k_0^\tau(j))$ and $\log(\xi^\tau(j))$ asymptotically to zero.

Since we are creating our own synthetic market data, we will set them to 1 *prior* to the start of the evolution and for all future dates. In our experiment no deterministic jump will take place since the initial portfolios will already be valued according to model prices computed under $\log(k_0^\tau(0)) = \log(\xi^\tau(0)) = 1$.

4.3.5 Evolution of SABR ρ

When modeling the SABR ρ we need to keep in mind that it represents a correlation, and any approach we follow will have to make sure ρ stays within the interval $[-1, 1]$.

Andreichenko proposes to evolve Fisher z -transformed values of ρ :

$$z(\rho) = \frac{1}{2} \ln \left(\frac{1 + \rho}{1 - \rho} \right).$$

The function z is a bijection from the interval $(-1, 1)$ to the real numbers \mathbb{R} (see fig. 4.6). The inverse is

$$z^{-1}(u) = \frac{\exp(2u) - 1}{\exp(2u) + 1}.$$

The idea behind this transformation is explained in Fisher (1934). If pairs

$$(x_i, y_i), \quad i = 1, 2, \dots, N,$$




Figure 4.6: Fisher z -function.

have been drawn from a bivariate normal distribution, then the empirical correlation coefficient r is not normally distributed (neither exactly, nor approximately, nor in the limit as $N \rightarrow \infty$). Furthermore, its standard deviation depends on the true value of ρ :

$$\sigma_r = \frac{1 - \rho^2}{\sqrt{N - 1}}.$$

As Fisher (1934) shows, for an increasing sample size N , the z -transformed value $z(r)$ rapidly tends to a normal distribution, and its standard deviation is independent of ρ :

$$\sigma_{z(r)} = \frac{1}{\sqrt{n - 3}}.$$

This allows for straightforward hypothesis testing and the computation of confidence intervals.

These properties are also beneficial for the purpose of our evolution model. Especially the fact that empirical correlations are level-independent motivates (and allows) to use a level-independent sampling mechanism.

Andreichenko (2011) evolves the SABR ρ by z -transforming the time series of past values, extracting the daily changes, re-applying them in the context of the sampling model and transforming the results back via z^{-1} .

This rather simple approach has two drawbacks: Without applying a correction term, the historical residuals have an inherent drift so that the z -transformed values will tend towards $+\infty$ or $-\infty$, and hence the SABR ρ will tend towards $+1$ or -1 . This effect can of course be remedied by applying a drift correction term. But the more fundamental problem is that despite a drift correction of this kind, in the long run the simulated ρ will happen to be in the vicinity of 1 in absolute value for the majority of the time.

To understand why this is the case, consider the sum S of W consecutive increments, i.e., a full window of length W resulting from applying the drift-corrected sampling mechanism described above. S is drawn from the discrete uniform distribution on the set

$$\{z(\rho(k + W \bmod N)) - z(\rho(k)) \mid k \in 0, 1, \dots, N - 1\}$$

where

$$k + W \bmod N := \begin{cases} k + W & \text{if } k + W \leq N - 1 \\ k + W - N & \text{otherwise.} \end{cases}$$

Due to the central limit theorem, the sum of a large number of independent draws from this set,

$$S_1, S_2, \dots, S_m,$$

will be approximately normally distributed

$$S_1 + S_2 + \dots + S_m \stackrel{\text{approx.}}{\sim} m\mathcal{N}(0, \mathbb{V}[S]).$$

Even though the process of z -transformed values will almost surely return to its point of origin at some point in time, the variance of a particular realization will rapidly grow as the process evolves, which means the likelihood of finding a future simulated z -transformed ρ within any given interval around zero, say $(-0.99, 0.99)$, will tend towards zero.

This is a very undesirable feature if one plans on performing a long-term evolution (which we do). Therefore, we cannot do without including mean reversion in the approach. Once again, we make use of the discrete variation of an Ornstein-Uhlenbeck process to evolve each of the z -transformed SABR $\tilde{\rho}^\tau$ as

$$z(\tilde{\rho}^\tau(j + 1)) - z(\tilde{\rho}^\tau(j)) = a_\rho^\tau(b_\rho^\tau - z(\tilde{\rho}^\tau(j))) + v_\rho^\tau \Delta_\rho^\tau(j), \quad \tau \in \{1, 2, \dots, 10\}, \quad (4.24)$$

where the volatility adjustment parameters v_ρ^τ are defined in the same way as before and a_ρ^τ as well as b_ρ^τ are constants yet to be determined.

The historical samples Δ_ρ^τ are obtained as described in section 4.2.5.

4.3.6 Obtaining Synthetic Implied Volatilities

To translate the simulated set of parameters and state variables into implied volatilities, we will make use of the LMM-SABR. For any given date j we have a set of synthetic values

$$\begin{aligned} \tilde{\sigma}_0^\tau(j), \quad \tau \in \{0.5, 1, 2, 5, 10\}, \\ \tilde{\nu}^\tau(j), \quad \tau \in \{0.5, 1, 2, 5, 10\}, \\ \tilde{\rho}^\tau(j), \quad \tau \in \{1, 2, \dots, 10\}. \end{aligned}$$

We will use these to calibrate the LMM-SABR functions g and h according to the approach described in section 3.4.2.

We will impose certain restrictions on the admissible sets of $abcd$ -parameters for both g and h in order to obtain plausible term structures of volatility.

To gain an intuition with regard to the shapes of these functions, we examine the historical time series of σ_0^τ and ν^τ . For this, we perform a SABR calibration to the implied volatilities

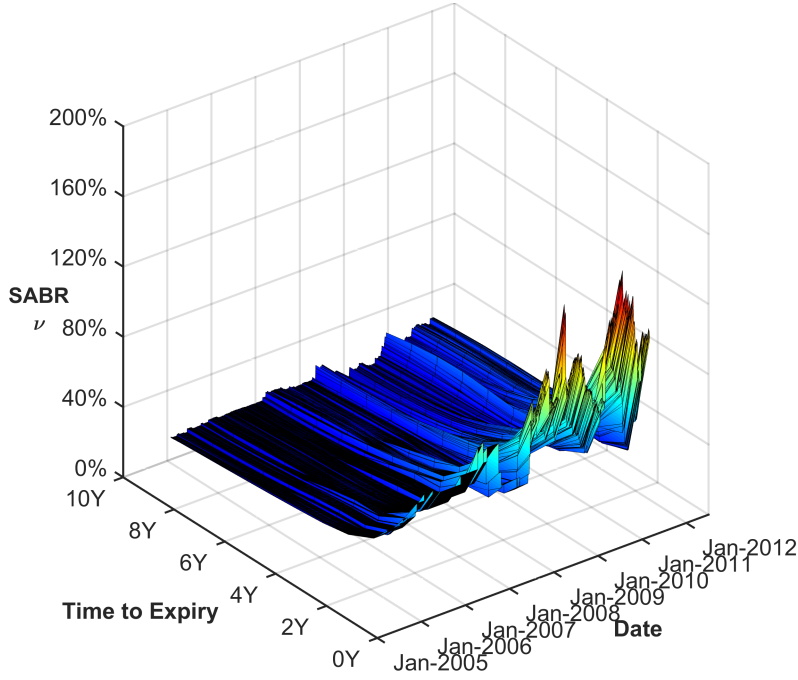


Figure 4.7: Evolution of SABR ν obtained by calibrating the SABR model to market prices of $n\text{Yr} \rightarrow 1\text{Yr}$ swaptions ($\beta = 0.5$).

of one-period swaptions over the course of our entire history of data. The resulting ν^τ are shown in fig. 4.7. As we can see, the SABR ν decreases monotonically with increasing expiry, except for very few dates at the end of our history.

As explained in Rebonato et al. (2009), this is a common observation in most market states. Therefore, we will impose the condition

$$\beta = 0,$$

and thus force the synthetic volatility of volatility function h into a monotonically decreasing shape.

In addition to letting $\beta = 0$, Andreichenko (2011) sets γ to a constant value. While the remaining parameters α and δ determine the values $h(0)$ and $\lim_{t \rightarrow \infty} h(t)$, one has no control over the speed of decay without the ability to adjust γ . Since our market data covers “normal” as well as very “excited” states during the financial crisis of 2007–2008 and the European sovereign debt crisis, this is not sufficient for our purposes.

Similarly, an examination of the SABR σ_0 parameters provides us with information about which shape we can expect for the deterministic part of the volatility, $g(t)$. Figure 4.8 shows that we can expect both humped shapes (under “normal” market conditions) as well as monotonically decreasing shapes (under “excited” market conditions) for g .




Figure 4.8: Evolution of SABR σ_0 obtained by calibrating the SABR model to market prices of $n\text{Yr} \rightarrow 1\text{Yr}$ swaptions ($\beta = 0.5$).

Therefore, we will work with the full parameter set for g . Hence, the LMM-SABR calibration provides us with the set

$$\{a, b, c, d, \alpha, \gamma, \delta\}$$

for each particular day.

As described in section 4.3.4, the resulting initial loadings and correction terms k_0^τ and ξ^τ will be set to 1. Hence, the calibration method applied here will yield relatively stable and plausible term structures of volatility, even if the individual sets of σ_0^τ and ν^τ show certain kinks and fluctuations in the expiry dimension.

In combination with the yield curve and the evolved SABR ρ^τ , these parameters allow us to recover the implied volatility smile for each maturity $\tau \in \{1, 2, \dots, 10\}$ via Hagan's formula eq. (3.27). By interpolating the parameters with regard to the time of maturity we will obtain the full implied volatility surface. The interpolation will be performed according to the method described in section 5.1.1.

Chapter 5

Hedging

5.1 Methodology

We will measure the delta hedging performance of our term structure models in two separate experiments. As a preliminary study we hedge within the framework of the historical data set. In particular, we examine how well each model would have performed had the hedging strategies it implies actually be used in the historical time span covered by our data set. To keep the setup realistic, it is of course essential that on any given date only information that was available before and up to this date is used. In our experiment this requirement will be fulfilled, since on each day we will calibrate the models to that day's market data only.

