Copenhagen Business School

Department of Finance

The SABR model
– theory and application

Thesis for M.Sc. in Business Administration and Management Science
(cand.merc.mat)

Author:
Søren Skov Hansen
Cpr: xxxxxx–xxxx

Thesis supervisor:
Mads Stenbo-Nielsen

Thesis submitted on April 7th 2011

Abstract

This thesis is based on the market observable volatility smiles for swaptions. We present different ways of handling these smiles and discuss the models as well as their implications.

Initially, we introduce the financial theory that will be fundamental to our thesis. Following, we move on to a brief empirical study of volatility, where we show that volatility, contrary to the assumption made in the classical Black-Scholes setup, is constant neither in time nor in strike price.

The first model we analyze is the *the local volatility model*. We discuss the model and give a numerical example on how the model can be calibrated to fit an observed volatility smile. Ultimately, we look into the inherent dynamics of the local volatility model, and we find that it has the counterintuitive property of shifting the volatility smile in the *opposite* direction of the price of the underlying asset when this shifts.

Our second and main model is *the SABR model*. We present the model, and examine how its parameters influence the shape of the fitted volatility smile. Following, we investigate the SABR model's ability to fit a volatility smile using different methods of estimation and parametrization. We find that the SABR model is very capable of fitting an observed volatility smile, seemingly regardless of choice of estimation and parametrization method. However, subsequently we note that the Δ risk measure that arises from the SABR model is very much dependent on the parametrization. We analyze this problem and give a correction to the Δ measure, which marginalizes the effect of the choice of parametrization, thus causing the SABR model to yield fairly similar Δ measures for different parametrizations.

Finally, we show how the SABR model's ability to inter- and extrapolate a volatility smile can be utilized in a pricing scenario to price a *constant maturity swap*. Initially, we explain the theory behind the pricing, and following we present a detailed numerical example using market data and comparing our findings to market prices.

Preface

I would like to thank my thesis supervisor, Mads Stenbo-Nielsen, for his qualified guidance, the motivating talks and his flexibility throughout the process. I would also like to thank Martin Linderstrøm for his help on the choice of subject as well as for his help on data collection and empirical discussions. Finally, I owe thanks to all those who have given me the occasional push in the right direction.

<div align="right">

Copenhagen, April 2011

Søren Skov Hansen

</div>

Contents

Chapter 1

Introduction

In financial terms, an *option* is the right but not the obligation to buy or sell a certain asset at a certain price at a certain time. Hence, an option is a key derivative when it comes to securing a price at a certain point in time whether that price is a selling price or a buying price. Another financial derivative commonly used as a means of securing a certain price is a *swap*. A basic description of a swap is: an exchange of a series of payments for a different series of payments. The two derivatives can be combined into what is known as a *swaption*. A swaption is the right but not the obligation to enter into a swap at a certain time. The nature of the swap, with its series of exchanges of payments, makes it an obvious choice for locking down the price of an asset needed on a regular basis. Examples could be natural gas, a foreign exchange rate or an interest rate. The latter is an especially interesting case, since hardly any business entity exists that does not have at least some exposure to some interest rate. Numbers from the *Bank for International Settlements* show that interest rate risk is indeed a major concern in contemporary society. No less than three quarters of total amounts outstanding on standard over-the-counter options[1] are linked to interest rate options (swaptions).[2] Therefore, we have chosen the swaption as our main derivative for this thesis.

In the process of pricing an option, an important piece of information is some kind of measure of the uncertainty of the price of the asset underlying the option. For most pricing methods, some a priori assumption regarding the development of the underlying asset is required along with an estimate of how much the price process will fluctuate. Empirically, estimates of volatility are not easily deter-

---

[1]Divided into options on foreign exchange, interest rates, equity and commodities.
[2]Actual numbers are presented in section 2.2.

mined. Estimates tend to depend heavily on the length of the (historical) time series used for calculations. Further, in most cases when looking at asset returns series, it becomes clear that the returns' fluctuations are not of a constant magnitude. Rather, they seem to continuously change over time. Despite that fact that volatility might not be constant, one of the classics within the field of option pricing, the Black-Scholes model, has as a key assumption that volatility is indeed constant. While this might not be the case in real life, the model still has the attractive feature that it implies a one-to-one relationship between a certain volatility level and the price of an option. This feature allows for a unique way to quote option prices as *implied volatilities*—the level of volatility that, when used in the classic Black-Scholes model, yields the observed option price.

However, when looking at implied volatility quotes, empirically these have a tendency to change significantly with the strike level of the option. In general, the implied volatility as a function of strike price will have the shape of a smile, thus creating what is known as *the volatility smile*. The main objective of our thesis is to show how this volatility behavior can be modeled, and subsequently how the modeling framework can be utilized for risk management purposes, and for the pricing of more complex derivatives. Initially, we will explore the properties of *the local volatility model* of Derman and Kani (1994). Following, we work our way on to the more contemporary *SABR model* of Hagan et al. (2002), which will be our main emphasis for this thesis.

The thesis is structured in the following manner:
We start out in Chapter 2 by laying down the theoretical foundation on which the thesis will be built. The fundamentals include definitions of interest rates, various derivatives and some basic techniques required for calculations. Following the fundamentals, in Chapter 3 we give an empirical motivation for the concept of non-constant volatility, and how this can present several issues within risk management and derivatives pricing. In Chapter 4 we introduce the local volatility model. We explain the principles of the model and present an implementation of the model in a discrete time setup using binomial tree pricing techniques. Finally, we explore the dynamics implied by the local volatility model. Chapter 5 is where we introduce the SABR model. The model and its components are explained in detail, and a closed-form approximation of the model-implied volatility is presented. Following, various estimation techniques for estimating the model's parameters are discussed and implemented using market data. As a closing section of Chapter 5, we present the SABR model's implications to different measures of risk, and show how these differ from those of the Black-Scholes setup. In Chapter 6 we show how the SABR

model can be applied beyond the scope of merely modeling a volatility smile. We introduce the concept of *constant maturity swaps* and show how the SABR model, in combination with pricing frameworks established previously in the thesis, can be used to price these more complex derivatives. Initially, we go through the theory of the pricing, and subsequently we apply the theory (together with results from Chapters 2 and 5) to market data and give a discussion of our findings.

Part I

Theoretical background

Chapter 2

The fundamentals

2.1 xIBOR rates

When using the term *xIBOR* or *xIBOR rates* in this paper this will be a reference to the *x InterBank Offered Rate*. The *x* refers to the entity that fixes the rate, a major one being the *British Bankers Association* (BBA) who sets the *London InterBank Offered Rate* (LIBOR) fixings. The LIBOR gives official fixings on interest rates with maturities ranging from a single business day to 12 months and are set each day at 11:00 GMT. In addition to the BBA, other fixing entities also exist, such as the *European Banking Federation* who sets the EURIBOR fixings and the *Danish Central Bank* who sets the CIBOR fixings.

The ways the various xIBOR rates are set differ, but they are all quoted using the *money market convention*. This means that the interest paid on a loan is simply calculated as $N\,\delta\,L$ where N is the notional, δ is the coverage[1] and L is the xIBOR rate. Letting $D(t,T)$ denote today's (time 0) price of a zero coupon bond purchased at t and maturing at T, and letting $F(0,T,T+\delta)$ denote the time 0 *forward xIBOR rate* between time T and $T+\delta$, then by argument of no arbitrage we can derive an expression for calculating the forward rate as

$$1 + \delta F(0,T,T+\delta) = \frac{1}{D(T,T+\delta)} \Leftrightarrow$$

$$1 + \delta F(0,T,T+\delta) = \frac{1}{\frac{D(0,T+\delta)}{D(0,T)}} \Leftrightarrow$$

$$F(0,T,T+\delta) = \frac{1}{\delta}\left(\frac{D(0,T)}{D(0,T+\delta)} - 1\right) \tag{2.1}$$

---

[1]The *coverage* is the length of the interest rate accrual period of the loan expressed in years. Hence, coverage will depend on the applied day count convention, and will as such also be a function of T. However, we will think of it as a constant given a day count convention and a maturity.

Swaps and forward contracts (bn USD)			
Foreign exchange	Interest rate	Equity	Commodity
39,638	400,985	1,830	1,675
8.9%	*90.3%*	*0.4%*	*0.4%*

Options (swaptions) (bn USD)			
Foreign exchange	Interest rate	Equity	Commodity
9,558	48,808	4,762	846
14.9%	*76.3%*	*7.4%*	*1.3%*

Table 2.1: *Amounts outstanding of OTC derivatives in bn USD (ultimo 2009).*

2.2 Interest rate swaps

In this section we will look at what characterizes a plain vanilla interest rate swap. We will look at the contract from a practitioners point of view, with a brief introduction to how interest rate swaps are treated in real life. After that, we will look into the pricing of the interest rate swap, still keeping a practical focus.

2.2.1 Basics of the interest rate swap

One of the most commonly used financial instruments is the *plain vanilla interest rate swap* (IRS). Table 2.1 shows some over-the-counter (OTC) derivatives statistics obtained from the Bank for International Settlements.[2] The numbers clearly show that interest rate derivatives are by far the dominating type compared to OTC derivatives on foreign exchange, equity and commodities—even when combined. In an IRS two counterparties agree to exchange two series of payments. In the following we will assume, that a series of fixed payments are exchanged for a series of floating payments. These two components of the IRS are referred to as the *fixed leg* and the *floating leg* respectively. Each party's position in the swap is named relative to the fixed leg. Hence, the person paying the fixed rate has entered in to a *payer swap* and the person paying the floating rate (and thus receiving the fixed rate) has entered in to a *receiver swap*.

Focusing first on the floating leg, this is linked to a specific interest rate, e.g. the LIBOR rate with a tenor of 6 months (LIBOR6M). The floating rate is set in advance (*before* each accrual period) and paid in arrears (*at the end* of each accrual period) using a suitable coverage depending on the day count convention. The standard market conventions for plain vanilla IRSs in major currencies and

---

[2]Data available at www.bis.org/statistics/extderiv.htm.

6

DKK are presented in Table 2.2. A brief explanation of some of the terms in Table 2.2 is given below:

Spot start is denominated in business days (B). This means, that if one enters into a *spot starting* EUR IRS on a Thursday,[3] then "spot start" does not mean that the IRS will actually start on that Thursday, but that it starts two business days later on Monday.

Term indicates which interest rate curve is the standard for the given currency. All of the listed currencies use the interest rate curves based on semi-annual payments except for USD denominated plain vanilla IRSs that are based on quarterly interest rate payments. As a consequence of this, the floating legs all have semi-annual payments except for the USD IRSs that have quarterly floating leg payments.

Day count conventions determine how the coverage (δ) is calculated. In Table 2.2 there are two different day count conventions:

Act365: This day count convention assumes that every year has 365 days. Thus, the coverage for a period is calculated by taking the actual number of days in a period and dividing it by 365

$$\delta_{\text{Act365}} = \frac{\text{Days}}{365}$$

30/360: In this day count convention all months are assumed to have 30 days, resulting in a 360-day year. The expression for calculating coverage according to the 30/360 day count convention is

$$\delta_{30/360} = \frac{\text{Years} \times 360 + \text{Months} \times 30 + \min[\text{Days}, 30]}{360}$$

Note that for most plain vanilla IRSs the day count convention for the floating leg differs from that of the fixed leg. Also, only GBP and JPY have the same payment frequencies (semi–annual) for the two legs.[4]

2.2.2 Valuation of the interest rate swap

Having established the basic set of rules for the plain vanilla IRS, we now turn to the valuation. Beginning with the floating leg, given a day count convention and

---

[3]Assuming we are in a period with no holidays.

[4]The payment frequency on the floating leg matches the term, while the fixed leg does not necessarily have that property.

		Floating leg			Fixed leg	
Currency	Spot start	Term	Freq	Day count	Freq	Day count
EUR	2B	6M	S	Act360	A	30/360
USD	2B	3M	Q	Act360	S	30/360
GBP	0B	6M	S	Act365	S	Act365
JPY	2B	6M	S	Act360	S	Act365
DKK	2B	6M	S	Act360	A	30/360

Table 2.2: *Plain vanilla IRS conventions.*

knowing the start and end date of an IRS, denoted by T_S and T_E respectively, we are able to calculate a set of coverages, $\delta_{S+1}^{\text{float}}, \ldots, \delta_E^{\text{float}}$. The no-arbitrage principle tells us that the floating interest rate for each period must be equal to the implied forward rate that is consistent with market-observable rates. Using the already established framework for the forward rates and zero coupon prices we can calculate the unit value of the floating leg as the sum of discounted forward rate payments

$$PV^{\text{float}} = \sum_{i=S+1}^{E} \delta_i^{\text{float}} F(0, T_{i-1}, T_i) D(0, T_i) \qquad (2.2)$$

Assuming that forward rates and zero coupon prices are obtained from the same yield curve, we can actually simplify (2.2) further. Substituting (2.1) from page 5 for $F(0, T_{i-1}, T_i)$ in (2.2) we come up with the simple expression

$$PV^{\text{float}} = D(0, T_S) - D(0, T_E) \qquad (2.3)$$

This simplification, however, is not empirically valid since the forward rates and zero coupon prices are *not* based on the same yield curve.[5] Therefore, we will continue thinking of (2.2) as the price of the floating leg. The difference between the yield curves is caused by the implicit risk element built into the forward rates. This risk element "quantifies" default risk—the case when the borrower cannot pay the loaner. Therefore, empirically, it is cheaper to fund on a 3 months rate than a 6 months rate, meaning that the effective annual interest rate on a loan with quarterly payments is smaller than that of a loan with semi-annual payments. Figure 2.1 shows the swap curve against EURIBOR3M together with the swap curve against EURIBOR6M. We see that the swap curve with the 3M tenor consistently lies under that with the 6M tenor.

Now turning to the valuation of the fixed leg of the IRS, we have already seen in Table 2.2 that the day count convention and the payment frequency of the fixed

---

[5] For a text on the matter of several simultaneous yield curves we refer to Fujii et al. (2010).
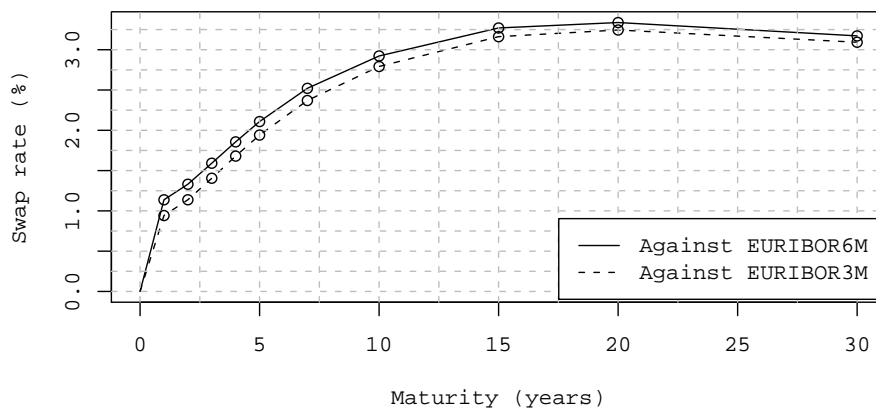
Figure 2.1: *Swap curves per June 1st 2010. Notice how the swap curve against EURI-BOR3M is consistently lower than the swap curve against EURIBOR6M.*

leg do not necessarily match those of the floating leg. However, we will continue to use the notation $\sum_{i=S+1}^{E}$ as a way of saying "summing over all i's (periods) between $S + 1$ and E". The start and end date will not change, but for the fixed leg we calculate a new set of coverages, $\delta_{S+1}^{\text{fix}}, \ldots, \delta_{E}^{\text{fix}}$. Letting K denote the fixed rate paid in the swap, the present value of the fixed leg can be obtained by

$$PV^{\text{fix}} = \sum_{i=S+1}^{E} \delta_i^{\text{fix}} \, K \, D(0, T_i) \tag{2.4}$$

As shown, the price of the fixed leg is merely the discounted fixed rate payments.

We have now calculated the present values of each leg in the swap and we can now calculate the value of the full swap by combining the prices of the two legs correctly. Remembering that the name of the swap refers to the *fixed* leg, we will use the term *payer swap* for a swap that pays a fixed rate and receives a floating rate. The value of the payer swap is

$$PV^{\text{pay}} = \sum_{i=S+1}^{E} \delta_i^{\text{float}} \, F(0, T_{i-1}, T_i) \, D(0, T_i) - \sum_{i=S+1}^{E} \delta_i^{\text{fix}} \, K \, D(0, T_i) \tag{2.5}$$

Since the positions are simply reversed for the receiver swap, the price of the receiver swap can found as $PV^{\text{rec}} = -PV^{\text{pay}}$.

When swaps are entered into, it is customary to do so at a present value of 0. Hence, we need to set $PV^{\text{float}} = PV^{\text{fix}}$ and solve for the appropriate fixed rate. Setting (2.5) equal to 0 we isolate the fixed rate that results in a present value of

9

0 of the IRS. This fixed rate is called the *par swap rate*, and we will denote it by $R(0, T_S, T_E)$, so

$$R(0, T_S, T_E) = \frac{\sum_{i=S+1}^{E} \delta_i^{\text{float}} F(0, T_{i-1}, T_i) D(0, T_i)}{\sum_{i=S+1}^{E} \delta_i^{\text{fix}} D(0, T_i)} \tag{2.6}$$

Under the theoretical assumption, that there is indeed only *one* yield curve from which both forward rates and discount factors are calculated, using the expression for the forward rate from (2.1) we can simplify (2.6) further writing

$$R(0, T_S, T_E) = \frac{D(0, T_S) - D(0, T_E)}{\sum_{i=S+1}^{E} \delta_i^{\text{fix}} D(0, T_i)} \tag{2.7}$$

Having defined the par swap rate, we are now able to express the present value of the IRS in a different manner. Assume that we have entered into a payer swap at time T_S and are thus paying the fixed rate K and receiving the floating xIBOR rate until the IRS expires at T_E. Should we, at some time t, enter into a matching receiver swap implying that we are receiving the fixed par swap rate $R(t, T_S, T_E)$ and paying the floating xIBOR rate, the floating rates will cancel out and leave us with a net periodical (fixed) cash flow of $R(t, T_S, T_E) - K$. Defining the term $A(\cdot)$ as the time t sum of discounted fixed coverages, $A(t, T_S, T_E) = \sum_{i=S+1}^{E} \delta_i^{\text{fix}} D(t, T_i)$, the time t present value of our initial payer swap can be written as[6]

$$PV_t^{\text{pay}} = A(t, T_S, T_E) \left(R(t, T_S, T_E) - K \right) \tag{2.8}$$

This is a very convenient way of determining the value of a swap position. Assuming that we are currently paying the fixed rate K, all we need to know to find the value of our current position is $A(\cdot)$ and the current par swap rate $R(\cdot)$. If we are paying a fixed rate that is *higher* than the time t par swap rate, then our position has a negative value, but if the opposite is the case, we have a positive value and we can offset[7] our swap position at a gain.

Differentiating (2.8) with respect to the par swap rate we get exactly $A(\cdot)$ and therefore we might think of $A(\cdot)$ as being the value of receiving 1 interest rate basis point (bp) over a period of $T_E - T_S$, or we might think of it as being the payer IRS's sensitivity towards the par swap rate.[8] This also referred to as *the*

---

[6]For $t > T_S$, replace T_S with t (+ *spot start* business days according to Table 2.2).

[7]Instead of an actual "termination" of a swap contract it is common practice to offset it by taking on an opposite position in yet another swap contract, thus canceling out the swap payments. At the time of the "termination", however, the swap contract rarely has a present value of 0 and an up-front fee is paid or received on the swap contract used to offset the initial swap contract—hence the gain (or loss).

[8]Correspondingly, the sensitivity of the receiver IRS is $-A(\cdot)$.

annuity of the swap.

2.3 Option pricing

This section will start off with a brief recapitulation of some of the theory behind option pricing. Once the necessary steps have been taken, the option pricing theory will be extended to cover options written on the IRSs introduced in section 2.2.

2.3.1 Building blocks

First, as a beginning to our option pricing theory, we start out by defining the *European call option* as the right, but not the obligation, to *buy* the underlying asset (UA) at a fixed strike price K at a specified time T. Letting V_t denote the value of UA at time t, the call option will have a time T payoff of $\max[V_T - K, 0]$. For the remainder of this paper we will use the short hand notation $(V_T - K)^+$ for payoffs of this type. Similarly, we define the *European put option* as the right, but not the obligation, to *sell* UA at strike K at time T yielding a payoff of $(K - V_T)^+$. The values of these options are calculated by discounting their expected (stochastic) time T payoffs under the risk neutral probability measure Q.[9]

Second, we introduce the concept of *martingales*. A stochastic process b_t is a martingale under the probability measure Q if it holds that $\mathbb{E}_t^Q[b_T] = b_t$, $\forall t < T$.[10] Essentially, this implies that the process has no drift, and therefore its expected future value must equal its present value. As an example, we observe an asset with a price process V_t. Let r_t denote the process of the risk free interest rate, then under the risk neutral probability measure Q the dynamics of V_t are $dV_t = r_t V_t dt + \sigma_{V_t} dW_t^Q$ where W_t^Q is a Brownian motion and σ_V is some volatility parameter. Now, V_t obviously is not a martingale, but the discounted value is. We can (under the risk neutral probability measure) transform our asset process V_t into a martingale by using $g_t = e^{\int_0^t r_s ds}$ as the numeraire. Mathematically expressed we have

$$\mathbb{E}_t^Q\left[\frac{V_T}{g_T}\right] = \frac{V_t}{g_t} \quad \left(\text{and especially} \quad \mathbb{E}_0^Q\left[\frac{V_T}{g_T}\right] = V_0\right)$$

This leads us to Theorem 2.1

---

[9]Unless anything else is stated, we will assume that processes and expectations are under the risk-neutral probability measure Q.

[10]Letting $\mathbb{E}_t^Q$ denote the time t expectation under the probability measure Q.

Theorem 2.1. (FIRST FUNDAMENTAL THEOREM OF ASSET PRICING)[11]
Given a time horizon T, a risky asset with price process V_t and a risk-free asset with price process g_t, a market is arbitrage-free (under the probability measure P) if and only if there exists an equivalent probability measure Q such that the discounted price process, $\left(\frac{V}{g}\right)_t$, is a martingale.

This tells us, that if we indeed are in a market with no arbitrage, we calculate prices by finding expected (discounted) values. Especially, this is useful in mathematical finance together with the concept of *change of numeraire*. Changing the numeraire is a technique that essentially allows us to use any traded asset N_t as our numeraire. One can show that if the discounted numeraire itself is a martingale under the risk neutral measure Q, then there exists an equivalent probability measure Q^N such that $\mathbb{E}_t^{Q^N}\left[\frac{V_T}{N_T}\right] = \frac{V_t}{N_t}$ for any asset with a price process V_t.

Having established the concept of a martingale process we turn to the *martingale representation theorem*

Theorem 2.2. (MARTINGALE REPRESENTATION THEOREM)[12]
Let W_t be a standard Brownian motion, and let $(m_t)_{0 \leq t < \infty}$ be a martingale process. Then there exists a uniquely determined (possibly stochastic) process $c(\cdot)$ such that

$$m_t = m_0 + \int_0^t c(\cdot)_s \, dW_s \quad \text{or (equivalently)} \quad dm_t = c(\cdot)_t \, dW_t \tag{2.9}$$

This tells us, that if we have a martingale process m_t, then the uncertainty regarding its development arises from a Brownian motion multiplied by some process $c(\cdot)$.

2.3.2 The Black-Scholes result

The next step naturally would be to come up with a model for $c(\cdot)$. A particularly well-known model is that of Black (1976). Black assumes $c(t, m_t)_t = \sigma \, m_t$ so that (2.9) becomes

$$dm_t = \sigma \, m_t \, dW_t \tag{2.10}$$

This model implies that $\log(m_t)$ is normally distributed with a standard deviation of $\sigma\sqrt{t}$ (and a mean of m_0). In order to take the model further we utilize the Black-Scholes formula (Black and Scholes, 1973). The BS formula in general considers an asset with a price process V_t with lognormally distributed prices and log returns with a standard deviation of ω. Adapting this notation and letting $\Phi(x)$ denote

---

[11]Cf. Björk (2004).
[12]Ibid.

the cumulative standard normal distribution the BS formula says

$$\mathbb{E}_t[(V_T - K)^+] = \mathbb{E}_t[V_T]\,\Phi(d_1) - K\,\Phi(d_2) \tag{2.11}$$

$$d_1 = \frac{\log(\mathbb{E}_t[V_T]/K) + \frac{1}{2}\omega^2}{\omega}$$

$$d_2 = d_1 - \omega$$

$\mathbb{E}_t[V_T]$ is the forward price—the time t expectation of a price at time $T > t$. Hence, by incorporating Black's assumption regarding the standard deviation $\omega = \sigma\sqrt{T-t}$ into (2.11) and discounting to present value by a risk free rate r_f we are left with what is known as *the Black '76 formula*

$$\text{B76 Call}_t = e^{-r_f\,(T-t)}[f_t\,\Phi(d_1) - K\,\Phi(d_2)] \tag{2.12}$$

$$d_1 = \frac{\log\left(\frac{f_t}{K}\right) + \frac{\sigma^2}{2}(T-t)}{\sigma\sqrt{T-t}}$$

$$d_2 = d_1 - \sigma\sqrt{T-t}$$

Where f_t denotes the the time t forward price of an asset at time $T > t$. The Black '76 formula is a special case of the Black-Scholes formula for the price of European options. The difference is that in the original Black-Scholes setting, the *spot price* is used instead of the forward price.

Since the Black '76 formula uses forward prices rather than spot prices it is especially useful in pricing bond options, caps, floors and swaptions—the latter being a key derivative in this paper.

2.4 Swaptions

Now that we are familiar with both swaps and options we will introduce the concept of *swaptions*. We will start out by explaining what constitutes a swaption and then we will move on to the pricing of the swaption using the tools and methods described in the previous sections.

2.4.1 Swaption basics

A swaption is the right, but not the obligation, to enter into an IRS[13] starting at T_S and maturing at T_E. Assuming that $t = 0$, standard notation for a swaption of this type is "T_S into $(T_E - T_S)$ payer/receiver swaption". For example, if we have $T_S = 1Y$ and $T_E = 6Y$ on an agreement in which we have the right to receive the

---

[13]Or any other type of swap.

fixed rate K in exchange for paying some floating xIBOR rate we would be entering into a *1Y into 5Y (1Y5Y) receiver swaption.* When the swaption is entered into, the two parties must agree on the type of settlement used. There are two types of settlement:

1. *Physical settlement* which means that the swap is entered into in a normal fashion where the actual exchanges of cash flows according to the underlying swap take place. Unless otherwise stated a swaption has physical settlement.

2. *Cash settlement* where—as the name implies—the swaption is settled through an exchange of cash corresponding to the value of the swap at the time of exercise.

2.4.2 Swaption pricing

As described in the previous section, the swaption can be settled in two different ways. The choice of settlement does not only cause a difference in the cash flows, but also in the actual valuation of the swaption. For the pricing analysis we will assume the option to be on a *payer* swaption.

Physical settlement In the case of physical settlement the option holder simply has the possibility of entering into the swap on the previously established conditions. This corresponds to (2.8) with the addition of a possibility to opt out of the swap. Therefore we get

$$\text{Payer swaption } PV_t^{\text{phys}} = A(t, T_S, T_E) \left[(R(t, T_S, T_E) - K)^+ \right] \qquad (2.13)$$

$$\text{where} \quad A(t, T_S, T_E) = \sum_{i=S+1}^{E} \delta_i^{\text{fix}} D(t, T_i)$$

Cash settlement For the cash settlement, the option holder receives the potential positive present value of the swap at the time of exercise, but in this case the par swap rate is used as the discount rate. In a theoretical world with only one yield curve this would have no effect, but, as we have seen in Figure 2.1 on page 9, the real-life market operates with several yield curves, which in turn leads to varying prices depending on what yield curve is used to calculate swap rates and discount factors. We get

$$\text{Payer swaption } PV_t^{\text{cash}} = \tilde{A}(t, T_S, T_E) \left[(R(t, T_S, T_E) - K)^+ \right] \qquad (2.14)$$

$$\text{where} \quad \tilde{A}(t, T_S, T_E) = \sum_{i=S+1}^{E} \frac{\delta_i^{\text{fix}}}{(1 + \delta_i^{\text{fix}} R(t, T_S, T_E))^{T_i - t}}$$

As it turns out, the main difference between valuing a cash or physically settled swaption is in the discounting. For ease of notation we will therefore use $A(\cdot)$ in the following as a reference to both cases. Having already established the Black-Scholes framework in (2.11) we will now use this to come up with a closed form solution to the value of a swaption. Thinking of the par swap rate $R(\cdot)$ as the underlying asset the similarity to the standard European call option on a stock becomes apparent. The main difference in the pricing, is that while the underlying asset in the standard Black-Scholes setting appreciates with the risk free rate r_f under the risk neutral measure

$$\mathbb{E}_0^Q[S_T] = e^{r_f T} S_0$$

making the present value of the asset a martingale when using $g_t = (e^{r_f t})_t$ as a numeraire, the par swap rate (the underlying asset in a swaption) will need another numeraire, namely the $A(\cdot)$ function.[14] We will not go into the derivations, but merely state the result

$$\text{Payer swaption } PV_t = A(t, T_S, T_E)[R(t, T_S, T_E)\Phi(d_1) - K\,\Phi(d_2)] \qquad (2.15a)$$

$$d_1 = \frac{\log\left(\frac{R(t,T_S,T_E)}{K}\right) + \frac{\sigma^2}{2}(T-t)}{\sigma\sqrt{T-t}} \qquad (2.15b)$$

$$d_2 = d_1 - \sigma\sqrt{T-t} \qquad (2.15c)$$

The price of the opposite swaption—the *receiver swaption*—can be obtained through a parity relationship[15] similar to the put-call-parity for standard European plain vanille options. All contracts having the same strike rate K and the same time dimensions[16] the parity states

$$\text{Swaption}^{\text{rec}} = \text{Swaption}^{\text{pay}} + \text{Forward starting swap}^{\text{rec}} \qquad (2.16)$$

Rearranging (2.16) we see that a forward starting receiver swap can be replicated by a long position in a receiver swaption and a short position in a payer swaption. To validate this claim we will look at the payoffs. We assume that we are at time 0 and that the swaptions both have the same starting date T_S and the same end date T_E. Further, we assume that the swaptions are struck at the same strike K. Now, obviously there are no payoffs up until T_S. After this point the payoffs can go two ways depending on whether the par swap rate at T_S is smaller or greater

---

[14]Cf. Schrager and Pelsser (2006).
[15]Cf. Sundaresan (2009).
[16]All entered into at t with start date T_S and end date T_E.

$t = T_S$	$K > R(T_S, T_S, T_E)$	$K < R(T_S, T_S, T_E)$
Long receiver swaption:	$K - \text{xIBOR}$	0
Short payer swaption:	0	$-(\text{xIBOR} - K)$
Net	$K - \text{xIBOR}$	$K - \text{xIBOR}$

Table 2.3: *Example of the put-call parity for swaptions. Net payoffs equal those of a receiver swap.*

than K. The payoffs are shown in Table 2.3. We see that the net payoff from being long a receiver swaption and short a payer swaption does indeed equal that of being long a (forward starting) receiver swap.

Part II

Volatility models

Chapter 3

Motivation

3.1 Constant volatility

In the standard Black-Scholes formula, the volatility of an asset's returns is assumed to be constant. The Black-Scholes formula provides a 1-to-1 relationship between an option's price and its volatility, and therefore prices are often quoted by stating the *implied volatility*. Given an option price C^* and letting $C(\cdot)$ denote the Black-Scholes price for a European plain vanilla call option

$$\text{BS Call } PV_t = V_t\,\Phi(d_1) - e^{-r_f(T-t)}K\,\Phi(d_2) \tag{3.1a}$$

$$\text{where} \quad d_1 = \frac{\log(V_t/K) + (r_f + \frac{\sigma^2}{2})(T-t)}{\sigma\sqrt{T-t}} \tag{3.1b}$$

$$\text{and} \quad d_2 = d_1 - \sigma\sqrt{T-t} \tag{3.1c}$$

then the implied volatility is the σ_B that sets $C^* = C(V, K, \sigma_B, r_f, T)$. The function cannot be inverted in closed form, and thus σ_B must be solved for numerically. However, this is a trivial problem that can easily be solved by applying for example the Newton-Raphson method.

The idea of a single and constant volatility might seem appealing, but will empirical studies support it? Figure 3.1 shows daily log-returns for the Carlsberg A/S B stock over the past seven years, together with 95% confidence intervals.[1] From just looking at the plot it becomes apparent, that the return volatility of Carlsberg is not constant. Further evidence of time-varying volatility is found in Figure 3.2 which shows the empirical distribution of the returns together with the best fitting normal distribution. The returns seem to have a tendency to have a higher peak and slightly fatter tails than expected under the normal distribution

---

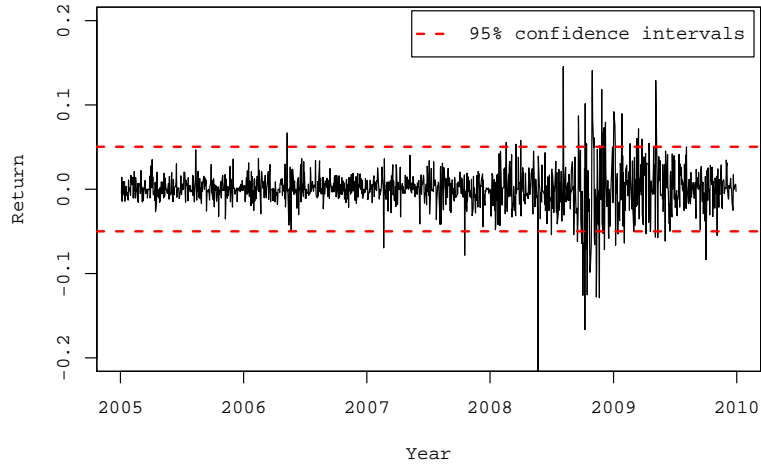[1]Calculated as mean(returns) $\pm\,\Phi^{-1}(97.5\%) \times$ StdDev(returns).