In our second and more important experiment, we produce synthetic evolutions of the yield curve and volatility surface via the approach developed in chapter 4. We will then apply the hedging strategies in these synthetic scenarios and again compare their performances.

In the synthetic evolution case, the hedging performance will be measured with the help of three metrics, relating to the daily hedging slippages, the maximum drawdown and the terminal portfolio values. Due to the limited amount of data, the historical hedging performance will only be analyzed according to the daily slippages.

For each of the models we consider two hedging strategies. In the first variant a delta hedge will be constructed by shifting the entire forward rate curve in parallel, and a single hedging instrument will be used to obtain a delta neutral portfolio. In the second variant, each relevant forward rate will be shifted individually, so that the hedged portfolio will contain multiple hedging instruments.

These two methods relate to the idea of *factor hedging* and *bucket hedging*. From a mathematical point of view, in a single-factor model with a single source of uncertainty, one requires only a single hedging instrument to obtain a perfect delta hedge. This method is called *factor hedging*. Practitioners on the other hand usually perform so-called *bucket hedging* (cf. Pietersz and Pelsser (2010)). In this case one uses more hedging instruments than there are factors in the model. When hedging a cap, e.g., this could mean that one uses an assortment of bonds whose maturities correspond to all cash flow dates of the option (see Driessen et al. (2003)).

Driessen et al. (2003) as well as Fan et al. (2007) find that bucket hedging outperforms factor-hedging in the context of their studies which deal with the pricing and hedging performance of term structure models in the cap and swaption markets, respectively.

We note that the one-factor hedge we perform is not exactly a “factor-hedge” for the SABR model since the model contains two stochastic drivers.

5.1.1 Data Selection and Processing

We have at our disposal time series of daily SABR parameters ranging from July 2005 through June 2011, calibrated to the USD swaption smile on each date and for a number of tenors. The available data encompasses swaptions with expiry after 6 months, 1, 2, 5 and 10 years, written on swaps of tenors 1, 2, 3, 5, 7 and 10 years, so that in total we have 30 sets of swaption data.

For the hedging experiment we will only work with a selection of $n\text{Yr} \rightarrow 1\text{Yr}$ swaptions but we will use some of the remaining data to bootstrap the full set of spanning $n\text{Yr}$ into 1Yr forward swap rates. These can also be interpreted as $n\text{Yr}$ into 1Yr forward LIBOR rates if we assume that the underlying swap has a one year tenor structure. This is of course a simplified interpretation of the data since in reality most interest rate swaps have an underlying tenor structure of 3 months or 6 months. The reason for making this simplification is two-fold. First, it will help lower the computational burden since it allows us to significantly reduce the number of parameters to be simulated compared to, say, working with a 3 months tenor structure. Second, it is a matter of availability of data. The above interpretation allows us to process the available swap rates more easily into forward LIBOR rates without having to rely too much on a potentially unstable interpolation scheme, and it allows us to interpret the $n\text{Yr} \rightarrow 1\text{Yr}$ swaptions as one-period swaptions and hence as caplets. Since we are using pre-calibrated SABR systems as inputs rather than raw option prices, we can be confident that the simplification will not affect the results on a qualitative level.

Bootstrapping of Forward LIBOR Rates

Our model will be based on a tenor structure of

$$T_0 = 0 < \dots < T_{11} = 11.$$

Hence, we need to obtain time series of forward rates f_i spanning the intervals $[T_i, T_{i+1}]$, $i = 0, 1, \dots, 10$. Furthermore, we compute the 19 year rate f_{19} to be used as the long rate in the yield curve evolution model.

Since the swaption data contains no information on rates spanning the interval $[0, 0.5]$, we use the 6 months US dollar treasury constant maturity rate obtained from the website of the Federal Reserve Bank of St. Louis. For each date this provides us with the price of the zero-coupon bond $P(0, 0.5)$.

Furthermore, we interpolate

$$P(0, 1) = \frac{P(0, 0.5)}{1 + 0.5S^{0.5, 1.5}(0)}$$

where $S^{0.5,1.5}(0)$ can be interpreted as the 1 year forward LIBOR rate spanning the interval $[0.5, 1.5]$. To obtain the remaining rates we make use of the formula

$$S^{m,n}(t) = \frac{P(t, T_m) - P(t, T_n)}{\sum_{i=m}^{n-1} \tau_i P(t, T_{i+1})},$$

cf. eq. (2.8), which allows us to solve iteratively for all remaining bond prices

$$P(0, 2), \dots, P(0, 10).$$

E.g., we can compute

$$P(0, 2) = \frac{P(0, 1)}{1 + S^{1,2}(0)}.$$

After obtaining $P(0, 3)$ in the same way, we make use of the 2Yr \rightarrow 2Yr forward swap rate to compute

$$P(0, 4) = \frac{P(0, 2) - P(0, 3)S^{2,4}(0)}{1 + S^{2,4}(0)},$$

and so forth.

The resulting bond prices can then be translated into the desired forward LIBOR rates via eq. (2.3).

To obtain the 19 year rate f_{19} , we continue the above process up to the discount bond $P(0, 20)$. This requires knowledge of further forward swap rates. In particular, we linearly interpolate the rates

$$S^{10,n}, \quad n \in \{14, 16, 18, 19\},$$

from the available 10 year forward swap rates.

Interpolation and Processing of Data

We wish to work with market data for each non-weekend day, so that each week contains exactly five business days and hence each set of five contiguous one-day increments in our data set will always span a full week and include a weekend effect.

We will therefore linearly interpolate all values we are lacking (due to bank holidays, e.g.) for each of the time series of SABR parameters and forward swap rates.

Furthermore, we wish to consistently work with a SABR β of 0.5. The original data contains different betas for certain dates. We will make the transition by transforming the SABR parameters into Black implied volatilities via Hagan's formula (eq. (3.27)) for a number of strikes ranging from $\frac{2}{3}f_{\text{atm}}$ to $\frac{3}{2}f_{\text{atm}}$ and obtaining new SABR parameters by re-calibrating to these implied volatilities a SABR system under $\beta = 0.5$.

Later on, we will require LMM-SABR correlations between each forward rate and its own volatility. As we have seen in section 3.4.2, they are retrieved by calibrating a SABR system to caplets and using the resulting SABR correlations. In order to recover a full set of them, i.e., one SABR ρ for each of the nine stochastic forward rates, it will be helpful

to have SABR data for all n Yr into 1Yr swaptions, $n = 1, 2, \dots, 10$. To obtain the missing parameters — for $n \in \{3, 4, 6, 7, 8, 9\}$ — we will apply a particular interpolation scheme in the maturity dimension.

Rather than interpolate the SABR parameters directly, we wish to interpolate implied volatilities since they have a more straightforward economic interpretation. And rather than work with absolute strikes, we will group implied volatilities of the same moneyness since they are more closely related. Assume we were instead to take the available implied volatilities for all swaptions with a fixed strike of 2%. On many dates the corresponding swaptions would range from deep in the money to deep out of the money depending on their maturity, and the implied volatilities would have very different positions on their respective smiles.

Instead, we will use a number of factors $x \in [2/3, 3/2]$, take the available implied volatilities corresponding to the strikes

$$K = xf,$$

where f is the ATM strike for each of the respective swaptions, and linearly interpolate the missing values. Afterwards, we will calibrate a SABR system to the interpolated values which will provide us with the desired SABR parameters.

Interpolation of Prices and Implied Volatilities

When hedging a certain swaption, we need to observe how its price evolves over time. Unfortunately, a time series of market quotes over the lifetime of a specific option is, in general, unavailable.

Consider the following example: In today's market data we find the implied volatility and price of an ATM cap expiring in 1 year. Say we construct a hedging strategy designed to delta hedge this cap, and assume that tomorrow we will try to evaluate the hedging slippage of our strategy. In order to determine tomorrow's value of our portfolio, we will need the price of the very cap we added today, which tomorrow will be a cap expiring in 1 year less 1 day. But in tomorrow's market data we will only find the price of a cap expiring in 1 year. Moreover, due to a change in the underlying forward rate, it is likely that the new quotes will not even contain a 1 year cap of the same strike.

Therefore, we will once again apply an interpolation scheme. Note that in our experiment we are only considering swaptions written on the particular swaps with tenors as given in our market data. This means we will use swaps with a tail of 1, 2, 3, 5, 7 or 10 years. Interpolation is necessary to value these swaptions at different points in time. Let $m \in \{1, 2, 3, 5, 7, 10\}$, and assume that at time $t_0 < T$ we wish to value a $T \rightarrow m$ swaption, i.e., a swaption expiring at time T where the underlying swap stretches over m one-year periods starting at T . Let K denote the fixed rate of this swaption.

Our historical market data contains SABR parameters for the cases

$$T \in \{0.5, 1, 2, 5, 10\}.$$

We translate these into implied volatilities $\hat{\sigma}_{0.5}, \hat{\sigma}_1, \hat{\sigma}_2, \hat{\sigma}_5$ and $\hat{\sigma}_{10}$ via Hagan's formula eq. (3.27). We use these values to linearly interpolate the desired implied volatility $\hat{\sigma}_T$.

The interpolated price of the swaption is obtained by inserting $\hat{\sigma}_T$ into Black's formula eq. (2.18). Note that this requires knowledge of the associated swap rate $S^{T, T+m}(t_0)$

which cannot be computed directly from the discrete set of forward rates

$$f_i(t_0) = f(t_0, t_0 + i, t_0 + i + 1), \quad i = 0, \dots, 10.$$

Hence, we construct an interpolated discount curve

$$t \mapsto P(t_0, t)$$

by linearly interpolating between all known values for $P(t_0, t)$. The discount curve can then be translated into forward rates $f(t_0, t, t + 1)$ for general expiries $t > t_0$ via eq. (2.3). These forward rates can in turn be used to compute $S^{T, T+m}(t_0)$.