Figure 3.1: *Daily returns for Carlsberg A/S B stock between January 1st 2005 and January 1st 2010.*
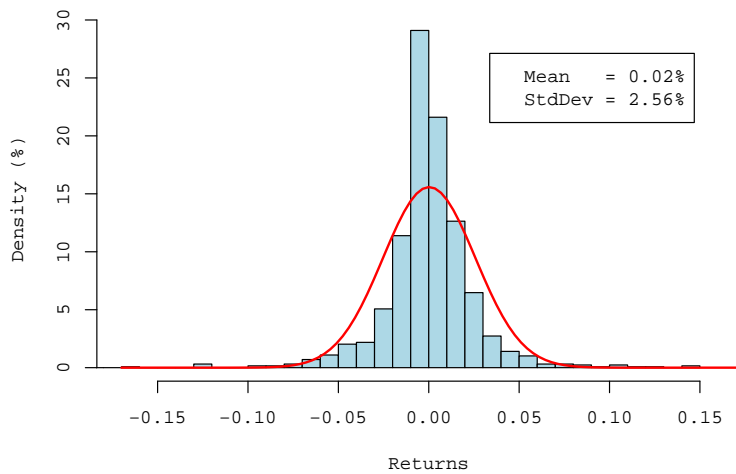

Figure 3.2: *Distribution of daily returns for Carlsberg A/S B stock between January 1st 2005 and January 1st 2010. The red line represents the best fitting normal distribution.*

19

assumption. This is confirmed when looking at the skewness and kurtosis

$$\text{Kurtosis} \equiv \frac{\frac{1}{n}\sum_{i=1}^{n}(x_i - \mu)^4}{\sigma^4} \quad \text{and} \quad \text{Skewness} \equiv \frac{\frac{1}{n}\sum_{i=1}^{n}(x_i - \mu)^3}{\sigma^3}$$

If the returns were indeed normally distributed they should have a kurtosis ≈ 3 and a skewness ≈ 0. However, we find that the kurtosis is 18.03 and skewness is -0.96. The kurtosis tells us, that the returns are indeed excessively peaked, and the negative skewness indicates, that the left tail is "longer", meaning that we are more prone to see large *negative* returns compared to large *positive* returns. According to Gatheral (2002), high peaks and the fat tails are typical traits of mixtures of distributions with different variances, which again is an indicator that the Black-Scholes assumption regarding constant volatility is violated.

Now that we have convinced ourselves that volatility is indeed time-varying, the concept of implied volatilities varying with maturities of options seems reasonable. However, as counterintuitive as it might seem, option market data suggests yet another anomaly in the volatility structure. When inferring volatilities from a specific option, a common pattern appears. In general, the implied volatility is lower for *at-the-money* (ATM) options, and higher for *in-the-money* (ITM) and *out-of-the-money* (OTM) options. This is known as the volatility smile. An example of a volatility smile is presented in Figure 3.3. Hagan et al. (2002) point out that not only is the smile present but it also seems to be moving as the price of the underlying asset shifts. These dynamics present market participants with a number of issues.

First, we note that the task of pricing exotic derivatives[2] becomes increasingly non-trivial since a single volatility no longer suffices. Take for example barrier options.[3] Imagine a call option with $V_0 = 105$ and $K = 110$, but with the built-in condition that it does not kick in unless the price of the underlying asset drops to 100 at some time $t > 0$ before a maturity $T > t$

$$\text{Exotic barrier call payoff}_T = \mathbf{1}_{\{V_t \leq 100 | 0 < t < T\}}(V_T - 110)^+$$

We are now faced with the question of which volatility to use. The volatility at a strike of 100? The volatility at a strike of 110? A mix of the two maybe? Furthermore, we might find ourselves in a situation where we need a quote for an implied volatility that we cannot observe directly in the market, e.g. we might need an implied volatility for an option that is far OTM. Therefore, we must implement

---

[2]More complicated derivatives—as opposed to plain vanilla derivatives.

[3]Options whose value is dependent on the price of the underlying asset hitting a certain barrier.
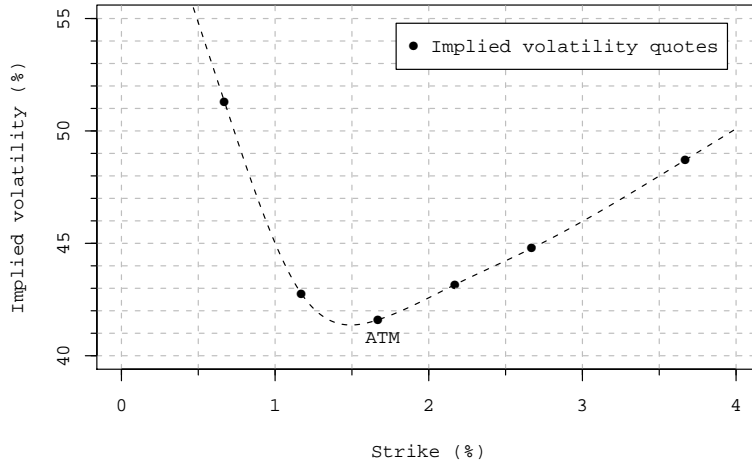
Figure 3.3: *Volatility smile for the 3M2Y swaption against EURIBOR6M as of November 1st 2010. The dashed line represents a cubic spline interpolation.*

some kind of rule for inter- and extrapolating the observed points of the volatility smile.

Secondly, in addition to correct pricing, financial institutions must also be able to manage the risk entailed by their positions in derivatives. Since the implied volatility itself changes as the strike price differs and since the smile shifts as the price of the underlying asset fluctuates, we need to come up with a way to manage the risks implied by these dynamics. In classical terms we want to be able to correctly identify and isolate delta (Δ) and vega ($\mathcal{V}$) risk from one another.[4]

What we need is a single, self-consistent model for the implied volatility from which we can extract volatilities for any (reasonable) strike price.

---

[4]Δ risk and $\mathcal{V}$ risk are risks associated with a change in the price of the underlying asset and with a change in implied volatility respectively.

Chapter 4

The local volatility model

4.1 Local volatility

Some of the first steps taken towards taming the volatility smile were taken by Derman and Kani (1994) with their *local volatility model*. Using notation from (2.9) Derman and Kani proposed a model in which $c(\cdot) = \sigma_{\mathrm{loc}}(m_t, t) \, m_t$ so that

$$dm_t = \sigma_{\mathrm{loc}}(m_t, t) \, m_t \, dW_t \tag{4.1}$$

The model is best described in its discretized version, in which Derman and Kani use the technique of *binomial option pricing* (Cox et al., 1979) to obtain a grid of volatilities that causes the binomial tree's prices to be consistent with market data.

    The standard setting for binomial pricing uses a start value for the underlying asset, a risk-free interest rate and assumes a constant volatility. From these parameters, and given a time frame, a grid of asset prices is laid out, and this grid is then used to price claims by calculating and discounting the expected value of the claims using risk-neutral transition probabilities. Derman and Kani, however, reverse the process. Instead of pricing claims based on the standard inputs, Derman and Kani use market quotes on option prices[1] together with forward prices to come up with the *local volatilities* for each node in the binomial tree. Together, these local volatilities make up *the local volatility surface*, and they are calculated in a way so that the resulting binomial tree becomes consistent with what is observed in the market. Hence, the procedure results in a local volatility grid, that can be used, together with the grid describing the development of the underlying asset, to price various more complex products in a manner that is self-consistent
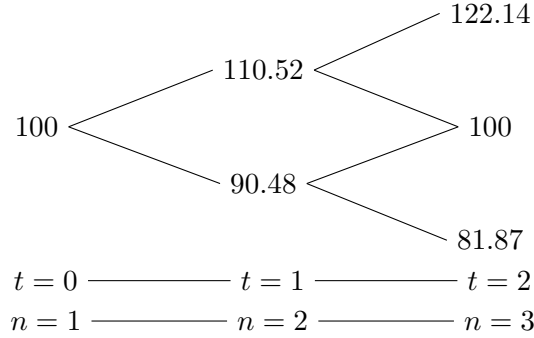
---

[1]Implied volatilities.

Figure 4.1: *Development of S_t under the assumption of constant volatility. Nodegroup denoted by n and time denoted by t.*

and in line with the market's belief regarding volatilities.

4.2 Example of discrete time local volatility

To demonstrate the idea of local volatility, we will apply it to a simple two-period binomial tree. We will let the underlying asset, S, have an initial value of $S_0 = 100$, the risk-free rate is set at $r_f = 3\%$, the (constant) volatility is $\sigma = 10\%$ and the time periods will be in years starting at year 0 and ending in year 2. In the usual constant volatility binomial setting, the asset prices when moving up or down in the grid are determined as $u \times S$ and $d \times S$ respectively, where $u = e^{\sigma\sqrt{h}}$ and $d = e^{-\sigma\sqrt{h}}$. This is done based on a belief that stock prices are lognormally distributed and that returns are normally distributed with a standard deviation of $\sigma\sqrt{h}$ where h is the length of the period (McDonald, 2006). We assume period lengths of 1 and therefore we get an up factor $u = e^{10\%} = 1.105$ and a down factor $d = e^{-10\%} = 0.905$. The risk-neutral probability p is calculated based on fairly priced forwards. Letting F_S denote the forward price of S and assuming a period length of 1, it must hold, that

$$p\,S\,u + (1-p)\,S\,d = F_S \equiv S\,(1+r_f) \quad \Leftrightarrow$$
$$p = \frac{(1+r_f) - d}{u - d}$$

For our standard binomial setup we find $p = \frac{(1+0.03)-0.905}{1.105-9.905} = 0.625$. The development of S_t under these assumptions is shown in Figure 4.1. However, before we turn to the example calculations we will explain the setup for the local volatility binomial model and build the required set of tools.

4.2.1 Preliminaries

The first thing to note regarding the local volatility model is that what we *know* is the implied volatility smile. This corresponds to us knowing European option prices for all strikes. We also know the risk-free interest rate r_f and since we are under the risk neutral probability measure, we also know that the (1-year) forward prices must satisfy $F = (1 + r_f)S$. Our objective is now to use this information to calculate the implied risk neutral transition probabilities of moving up or down in our tree, and to calculate the state values of these up and down states.

The setup

Before we start deriving formulae and doing actual calculations, we will go through the setup of our model. The first thing that needs to be explained is the time dimension. We will keep the standard notation in which the tree starts at time $t = 0$, but at the same time we will introduce the concept of a "nodegroup"[2] denoted by n (cf. Figure 4.1). The link between t and n is simply $n = t + 1$. This means that the very first group of nodes—the single node at $t = 0$—is nodegroup 1 ($n = 1$). This will also ease our notation since in any nodegroup n' there is exactly n' nodes. The local volatility model is solved *one nodegroup at the time*, starting with $n = 1$, $n = 2$, $n = 3$ and so on. Now that we have defined our horizontal dimension in the grid we will move on to the vertical dimension. We will index vertically by i, so that for every nodegroup n' the bottom node will have index $i = 1$ and the top node will have index $i = n'$. This further allows us to easily refer to any node in the grid by using the notation (n, i).

We will need several parameters for our local volatility model. To help illustrate the meaning of these parameters we refer to Figure 4.2. Below we will list the parameters and explain their roles.

S_0: The spot price of the underlying asset. This price is known at $t = 0$ and is a key value. S_0 is kept at the center of the grid, which means that whenever there *exists* a center node in a nodegroup (when n is odd), this node will have the state value S_0—illustrated by the horizontal dashed line in Figure 4.2a.

s_i: The small s will act as an indicator of a calculated (known) state value in the grid. From this known state value s_i, the tree can evolve either up to the unknown S_{i+1} (next item) or down to the unknown S_i (two items down).

---

[2]A node is defined as the point where edges meet in the tree.

(a) Implied grid around the center.   **(b)** Implied grid away from the center.
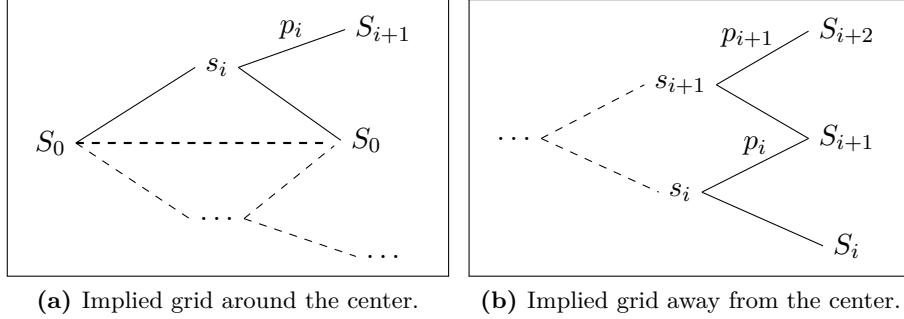
Figure 4.2: *Setup for the local volatility implied grid. To the left the grid around the center (S_0) is illustrated (the horizontal line is merely an illustration of the center—it is NOT a path). To the right we see the the setup for the remaining grid.*

S_{i+1}: This is the (unknown) state value reached after an upward move from s_i.

S_i: This is the (unknown) state value reached after a downward move from s_i.

p_i: The risk neutral probability of moving from s_i to S_{i+1}. Equivalently, the probability of moving from s_i to S_i is $1 - p_i$.

F_i: The one-period forward price at the node where s_i is the known spot price.

λ_i: This is more of a helper on the notation than an actual piece of the model. λ_i denotes the Arrow-Debreu price[3] of node i. Both s_i and λ_i will be "on the same node".

Having defined our setup and the parameters we are now ready to derive the tools needed to determine the unknown parameters through our knowledge of the spot price, the interest rate (and thereby iteratively the forward price) and the volatility smile.

The formulae

As a general reference to this section we give Derman and Kani (1994).

As mentioned, the unknowns that we are trying to calculate in this model are the transition probabilities p_i, and the future state values S_i. We *choose* to maintain the spot price S_0 as a central node throughout the tree. A center is chosen in order to uniquely determine the parameters. We no longer have an a priori structure for the development of S_t as we did in the standard binomial case. Instead we must

---

[3]The sum (over all possible paths to the node) of the product (of discounted transition probabilities). Note: since the transition probabilities change throughout the tree, the standard binomial model formula for the Arrow-Debreu prices $(1 + r_f)^t \binom{n}{i} p^i (1-p)^{n-i}$ cannot be used.

infer the market's take on the development of S_t through the observed volatility smile. At an arbitrary point in time $n > 0$ in developing the grid there are $2n + 1$ parameters perfectly describing the transition to time $n + 1$: the $n + 1$ new stock prices S_i and the n transition probabilities p_i.[4] However, at time n all the information we have is n forward prices and n European option prices.[5] To sum up we are left with the task of determining $2n + 1$ parameters from $2n$ equations (options and forward prices), and therefore we add the centering constraint to be able to uniquely determine all parameters. As mentioned we choose to center the nodes around S_0. This condition leaves us with an equal amount of parameters and equations.

We are now set up to stepwise determine the local volatility grid corresponding to the observed volatility smile. To do so, the first identity that must hold is

$$F_i = p_i S_{i+1} + (1 - p_i) S_i \qquad (4.2)$$

We remind ourselves, that i tells us where we are located *vertically* in the grid, while n tells us where we are *horizontally*. F_i is the known forward price while S_i denotes a downwards shift and S_{i+1} denotes an upwards shift. We now let $C(K, T)$ represent the present value of a call option with strike K and maturity T. Using the Arrow-Debreu prices λ_i, the value of the call option can be written as

$$C(K, n+1) = e^{-r_f \Delta t} \sum_{j=1}^{n} (\lambda_j \, p_j + \lambda_{j+1}(1 - p_{j+1})) \max[S_{j+1} - K, 0] \qquad (4.3)$$

where we use the notation $e^{-r_f \Delta t}$ to indicate a discounting all the way back to time 0. Setting $K = s_i$ allows us to rearrange and simplify the expression using the resulting fact that only up-moves will have a positive contribution while down-moves will have a value of 0. Also, we know that forwards must be priced according to (4.2). With this extra insight we rewrite (4.3) into

$$C(s_i, n+1) = e^{-r_f \Delta t} \left[\lambda_i \, p_i (S_{i+1} - s_i) + \sum_{j=i+1}^{n} \lambda_j (F_j - s_i) \right] \qquad (4.4)$$

One can think of the first term within the brackets in (4.4) as a "maybe-value"

---

[4]At any node, the probability of a down-move p_{down} can be parametrized as $1 - p_{\text{up}}$ and vice versa. Therefore, the probabilities of the total $2n$ possible movements from the nodes in nodegroup n are fully described by n parameters.

[5]Similar to the argument for the transition probabilities, one could be led to believe that we in fact know $2n$ option prices—the calls and the puts—but these are linked through the put-call parity, and thus one can be parametrized through the other.

(with probability p_i) and the second term as "sure-value" since the former only contributes with a postive payoff in the case of an up-move, and the latter will surely contribute with a positive payoff (with an expected value equal to the forward price as presented in (4.2)) regardless of the direction of the move.

Since a volatility smile—and thereby an option price—is given and the forward prices are known, we are left with an expression containing only two unknowns: the transition probability p_i and the value of the underlying asset after an upwards move S_{i+1}. By combining (4.2) and (4.4) one might say that we are adding an extra unknown variable S_i, but this brings us back to our choice of centering the grid around S_0. Having done so, we are able to start at the central node $(i = n/2 + \frac{1}{2})$ and work our way upwards. Solving (4.2) and (4.4) simultaneously yields

$$S_{i+1} = \frac{S_i \left[e^{r_f \Delta t} C(s_i, n+1) - \Sigma \right] - \lambda_i s_i (F_i - S_i)}{\left[e^{r_f \Delta t} C(s_i, n+1) - \Sigma \right] - \lambda_i (F_i - S_i)} \tag{4.5}$$

$$p_i = \frac{F_i - S_i}{S_{i+1} - S_i} \tag{4.6}$$

$$\text{where } \Sigma = \sum_{j=i+1}^{n} \lambda_j (F_j - s_i)$$

As mentioned, one of the key elements to this method of iteratively calculating state values and transition probabilities is knowing the value of the central node. So, what happens when we are in a position where there *is* no central node— e.g. moving from $n = 1$ to $n = 2$? Obviously, when n is odd (meaning that we are moving *to* an even n) there will be no central node. Derman and Kani solve this problem by setting the mean of the natural logarithms of the two central nodes equal to the natural logarithm of todays spot price

$$\log(S_0) = \frac{\log(S_{i+1}) + \log(S_i)}{2} \quad \text{or equivalently}$$

$$S_i = S_0^2 / S_{i+1} \tag{4.7}$$

Substituting (4.7) into (4.5) and rearranging we find the following expression which we will use to determine the state value of the first node above the center of the grid when moving *from* an odd n

$$S_{i+1} = \frac{S_0 \left[e^{r_f \Delta t} C(S_0, n+1) + \lambda_i S_0 - \Sigma \right]}{\lambda_i F_i - e^{r_f \Delta t} C(S_0, n+1) + \Sigma} \tag{4.8}$$

Calculations similar to those just performed can be done using put options instead

of call options to obtain the lower half of the grid. We will, however, merely refer
to Derman and Kani (1994), and simply state the result for the formula used to
calculate the state values iteratively for the nodes below the central node. Letting
$P(K,T)$ denote the present value of the put option struck at K and maturing at
T, the expression is

$$S_i = \frac{S_{i+1}\left[e^{r_f \Delta t}P(s_i, n+1) - \Sigma\right] + \lambda_i\, s_i(F_i - S_{i+1})}{\left[e^{r_f \Delta t}P(s_i, n+1) - \Sigma\right] + \lambda_i(F_i - S_{i+1})} \tag{4.9}$$

$$\text{where } \Sigma = \sum_{j=1}^{i-1} \lambda_j(s_i - F_j)$$

The transition probability p_i is calculated as given in (4.6).