We will employ the same approach to interpolate implied volatilities and prices in the evolution framework. For a given tail m , we compute the available synthetic implied volatilities for the available swaptions expiring at

$$t_0 + i, \quad i = 1, 2, \dots,$$

via the method described in section 3.4.3. Then we linearly interpolate the required value and insert it into Black's formula alongside the (interpolated) forward swap rate.

5.1.2 Construction of a Delta Hedged Portfolio

Both in the historical as well as the evolution framework, we proceed as follows: On each given date we construct a portfolio comprising a \$10,000 notional one-period swaption, an adequate position in hedging instruments to make the entire portfolio delta neutral and an amount of cash so that the overall portfolio has zero value. We perform one-factor and multi-factor hedges.

Consider a particular $n\text{Yr} \rightarrow 1\text{Yr}$ swaption coming into life at t_0 which is set to expire on the fixed date T . The underlying one-period swap will thus pay at $T + 1$.

For the one-factor hedge we always use as hedging instrument the very same zero-coupon bond

$$P(t_j, T + 1),$$

since it pays at the same date as the underlying swap. For the multi-factor hedge we use the discount bonds

$$P(t_j, t_j + i), \quad i = 1, \dots, n + 1.$$

In all cases we compute the deltas by a numerical finite difference method implemented in the commercial *Extreme Optimization Numerical Libraries for .NET*. Let

$$C_{n\text{Yr} \rightarrow 1\text{Yr}}(t_j)$$

denote the time t_j price of the $n\text{Yr} \rightarrow 1\text{Yr}$ swaption under the respective model.

For a one-factor hedge, we obtain the model delta

$$\Delta^{\text{Swptn}}(t_j) = \frac{\partial C_{n\text{Yr} \rightarrow 1\text{Yr}}(t_j)}{\partial f(t_j)}$$

by parallel shifting the entire yield curve upwards / downwards and re-computing the swaption price under the given model (two-point finite difference method).

For the multi-factor hedge, we compute the model deltas by shifting each forward rate as well as the spot rate individually and re-valuing the swaption under the respective model. Since the price of an $n\text{Yr} \rightarrow 1\text{Yr}$ swaption depends only on the first n forward rates and the spot rate, we obtain $n + 1$ non-zero deltas,

$$(\Delta_0^{\text{Swptn}}(t_j), \dots, \Delta_n^{\text{Swptn}}(t_j)),$$

where

$$\Delta_i^{\text{Swptn}}(t_j) = \frac{\partial C_{n\text{Yr} \rightarrow 1\text{Yr}}(t_j)}{\partial f_i(t_j)}, \quad i = 0, \dots, n.$$

A more detailed description of the exact computations required for each model will be given further down this section.

In the single-factor case we compute the relevant bond delta

$$\Delta^{\text{Bond}}(t_j) = \frac{\partial P(t_j, T + 1)}{\partial f(t_j)}$$

by performing a parallel shift of the entire yield curve. The hedged portfolio will include

$$p(t_j) = -\frac{\Delta^{\text{Swptn}}(t_j)}{\Delta^{\text{Bond}}(t_j)}$$

zero-coupon bonds $P(t_j, T + 1)$ so that the overall portfolio delta is zero.

In the multi-factor case we compute a full matrix comprising the deltas for each individual bond with regard to each of the forward rates

$$(\Delta_{ik}^{\text{Bond}})_{i,k=0,\dots,9}(t_j) = \left(\frac{\partial P(t_j, t_j + k + 1)}{\partial f_i(t_j)} \right)_{i,k=0,\dots,9}.$$

Note that only the upper right triangle $i \geq k$ will contain non-zero entries and that these bond deltas are model-independent.

To construct a delta neutral portfolio, we solve the system

$$(\Delta_{ik}^{\text{Bond}})_{i,k=0,\dots,9}(t_j) \begin{pmatrix} p_0(t_j) \\ \vdots \\ p_n(t_j) \end{pmatrix} = - \begin{pmatrix} \Delta_0^{\text{Swptn}}(t_j) \\ \vdots \\ \Delta_n^{\text{Swptn}}(t_j) \end{pmatrix}$$

and take a position of $p_i(t_j)$ discount bonds $P(t_j, t_j + i + 1)$ for each $i = 0, \dots, n$.

Furthermore, a cash position of

$$B(t_0) = \begin{cases} -C_{n\text{Yr} \rightarrow 1\text{Yr}}(t_0) - pP(t_0, T + 1) & \text{in the one-factor case or} \\ -C_{n\text{Yr} \rightarrow 1\text{Yr}}(t_0) - \sum_{i=0}^n p_i P(t_0, t_0 + i + 1) & \text{in the multi-factor case} \end{cases}$$

will be entered into. It will be invested into a zero-coupon bond expiring on the next (business) day. Thus, it will earn interest — or cost interest if the position is negative — and will not be exposed to delta risk.

Overall, the hedged portfolio contains:

- A \$10,000 notional one-period swaption worth $C_{nYr \rightarrow 1Yr}(t_0)$,
- $\begin{cases} p(t_0) \text{ discount bonds } P(t_0, T+1) \text{ in the one-factor case or} \\ p_i(t_0) \text{ bonds } P(t_0, t_0 + i + 1), i = 0, \dots, n, \text{ in the multi-factor case,} \end{cases}$
- $B(t_0)$ invested into the discount bond expiring on the next (business) day.

We will re-hedge the portfolio daily. On each day t_j we will re-compute the hedge ratios p or p_0, \dots, p_n for each of the models and update the positions with regard to the respective hedging instruments.

We will now see how the particular deltas are computed for each of the models. On each day t_j the discount curve

$$P(t_j, T), \quad T > t_j,$$

is constructed by applying the interpolation method described in section 5.1.1. All relevant forward rates and forward swap rates are computed via this interpolated curve.

Delta Hedging under the Black Model

Under the Black model, the time t price of a swaption with fixed rate K and underlying tenor

$$T_m < T_m + 1 < \dots < T_n$$

is given by

$$V_{\text{PSwaption}}(t) = A(t) \text{Bl}^+(K, S^{m,n}(t), (T_m - t), \sigma), \quad (5.1)$$

where σ denotes the Black volatility.

The only free parameter in the Black model, i.e., the only variable that is not given by a swaption's observable characteristics, is the volatility σ .

We might wonder which exact volatility we are to use in order to compute the hedge ratio. In the absence of smiles and under the assumption of constant volatility, where a single σ would produce correct swaption prices for all strikes, the answer would be unambiguous.

In the presence of smiles, on the other hand, we might ask ourselves whether we should compensate for the obvious inadequacies of the Black formula in pricing ITM / OTM options. Rebonato (2004) makes the following consideration. Assume an at-the-money option has an implied volatility of 20%, and a 30-delta option¹ has an implied volatility of 24%. We wish to hedge the 30-delta option. Should we use the Black formula with an input volatility of 20% or would the implied volatility of our particular option, 24%, provide for a better hedge? Rebonato (2004) points out that the implied volatility supplies no additional information about the process of the underlying forward rate. The presence of smiles simply shows that the assumptions of the Black framework are strongly violated in reality, and the strike-dependent implied volatility merely becomes “*the wrong number*”

¹The term “30-delta option” has the following meaning [Rebonato (2004)]: Consider the implied volatility of an option with a particular strike. Inputting this number into the Black delta formula, $N(d_1)$ [cf. eq. (2.10)], and multiplying the result by 100 yields the market quote for the *delta* of this option. In the presence of smiles, the Black formula does not apply, and this procedure merely provides a way to quote the strike of an option after performing a certain normalization with regard to the option's moneyness.

to put in the wrong formula to get the right price of plain-vanilla options” [Rebonato (2004)].

Therefore, we will “calibrate” the Black model to the corresponding same-expiry ATM swaption, i.e., we will compute our hedge ratios by using the Black formula, eq. (5.1), in conjunction with the ATM swaption’s implied volatility.

In particular, the Black delta will be given by

$$\frac{dA}{df} \text{Bl}^+ + A \frac{d\text{Bl}^+}{dS} \frac{dS}{df},$$

where $\frac{d}{df}$ refers to the differential with regard to a shift in the forward rates f_0, \dots, f_{10} .

In the multi-factor case we compute all deltas by shifting the forward LIBOR rates one at a time while leaving the remaining rates fixed, i.e., the terms $\frac{d}{df}$ in the above equations relate to a change with regard to a single forward rate. This will provide us with a number of deltas, one for each forward rate we are shifting. In the single-factor case, a single delta is computed by shifting all forward rates in unison.

Remark: Even when computing $S^{m,n}(t)$ from interpolated forward rates [cf. section 5.1.1], we will not shift these interpolated values to compute the above derivatives but rather the original input forward rates $f_i = (t, t + i, t + i + 1)$. This goes for all of the subsequent models as well.

Delta Hedging under the Hull-White Model

Prior to computing any hedge ratios, we need to calibrate the Hull-White model. This will be done by the method described in section 3.2.2. For the calibration we use the prices of one-period ATM swaptions (historical or synthetic).

In all cases the Hull-White model will be calibrated to prices of one-period ATM swaptions expiring in one, three and five years’ time. Since we are working with option dollar prices instead of implied volatilities, we choose to perform the calibration by minimizing the mean-squared *relative* distances between market and model prices [cf. eq. (3.20)].

When computing market instantaneous forward rates, we interpolate the discount curve by a cubic spline (using the open-source library *Math.NET Numerics*).

Similar to the Black model, we obtain the Hull-White delta by numerically differentiating the swaption pricing equation derived in corollary 3.2.11 with regard to a shift of each individual forward rate (multi-factor) or a simultaneous parallel shift of all forward rates (single-factor).

Delta Hedging under the SABR Model

Consider again the payer swaption described above. For its treatment we consider a SABR system modeling the forward swap rate $S^{m,n}(t)$ under the swap measure, i.e., under the numéraire $A^{m,n}(t)$.

We calibrate the SABR model to the implied volatility smile of swaptions expiring at T_m . The required implied volatilities for a set of strikes are interpolated via the approach described in section 5.1.1. The SABR system is calibrated according to the method described in section 3.3.2.

The Hagan formula eq. (3.27) gives the (approximate) Black implied volatility of a European call or put option under the SABR model.

Consider a SABR system modeling the forward swap rate $S^{m,n}(t)$ under the swap measure.

As we have seen in section 3.3.1, a payer swaption with fixed rate K and expiry date T_m can be approximately valued under the SABR model as

$$V_{\text{PSwaption}}(t) = A(t)\text{Bl}^+(K, S^{m,n}(t), (T_m - t), \hat{\sigma}), \quad (5.2)$$

where $\hat{\sigma}$ denotes the Black implied volatility computed via Hagan's formula. Note that $A(t)$ as well as $S^{m,n}(t)$ and $\hat{\sigma}$ depend on the forward rates f_0, \dots, f_{m-1} .

The SABR model has two driving Brownian motions. They cannot both be neutralized by taking a Δ position in the forward swap rate alone.

In their original paper, Hagan et al. (2002) compute the SABR delta for a European call option as follows. Consider a call option on the forward rate f . Let $C(f, \sigma)$ denote the value of this option under the Black formula using the volatility σ . By setting $\sigma = \hat{\sigma}(f_t, \sigma_t)$, where $\hat{\sigma}$ is computed via Hagan's formula and where σ_t is the SABR volatility, one obtains the (approximate) SABR price of the call option as

$$V_{\text{Call}}(t) = C(f_t, \hat{\sigma}(f_t, \sigma_t)).$$

Hagan et al. (2002) arrive at a SABR delta of

$$\begin{aligned} \Delta_{\text{Hagan}} &= \frac{\partial V_{\text{Call}}}{\partial f} \\ &= \frac{\partial C}{\partial f} + \frac{\partial C}{\partial \hat{\sigma}} \frac{\partial \hat{\sigma}}{\partial f}. \end{aligned} \quad (5.3)$$

Note that $\frac{\partial C}{\partial f}$ and $\frac{\partial C}{\partial \hat{\sigma}}$ are, respectively, merely the Black delta and Black vega.

Bartlett (2006) argued that the delta risk can be hedged more precisely by computing

$$\Delta_{\text{Bartlett}} = \frac{\partial C}{\partial f} + \frac{\partial C}{\partial \hat{\sigma}} \left(\frac{\partial \hat{\sigma}}{\partial f} + \frac{\partial \hat{\sigma}}{\partial \sigma} \frac{\rho \nu}{f^\beta} \right). \quad (5.4)$$

The intuition behind eq. (5.4) is that since the forward rate f and the SABR volatility σ are correlated, a change in the forward rate f by Δf will, on average, be accompanied by a change in the volatility σ by $\frac{\rho \nu}{f^\beta} \Delta f$. Accounting for this effect reduces the variance of the hedged portfolio compared to eq. (5.3).

Rebonato et al. (2009) (citing Pogudin (2008)) prove that eq. (5.4) in fact produces the minimum-variance delta hedge.

In our experiment, we will compute all hedge ratios numerically. We wish to hedge the entire risk relating to a change in the yield curve, including the resulting change in $A(t)$. Furthermore, we wish to bump each forward rate individually when performing a multi-factor hedge.

For the single-factor hedge, we compute the Hagan delta as

$$\begin{aligned}\Delta_{\text{Hagan}} &= \frac{dA}{df} \text{Bl}^+ + A \frac{d\text{Bl}^+}{df} \\ &= \frac{dA}{df} \text{Bl}^+ + A \left(\frac{\partial \text{Bl}^+}{\partial S} \frac{dS}{df} + \frac{\partial \text{Bl}^+}{\partial \hat{\sigma}} \frac{d\hat{\sigma}}{df} \right).\end{aligned}$$

where $\frac{d}{df}$ denotes the differential with regard to a parallel shift of all forward rates and will be computed by a numerical two-point finite difference method.

We will compute the Bartlett delta as

$$\Delta_{\text{Bartlett}} = \frac{dA}{df} \text{Bl}^+ + A \left[\frac{\partial \text{Bl}^+}{\partial S} \frac{dS}{df} + \frac{\partial \text{Bl}^+}{\partial \hat{\sigma}} \left(\frac{d\hat{\sigma}}{df} + \frac{\partial \hat{\sigma}}{\partial \sigma} \frac{\rho \nu}{S^\beta} \frac{dS}{df} \right) \right].$$

In both cases the formulas have been adapted to account for the change in S through a parallel shift in each single forward rate. Consider, e.g., the term

$$\frac{\partial \hat{\sigma}}{\partial \sigma} \frac{\rho \nu}{S^\beta} \frac{dS}{df}.$$

It reflects the idea behind Bartlett's approach that a bump of

$$f_i \rightarrow f_i + \varepsilon, \quad i = 0, \dots, m-1,$$

will locally have the effect

$$S \rightarrow S + \frac{dS}{df} \varepsilon$$

which will, on average, result in

$$\sigma \rightarrow \sigma + \frac{\rho \nu}{S^\beta} \frac{dS}{df} \varepsilon.$$

Again, in the multi-factor case each forward LIBOR rate is shifted individually, and in the single-factor case all forward rates are shifted in parallel. We compute both types of delta, with and without the Bartlett correction term.

5.1.3 Measurement of the Hedging Performance

We require that our hedging strategy is self-financing, i.e., no cash injections or withdrawals are allowed. For every change we make in the position of the hedging instruments, we will make an according adjustment to the cash account.

Let $H(t_j)$ denote the time t_j value of the hedged portfolio.

According to the model used to construct the above portfolio, the overall position will be locally riskless with regard to a change in the term structure of interest rates. In reality, a certain slippage will of course occur. Consider the one-factor case. On the second day, t_1 , of the experiment the value of the above portfolio will be

$$H(t_1) = C_{nYr \rightarrow 1Yr}(t_1) + p(t_0)P(t_1, T + 1) + \frac{B(t_0)}{P(t_0, t_1)}.$$

Let

$$S_j = H(t_{j+1}) - H(t_j)$$

denote the one-day hedging slippage between the dates t_j and t_{j+1} . We will compare these one-day movements of the hedging portfolio to the one-day movements of the unhedged position (the naked swaption).

If

$$t_0 < t_1 < \dots < t_N = T$$

are all business dates during the life of the swaption, we compute the empirical variance of the sequence

$$(C_{nYr \rightarrow 1Yr}(t_{j+1}) - C_{nYr \rightarrow 1Yr}(t_j))_{j=0, \dots, n-1}$$

by using the unbiased estimator, i.e., by applying the formula

$$\hat{\sigma}_Y^2 = \frac{1}{m-1} \sum_{i=1}^m (y_i - \bar{y})^2,$$

where y_1, \dots, y_m is a sample drawn from a general distribution Y , and where \bar{y} denotes the sample mean.

Let $\mathbb{V}[\Delta_{\text{Swtpn}}]$ denote the empirical variance of the one-day movements in the unhedged position and let $\mathbb{V}[\Delta_X]$ denote the empirical variance of the sequence $(S_j)_{j=0, \dots, n-1}$ where the hedging portfolio has been constructed under the model “ X ”.

The first metric we consider will be the *hedge variance ratio*

$$\text{HVR}_X = 1 - \frac{\mathbb{V}[\Delta_X]}{\mathbb{V}[\Delta_{\text{Swtpn}}]}, \quad (5.5)$$

which describes the percentage variance reduction achieved by the hedging strategy compared to the unhedged position. The higher that HVR_X is, the better the hedging performance of the model X will be rated in the experiment.

Note: In our experiment the respective variances will be computed with regard to the one-day movements of all individual swaptions and simulated paths combined.

The same metric has, e.g., be used by Driessen et al. (2003), Attouli (2011) and Li and Zhao (2006) to evaluate a hedging strategy in their respective studies. As Li and Zhao (2006) point out, this measure is similar in spirit to R^2 in linear regression.

The second metric we consider will be the average maximum drawdown.

Definition 5.1.1 (Maximum drawdown). If $(X_t)_{t \geq 0}$ is a stochastic process, the *drawdown* at time T is defined as

$$D(T) = \left(\max_{t \in (0, T)} X(t) - X(T) \right)^+.$$

The *maximum drawdown* at time T is defined as

$$\begin{aligned} \text{MDD}(T) &= \max_{s \in (0, T)} D(s) \\ &= \max_{s \in (0, T)} \left(\max_{t \in (0, s)} X(t) - X(s) \right). \end{aligned}$$

Applied to our context, the drawdown of a portfolio measures the loss (if positive) an investor would have incurred until time T who bought the portfolio when it was most expensive. The maximum drawdown $\text{MDD}(T)$ describes the loss an investor would have taken who bought the portfolio at its highest and sold it at its lowest price.

The maximum drawdown is an informative statistic to consider when examining the performance of a hedging strategy. It is particularly important since we hedge each swaption all the way from inception to expiry. In practice, an investor might carry an option for only part of its life. The maximum drawdown gives an indication of the *worst case* one can *expect* in over a single path.

In particular, we will analyze the average maximum drawdown of each hedging strategy, i.e., we compute the maximum drawdown over the lifetime of each individual swaption in each simulated path and average over the results.

As with the hedge variance ratio, we will also present the result in terms of the percentage reduction achieved in comparison to the unhedged position, i.e., we consider

$$\text{AMDR}_X = 1 - \frac{\text{AMD}_X}{\text{AMD}_{\text{Swtpn}}}, \quad (5.6)$$

where AMD_X denotes the average maximum drawdown of the hedging strategy “ X ”.