For both the lower and the upper half of the implied grid, the implied volatility
for each node based on the possible state values and the transition probability is
calculated as

$$\sigma_i = \sqrt{p_i\,(1-p_i)}\log(S_{i+1}/S_i) \tag{4.10}$$

To see how (4.10) is valid, consider a stochastic variable Y that follows a two-point
distribution, so that $\Pr(Y = a) = p$ and $\Pr(Y = b) = 1 - p$. Now, the variance of
Y is

$$\begin{aligned}
\mathrm{Var}[Y] &= \mathbb{E}[Y^2] - \mathbb{E}[Y]^2 \\
&= p\,a^2 + (1-p)b^2 - (p\,a + (1-p)b)^2 \\
&= p(1-p)(a-b)^2 \tag{4.11}
\end{aligned}$$

Since the only thing we changed, going the Black-Scholes setup to the local volatil-
ity setup, was the nature of the volatility, we still assume the price of the underlying
asset to follow a geometric Brownian motion

$$\frac{dS_t}{S_t} = \mu\, dt + \sigma(S_t, t)\, dW_t$$

Using results from Björk (2004) obtained through application of Itô's lemma the
dynamics of $X_t = \log(S_t)$ can be written as

$$dX_t = \left(\mu - \frac{\sigma(S_t, t)^2}{2}\right) dt + \sigma(S_t, t)\, dW_t$$

Assuming that the volatility function is deterministic, and therefore does not de-

pend on S_t, the variance of X_t is

$$\text{Var}[X_t] = \int_0^t \sigma(s)^2 ds \tag{4.12}$$

We consider the volatility for a single node at the time, and as the name of the model implies, the volatility is assumed locally constant. Therefore we can write (assuming a single period $t \in [0, 1]$)

$$\begin{aligned} \sigma^2 &= \int_0^1 \sigma(s)^2 ds \\ &= \text{Var}[\log(S_1)] \end{aligned} \tag{4.13}$$

From the binomial model, we know that S_1 follows a two-point distribution—it can either go to S_{up} or S_{down}. Combining (4.11) and (4.13) we now find

$$\sigma = \sqrt{p\,(1-p)}\log(S_{\text{up}}/S_{\text{down}})$$

which corresponds to (4.10).

4.2.2 Example calculations

We now move on to a numerical example of the local volatility model. Assuming that the ATM volatility is still 10%, we define a simple volatility "smile" by letting the volatility decrease by 0.5 percentage point for every 10 units the strike increases.[6]

We now have what we need to go ahead and obtain the state value tree and the transition probabilities implied by the volatility smile. We remind ourselves of our setting

- We are looking at a two-period binomial grid (time $t \in \{0, 1, 2\}$, and corresponding nodegroups $n \in \{1, 2, 3\}$—see Figure 4.1 on page 23 for a graphical explanation of the time dimension).

- The underlying asset has a spot price of $S_0 = 100$.

- A risk-free interest rate of 3% per period is assumed for all periods.

- Our volatility smile is defined as having an ATM volatility of 10% and a 0.5 percentage point in-/decrease for every 10 unit de-/increase in the strike

---

[6]Indeed, this describes a linear, and not very "smile-like", smile, but for our purposes this is of no importance.

price (starting at $K = S_0$). Mathematically

$$\sigma_{\text{imp}}(K) = 10\% - 0.5\%(K - S_0)/10 \tag{4.14}$$

To clarify which node we are talking about we remind ourselves of the previously established notation (n, i) meaning "node group n, node i", $i = 1$ indicating the very bottom node. For example, looking at Figure 4.1 on page 23 node $(2, 1)$ has the value 90.48 while node $(3, 3)$ has the value 122.14.

To begin building the implied tree, the first thing we must do is to determine the Arrow-Debreu price for node $(1, 1)$. Since this is the spot node, the Arrow-Debreu price is simply 1. We write $\lambda_{(1,1)} = 1$, where the subscript corresponds to our node notation. Moving on to node $(2, 2)$ we note that we are in the special case with no central node. Hence, we must use (4.8) to find $S_{(2,2)}$, but in order to do so, this requires us to determine $C(100, 1)$ using the volatility implied by the smile. In this case, since the option has strike 100, we know from (4.14) to use $\sigma_{\text{imp}} = 10\%$. In the standard binomial setting using $\sigma = \sigma_{\text{imp}} = 10\%$ we find $C(100, 1) = 6.38$. Since there are no nodes above $(2, 2)$ the Σ-term is equal to 0 and by (4.8) we find

$$
\begin{aligned}
S_{(2,2)} &= \frac{100 \left[e^{3\%} 6.38 + 1 \times 100 - 0\right]}{1 \times 100(1 + 3\%) - e^{3\%} 6.38 + 0} \\
&= 110.52
\end{aligned}
$$

The price in the lower node of nodegroup 2 can now be found using our centering condition (4.7)

$$
\begin{aligned}
S_{(2,1)} &= 100^2/110.52 \\
&= 90.48
\end{aligned}
$$

and by (4.6) we find the transition probability

$$
\begin{aligned}
p_{(1,1)} &= \frac{100 \times 1.03 - 90.48}{110.52 - 90.48} \\
&= 0.625
\end{aligned}
$$

We now have the necessary parameters to determine the grid implied volatility at

(a) State values - $s_{(n,i)}$.

(b) Grid implied volatilities - $\sigma_{(n,i)}$.

(c) Transition probabilities - $p_{(n,i)}$.

(d) Arrow-Debreu prices - $\lambda_{(n,i)}$
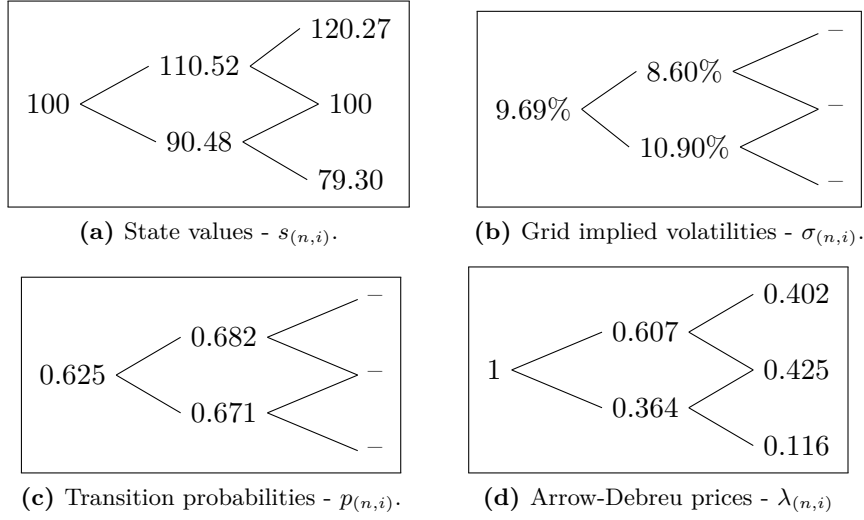
Figure 4.3: *Results from binomial local volatility modeling.*

node $(1, 1)$ using (4.10)

$$\sigma_{(1,1)} = \sqrt{0.625\,(1 - 0.625)}\log(110.52/90.48)$$
$$= 9.69\%$$

$p_{(1,1)}$ together with r_f gives ud the Arrow-Debreu price for node $(2, 2)$, $\lambda_{(2,2)} = 0.625 \times 1.03^{-1} = 0.607$. All we need now to be able to use (4.5) is the value of the call option $C(110.52, 2)$, which we must calculate at the implied volatility dictated by our smile definition

$$\sigma_{imp}(110.52) = 10\% - 0.5\%(110.52 - 100)/10 = 9.474\%$$

With this volatility, using the standard binomial pricing scheme we find a call option price of $C(110.52, 2) = 3.92$. The forward price must be $F_{(2,2)} = 110.52 \times 1.03 = 113.84$, and entering this information into (4.5) (recalling that we chose to center our grid around $S_0 = 100$, so that $S_{(3,2)} = 100$) we find $S_{(3,3)} = 120.27$ with $p_{(2,2)} = 0.682$ yielding a grid implied volatility of $\sigma_{(2,2)} = 8.60\%$. The final task of finding $S_{(3,3)}$ using (4.9), along with $p_{(2,1)}$ and $\sigma_{(2,1)}$ is now trivial, and we will merely state the results. Figure 4.3 presents our findings from the application of the local volatility model.

4.3 Dynamics of the local volatility model

Having seen a numerical example of how the local volatility model can be calibrated to a volatility smile, we will move on to examining the dynamics resulting from this model. In order to do so, we will simplify our initial setup by removing the time dimension from (4.1) on page 22, leaving us with

$$dm_t = \sigma_{\text{loc}}(m_t)\, m_t \, dW_t \tag{4.15}$$

This particular model is the main focus in an article by Hagan and Woodward. Hagan and Woodward (1999) give a closed form approximation for the implied volatility to use with Black's formula as stated in (2.12) on page 13. Letting f denote the forward price, the approximation is

$$\sigma_B(K, f) = \sigma_{\text{loc}}\left(\frac{f+K}{2}\right)\left[1 + \frac{1}{24}\frac{\frac{d^2\, \sigma_{\text{loc}}}{d\, f^2}\left(\frac{f+K}{2}\right)}{\sigma_{\text{loc}}\left(\frac{f+K}{2}\right)}(f-K)^2 + \cdots\right] \tag{4.16}$$

Especially, we note the following regarding this approximation:

1. The local volatility function and its derivatives are evaluated at the average point of the forward price and the strike.

2. Intuitively, the first term of (4.16) ($\times 1$) dominates the second term ($\times \frac{1}{24}$) for options not too far ITM or OTM.

3. The contribution of the omitted terms to the correction is usually less than 1% of that of the first term.[7]

As a result of the second point, we will analyze the dynamics of the implied volatility arising from the use of the local volatility model, by focusing on the first term.

Ignoring everything but the first term, we are left with

$$\sigma_B(K, f) = \sigma_{\text{loc}}\left(\frac{f+K}{2}\right) \tag{4.17}$$

Now, suppose that today we observe two things: a forward price f_0 and a smile of implied volatilities which is dependent on the strike price K, takes f_0 as a given and is denoted by $\sigma_B^0(K; f_0)$. To be in line with the observed implied volatilities

---

[7]Cf. Hagan et al. (2002).

$\sigma_B^0(\cdot)$ and (4.17) we note that it must hold that $\sigma_{\text{loc}}(f) = \sigma_B^0(2f - f_0)$ since

$$\sigma_B^0(2f - f_0) = \sigma_{\text{loc}}\left(\frac{f_0 + [2f - f_0]}{2}\right) = \sigma_{\text{loc}}(f) \tag{4.18}$$

However, if the current forward price f_0 shifts to another forward level f, then (4.17) together with (4.18) imply that

$$\begin{aligned}
\sigma_B(K, f) &= \sigma_{\text{loc}}\left(\frac{f + K}{2}\right) \\
&= \sigma_B^0\left(2\frac{f + K}{2} - f_0\right) \\
&= \sigma_B^0(K + [f - f_0]) \tag{4.19}
\end{aligned}$$

This tells us, that the local volatility model has some peculiar properties. Consider an initial implied volatility smile $\sigma_B^0(K + [f - f_0])$ with the shape of a parabola. Given an initial forward level f_0, assume that an implied volatility of $\sigma_B(K, f_0) = \sigma_B^0(K + 0)$ is observed. Let the forward level increase to $f > f_0$. Now the implied volatility is is given as $\sigma_B(K, f) = \sigma_B^0(K + \underbrace{[f - f_0]}_{>0})$. Intuitively, an increase in the forward price should result in a shift *to the right* of the volatility smile. However, the local volatility model now predicts that the volatility smile shifts *opposite* the direction of the shift in the forward price. Figure 4.4 shows a graphical rendering of the argument.
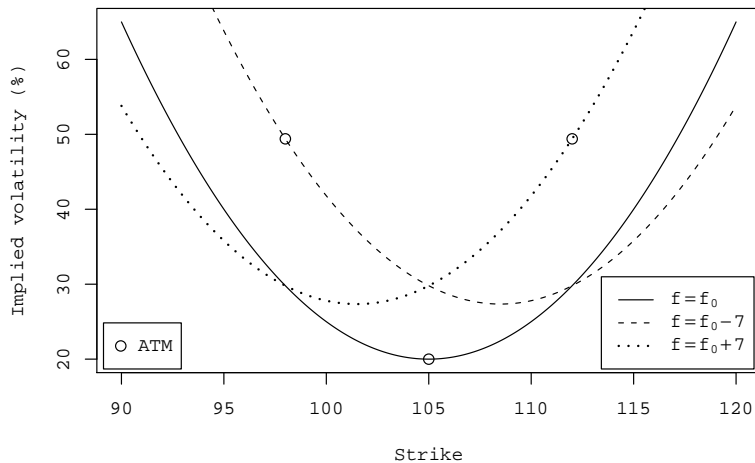

Figure 4.4: *Volatility smile dynamics from local volatility.*

This inconsistency obviously has various implications. The main implication being, that when calculating risk measures under the local volatility model, the resulting hedging decisions might not hedge risk exposure at all.

Let $C(K, f, \sigma_B(K, f))$ denote the value of the call option with strike K, forward price on the underlying asset f and a volatility calculated from the local volatility model $\sigma_B(K, f)$—the interest rate and maturity are arbitrary. Now, calculating the Δ risk for this option requires finding the partial derivative of $C(K, f, \sigma_B(K, f))$ with respect to f. Using standard rules for partial derivatives the Δ risk is

$$\Delta_C = \frac{\partial C(K, f, \sigma_B(K, f))}{\partial f} + \underbrace{\frac{\partial C(K, f, \sigma_B(K, f))}{\partial \sigma_B} \frac{\partial \sigma_B(K, f)}{\partial f}}_{\text{Correction term}}$$

Where the first term is the standard Δ risk from the constant volatility Black-Scholes setup and the correction term is a result of the volatility itself being a function of f. As we saw in Figure 4.4 the change in volatility when the underlying forward price shifts is exactly opposite to market dynamics and therefore our Δ_C measure will be incorrect. For $f > f_0$ the local volatility model tells us that the correction term is negative even though market observations (and our intuition) tell us that it should be positive.

4.4 Wrapping up local volatility

Concluding on our analysis of the local volatility model we note several features—some more attractive than others. The local volatility model assigns different volatilities to different points in the market, and by doing so essentially creates a deterministic volatility function. The volatility function's values can be determined by numerical methods as shown by example in section 4.2. Using quotes on different liquid options, the local volatility model can provide a framework for consistent pricing of exotic and path-dependent derivatives. However, as we have shown in section 4.3, the model is not very well suited as a risk management tool, since the dynamics of the volatility smile implied by the model are opposite to the market observable (and intuitive) dynamics.

Before one completely writes off local volatility as a failed attempt on a model for volatility, due to its short-comings as a risk management tool, it should be mentioned that some sources suggest, that the model might never actually have been thought up for purposes of risk management.

It is unlikely that Dupire, Derman and Kani ever thought of local

*volatility as representing a model of how volatilities actually evolve
(...) the idea is more to make a simplifying assumption that allows
practitioners to price exotic options consistently with the known prices
of vanilla options.*

– Jim Gatheral, (Gatheral, 2002, page 6)

Chapter 5

The SABR model

In our search for a volatility model that will provide both a consistent base for pricing options at different strikes as well as a satisfying quantification of risk we now turn our attention to *the SABR model.*

5.1 Specification of the SABR model

The SABR model differs from the Black-Scholes and the local volatility approach in its assumption about the dynamics of the volatility of the underlying asset. Where in the Black-Scholes setup and the local volatility model, the volatility is assumed to be constant respectively locally constant, the SABR model allows the volatility to evolve as a function of time, t, the strike price, K, and the current forward price, f_t. Furthermore, acknowledging the fact that volatility sometimes behaves relatively calm but also sometimes varies a lot, the SABR model allows for the volatility itself to be random over time. Going back to (2.9) on page 12, the SABR model assumes $c(\cdot) = \alpha_t m_t^\beta$. Where the local volatility model's $c(\cdot)$ function only depended on time and the level of the forward price itself, the SABR model's $c(\cdot)$ includes an extra source of randomness by also depending on random development of α_t. This extra randomness is scaled through the inclusion of a "volatility of volatility" parameter ν. Mathematically, letting f_t denote the forward price process, the SABR model is formulated as

$$df_t = \alpha_t \, f_t^\beta \, dW_t^1 \tag{5.1a}$$

$$d\alpha_t = \nu \, \alpha_t \, dW_t^2 \tag{5.1b}$$

where W_t^1 and W_t^2 are correlated Brownian motions with correlation coefficient ρ so that

$$dW_t^1 dW_t^2 = \rho \, dt \tag{5.1c}$$

In total, the SABR model is described by the stochastic α_t process, the β parameter and the correlation coefficient ρ. This is also where the model takes its name: **S**tochastic **A**lpha **B**eta **R**ho. We note, that as a special case of the SABR model, setting $\beta = 1$ and $\nu = 0$ leaves us with the original Black-Scholes setup, since this particular combination of parameters results in a constant volatility α_0 and a forward process with returns that are normally distributed with a mean of 0 and standard devaiation of $\alpha_0 \sqrt{t}$.

As we will show, the SABR model as stated in (5.1) can be calibrated to fit a given implied volatility smile very well for an asset with a single time to maturity.[1] Swaptions are traded liquidly over different "standardized strikes" (ATM $+ x$ bps, $x \in \{\pm 200, \pm 100, \pm 50, 0\}$). At the same time swaptions are also traded on "standardized maturities and tenors" (6M2Y, 1Y2Y, 5Y2Y etc.). The fact that these products are both highly liquid *and* are trading at a range of strikes for single maturities makes the SABR model an obvious choice for fitting implied volatility smiles for swaptions. Therefore, our analysis of the SABR model will be focused on this specific type of derivatives.

5.2 Solving the SABR model

Before moving on with our analysis of the SABR model, we need a basis for calculations. The model can be solved using Monte Carlo techniques, but this would be quite cumbersome and time consuming for our purposes. Instead we choose to follow in the footsteps of Hagan et al. (2002). In the article from 2002, the following approximation to the implied volatility resulting from the SABR model is presented

$$\sigma_B(K, f) = \frac{\alpha_0 \left\{ 1 + \left[\frac{(1-\beta)^2}{24} \frac{\alpha_0^2}{(f\,K)^{1-\beta}} + \frac{\rho \, \beta \, \nu \, \alpha_0}{4(f\,K)^{(1-\beta)/2}} + \frac{2-3\rho^2}{24} \nu^2 \right] T \right\}}{(f\,K)^{(1-\beta)/2} \left[1 + \frac{(1-\beta)^2}{24} \log^2\left(\frac{f}{K}\right) + \frac{(1-\beta)^4}{1920} \log^4\left(\frac{f}{K}\right) \right]} \times \frac{z}{x(z)} \tag{5.2}$$

$$\text{where } z = \frac{\nu}{\alpha_0}(f\,K)^{(1-\beta)/2} \log\left(\frac{f}{K}\right)$$

$$\text{and } x(z) = \log\left(\frac{\sqrt{1 - 2\rho z + z^2} + z - \rho}{1 - \rho} \right)$$

---

[1]The SABR model can also be extended to cover an entire volatility surface rather than just a smile for a single maturity, but that is beyond the scope of this paper. Se appendix on *The Dynamic SABR Model* (Hagan et al., 2002).

For the ATM case when $K = f$, the formula is drastically simplified as a result of $\log \frac{f}{K}$ evaulating to 0.

$$\sigma_B(f,f) = \frac{\alpha_0}{f^{(1-\beta)}} \left\{ 1 + \left[\frac{(1-\beta)^2}{24} \frac{\alpha_0^2}{f^{(2-2\beta)}} + \frac{\rho\,\beta\,\nu\,\alpha_0}{4\,f^{(1-\beta)}} + \frac{2-3\rho^2}{24} \nu^2 \right] T \right\} \quad (5.3)$$

The approximation is carried out by initially applying singular perturbation techniques[2] to obtain prices of European options, and from these the implied volatilities $\sigma_B(K, f)$ are then inferred.

Having established a closed-form approximation of the implied volatility resulting from the SABR model, we can begin our analysis of the model.

5.3 Elements of the SABR model

In this section we will investigate the various parts of the SABR model. We will take the model apart in order to examine the different parameters's effects on the dynamics of the model, and how changing parameter values can alter the shape of the predicted volatility smile. As a basis for our analysis we will use (5.2) and (5.3).

In the following sections we will use an initial smile setup with a set of more or less randomly chosen parameters ν, α_0, ρ, β and f. Each parameter is then shifted in order to assess its impact on the shape of the smile. As our basic smile we define the smile described by (5.2) using the following parameters:

$$\nu = 0.30 \quad \alpha_0 = 0.03 \quad \rho = -0.30 \quad \beta = 0.60 \quad f = 0.025 \quad (T = 1)$$

f To start out, we take a look at what happens to the volatility smile when the forward price f shifts away from its initial state. The problem we faced with the local volatility model was that the volatility smile shifted opposite the shift of the forward price. Looking at Figure 5.1, we see that the SABR model remedies this issue. Clearly, when the forward price shifts to a higher level, we see that the volatility smile shifts to the right, in line with as well our intuition as with market observable dynamics.

ν The ν parameter is initially set at 0.30 and is subsequently shifted by adding 0.20 causing a $\nu_{\text{new}} =$ of 0.50. Figure 5.2 reveals that ν governs the curvature of the smile. Ceteris paribus, increasing ν causes $\sigma_B(K)$ to increase for OTM

---

[2]Perturbation techniques are beyond the scope of this paper. For details on these we refer to Hagan et al. (2002).
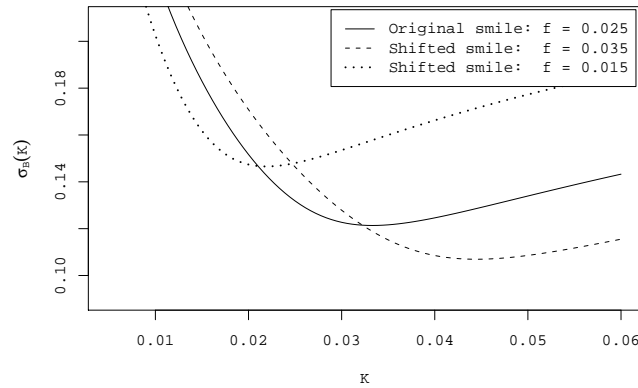
Figure 5.1: *Shifting f.*

strikes and for reasonably far ITM strikes. The volatility smile changes around the ATM point resulting in a smile that is more convex than the initial one.
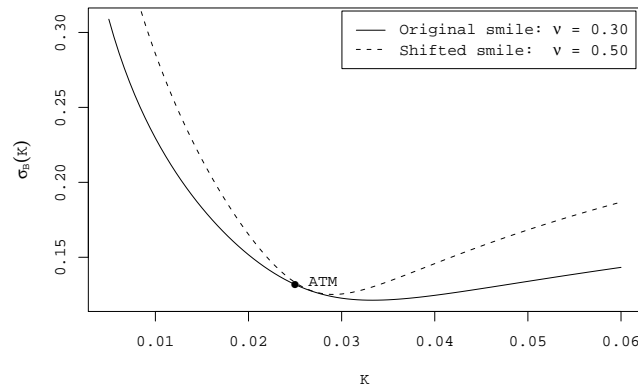
Figure 5.2: *Shifting ν.*

α_0 We choose to shift the α_0 parameter both up and down from its initial state 0.03. We shift it to values of 0.01 and 0.05. The α_0 parameter can be interpreted as "the initial volatility" since it is from this point that the stochastic volatility process takes its origin. With that in mind, the smiles resulting from the shifts are not surprising. Figure 5.3 shows that by shifting α_0 up or down we essentially shift the entire smile up or down. Apparently, α_0 does not influence (notably) on the shape of the smile, but rather on the vertical location.

ρ Initially we assign ρ the value of -0.30. Obviously, since ρ is the correlation between the two Brownian motions governing the development of the SABR model, it is bounded: $\rho \in [-1, 1]$. However, our choice of a negative value for the correlation coefficient is not a mere coincidence. In general, the volatility is lower for higher (ITM) strike prices and vice versa. This inverse relation
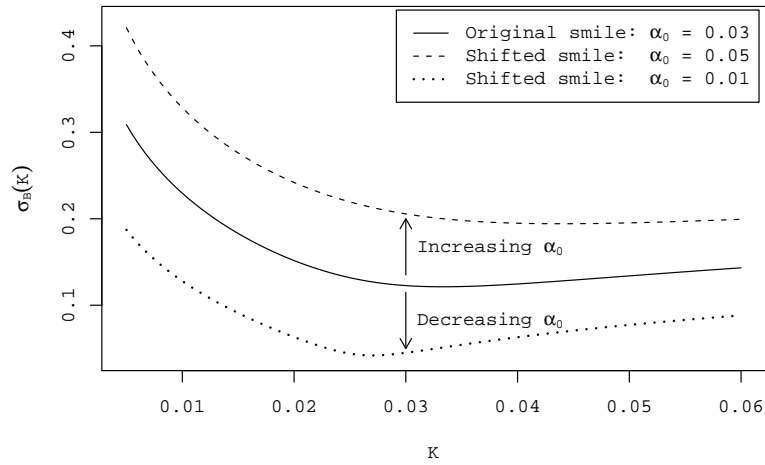
Figure 5.3: *Shifting α_0.*

is imposed on the SABR model by setting $\rho < 0$.

Figure 5.4 shows the effect of ρ on the shape of the volatility smile. Changing the level of correlation seems to cause the smile to "rotate" around the ATM point. As we would expect after having discussed the intuition behind negative correlation, we see that decreasing ρ (making it more negative) to a level of -0.80 causes a steeper smile, while increasing the correlation to $+0.80$ somewhat flattens the smile and gives it an opposite relation to the forward price compared to the -0.80 case.
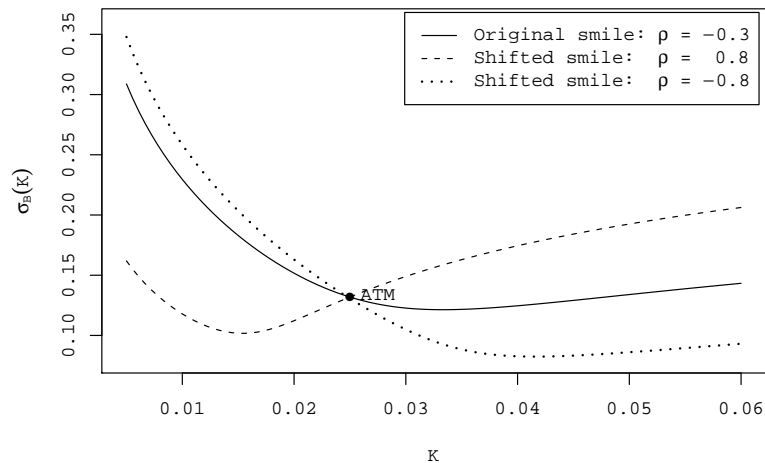
Figure 5.4: *Shifting ρ.*

β In Figure 5.5 we shift β. We start out with a β of 0.60 and then shift this value up and down respectively by 0.20. The β parameter is not directly constrained in the original article by Hagan et al. (2002), but we will, however, impose the constraint $\beta \in [0, 1]$. Clearly, choosing a $\beta < 0$ would

result in a case, where greater forward prices f would imply smaller relative change in the process f_t. This is by no means logical, and therefore we impose the lower bound $\beta \geq 0$. Similarly, choosing a $\beta > 1$ would mean that the expected deviation from the current state of f_t would be greater than the volatility times the current forward level (times $\sqrt{t}$) which is also undesirable. Hence, we impose the upper bound $\beta \leq 1$.

The numerical manipulation of the β shows us, that it—like the ν parameter—influences some on the curvature of the volatility smile. However, the impact of β is somewhat different from that of ν. Where ν seems to curve the entire smile, β apparently has its greater effect on the side of the smile that is left of ATM. The lower we set β, the less curvature in the volatility smile, and at the same time the entire smile is shifted in the same direction as the β—a property similar to that of the α_0 parameter.
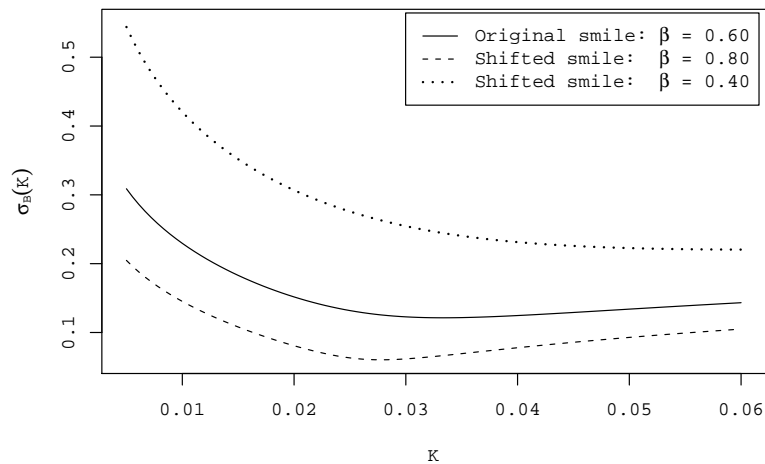
Figure 5.5: *Shifting β.*

In conclusion, we note that several parameters seem to have similar properties or a mix of properties from other parameters. The ρ and the β both seem to govern the curvature of the smile. At the same time, β and α_0 both help shifting the location of the smile up and down. These similarities and connections between the parameters of the SABR model will be further dealt with in the following section where we will address the issues of calibrating the SABR model to real-life data.

5.4 Estimating parameters in the SABR model

In this section we will discuss different ways of calibrating the necessary parameters for a SABR single-maturity volatility smile. Further, we will investigate the uniqueness of the parameters that we obtain from the calibrations. The data that will be used in this section is swaption data[3] as of June 1st 2010 and November 1st 2010. The data is in the form of a smile of Black-Scholes implied volatilities for a single swaption—implied volatilities are given for ATM as well as ATM $+x$ bps, $x \in \{\pm 200, \pm 100, \pm 50\}$.[4]

In general, what we want out of our estimation is to minimize the error between the points we regard as "true" or "observed" implied volatilities, and the points that are fitted by our SABR model for corresponding strikes.

To be more specific, in our case we will be minimizing the sum of squared differences between observed and fitted volatilities. Hence, our general problem is formulated as

$$\min_{\nu,\alpha_0,\rho,\beta} \sum_i \left(\bar{\sigma}_i - \sigma_B(\nu, \alpha_0, \rho, \beta; K_i, f)\right)^2 \tag{5.4a}$$

where $\sigma_B(\nu, \alpha_0, \rho, \beta; K_i, f)$ are the SABR implied volatilities as a function of the SABR parameters and given the strikes K_i and an ATM forward level f, and $\bar{\sigma}_i$ are the market observed implied volatilities. In some cases different object functions might be considered. As an example, one might choose to weight by "traded amount", "trade frequency" or some other liquidity measure. However, given the relative few data points available for our estimation we will not apply any weights, and as we will show, the SABR model is indeed capable of producing very nice fits to our volatility smiles. We impose the following restrictions on the optimization problem in (5.4a)

$$\rho \in [-1; 1] \tag{5.4b}$$

$$\beta \in [0; 1] \tag{5.4c}$$

$$\nu \geq 0 \tag{5.4d}$$

$$\alpha_0 > 0 \tag{5.4e}$$

Eventually, we are faced with the task of estimating the four SABR parameters based on six or seven market quotes. Obviously, this is possible, but the smaller the ratio $\frac{\text{parameters to estimate}}{\text{available market quotes}}$ the better. Assuming that we cannot increase the number of available market quotes, we turn to the numerator.

---

[3]Courtesy of Danske Bank.

[4]Note: on several occasions the ATM forward rate is below 2% bps. In these cases the quote for ATM $- 200$ bps is discarded since this implies a swaption with a negative strike rate.

In section 5.3 we analyzed the various parameters in the SABR model, and we found that several had somewhat similar effects on the shape of the resulting estimated volatility smile. In the following section we will explore the possibility and implications of fixing or pre-estimating one of the SABR parameters prior to the actual calibration to market data. The parameter on which we will focus is the β.

5.4.1 Fixing the β

In this section we will explore the possibilities of pre-determining the β. As one sees from (5.1a) on page 36, the β is a key determinant for the forward process of the underlying asset under the SABR model's assumptions. As such, the choice of β might depend on an a priori belief about the forward process. We will divide these beliefs into three common models: $\beta \in \{0, \frac{1}{2}, 1\}$

$\beta = 0$ — the stochastic normal model

Setting $\beta = 0$ results in a forward process that looks like

$$df_t = \alpha_t \, dW_t$$

The process describes a forward price whose increments are stochastic normally distributed. "Stochastic" in the sense that they are normally distributed with a mean of 0 and a stochastic standard deviation that is lognormally distributed. This essentially will enable the forward process to be negative, and for most practical purposes this is probably not a desirable feature.

$\beta = \frac{1}{2}$ — the stochastic CIR model

This "stochastic CIR model" takes its name from the short term interest rate process suggested by Cox, Ingersoll and Ross (Cox et al., 1985), in which the volatility term is multiplied by the square root of the current level of the process. In our case, this occurs as the forward process takes the form

$$df_t = \alpha_t \, f_t^{\frac{1}{2}} \, dW_t = \alpha_t \sqrt{f_t} \, dW_t$$

The "$\sqrt{f_t}$"-term forces the process to stay above 0, and thereby preventing negative forward rates.

43

$\beta = 1$ — the stochastic lognormal model

Setting $\beta = 1$ takes us to the lognormal case. For this value, the forward proces is

$$df_t = \alpha_t \, f_t \, dW_t$$

which is almost similar to the standard Black-Scholes setup where the value of the underlying asset follows a geometric Brownian motion. The difference between Black-Scholes and the stochastic lognormal model is, that in the latter case, the volatility is a stochastic process itself, where Black and Scholes assumed the volatility to be constant. The SABR model with a β of 1 implies that forward rates are lognormally distributed and thus, the non-negativity property from the CIR case also applies to the lognormal model.