The third metric we compute will be the variance of the terminal portfolio values. Hence, we compute the variance $\mathbb{V}[\text{PO}_{\text{Swtpn}}]$ of the set of terminal payoffs of all unhedged swaptions as well as the variance $\mathbb{V}[\text{PO}_X]$ of the set of terminal values of all hedged portfolios for each model X . As in eq. (5.5) and eq. (5.6), we consider the variance reduction achieved by each model:

$$\text{VPOR}_X = 1 - \frac{\mathbb{V}[\text{PO}_X]}{\mathbb{V}[\text{PO}_{\text{Swtpn}}]}. \quad (5.7)$$

Note that the terminal payoff can also be interpreted in a different way. If, instead of setting up a delta neutral zero-value portfolio, we constructed a “replicating” portfolio² as

²For simplicity, we use the term replicating portfolio. Of course the portfolio will only be imperfectly replicating due to model misspecification, market frictions and discrete-time trading. As Rebonato (2004) points out, even if markets were frictionless and trading were continuous, perfect replication would only be possible under strong conditions. The process for the underlying must be a diffusion and have a volatility either stochastic or functionally dependent on the underlying itself. For the SABR model, e.g., this is not the case.

a portfolio comprising only zero-coupon bonds and a cash position so that this portfolio has the same delta and initial value as the original swaption, then the value PO_X would constitute the terminal slippage of our replicating strategy, i.e., the difference between the option payoff and the terminal value of the “replicating” portfolio.

Schröter et al. (2012) point out that if one knows the *true* (market) price of an option at the time of inception, the terminal hedging error is a measure for the pure hedging performance of the model. If instead one only knows the option price under the hedging model, the terminal hedging error reflects the model’s joint pricing and hedging performance. The latter would, e.g., occur if a bank sold an exotic option for which there is no market price available so that the bank would have to rely on the model for both the pricing and the hedging of the option. In our case, however, all market prices are exogenously given as historical or synthetic inputs. Therefore, the terminal payoffs are indicative of the models’ pure hedging performances.

5.2 Hedging Performance in the Historical Framework

In a first test, we study the hedging performances of the models in the regular real-world historical data, i.e., we make use of historical prices to analyze how well the respective models would have hedged had they been used to construct hedging portfolios at the time.

We wish to only consider options with start and end dates within our available set of market data so that we can follow each swaption from the day of its inception until its expiry date. Therefore, we will consider the swaptions

- 1Yr \rightarrow 1Yr,
- 2Yr \rightarrow 1Yr and
- 3Yr \rightarrow 1Yr.

We illustrate the method by means of a 2Yr \rightarrow 1Yr swaption. For each date t_j in our history we look ahead by one year. If the resulting date is not a business day, the swaption is assumed to expire on the next business day. We compute the set of inception dates t_0, \dots, t_K for which the respective end dates lie within the data set. For the 2Yr \rightarrow 1Yr swaption, e.g., this set encompasses 1033 dates.

For each of these dates, we consider the associated ATM swaption and construct a hedged portfolio which will be re-hedged daily according to the method described in section 5.1.2. In the case of the 2Yr \rightarrow 1Yr swaption, each of the swaptions will be alive for about 500 business days, so that we will obtain a total of roughly 500,000 data points for each hedging approach / model and for the unhedged swaption.

We note that the movements of the individual swaptions are of course highly correlated so that the results will convey significantly less information than the sheer number of data points suggests. Furthermore, in the time span we consider, the term structure saw a decrease in almost any two-year interval, so that basically all swaptions end up out of the money.

Therefore, it makes little sense to rate the hedging performance in terms of the terminal payoffs of each strategy. The unhedged portfolio would produce a zero-variance payoff of 0 and seemingly dominate all hedging strategies.

Approach	1Yr		2Yr		3Yr	
	$\mathbb{V}[\Delta]$	HVR	$\mathbb{V}[\Delta]$	HVR	$\mathbb{V}[\Delta]$	HVR
Unhedged swaption	7.6897	-	6.2012	-	4.5720	-
Black 1-fac	1.2076	84.30%	1.3236	78.66%	1.5712	65.63%
HW 1-fac	1.4582	81.04%	1.2095	80.50%	1.1806	74.18%
SABR Hagan 1-fac	1.3323	82.67%	1.1855	80.88%	1.2811	71.98%
SABR Bartlett 1-fac	1.2617	83.59%	1.1550	81.37%	1.1952	73.86%
Black n -fac	0.8907	88.42%	1.4375	76.82%	1.5633	65.81%
HW n -fac	0.9126	88.13%	0.8265	86.67%	0.8732	80.90%
SABR Hagan n -fac	0.8687	88.70%	1.0249	83.47%	1.1021	75.89%
SABR Bartlett n -fac	0.8602	88.81%	0.9772	84.24%	0.9987	78.16%

Table 5.1: Variance of historical one-day-movements in the unhedged swaptions and the hedged portfolios for each of the hedging models, as well as the hedge variance ratios, i.e., the percentage reductions in this particular variance statistic achieved by each of the hedging strategies compared to the unhedged positions, cf. eq. (5.5). The results have been generated in the *historical framework*.

Similarly, due to the high correlation between individual swaptions, the maximum draw-down contains too little information to be considered here.

Instead, we focus on the hedge variance ratios.

The results are shown in table 5.1. For the Hull-White and SABR model, the n -factor hedge is clearly superior to the one-factor hedge. For the Black model, the two strategies come slightly closer in their performances. In fact, for the 2-year swaption a one-factor Black hedge shows lower variance than an n -factor Black hedge. In all cases, the SABR model performs better when the Bartlett correction is applied (as it should).

Overall, the SABR and Hull-White models produce better hedges than the Black model. The SABR n -factor hedge is the most successful strategy for the 1-year swaption, and the Hull-White n -factor hedge comes in first in the other two cases.

5.3 Hedging Performance in the Synthetic Evolution Model

We will simulate three individual swaptions:

- 1Yr \rightarrow 1Yr,
- 3Yr \rightarrow 1Yr and
- 5Yr \rightarrow 1Yr.

Note that two of these swaptions have the same structure as those examined in the historical experiment in section 5.2. This allows us to make a comparison between the hedging performances in the evolution framework and in the historical data set.

However, we note that the mere presence of a certain disparity between the results under both frameworks would not be evidence of a deficient evolution framework. The reason for this is two-fold. First, as explained in section 5.2, the historical data set covers only

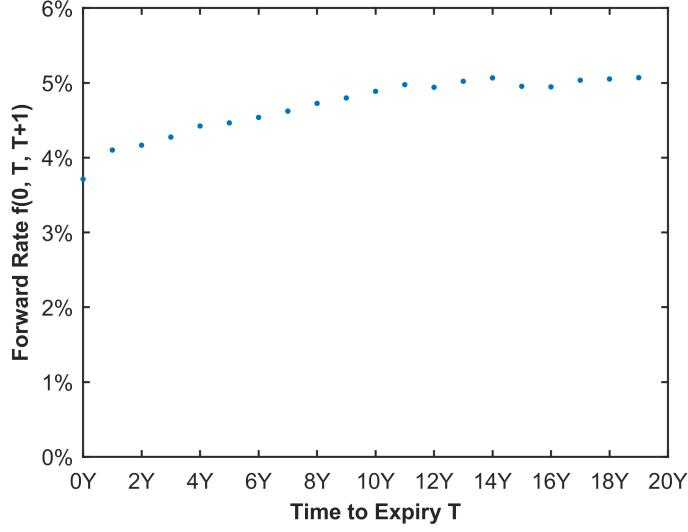


Figure 5.1: Market forward rate curve, $T \mapsto f(0, T, T + 1)$, as observed on July 1, 2005. These rates are assumed to be representative of “normal” market conditions and are used as initial inputs to the yield curve evolution in the first part of the experiment.

a very small number of non-overlapping “lifetimes” for the swaptions we consider, so that the statistical content of the historical experiment is very limited. Second, while in the historical time span virtually all swaptions were out-of-the-money for most of their lives up until their expiry dates, the synthetic framework will produce options both in and out of the money. This might have a significant impact on the hedging performance of the various models.

The evolution approach will be applied as follows:

We assume that each year consists of 250 evenly spaced business days. We choose a window length of $W = 80$ for the sampling mechanism. This has been chosen as a multiple of 5 which will help recover all weekend effects possibly present in the historical data. All mean reversion speeds have been set to 0.5% per day.

We will generate two sets of data. For the first set, we take as initial values the parameters observed on July 1, 2005. At this time, “normal” market conditions prevailed, with interest rates ranging from 3.5% to 5.5% (see fig. 5.1). The second simulation will be initialized using data from July 1, 2010. At this time the market was “excited”. Short term rates were as low as 0.5% (see fig. 5.2).

Figure 5.3 shows the root-mean-squared volatility function

$$\hat{g}(T) = \sqrt{\frac{1}{T} \int_0^T g(T - u)^2 du}$$

resulting from calibrating the g -function to implied volatilities of one-period swaptions on each of the given dates. As a result of the “excitedness” present in the market at the time, on July 1, 2010, the volatility was significantly higher than it was five years prior.




Figure 5.2: Market forward rate curve, $T \mapsto f(0, T, T+1)$, as observed on July 1, 2010. These rates are assumed to be representative of low-rate “excited” market conditions and are used as initial inputs to the yield curve evolution in the second part of the experiment.