We have now presented three different choices of β that all yield some properties similar to those of pre-existing modeling frameworks. However, the SABR model does not *require* the user to predetermine a β based on a (biased) belief regarding how the process should evolve according to some textbook example. In the following section we will loosen the grip on the β and look a bit further into its role in the SABR framework.

5.4.2 Fitting the β

As we have already seen, the β parameter in the SABR model can be pre-specified to match a certain class of process. In this section, we will look into the possibility of inferring β from our data rather than fixing it. To establish methods for estimating β in the SABR model setup, we will turn our attention to the forward process.

The forward process in the SABR framework is formulated as we have seen in (5.1a) on page 36. This implies that the relative change in the forward price ($\frac{df_t}{f_t}$) has a standard deviation of

$$\begin{aligned} \text{StDev}\left[\frac{df_t}{f_t}\right] &= \alpha_t \frac{f_t^{\beta}}{f_t} \\ &= \alpha_t \, f_t^{\beta-1} \end{aligned} \tag{5.5}$$

Restraining β to the interval $[0, 1[$ we see that the instantaneous standard deviation has a direct inverse relation to the forward price. For the case of stock prices Beckers (1980) presents an economic reasoning behind this phenomenon. He argues, that as the stock price of a firm falls, the market value of its equity tends
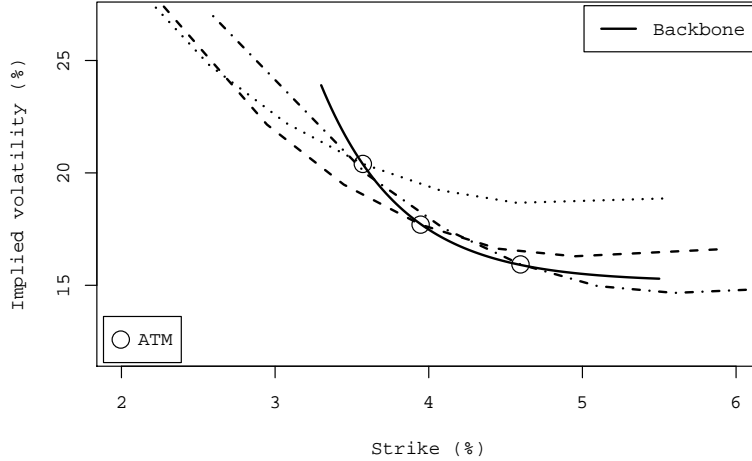
Figure 5.6: *Example of a backbone for 10Y10Y swaptions against EURIBOR6M for different ATM forward levels.*

to fall more rapidly than that of its debt (or at least its fixed costs, in case of no debt). This will cause the company to become increasingly risky, and thereby increase the volatility. Though this argument is not directly transferable to forward interest rates it is still a nice bit of intuition.

Seeing how the standard deviation in (5.5) depends negatively on the current level of f_t, what does this imply? To move into the implications of this inverse relationship, we introduce the concept of *the backbone.*

The backbone is the curve traced by the ATM implied volatilities over various ATM levels of the forward process. An example of a backbone is shown in Figure 5.6. Rewriting the expression for the instantaneous standard deviation of the forward process from (5.5) in log terms we obtain

$$\log\left(\text{StDev}\left[\frac{df_t}{f_t}\right]\right) = \log(\alpha_t) + (\beta - 1)\log(f_t)$$

$$\Leftrightarrow \quad \log(\sigma_{ATM}) = \log(\alpha_t) + (\beta - 1)\log(f_t) \tag{5.6}$$

Also, as we have already shown in (5.3) on page 38 the SABR ATM volatility can be approximated by the closed form expression

$$\sigma_B(f, f) = \frac{\alpha_0}{f^{(1-\beta)}} \left\{ 1 + \left[\frac{(1-\beta)^2}{24} \frac{\alpha_0^2}{f^{(2-2\beta)}} + \frac{\rho \beta \nu \alpha_0}{4 f^{(1-\beta)}} + \frac{2 - 3\rho^2}{24} \nu^2 \right] T \right\}$$

This expression is fairly complicated in its present form. However, the main con-

tribution to σ_B comes from the first term within the curly brackets. According to Hagan et al. (2002), omitting the last term will result in a relative error that in extreme cases will exceed three percent. Hence, for pricing purposes the full expression should be used, but for analytical purposes, we will keep only the first term. By doing so, we are left with

$$\sigma_B(f, f) \approx \frac{\alpha_0}{f^{(1-\beta)}}$$

Rearranging using log terms, we are left with an expression similar to (5.6). We have now shown in two ways, that we can estimate the SABR β parameter by performing a linear regression on sets of $(\log(\sigma_{ATM}), \log(f_{ATM}))$ for different ATM levels. The β is estimated by the slope of the line $+ 1$.

5.4.3   Parameterization in the SABR model

In the previous sections we have discussed different possibilities for pre-specifying β based on an a priori belief regarding the dynamics of the process of f_t, or estimating the β parameter using a linear regression technique. In the following we will assume that β has already been estimated. Hence, we are left with the task of estimating α_0, ρ and ν in order to have a fully calibrated SABR model. We will present two different ways of estimating these three remaining parameters.

Estimating α_0, ρ and ν

The first parametrization scheme we will present is very intuitive and simple. Reminding ourselves of our general optimization problem from (5.4a)

$$\min_{\nu,\alpha_0,\rho} \sum_i \left(\bar{\sigma}_i - \sigma_B(\nu, \alpha_0, \rho; K_i, f, \beta)\right)^2$$

an obvious solution is to simply find the parametrization $\{\nu, \alpha_0, \rho\}$ that minimizes the sum of the square errors. This can be done using standard optimization techniques such as the Newton-Raphson method.

Estimating ρ and ν

For the second parametrization, we will estimate ρ and ν, and from those we will infer an α_0 using the approximation for ATM volatility given by (5.3) on page 38 along with the known ATM implied volatility σ_{ATM}. Knowing σ_{ATM} we are able

to invert (5.3) to solve for α_0. We find

$$\sigma_{ATM} = \frac{\alpha_0}{f^{(1-\beta)}}\left\{1 + \left[\frac{(1-\beta)^2}{24}\frac{\alpha_0^2}{f^{(2-2\beta)}} + \frac{\rho\,\beta\,\nu\,\alpha_0}{4\,f^{(1-\beta)}} + \frac{2-3\rho^2}{24}\nu^2\right]T\right\}$$

$$\Leftrightarrow 0 = A\,\alpha_0^3 + B\,\alpha_0^2 + C\,\alpha_0 - \sigma_{ATM}f^{(1-\beta)} \tag{5.7}$$

where $A = \left[\frac{(1-\beta)^2 T}{24\,f^{(2-2\beta)}}\right]$, $B = \left[\frac{\rho\beta\nu T}{4\,f^{(1-\beta)}}\right]$ and $C = \left[1 + \frac{2-3\rho^2}{24}\nu^2 T\right]$. This cubic can have up to three real roots. However, typically there will only be one real root.[5] Should there be more than one real root, one should seek the smallest positive root—in cases with three real roots these will be of magnitude -1000, 1000 and 1, and we would choose the latter.[6]

The steps described above, can be comprised into the following optimization algorithm:

   i. Assign initial values to ρ and ν.

   ii. Solve (5.7) to obtain an estimate of α_0.

   iii. Minimize the sum of squared errors with regards to ρ and ν.

   iv. Repeat [ii.]

   v. If the sum of squared errors exceeds a prespecified tolerance level go back to [iii.].

Stated mathematically, the problem is now to minimize the following expression

$$\min_{\nu,\rho}\sum_i\left(\bar\sigma_i - \sigma_B(\nu,\rho,\alpha_0(\rho,\nu;\sigma_{ATM});K_i,f,\beta)\right)^2$$

This means that we now implicitly determine α_0 using σ_{ATM} as an "extra" input compared to the case where we directly estimate α_0, ρ and ν.

5.5 Fitting a smile

In this section we will apply the various methods of fitting the SABR model to an actual volatility smile. For the SABR modeling we will be using the volatility smiles graphed in Figure 5.6 on page 45. The data presented is a 10Y10Y swaption against EURIBOR6M, and can be seen in detail in Table 5.1. The table shows seven forward levels (ATM and ATM $\pm\,200, 100, 50$ bps) along with their corresponding Black-Scholes implied volatilities.

---

[5]Cf. West (2005).

[6]Ibid.

	December 1st 2010						
σ_{BS}	32.15	24.80	22.22	**20.40**	19.23	18.67	18.87
f	1.571	2.571	3.071	**3.571**	4.071	4.571	5.571

	September 1st 2010						
σ_{BS}	26.97	20.16	17.63	**16.11**	14.80	14.66	14.99
f	2.597	3.597	4.097	**4.597**	5.097	5.597	6.597

	June 1st 2010						
σ_{BS}	29.97	22.15	19.48	**17.70**	16.63	16.29	16.62
f	1.947	2.947	3.447	**3.947**	4.447	4.947	5.947

Table 5.1: *10Y10Y EURIBOR6M swaption data. All numbers are in percentages. ATM forwards and volatilities are typeset in **bold**.*

5.5.1 Fitting a smile with a prespecified β

As previously mentioned, the first task when calibrating a SABR model is that of choosing a suitable value for the β parameter. In our first attempt of a calibration, we will proceed as described in section 5.4.1 and simply *choose* a β value. Since we are dealing with interest rates we somewhat arbitrarily choose to fix $\beta = \frac{1}{2}$. Hence, our SABR model will be of the stochastic CIR type. Since this type of modeling does not require any fitting of the β, we will only need to focus on one of the three smiles presented. We choose to focus on the the newest available data: the volatility smile per December 1st 2010.

Estimating α_0, ρ and ν

Initially, we will calibrate a SABR model to the volatility smile, by predetermining a $\beta = \frac{1}{2}$ and simply estimate the remaining parameters α_0, ρ and ν by solving

$$\min_{\nu,\alpha_0,\rho} \sum_i \left(\bar{\sigma}_i - \sigma_B(\nu,\alpha_0,\rho;K_i,f,\beta) \right)^2$$

Setting $\beta = \frac{1}{2}$ and estimating the remaining parameters result in a parametrization $\{\alpha_0,\rho,\nu\} = \{3.574\%, -24.862\%, 35.950\%\}$ with a square error sum of approximately 1.22E−5. Figure 5.7 shows the Black-Scholes volatilities implied by the market prices along with those inferred by our SABR model (for $\beta = \frac{1}{2}$). We see that the SABR model is indeed capable of producing a highly convincing fit to market data. The observed volatilities and the SABR volatilities coincide
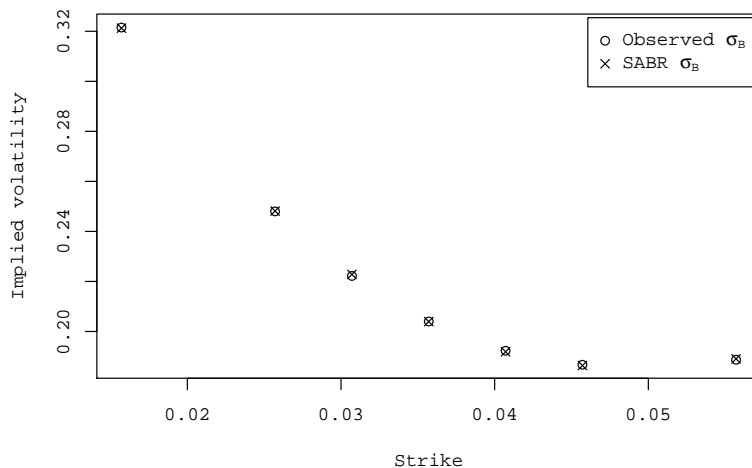
Figure 5.7: *SABR model fitted to 10Y10Y EURIBOR6M swaption data as per December 1st 2010. The β is prespecified at $\frac{1}{2}$ and the remaining parameters are estimated: $\{\alpha_0, \rho, \nu\} = \{3.574\%, -24.862\%, 35.950\%\}$*

β	α_0	ρ	ν	Error
$\frac{1}{2}$	3.574%	-24.862%	35.950%	1.22E-5
1	20.226%	-47.301%	46.442%	8.45E-7

Table 5.2: *SABR parameters fitted to 10Y10Y EURIBOR6M swaption data as per December 1st 2010. Depending on the set value of β, the parameter estimates vary, but both models fit market data very well.*

almost perfectly. However, we have only shown this for our choice of $\beta = \frac{1}{2}$. We will now set $\beta = 1$ and recalibrate the model to see the effects of such a change to the model—and especially the implicit change to the a priori belief on the process of the forward rates.[7] Setting $\beta = 1$ and applying our numerical minimization of the squared error terms leaves us with a set of parameters $\{\alpha_0, \rho, \nu\} = \{20.226\%, -47.301\%, 46.442\%\}$ with a square error sum of approximately 8.45E-7. We note that with $\beta = 1$ we are left with an even smaller error compared to the case of $\beta = \frac{1}{2}$ which we have seen in Figure 5.7 already fits market data very accurate. The two sets of parameter estimates are shown in Table 5.2.

---

[7]We remind ourselves, that the special cases of $\beta = 0$, $\beta = \frac{1}{2}$ and $\beta = 1$ imply a Vasicek, Cox-Ingersoll-Ross and Black-Scholes model respectively.

β	α_0	ρ	ν	Error
$\frac{1}{2}$	3.564%	-24.696%	36.142%	$1.25\mathrm{E}{-5}$
1	20.239%	-47.343%	46.416%	$8.56\mathrm{E}{-7}$

Table 5.3: *SABR parameters fitted to 10Y10Y EURIBOR6M swaption data as per December 1st 2010. The β's are predetermined while α_0 is determined based on σ_{ATM}, ρ and ν as the $\alpha_0^* \in [0;1]$ that solves (5.8).*

Estimating ρ and ν

As shown in section 5.4.3, the α_0 parameter does not necessarily need to be estimated through the numerical minimization of the error terms. Instead, α_0 can be expressed as a root of a cubic involving ρ, ν and the ATM implied volatility σ_{ATM}

$$0 = A\,\alpha_0^3 + B\,\alpha_0^2 + C\,\alpha_0 - \sigma_{ATM} f^{(1-\beta)} \qquad (5.8)$$

where $A = \left[\frac{(1-\beta)^2 T}{24\, f^{(2-2\beta)}} \right]$, $B = \left[\frac{\rho\beta\nu T}{4\, f^{(1-\beta)}} \right]$ and $C = \left[1 + \frac{2-3\rho^2}{24}\nu^2 T \right]$.

Using this parametrization of α_0, we repeat our calibration of the SABR model with a prespecified $\beta = \frac{1}{2}$ and with a $\beta = 1$. The results are shown in Table 5.3. We note, that the parameter estimates change very little compared to our initial fitting of the SABR model, where none of the three parameters (α_0, ρ and ν) were bounded by the values of the others. Further, we see that the error term does not change much either. The error term does increase a little bit, but that was to be expected since we are now expressing one of the three variables through the other two, thus removing one degree of freedom.

Based on the results in Table 5.3 we conclude that parameterizing α_0 by ρ, ν and σ_{ATM} does not diminish our models ability to fit market data. Rather, based on the belief that obtaining α_0 through other parameter estimates in accordance with the model setup will lead to a more self-consistent model, we tend to prefer this parametrization approach. Our preference is further supported by the fact that when comparing the parameter estimates in Table 5.2 and Table 5.3 there seems to be little deviance.

5.5.2 Fitting a smile—estimating the β

We saw in section 5.5.1 that our choice of β has little effect on the ability of the SABR model to fit to market data. Whether we specified a β of $\frac{1}{2}$ or a β of 1, the remaining error was very small. However, in this section we will look into fitting the β parameter based on observed movements of the volatility smiles. To be more

specific, we will estimate β based on the backbone of the volatility curves.[8]

In section 5.4.1 we showed, that expressed in log terms, the ATM volatility is a linear function of the ATM forward level

$$\log(\sigma_{ATM}) = \log(\alpha_t) + (\beta - 1)\log(f_t)$$

For our estimation of the the β we will use the data presented in Table 5.1 on page 48. Taking logs of the three ATM levels and running a standard linear regression yields the following ANOVA table

	Estimate	Std. Error	t value	Pr($>$|t|)
(Intercept)	4.1530	0.3041	13.66	0.0465
$\log f_t$	-0.9087	0.2181	-4.17	0.1500

The first thing we note is that the slope estimate of approximately -0.91 shows a low degree of significance with a *p*-value of 0.15. However, we only have three observations and therefore—despite its insignificance—we choose to move on with this estimate.

A slope estimate of -0.91 tells us, that the appropriate β estimate is $-0.91 + 1 = 0.09$. Before we move on with this number, we note that this almost puts in the class of normally distributed forward rates, implying that the forward rate has a positive probability of being negative. Obviously, this is not an attractive feature, since this would mean that one could make a profit borrowing money.[9]

Estimating ρ and ν

Having obtained our β estimate from the linear regression technique, we will now fit the SABR model using this β value. In section 5.5.1 two ways of estimating the SABR parameters were presented: the free approach and the parametrization approach where α_0 is a function of other parameters. We took a preference to the parametrization approach, and therefore, we will move on using this method for estimating the SABR parameters.

Fitting the SABR model using a β of 0.09 yields the parameter estimates displayed in Table 5.4. Looking at the error term it quickly becomes apparent, that the β estimate of 0.09, calculated through the linear regression, actually leaves us with the worst fitting model so far. The models previously calibrated for βs

---

[8]See section 5.4.2 for an explanation of *the backbone* of a volatility smile.