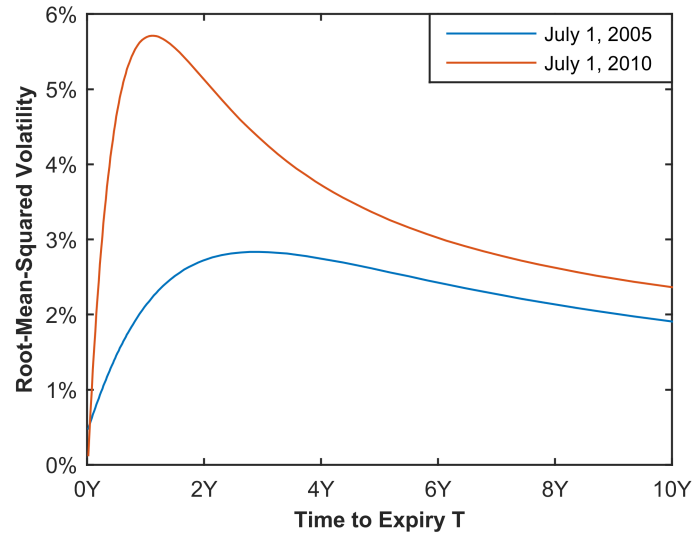


Figure 5.3: Root-mean-squared volatility functions \hat{g} calibrated to one-period swaptions on each of the two dates. Note that the calibration has been performed using a CEV-type exponent β of 0.5 so that the resulting volatilities are of a different magnitude than lognormal Black implied volatilities

For each of these dates, we generate 5000 evolution paths. To be precise, for each index $j \in 1, 2, \dots, 5000$, we take as initial values

$$\begin{aligned} y^\tau(0) &= f_\tau, \quad \tau \in \{0, 1, \dots, 10\} \cup \{\tau^l\} \\ \tilde{\sigma}_0^\tau(0) &= \sigma_0^\tau, \quad \tau \in \{0.5, 1, 2, 5, 10\}, \\ \tilde{\nu}^\tau(0) &= \nu^\tau, \quad \tau \in \{0.5, 1, 2, 5, 10\}, \\ \tilde{\rho}^\tau(0) &= \rho^\tau, \quad \tau \in \{1, 2, \dots, 10\} \end{aligned}$$

where f_τ , σ_0^τ , ν^τ and ρ^τ are the parameters observed on either July 1, 2005 (for the first simulation), or July 1, 2010 (for the second simulation).

Note that the interest rate processes y^τ in the evolution model have been chosen to constitute forward rates f_τ .

To keep the process tractable, we perform the simulation with a horizon of five years.

As for the mean reversion levels of each process, we proceed as follows:

For the time series of $\tilde{\sigma}_0^\tau$, $\tilde{\nu}^\tau$ and $\tilde{\rho}^\tau$, we assume that each process reverts to its initial value. The mean reversion levels of the rates y^τ are of particular importance and will be chosen slightly differently. In particular, we mildly follow the assumption that today's n year forward rate is a prediction of which spot rate we expect n years into the future. For the additive world long rate \bar{y}^{τ_l} we take as mean reversion level its initial value $\Sigma(f_\tau(j))$.

For the spreads of the remaining rates

$$\bar{y}^\tau - \bar{y}^{\tau_l}$$

we assume that the initial spreads

$$\Sigma(f_{\tau+5}) - \Sigma(f_{19})$$

are our expectation of the simulated spreads

$$\bar{y}^\tau - \bar{y}^{\tau_l}$$

at time $t = 5$ years.

In eq. (4.15) we have estimated this expectation as

$$(\bar{y}^\tau(0) - \bar{y}^{\tau_l}(0)) e^{-5a^{\tau*}} + \bar{b}^\tau(1 - e^{-5a^*}).$$

Hence, we choose \bar{b}^τ as

$$\bar{b}^\tau = \frac{(\Sigma(f_{\tau+5}) - \Sigma(f_{19})) - (\Sigma(f_\tau) - \Sigma(f_{\tau_l})) e^{-5a^{\tau*}}}{1 - e^{-5a^*}},$$

since $\Sigma(f_\tau) = \bar{y}^\tau(0)$. With this, the setup is complete and the simulation can be performed. The results of all simulations are shown in tables 5.2 to 5.7.

Interestingly, the SABR model performs significantly worse when the Bartlett correction is not applied, for all of the three statistics, more so than in the historical framework.

Approach	1Yr		3Yr		5Yr	
	$\mathbb{V}[\Delta]$	HVR	$\mathbb{V}[\Delta]$	HVR	$\mathbb{V}[\Delta]$	HVR
Unhedged swaption	11.5620	-	14.8051	-	13.4355	-
Black 1-fac	0.8970	92.24%	1.9151	87.06%	2.6346	80.39%
HW 1-fac	0.7664	93.37%	1.0491	92.91%	1.1133	91.71%
SABR Hagan 1-fac	0.9077	92.15%	1.9955	86.52%	2.4779	81.56%
SABR Bartlett 1-fac	0.7721	93.32%	1.0560	92.87%	1.1121	91.72%
Black n -fac	0.9214	92.03%	1.8546	87.47%	2.5387	81.10%
HW n -fac	0.8461	92.68%	1.0929	92.62%	1.3108	90.24%
SABR Hagan n -fac	1.0085	91.28%	1.8989	87.17%	2.5464	81.05%
SABR Bartlett n -fac	0.7797	93.26%	1.2773	91.37%	1.0500	92.19%

Table 5.2: Variance of historical one-day-movements in the unhedged swaptions and the hedged portfolios for each of the hedging models, as well as the hedge variance ratios, i.e., the percentage reductions in this particular variance statistic achieved by each of the hedging strategies compared to the unhedged positions, cf. eq. (5.5). The results have been generated in the *synthetic evolution framework* initialized with market data prevailing on *July 1, 2005*.

Approach	1Yr		3Yr		5Yr	
	AMD	AMDR	AMD	AMDR	AMD	AMDR
Unhedged swaption	58.3205	-	121.6088	-	145.7820	-
Black 1-fac	15.9683	72.62%	41.4255	65.94%	54.0140	62.95%
HW 1-fac	13.2280	77.32%	25.3147	79.18%	28.7211	80.30%
SABR Hagan 1-fac	16.1843	72.25%	41.0922	66.21%	52.5908	63.93%
SABR Bartlett 1-fac	13.3885	77.04%	24.7360	79.66%	27.8835	80.87%
Black n -fac	16.5321	71.65%	39.9829	67.12%	55.5942	61.86%
HW n -fac	13.9776	76.03%	24.8559	79.56%	32.2612	77.87%
SABR Hagan n -fac	17.0727	70.73%	41.7871	65.64%	55.1432	62.17%
SABR Bartlett n -fac	13.6583	76.58%	27.8382	77.11%	29.4203	79.82%

Table 5.3: Average maximum drawdown of the unhedged swaptions and the hedged portfolios for each of the hedging models, as well as the percentage reductions in the AMD achieved by each of the hedging strategies compared to the unhedged positions, cf. eq. (5.6). The results have been generated in the *synthetic evolution framework* initialized with market data prevailing on *July 1, 2005*.

Approach	1Yr		3Yr		5Yr	
	$\mathbb{V}[\text{PO}]$	VPOR	$\mathbb{V}[\text{PO}]$	VPOR	$\mathbb{V}[\text{PO}]$	VPOR
Unhedged swaption	1209.7995	-	1346.8037	-	983.6205	-
Black 1-fac	62.3006	94.85%	213.5472	84.14%	499.6651	49.20%
HW 1-fac	88.4830	92.69%	175.9876	86.93%	493.8816	49.79%
SABR Hagan 1-fac	61.1255	94.95%	265.2230	80.31%	309.5936	68.53%
SABR Bartlett 1-fac	86.1619	92.88%	257.4722	80.88%	296.7795	69.83%
Black n -fac	63.7531	94.73%	198.6942	85.25%	333.0726	66.14%
HW n -fac	90.6991	92.50%	211.1684	84.32%	237.4563	75.86%
SABR Hagan n -fac	55.0669	95.45%	206.5495	84.66%	329.0616	66.55%
SABR Bartlett n -fac	69.9802	94.22%	185.9767	86.19%	234.8335	76.13%

Table 5.4: Variance of terminal payoffs of the unhedged swaptions and of the hedged portfolios for each of the hedging models, as well as the percentage reductions in this particular variance statistic achieved by each of the hedging strategies compared to the unhedged positions, cf. eq. (5.7). The results have been generated in the *synthetic evolution framework* initialized with market data prevailing on *July 1, 2005*.

Approach	1Yr		3Yr		5Yr	
	$\mathbb{V}[\Delta]$	HVR	$\mathbb{V}[\Delta]$	HVR	$\mathbb{V}[\Delta]$	HVR
Unhedged swaption	23.3909	-	20.2739	-	12.6971	-
Black 1-fac	1.0036	95.71%	2.0892	89.69%	2.4745	80.51%
HW 1-fac	0.2489	98.94%	0.5981	97.05%	0.8909	92.98%
SABR Hagan 1-fac	0.6215	97.34%	1.9041	90.61%	2.6103	79.44%
SABR Bartlett 1-fac	0.4419	98.11%	1.2220	93.97%	1.6997	86.61%
Black n -fac	0.6376	97.27%	1.8595	90.83%	2.5023	80.29%
HW n -fac	0.4854	97.92%	1.0749	94.70%	1.4485	88.59%
SABR Hagan n -fac	0.7228	96.91%	1.8521	90.86%	2.3753	81.29%
SABR Bartlett n -fac	0.2037	99.13%	0.6610	96.74%	0.9870	92.23%

Table 5.5: Variance of historical one-day-movements in the unhedged swaptions and the hedged portfolios for each of the hedging models, as well as the hedge variance ratios, i.e., the percentage reductions in this particular variance statistic achieved by each of the hedging strategies compared to the unhedged positions, cf. eq. (5.5). The results have been generated in the *synthetic evolution framework* initialized with market data prevailing on *July 1, 2010*.

Approach	1Yr		3Yr		5Yr	
	AMD	AMDR	AMD	AMDR	AMD	AMDR
Unhedged swaption	53.6589	-	131.7806	-	156.1591	-
Black 1-fac	17.7668	66.89%	65.0915	50.61%	80.9125	48.19%
HW 1-fac	7.2561	86.48%	26.5690	79.84%	44.0949	71.76%
SABR Hagan 1-fac	26.1029	51.35%	80.5888	38.85%	95.6315	38.76%
SABR Bartlett 1-fac	17.2386	67.87%	37.9717	71.19%	54.5557	65.06%
Black n -fac	26.1039	51.35%	76.2900	42.11%	89.6476	42.59%
HW n -fac	17.4868	67.41%	35.8659	72.78%	51.7319	66.87%
SABR Hagan n -fac	20.5050	61.79%	67.3065	48.93%	80.4639	48.47%
SABR Bartlett n -fac	9.9985	81.37%	26.8681	79.61%	42.6382	72.70%

Table 5.6: Average maximum drawdown of the unhedged swaptions and the hedged portfolios for each of the hedging models, as well as the percentage reductions in the AMD achieved by each of the hedging strategies compared to the unhedged positions, cf. eq. (5.6). The results have been generated in the *synthetic evolution framework* initialized with market data prevailing on *July 1, 2010*.