[9]More accurately, since we are dealing with *forward rates* this implies that one can lock in a *future* profit by agreeing to borrow money in the future. Either way, a negative rate—instantaneous or forward—assigns negative value to holding money which is very rarely the case.

β	α_0	ρ	ν	Error
0.09	0.901%	3.916%	30.801%	5.67E−5

Table 5.4: *SABR parameters fitted to 10Y10Y EURIBOR6M swaption data as per December 1st 2010. The $\beta = 0.09$ is calculated using linear regression.*
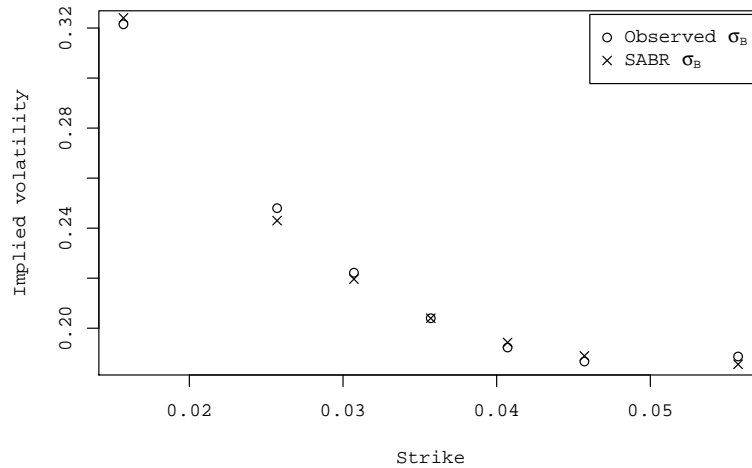

Figure 5.8: *SABR model fitted to 10Y10Y EURIBOR6M swaption data as per December 1st 2010. The $\beta = 0.09$ is calculated using the linear regression technique. The error term is $5.67E-5$.*

(semi-)arbitrarily chosen at $\frac{1}{2}$ and 1 all have much smaller error terms. For $\beta = \frac{1}{2}$ the error is 1.25E−5 and for $\beta = 1$ it is as small as 8.56E−7. Hence, we do not gain anything estimating β using the linear regression technique—*au contraire*. However, even though we see that the linear regression technique performs the worst so far, we still note the SABR model's ability to fit a market smile regardless of the choice of β. Further, we also recognize that our fits are all based on a single point in time. Indeed the model with a β of 0.09 might be the better model over time as the volatility smiles shift as a result of shifts to ATM forward price. However, the analysis of the time-dynamic SABR model is beyond the scope of this thesis.

Figure 5.8 shows market implied volatilities along with the SABR implied volatilities from the model with $\beta = 0.09$. As we see, the model still fits very well, despite the tendency we have seen so far, that the error term is smaller for β's near 1 and grows larger as β approaches to 0. It seems as if a conscious choice of a β based on an a priori belief regarding the dynamics of the underlying asset is plenty. Also, in our case it should be kept in mind that we only had three data

points available for the estimation of β, which is hardly a sufficient number of data points for convincing results or stable conclusions.

5.6 Risk measures under the SABR assumptions

In this section we will analyse the concepts of delta (Δ) and vega ($\mathcal{V}$) risk when hedging in a SABR framework. Initially, we will go over these risk measures in the standard Black-Scholes setting with a contant volatility σ_B. Once the basics have been presented in the Black-Scholes setup we will extend the concepts to cover the SABR model in which volatility is no longer constant but instead is dependent on a range of parameters $\sigma_B(\cdot)$.

5.6.1 Δ and $\mathcal{V}$ in the Black-Scholes setting

We remind ourselves that the time 0 price of a European plain vanilla call option on an underlying asset with a price process S_t under the Black-Scholes assumptions is

$$C(S_0, K, r_f, \sigma_B, T) = S_0 \, \Phi(d_1) - e^{-r_f \, T} K \, \Phi(d_2) \qquad (5.9)$$

where $d_1 = \frac{\log(S_0/K) + (r_f - \frac{1}{2}\sigma_B^2)T}{\sigma_B\sqrt{T}}$ and $d_2 = d_1 - \sigma_B\sqrt{T}$ and $\Phi(\cdot)$ is the cumulative standard normal distribution function.

The value of the option is obviously dependent on a lot of parameters. All of which are subject to change over time. With the exception of the time to maturity T, these changes are stochastic and as such they represent risks to the value of the option. These are the risks option traders try to manage. The two risk measures on which we will focus are:

Δ (Delta) the the risk towards a change in the value of the underlying asset f_t.

$\mathcal{V}$ (Vega) the risk towards a change in the volatility of the returns of the underlying asset.

As mentiones, the Δ risk is the sensitivity of the option price towards changes in the underlying asset

$$\Delta \equiv \frac{\partial C(\cdot)}{\partial S} \qquad (5.10)$$

For the Black-Scholes case presented in (5.9) the Δ can be expressed explicitly. There are more ways to derive the Δ. One way is to apply the chain rule of differentiation. This is straightforward, but it requires some tedious calculations. Instead, we will derive the Δ using a neat little trick.

First, we note that $C(\cdot)$ is homogenous of degree 1 in spot price of the underlying asset and strike price. Omitting other terms than S_0 and K we write

$$C(\lambda(S_0, K)) = \lambda^1 C(S_0, K)$$

Secondly, we use Euler's Theorem.

Theorem 5.1. (EULER'S HOMOGENOUS FUNCTION THEOREM)[10]
Let h be a continuous differentiable function of n variables with continuous partial derivatives on an open space $\mathcal{D}$ such that $\mathbf{x} = (x_1, x_2, \ldots, x_n)^t \in \mathcal{D}$ and $t > 0$ means that $t\,\mathbf{x} = (t\,x_1, t\,x_2, \ldots, t\,x_n)^t \in \mathcal{D}$. Now h is homogenous of degree λ if and only if, for all $\mathbf{x} \in \mathcal{D}$ it holds that:

$$\mathbf{x}^t \nabla h = \sum_{i=1}^{n} x_i h_i'(\mathbf{x}) = \lambda\, h(\mathbf{x}) \tag{5.11}$$

Since we have seen that $C(\cdot)$ is indeed homogenous of degree 1 in (S_0, K) we now find from Theorem 5.1 that $C(\cdot)$ must be of the form corresponding to the left-hand side of (5.11). This means that $C(\cdot)$ is written as a sum of S_0 and K weighted by their partial derivatives. Hence, the multiplication term on S_0 in (5.9) must be equal to $\frac{\partial C(\cdot)}{\partial S} \equiv \Delta$ and we find

$$\Delta = \Phi(d_1)$$

Δ tells us how much the value of the option will change given a 1 unit change in the underlying asset. We note that $\Delta \in [0;1]$.[11] A Δ near 0 is seen for far OTM call options, since these are practically worthless, and an increase in the underlying asset has close to no effect on the value of the option. Similarly, a far ITM call option will have a Δ near 1 since the unit increase in the value of the underlying asset will almost surely result in an extra unit payoff from the option.

The option Δ can also be used to create a portfolio that is neutral to changes in the value of the underlying asset. For example, say we buy 100 identical call options on Stock A each costing 1 EUR and each with a Δ_A of 0.40. This leaves our portfolio with a total Δ_{PF} of $100 \times 0.40 = 40$, implying that if the price of Stock A increases by 1 EUR our position will be worth 140 EUR, but if the price drops by 1 EUR it will be worth only 60 EUR. This risk can be hedged by a position in the underlying asset with an offsetting Δ. The underlying asset itself

---

[10]Cf. Sydsæter (2005).

[11]Opposite sign of the S_0 term in the Black-Scholes formula for plain vanilla European put options means that for the put option we find $\Delta_{put} = -\Phi(-d_1) = \Phi(d_1) - 1$ and therefore $\Delta_{put} \in [-1;0]$.

has a Δ of 1, and therefore a simple hedge strategy is to short (sell) $\Delta_{\mathrm{PF}} = 40$ units of the underlying asset leaving us with a net $\Delta = 0$.

Turning to the $\mathcal{V}$ risk of the option this is the option's sensitivity towards the uncertainty in the market—the volatility. Similar to the way we defined Δ in (5.10), we define $\mathcal{V}$

$$\mathcal{V} \equiv \frac{\partial C(\cdot)}{\partial \sigma_B} \tag{5.12}$$

There is no neat trick to come up with the partial derivative of $C(\cdot)$ with regards to σ_B, and instead of going through tedious calculations we will merely present the formula

$$\mathcal{V} = S_0 \sqrt{T} \, \Phi'(d_1)$$

$\mathcal{V}$ is the same for the put and the call, and we note that $\mathcal{V} \in [0; \infty[$. This implies that the call and the put option become more expensive when volatility rises, and is a result of the increasing uncertainty faced by the writers of the options.

5.6.2 Δ and $\mathcal{V}$ in the SABR model

We have now seen how Δ and $\mathcal{V}$ behave in the Black-Scholes model. As mentioned, Black and Scholes assumed the volatility to be constant, and as discussed in Chapter 3—and as seen in the various graphs depicting volatility smiles—this is hardly the case. This is the reason why we introduced the SABR model, which allows for a dynamic volatility $\sigma_B(\cdot)$ that has an approximate closed-form solution, and that can be used together with the well-known Black-Scholes formula.

In the SABR model, the σ_B is not constant, but instead it is assumed to be a function of several parameters. Among others it is a function of the current forward level f_t (roughly corresponding to S_0 in (5.9)) and α_t (which is the volatility term within the SABR model). Given the complexity of the closed form solution, a common way of calculating partial derivatives is simply to change the value of the differentiation variable a tiny bit and see what relative change this causes

$$\frac{\partial y(x, z)}{\partial x} \approx \frac{y(x + \epsilon, z) - y(x - \epsilon, z)}{2\,\epsilon} \tag{5.13}$$

We note that if we replace x with $x + \epsilon$ on the right hand side and take the limit for $\epsilon \to 0$ we have the definition of a partial derivative (Sydsæter, 2005).

The method described in (5.13) is what we will be using to calculate Δ and $\mathcal{V}$ risks in the SABR model, but first we will reconsider the definitions of these given in (5.10) and (5.12) respectively.

Starting out with the Δ risk, this is defined as the change in the option price caused by a unit change in the value of the underlying asset. In the Black-Scholes setup this is trivial, and can be done directly by applying the method from (5.13). However, reminding ourselves that in the SABR model, the volatility term $\sigma_B(\cdot)$ itself is a function of the underlying asset (the forward rate). Omitting other terms, we write the price of the European call option as $C(f_t, \sigma_B(f_t))$. In order to calculate the change in $C(\cdot)$ when the forward rate changes we need to apply *the chain rule of differentiation* stated in Theorem 5.2.

Theorem 5.2. (THE CHAIN RULE OF DIFFERENTIATION)[12]
If g is differentiable in x_0 and if h is differentiable in $u_0 = g(x_0)$ then $h(x) = h\left(g(x)\right)$ is differentiable in x_0 and

$$\left.\frac{dh}{dx}\right|_{x=x_0} = \left.\frac{dh}{du}\right|_{u_0=g(x_0)} \times \left.\frac{dg}{dx}\right|_{x=x_0}$$

Applying the chain rule to $C(f_t, \sigma_B(f_t))$ we find the SABR Δ as[13]

$$\frac{\partial C(f_t, \sigma_B(f_t))}{\partial f_t} = \frac{\partial C(f_t, \sigma_B(f_t))}{\partial f_t} + \frac{\partial C(f_t, \sigma_B(f_t))}{\partial \sigma_B(f_t)} \times \frac{\partial \sigma_B(f_t)}{\partial f_t} \tag{5.14}$$

We see in (5.14), that the SABR Δ has two terms. The first term corresponds to the regular Black-Scholes Δ and the second term is the correction for movements in $\sigma_B(\cdot)$ caused by the shift in f_t. With this representation of the SABR Δ risk, our discretization scheme for numerically approximating Δ becomes

$$\Delta_{\text{SABR}} \approx \frac{C(f_t + \epsilon, \sigma_B(f_t + \epsilon)) - C(f_t, \sigma_B(f_t))}{\epsilon} \tag{5.15}$$

Moving on to $\mathcal{V}$, our point of interest is now the effects on the call option price when the volatility of the underlying asset changes. We note that the volatility in question when it comes to estimating the SABR $\mathcal{V}$ is the α and not (directly) the $\sigma_B(\cdot)$. We remind ourselves, that in the SABR model, the price process evolves as $df_t = \alpha_t f_t^\beta\, dW_t^1$ where α_t is the (stochastic) volatility process. With this in mind and using Theorem 5.2 we are able to express the SABR $\mathcal{V}$ as

$$\frac{\partial C(\sigma_B(\alpha_t))}{\partial \alpha_t} = \frac{\partial C(\sigma_B(\alpha_t))}{\partial \sigma_B(\alpha_t)} \times \frac{\partial \sigma_B(\alpha_t)}{\partial \alpha_t} \tag{5.16}$$

---

[12]Cf. Sydsæter (2005).

[13]Here, we implicitly assume that the SABR model is parametrized in its most free form form $\sigma_B(\nu, \alpha_0, \rho; K_i, f, \beta)$. Had we chosen to parametrize α_0 through ρ and ν this should also had been factored in since f_t appears in the cubic in (5.8) solved to obtain α_0, and therefore a term $\frac{\partial C(\cdot)}{\partial \sigma_B(\cdot)} \times \frac{\partial \sigma_B(\cdot)}{\partial \alpha_0(\cdot)} \times \frac{\partial \alpha_0(\cdot)}{\partial f_t}$ should have been added to (5.14).
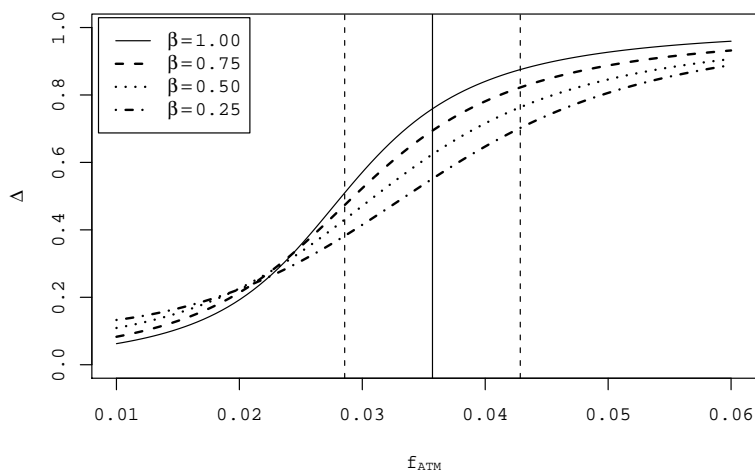
Figure 5.9: Δ *calculated for different values of β. The solid vertical line indicates the ATM forward level, and the dashed vertical lines indicate ATM $\pm$ 20%.*

β	α	ρ	ν
1.00	20.223%	-47.301%	46.442%
0.75	8.415%	-36.656%	40.719%
0.50	3.574%	-24.862%	35.950%
0.25	1.543%	-8.500%	32.063%

Table 5.5: *SABR parameters for different β values.*

Having now established the basics of the Δ and $\mathcal{V}$ risk measures in the SABR model, let us examine a plot of the former. Figure 5.9 shows Δ for the 10Y10Y EURIBOR6M swaption we have been working with so far. The swaption volatility skew has been fitted[14] for different levels of the β parameter. For parameter estimates see Table 5.5. After fitting the models the corresponding Δ figures have been estimated using the method described in (5.15). We see that the SABR setup yields significantly different Δ risk measures depending on which β is chosen. The Δ's obvious dependence on the value of β in conjunction with the fact, that we do not have any means of determining a "true" β imposes a problem: which Δ measure should we choose for hedging purposes?

---

[14]The models have been fitted freely, meaning that α_0 is *not* parametrized through ρ and ν.

5.6.3 Modifying the SABR risk measures

In this section we will build on our findings from section 5.6.2. We found, that using our derivations of the risk measures left us with risk measures dependent on the choice of β. Even though the SABR smiles resulting from various β estimates fit equally well, we see that the Δ measure varies substantially when changing the value of β. In the original article,[15] this issue is not adressed. However, Bartlett (2006) proposes a correction to the risk terms which we shall look into in this section. The main difference between the articles of Hagan et al. (2002) and Bartlett (2006) is that while the former specify the SABR model, the latter actually takes into account the proper dynamics of the model.

In the SABR model, the parameter ρ is used to indicate the level of correlation between the forward price process f_t and the volatility process α_t. This correlation is not accounted for in the Δ and $\mathcal{V}$ derivations shown earlier. Roughly speaking, we calculated our risk measures by shifting only one parameter. This means that for the Δ case the shift to f and α are

$$\Delta_{\text{Hagan}} : \quad \begin{aligned} f &\to f + df \\ \alpha &\to \alpha \end{aligned} \tag{5.17}$$

And for the the $\mathcal{V}$ case

$$\mathcal{V}_{\text{Hagan}} : \quad \begin{aligned} f &\to f \\ \alpha &\to \alpha + d\alpha \end{aligned} \tag{5.18}$$

Obviously, having specified a model with a correlated price and volatility process, these diagrams are erroneous. Whenever f changes, then—in the long run—α will change according to the correlation and vice versa.

Following Bartlett's approach we rewrite the SABR model as a model containing two *independent* Brownian motions[16]

$$df_t = \alpha_t f_t^\beta \, dW_t$$
$$d\alpha_t = \nu \alpha_t \left(\rho \, dW_t + \sqrt{1 - \rho^2} \, dZ_t \right)$$

The first expression can be rearranged to yield $dW_t = df_t \, (\alpha_t f_t^\beta)^{-1}$. Inserting this

---

[15](Hagan et al., 2002).