Approach	1Yr		3Yr		5Yr	
	$\mathbb{V}[\text{PO}]$	VPOR	$\mathbb{V}[\text{PO}]$	VPOR	$\mathbb{V}[\text{PO}]$	VPOR
Unhedged swaption	2940.1467	-	2740.0505	-	996.5144	-
Black 1-fac	31.7521	98.92%	221.7938	91.91%	281.8349	71.72%
HW 1-fac	48.8013	98.34%	98.7338	96.40%	162.3643	83.71%
SABR Hagan 1-fac	38.5353	98.69%	174.3219	93.64%	326.2759	67.26%
SABR Bartlett 1-fac	83.6857	97.15%	71.1734	97.40%	184.1804	81.52%
Black n -fac	25.4622	99.13%	173.5436	93.67%	351.2582	64.75%
HW n -fac	60.3565	97.95%	91.9346	96.64%	219.9288	77.93%
SABR Hagan n -fac	34.6657	98.82%	220.8070	91.94%	332.9613	66.59%
SABR Bartlett n -fac	75.7951	97.42%	119.4377	95.64%	188.7116	81.06%

Table 5.7: Variance of terminal payoffs of the unhedged swaptions and of the hedged portfolios for each of the hedging models, as well as the percentage reductions in this particular variance statistic achieved by each of the hedging strategies compared to the unhedged positions, cf. eq. (5.7). The results have been generated in the *synthetic evolution framework* initialized with market data prevailing on *July 1, 2010*.

Variance of One-Day-Movements

Contrary to the historical framework, the n -factor and one-factor strategies yield very similar results in the synthetic evolution framework. In several cases the one-factor hedge performs even (slightly) better than the n -factor hedge. Over all swaptions and approaches, the multi-factor delta hedges produce an average hedge variance ratio of 90.80%, compared to 90.70% for the one-factor hedges.

For all approaches the average reduction in variance is lower when the simulation is longer.

As in the historical framework, the Black model underperforms compared to the other models with regard to the variance of one-day-movements in both simulations. Overall, the SABR and Hull-White models yield similar results with the Hull-White model emerging as the close winner.

Average Maximum Drawdown

As with the hedge variance ratios, the reduction in maximum drawdown seems to depend little on whether a single- or a multi-factor hedge was used. The single-factor hedges lead to an average reduction in AMD by 67.30% while the n -factor hedges yielded a slightly lower reduction of 66.76%.

Using the Bartlett correction, the SABR approach performs slightly better than the Hull-White approach when the simulation was initialized with the normal rate environment of July 2005, and slightly worse when the simulation started with the 2010 data.

Overall, a large portion of unhedged drawdown risk appears to remain under any of the three approaches, at least in the case of a mere delta hedge.

Variance of Terminal Payoffs

As for the variance of terminal payoffs, when it comes to the 1Yr swaption, the Black model can keep up with the remaining approaches. However, for longer maturities the Black model's performance shows a steep decline compared to the Hull-White and SABR models. The latter two models seem to reduce the variance of terminal payoffs to a similar degree.

The multi-factor hedges achieve an average variance reduction of 85.98%, compared to a reduction of 84.24% for the single-factor hedges.

It appears that the hedge variance ratio and the variance of terminal payoffs can be reduced significantly with a delta hedge, especially when the evolution runs only for a single year.

Chapter 6

Conclusion and Outlook

We have introduced and refined a model for the evolution of the yield curve and the implied volatility surface. The model allows to generate artificial time series of market data using a window sampling mechanism that aims at recovering the characteristics of an original set of market data used as input to the approach. The method is particularly useful when only limited *relevant* market data is available.

Subsequently, we have presented a methodology for the evaluation of hedging strategies based on three statistics, namely the percentage reductions in the variance of one-day-movements, in the maximum drawdown and in the variance of terminal payoffs.

We have analyzed the hedging performance of the Black model, the Hull-White model and the SABR model. The models were used to generate single- and multi-factor hedges, both within the framework of actual historical data and within our synthetically generated market.

Due to the limited amount of historical data, we only considered the variance of one-day-movements for the historical study. Overall, in this study the SABR and Hull-White models were superior to the Black model, and the n -factor hedges were considerably better than the one-factor hedges.

Within the synthetic evolution framework on the other hand, there was only a small difference between single- and multi-factor hedges. Overall, the multi-factor hedges were slightly superior for the two variance-related statistics we analyzed, and slightly inferior with regard to the maximum drawdown.

Once again, the SABR and Hull-White models outperformed the Black model in most experimental setups.

Of course we cannot eliminate the possibility that any of the observations made in the synthetic market relate to specific traits of the evolution approach.

In this work, we did not engage in the hedging of any risk factor except for changes in the forward rate curve. Gupta and Subrahmanyam (2005) point out that with continuous trading and continuous state variable sample paths, the delta risk would be all that mattered. Of course, neither of these conditions is satisfied in practice. Furthermore, we only considered three interest rate models and only one-period plain-vanilla swaptions. Future

studies could employ the joint evolution model presented here to study the hedging performance of a variety of models. With sufficient computational resources (or time) and given one has sufficient market data at one's disposal, even complex derivatives could be dealt with by means of Monte Carlo simulation.

Appendix A

A.1 AR(1) Filters

Definition A.1.1 (AR(1) process). A *first order autoregressive process*, or AR(1) *process*, is a stochastic process X_t satisfying

$$X_t = c + \varphi X_{t-1} + \varepsilon_t, \quad (\text{A.1})$$

where φ and c are constants and ε_t is *white noise*, i.e., a sequence of independent random variables with zero mean and finite variance σ_ε^2 .

We will only consider the case $\varphi \in (0, 1)$ so that the process produces an Ornstein-Uhlenbeck-type mean reversion. In fact, if the white noise terms ε_t are normally distributed with zero mean and if $\varphi \in (0, 1)$, the AR(1) process defined by eq. (A.1) can be interpreted as the discretized version of an Ornstein-Uhlenbeck process. This can be seen by considering the exact solution of the continuous-time Ornstein-Uhlenbeck process

$$dX_t = a(b - X_t)dt + \sigma dW_t,$$

namely

$$X_t = X_0 e^{-at} + b(1 - e^{-at}) + Z_t, \quad (\text{A.2})$$

where

$$Z_t \sim \mathcal{N}\left(0, \frac{\sigma^2}{2a} (1 - e^{-2at})\right).$$

For time steps $\delta t = 1$, the discretized version of eq. (A.2) reads

$$X_t = b(1 - e^{-a}) + e^{-a} X_{t-1} + \varepsilon_t,$$

where

$$\varepsilon_t \sim \mathcal{N}\left(0, \frac{\sigma^2}{2a} (1 - e^{-2a})\right).$$

This corresponds to the AR(1) process defined in eq. (A.1), with parameters

$$\begin{aligned} c &= b(1 - e^{-a}) \\ \varphi &= e^{-a} \\ \varepsilon_t &\sim \mathcal{N}\left(0, \frac{\sigma^2}{2a} (1 - e^{-2a})\right). \end{aligned} \quad (\text{A.3})$$

A.1.1 Fitting the AR(1) Parameters

We will now derive a calibration mechanism for the AR(1) process. We assume that we have knowledge of a set of realizations

$$X_j = x_j, \quad j = 0, 1, \dots, N.$$

¹First write X_t as

$$X_t = \mu + X_t^* \tag{A.4}$$

where X_t^* is a latent AR(1) series of the type

$$X_t^* = \varphi X_{t-1}^* + \varepsilon_t. \tag{A.5}$$

Then X_t can be expressed as

$$X_t = c + \varphi X_{t-1} + \varepsilon_t, \quad c = \mu(1 - \varphi). \tag{A.6}$$

Rewriting the recursion in eq. (A.5) yields

$$\begin{aligned} X_t^* &= \varphi(\varphi X_{t-2}^* + \varepsilon_{t-1}) + \varepsilon_t \\ &= \varphi^2 X_{t-2}^* + \varepsilon_t + \varphi \varepsilon_{t-1} \\ &= \dots \\ &= \sum_{i=0}^{\infty} \varphi^i \varepsilon_{t-i}. \end{aligned}$$

Here we have made use of the fact that $|\varphi| < 1$ since we assumed $\varphi \in (0, 1)$. The above computation provides several insights. First, it shows that the influence of a past realization or shock tends towards zero as time passes. Assuming the process has an infinite history, we see that

$$\mathbb{E}[X_t^*] = 0$$

and thus

$$\mathbb{E}[X_t] = \mu.$$

This shows that $c/(1 - \varphi)$ is an estimate for $\mathbb{E}[X_t]$. Furthermore, the variance of X_t can be computed as

$$\begin{aligned} \mathbb{V}[X_t] &= \mathbb{V}[X_t^*] \\ &= \sum_{i=0}^{\infty} \varphi^{2i} \mathbb{V}[\varepsilon_{t-i}] \\ &= \frac{\sigma_\varepsilon^2}{1 - \varphi^2}. \end{aligned}$$

¹cf. (Martin et al., 2012, Example 7.2)

In order to estimate the parameters, we consider the case of independent and normally distributed white noise:

$$\varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2).$$

Knowing the distribution of ε_t allows us to construct a maximum likelihood estimator for the three unknown parameters $(c, \varphi, \sigma_\varepsilon^2)$. Since ε_t is normally distributed, X_0 is normally distributed as well. In particular,

$$X_0 \sim \mathcal{N}\left(\frac{c}{1-\varphi}, \frac{\sigma_\varepsilon^2}{1-\varphi^2}\right).$$

Thus, the probability density function of X_0 , conditional on the parameters $(c, \varphi, \sigma_\varepsilon^2)$, is

$$f_{X_0}(x_0; c, \varphi, \sigma_\varepsilon^2) = \frac{1}{\sqrt{2\pi\sigma_\varepsilon^2/(1-\varphi^2)}} \exp\left(-\frac{(x_0 - c/(1-\varphi))^2}{2\sigma_\varepsilon^2/(1-\varphi^2)}\right).$$

Next consider $j \geq 1$. The distribution of X_j depends on the value X_{j-1} observed at t_{j-1} . From eq. (A.6) it follows that

$$(X_j | X_{j-1} = x_{j-1}) \sim \mathcal{N}(c + \varphi x_{j-1}, \sigma_\varepsilon^2).$$

The conditional probability density of the observation $X_j = x_j$ is thus

$$f_{X_j}(x_j | X_{j-1} = x_{j-1}; c, \varphi, \sigma_\varepsilon^2) = \frac{1}{\sqrt{2\pi\sigma_\varepsilon^2}} \exp\left(-\frac{(x_j - c - \varphi x_{j-1})^2}{2\sigma_\varepsilon^2}\right).$$

In particular, for any $j \geq 1$ the distribution of X_j depends on past realizations of the process only through X_{j-1} . For a set of observations x_0, x_1, \dots, x_N , the joint probability density function — i.e., the likelihood function — is

$$\mathcal{L}(c, \varphi, \sigma_\varepsilon^2) = f_{X_0}(x_0; c, \varphi, \sigma_\varepsilon^2) \cdot \prod_{j=1}^N f_{X_j}(x_j | X_{j-1} = x_{j-1}; c, \varphi, \sigma_\varepsilon^2).$$

The log-likelihood function can be computed as

$$\begin{aligned} \ln \mathcal{L}(c, \varphi, \sigma_\varepsilon^2) &= -\frac{1}{2} \ln(2\pi\sigma_\varepsilon^2/(1-\varphi^2)) - \frac{(x_0 - c/(1-\varphi))^2}{2\sigma_\varepsilon^2/(1-\varphi^2)} \\ &\quad - \frac{N}{2} \ln(2\pi\sigma_\varepsilon^2) - \frac{1}{2\sigma_\varepsilon^2} \sum_{j=1}^N (x_j - c - \varphi x_{j-1})^2. \end{aligned} \tag{A.7}$$

The maximum-likelihood estimator (MLE) is obtained by maximizing this expression. From eq. (A.7) we also see that if we ignored the expression resulting from the first observation $X_0 = x_0$, i.e., calculated the MLE conditional on the initial observation, the

MLE for c and φ is merely the result of minimizing the sum of squared errors ε_t . This estimator is called the *conditional maximum-likelihood estimator*. When maximizing the full log-likelihood function, the resulting estimator is called the *exact maximum-likelihood estimator* (cf. Martin et al. (2012)).

In the sampling approach, one cannot expect the residuals ε_t to be normally distributed, but the above derivation motivates the estimation of c and φ via a least squares approach. Since we have a long history of data, the conditional MLE is clearly sufficient. It has the additional advantage that the result is independent of the variance σ_ε of the residuals which we do not need to estimate since new data is produced by sampling historical residuals rather than by drawing from a random distribution.

References

- Yunbi An and Wulin Suo. The compatibility of one-factor market models in caps and swaptions markets: Evidence from their dynamic hedging performance. *Journal of Futures Markets*, 28(2):109–130, 2008.
- Leif B.G. Andersen and Vladimir V. Piterbarg. *Interest Rate Modeling*, volume 1. Atlantic Financial Press, London, 2010.
- Pavel Andreichenko. A parsimonious model for the joint evolution of yield curves and the interest rate smile surface under the objective measure. MSc. thesis, University of Oxford, 2011. unpublished.
- Sami Attaoui. Hedging performance of the Libor market model: The cap market case. *Applied Financial Economics*, 21(16):1215–1223, 2011.
- Bruce Bartlett. Hedging under SABR model. *Wilmott Magazine*, pages 2–4, July/August 2006.
- Fischer Black. The pricing of commodity contracts. *Journal of Financial Economics*, 3: 167–179, 1976.
- Alan Brace, Dariusz Gątarek, and Marek Musiela. The market model of interest rate dynamics. *Mathematical Finance*, 7(2):127–147, April 1997.
- Damiano Brigo and Fabio Mercurio. *Interest Rate Models - Theory and Practice*. Springer-Verlag, Berlin, 2nd edition, 2006.
- Nick Deguillaume. Dependence of rate move magnitude on rate levels with an application to a real-world model. MSc. thesis, University of Oxford, 2009. unpublished.
- Nick Deguillaume, Riccardo Rebonato, and Andrey Pogudin. The nature of the dependence of the magnitude of rate moves on the rates levels: a universal relationship. *Quantitative Finance*, 13(3):351–367, 2013.
- Emanuel Derman and Iraj Kani. Riding on a smile. *Risk*, 7(2):139–145, February 1994.
- Joost Driessen, Pieter Klaassen, and Bertrand Melenberg. The performance of multi-factor term structure models for pricing and hedging caps and swaptions. *Journal of Financial and Quantitative Analysis*, 38(3):635–672, September 2003.
- Bruno Dupire. Pricing with a smile. *Risk*, 7(1):18–20, January 1994.
- Rong Fan, Anurag Gupta, and Peter Ritchken. On pricing and hedging in the swaption market: How many factors, really? *Journal of Derivatives*, 15(1):9–33, 2007.

- Ronald Fisher. *Statistical Methods for Research Workers*. Oliver and Boyd, Edinburgh, 5th edition, 1934.
- Anurag Gupta and Marti G. Subrahmanyam. Pricing and hedging interest rate options: Evidence from cap-floor markets. *Journal of Banking & Finance*, 29:701–733, 2005.
- Patrick S. Hagan, Deep Kumar, Andrew S. Lesniewski, and Diana E. Woodward. Managing smile risk. *Wilmott Magazine*, pages 84–108, September 2002.
- Marc Henrard. *Interest Rate Instruments and Market Conventions Guide*. OpenGamma, December 2013. <http://www.opengamma.com/sites/default/files/interest-rate-instruments-and-market-conventions.pdf>.
- John Hull and Alan White. Pricing interest-rate derivative securities. *Review of Financial Studies*, 3(4):573–592, Winter 1990.
- John C. Hull. *Options, Futures, and Other Derivatives*. Pearson Education Ltd., Essex, United Kingdom, 8th edition, 2012.
- Farshid Jamshidian. An exact bond option formula. *Journal of Finance*, 44(1):205–209, March 1989.
- Farshid Jamshidian. LIBOR and swap market models and measures. *Finance and Stochastics*, 1(4):293–330, September 1997.
- Robert Jarrow, Haitao Li, and Feng Zhao. Interest rate caps ‘smile’ too! But can the LIBOR market models capture the smile? *Journal of Finance*, 62(1):345–382, February 2007.
- Haitao Li and Feng Zhao. Unspanned stochastic volatility: Evidence from hedging interest rate derivatives. *The Journal of Finance*, 61(1):341–378, 2006.
- Vance Martin, Stan Hurn, and David Harris. *Econometric Modelling with Time Series: Specification, Estimation and Testing*. Themes in Modern Econometrics. Cambridge University Press, New York, 2012.
- Kristian R. Miltersen, Klaus Sandmann, and Dieter Sondermann. Closed form solutions for term structure derivatives with log-normal interest rates. *The Journal of Finance*, 52(1):409–430, March 1997.
- Ken Nyholm and Riccardo Rebonato. Long-horizon yield curve forecasts: Comparison of semi-parametric and parametric approaches. 2007. Working Paper.
- Jan Obłój. Fine-tune your smile: Correction to Hagan et al. *Wilmott Magazine*, May 2008.
- Louis Paulot. Asymptotic implied volatility at the second order with application to the SABR model. *SSRN eLibrary*, June 2009.
- Raoul Pietersz and Antoon Pelsser. A comparison of single-factor markov-functional and multi-factor market models. *Review of Derivatives Research*, 13(3), 2010.
- Andrey Pogudin. Theoretical and practical aspects of hedging within SABR and LMM-SABR models. MSc. thesis, University of Oxford, 2008. unpublished.

- Riccardo Rebonato. *Volatility and Correlation*. John Wiley, Chichester, United Kingdom, 2nd edition, 2004.
- Riccardo Rebonato. Forward-rate volatilities and the swaption matrix: Why neither time-homogeneity nor time-dependence are enough. *International Journal of Theoretical and Applied Finance*, 9(5):705–746, 2006.
- Riccardo Rebonato. A time-homogeneous, SABR-consistent extension of the LMM: Calibration and numerical results. *Risk*, December 2007.
- Riccardo Rebonato and Richard White. Linking caplets and swaptions prices in the LMM-SABR model. *Journal of Computational Finance*, 13(2):19–45, 2009.
- Riccardo Rebonato, Sukhdeep Mahal, Mark S. Joshi, Lars-Dierk Buchholz, and Ken Nyholm. Evolving yield curves in the real-world measures: A semi-parametric approach. *Journal of Risk*, 7(3):29–62, 2005.
- Riccardo Rebonato, Andrey Pogudin, and Richard White. Delta and vega hedging in the SABR and LMM-SABR models. *Risk*, December 2008.
- Riccardo Rebonato, Kenneth McKay, and Richard White. *The SABR/LIBOR Market Model*. John Wiley, Chichester, United Kingdom, 2009.
- Till Schröter, Michael Monoyios, Mario Rometsch, and Karsten Urban. Model uncertainty and the robustness of hedging models. November 2012. Working Paper.
- Oldrich Vasicek. An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5:177–188, 1977.