[16]We utilize, that if W_t^1 and Z_t are independent Browninan motions, then—applying the Cholesky decomposition—W_t^1 and $W_t^2 = \rho W_t^1 + \sqrt{1 - \rho^2} Z_t$ are correlated with correlation coefficient ρ (McDonald, 2006, page 657).

into the second expression and rearranging leaves us with

$$d\alpha_t = \frac{\rho\,\nu}{f_t^\beta} df_t + \nu\,\alpha_t \sqrt{1-\rho^2}\,dZ_t \tag{5.19}$$

It is now apparent that changes to α_t come from one of the two terms in (5.19). The first term is the systematic change occurring with a change in f_t and the second term is unsystematic change. Hence, we find that with a unit change in f_t the change in α_t due to a change in f_t must be $\frac{\rho\,\nu}{f_t^\beta}$. We write this as

$$\delta_f\alpha_t = \frac{\rho\,\nu}{f_t^\beta}\,df_t \tag{5.20}$$

Where we let $\delta_f\alpha_t$ denote "the change in α_t caused by a change in f_t".

In a similar way, we derive the implication for f_t when α_t changes. We find—using similar notation—the following link

$$\delta_\alpha f_t = \frac{\rho\,f_t^\beta}{\nu}\,d\alpha_t \tag{5.21}$$

Using our new findings regarding the correlation movements of α_t and f_t stated in (5.20) and (5.21) together with the Δ and $\mathcal{V}$ parameter change diagrams in (5.17) and (5.18) we can now write the basic dynamics involved in estimating Δ and $\mathcal{V}$ given Bartlett's correlation corrections

$$\Delta_{\text{Bartlett}}: \quad \begin{array}{c} f \to f + df \\ \alpha \to \alpha + \delta_f\alpha \end{array} \qquad \mathcal{V}_{\text{Bartlett}}: \quad \begin{array}{c} f \to f + \delta_\alpha f \\ \alpha \to \alpha + d\alpha \end{array} \tag{5.22}$$

Or stated mathematically, the new expressions for the Δ and $\mathcal{V}$ are[17]

$$\Delta = \underbrace{\frac{\partial C(\cdot)}{\partial f_t} + \frac{\partial C(\cdot)}{\partial \sigma_B(\cdot)} \times \frac{\partial \sigma_B(\cdot)}{\partial f_t}}_{\text{old } \Delta \text{ (5.14)}} + \underbrace{\frac{\partial C(\cdot)}{\partial \sigma_B(\cdot)} \times \frac{\partial \sigma_B(\cdot)}{\partial \alpha_t} \times \frac{\rho\,\nu}{f_t^\beta}}_{\text{Bartlett's correction}} \tag{5.23}$$

$$\mathcal{V} = \underbrace{\frac{\partial C(\cdot)}{\partial \sigma_B(\cdot)} \times \frac{\partial \sigma_B(\cdot)}{\partial \alpha_t}}_{\text{old } \mathcal{V} \text{ (5.16)}} + \underbrace{\left(\frac{\partial C(\cdot)}{\partial \sigma_B(\cdot)} \times \frac{\partial \sigma_B(\cdot)}{\partial f_t} + \frac{\partial C(\cdot)}{\partial f_t} \right) \frac{\rho\,f_t^\beta}{\nu}}_{\text{Bartlett's correction}} \tag{5.24}$$

Having now "calibrated" our risk measures for the inherent correlation in the SABR model, we once more calculate the Δ risk for various values of β. The new Δ's are shown in Figure 5.10. Wee see, that by applying Bartlett's correlation

---

[17]Note: in (Bartlett, 2006) the term $\frac{\partial C(\cdot)}{\partial f_t} \frac{\rho\,f_t^\beta}{\nu}$ in the expression for $\mathcal{V}$ (Bartlett's equation (19)) is (presumably) mistakenly left out.
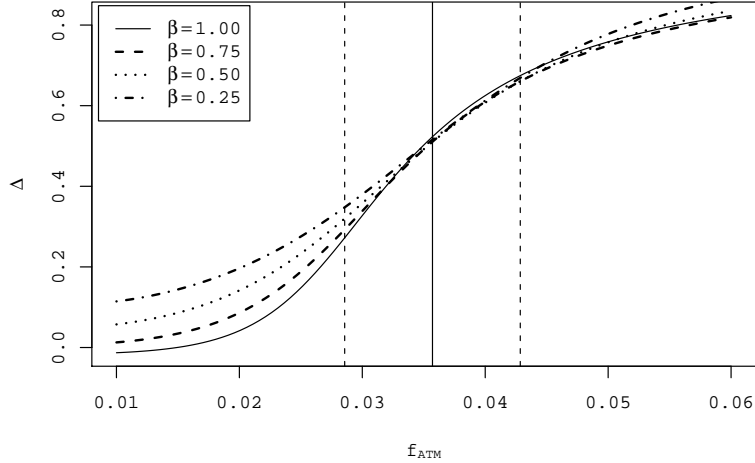
Figure 5.10: Δ *calculated and corrected cf. (Bartlett, 2006) for different values of β. The solid vertical line indicates the ATM forward level, and the dashed vertical lines indicate ATM $\pm$ 20%.*

correction to our Δ measures we obtain much more consistent measures. The Δ's within ATM $\pm$20% region are now very similar regardless of our choice of β. Comparing to Figure 5.9 the improvement is remarkable. Further, we note that in general, the Bartlett correction leaves us with smaller Δs. The reason for this reduction lies in the formulation of the Δ in (5.23). The correction term has three components.

The first component is $\partial C(\cdot)/\partial \sigma_B(\cdot)$. In its crude form, this is the standard Black-Scholes $\mathcal{V}$ measure from (5.12), which we have already seen is strictly positive.[18]

The second component is $\partial \sigma_B(\cdot)/\partial \alpha_t$, which tells us how the Black-Scholes volatility reacts to a change in the SABR volatility. Intuitively, this is also positive, which is confirmed when investigating our expression for $\sigma_B(\cdot)$ in (5.2) on page 37, remembering the result from Hagan et al. (2002) stated in section 5.4.2 that the second term in the curly brackets accounts for a very small part of the final $\sigma_B(\cdot)$. Ignoring the second term in the curly brackets leaves us with an expression that is clearly a strictly increasing function of α_t.

The third and final component of the Bartlett correction is $\rho \nu / f_t^{\beta}$. Obviously, the sign of this term is dependent on the values of ρ and ν. However, as discussed in section 5.3 the correlation between the underlying asset price and the volatility will in most cases be negative, and thereby we get $\rho < 0$. The ν parameter is the

---

[18]Assuming that the price of the underlying asset and time to maturity are both > 0.

"volatility of volatility", and as such it must hold that $\nu \geq 0$. This means that the third component in the Bartlett correction is negative.

Combining the signs of the three components, we now see why correcting the Hagan Δ causes a smaller Δ. Roughly speaking, the Hagan Δ fails to include the semi-deterministic movement in α_t caused by a shift in f_t, which in turn causes an extra (opposite[19]) shift in $C(\cdot)$. Bartlett's method corrects this.

So, does this mean, that by following Hagan's derivations one is not Δ-hedged? Yes and no. Indeed, a portfolio with a net Δ of 0 calculated based on the Δ measure presented in (5.14) on page 56 will not be completely Δ-hedged according to the dynamics of the SABR model. However, we note that in (5.23) the correction term is exactly the Hagan SABR $\mathcal{V}$. This implies that by using the derivations of Hagan et al. from (5.14) and (5.16) to hedge both Δ *and* $\mathcal{V}$ risk, one would actually—not knowing—have managed to correctly Δ-hedge the portfolio. However, the $\mathcal{V}$-hedge would still need to be corrected in accordance with (5.24).

---

[19]Assuming, that $\rho \in [-1; 0[$.

Part III

Application

Chapter 6

Applying the SABR model

6.1 SABR model for pricing

In the previous sections we have shown how the SABR model can effectively and self-consistently be used for purposes of risk management. By applying the SABR model, we are able to calculate Δ and $\mathcal{V}$ risk measures over different strike prices in a consistent manner and thereby more efficiently calculate the total exposure towards these risks of a portfolio of derivatives.

However, risk management is not the only benefit of adopting the SABR model. In the following sections we demonstrate how the SABR model can be utilized in a pricing scenario. We choose to base the application of the SABR model on *constant maturity swaps*. We start out with an introduction to constant maturity swaps, and after we have become familiar with the internal mechanisms of this class of derivatives we will move on to the actual pricing, where we will show how the problem of determining the value of a constant maturity swap can be first split up into smaller bits, and subsequently solved by applying the SABR model and our already established framework for pricing payer and receiver interest rate swaptions. Finally, we apply the theory to market data and give a discussion of our findings and their implications.

6.2 CMS products

The class of derivatives on which we will base our example of pricing using the SABR model is the *constant maturity swap* or CMS swap. The CMS swap is basically a swap contract like the ones presented in section 2.2. However, there are some differences that cause the CMS swap to be not nearly as well behaved a pricing problem as the plain vanilla IRS.
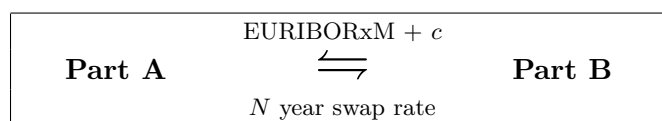
Figure 6.1: *An example diagram of a CMS swap.*

On a daily basis, the *International Swaps and Derivatives Association* (ISDA) publishes swap rate fixings for a range of swap maturities (1Y, 2Y, 5Y etc.). These published fixings are simply the daily par swap rates reported by swap dealers, and are all quoted for plain vanilla interest rate swaps using standard conventions.[1] Because the swap rates are always quoted for the same maturities, they are referred to as *constant maturity swap fixings*. As the name implies, the CMS swap is a swap contract written on such a rate.

6.3 Details of the CMS swap

The CMS swap is a contract where two parties agree on an exchange of a floating CMS rate (e.g. the 10Y swap rate) against (for example) a floating xIBOR rate (e.g. EURIBORxM). The standard for EUR contracts is a quarterly exchange of the floating EURIBOR3M plus a margin[2] c for some floating CMS rate over a predetermined period of time. The two floating rates are usually set in advance and paid in arrears. The additional margin c paid on the xIBOR leg can be either positive or negative depending on the shape of the yield curve. On an upward sloping yield curve the CMS spread is positive, since this curve implies that the short rate (the xIBORxM leg) is smaller than the longer rate (the N year par swap rate). For a downwards sloping yield curve the opposite is the case. An example of a CMS swap flow is depicted in Figure 6.1.

The problem with the CMS swap is the fact that the fixing rates are often accrued over non-matching periods of time. As an example, consider the EUR 10Y CMS swap on the 30Y fixing, which is the contract that specifies an exchange of the floating EURIBOR3M rate for the 30Y swap fixing over a period of ten years. Clearly, the 30Y fixing is mismatched since it is being paid for 10 years. In the pricing of the CMS swaplet[3], the short (floating) rate on which the xIBOR leg is priced has a different sensitivity to movements of the underlying swap curve compared to that of the leg paying the 30Y fixing rate. Graphically speaking, the two legs have different shapes when plotting price against rate, and this is what

---

[1]See Table 2.2 on page 8 for conventions.

[2]Also known as the *CMS spread*.

[3]A *swaplet* is merely a swap with only one exchange of payments.

Product	Payoff at t_p
CMS swap	R_p
CMS cap	$(R_p - K)^+$
CMS floor	$(K - R_p)^+$

Table 6.1: *Payoffs for CMS swap, cap and floor.*

we will refer to as *the convexity correction* that we need to take into account in order to price the CMS swaplet correctly.

6.4 Pricing the CMS swap

Pricing the CMS swap will require us to go through a few preliminary steps. We will break the pricing problem up into smaller bits and finally we will piece these bits together to obtain a pricing expression for a CMS swap.

To start out, we will consider three products: a CMS swap, a CMS cap and a CMS floor, all starting at t_0 with pay dates $t_1, \ldots, t_m$ where t_m is $t_0 + M$ years. The corresponding coverages are denoted $\delta_1^t, \ldots, \delta_m^t$.

The CMS swap pays the N year swap rate fixed on the previous fixing date. For example, at each $t_p \in \{t_1, \ldots, t_m\}$ the CMS swap pays out the par swap rate fixing set at $t_{p-1} - x$ business days.[4] For ease of notation we will define the fixing date as $\tau_p \equiv t_{p-1} - x$ business days, so that the par swap rate paid in the p'th period (paid at t_p) is fixed at τ_p.

The CMS cap pays out the difference between the par swap rate and some fixed rate K, if the par swap rate is greater than K.

The CMS floor is the opposite of the CMS cap.

The three products and their payoffs are summarized in Table 6.1.[5]

Clearly, these three products are comprised by a series of similar payments. Hence, to price either of the products all we need to be able to, is to calculate the value of a single payment of either of the types presented in Table 6.1 paid at some point in time $t_p \in \{t_1, \ldots, t_m\}$ and add them together. These single payment contracts are known as CMS swaplets, CMS caplets and CMS floorlets—meaning that for example a CMS cap running from t_0 to t_m is equivalent to m

---

[4]See Table 2.2 on page 8 for standard conventions.

[5]Note: we use the payoff R_p (denoting the CMS fixing rate) as time t_p payoff for the CMS swap for generality purposes (the CMS rate can be swapped against both a fixed rate or a floating xIBOR rate).
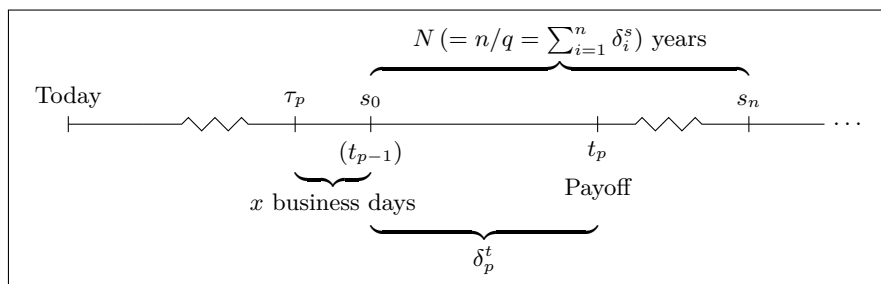
Figure 6.2: *The timeline of a CMS product.*

CMS caplets running from t_{i-1} to t_i where $i \in \{1, \ldots, m\}$. For generality, we adopt the notation $R_s(t)$ for the time t par swap rate for the N year swap starting at s_0 and ending N years later at s_n, and for ease of notation we introduce the notation $L(t)$ corresponding to the $A(\cdot)$ term presented in section 2.2.2, so

$$L(t) = A(t, s_1, s_n) = \sum_{i=1}^{n} \delta_i^s D(t, s_i) \tag{6.1}$$

where $s_1, \ldots, s_n$ are the pay dates for the swap that starts at s_0 and ends N years later at s_n and $\delta_1^s, \ldots, \delta_n^s$ are the corresponding coverages. A general timeline of a CMS product is presented in Figure 6.2.

Concluding on our preliminaries: $t_1, \ldots, t_m$ are the pay dates for the CMS product (with corresponding fixing dates $\tau_1, \ldots, \tau_m$), while $s_1, \ldots, s_n$ are the pay dates for what we will call *the reference swap*—e.g. the 10Y swap against EURIBOR6M. The par swap rate $R_s(t)$ is the time t par swap rate for the reference swap. When a CMS product is entered into, the time frame $\{t_0, \ldots, t_m\}$ is laid out and stays the same, but the time frame for the reference swap "shifts", so that a new time frame $\{s_0, \ldots, s_n\}$ is established at each fixing date $\tau_j \in \{\tau_1, \ldots, \tau_m\}$. In the following sections we will denote today as time 0.

6.4.1 The CMS caplet and floorlet

In order to price the CMS swaplet, we will initially focus on the CMS caplet and the CMS floorlet. We will use t_p as notation for the pay date for the caplet/floorlet and τ for the fixing date.[6] To price these we will need to use the change of numeraire technique introduced in section 2.3. This tells us that for any choice of a traded asset as our numeraire there exists a probability measure such that the value of

---

[6]As long as we are in the realm of the caplet, floorlet and swaplet we will relax our notation using τ rather than τ_p to denote the fixing date for the interest rate involved in the payoff at time t_p, since these products all have only one payment (at time t_p) and one fixing date.

any traded asset $V(t)$ divided by our numeraire is a martingale. Choosing $L(t)$—which essentially is a sum of zero coupon bonds, and therefore obviously a traded asset—as our numeraire we write this as

$$V(t) = L(t) \, \mathbb{E}_t^{Q^L} \left[\frac{V(T)}{L(T)} \right] \tag{6.2}$$

For our plain vanilla European payer swaption we know that at time T it has the value[7] $V_{sw}(T) = L(T) \, (R - K)^+$. Inserting this into (6.2) we end up with

$$V_{sw}(t) = L(t) \, \mathbb{E}_t^{Q^L} [(R - K)^+] \tag{6.3}$$

which is the price of our payer swaption under the Q^L measure, and a result we will use in the following sections.

The CMS caplet

Having established our framework for pricing swaptions we move on to the pricing of the CMS caplet. The CMS caplet has a payoff of $(R_s(\tau) - K)^+$ at time t_p. At the time of the swap rate fixing $\tau < t_p$, the value of the caplet is $D(\tau, t_p)(R_s(\tau) - K)^+$. Using this together with (6.2) we are able to write the time t value of the caplet as

$$V_{\text{caplet}}(t) = L(t) \, \mathbb{E}_t^{Q^L} \left[\frac{D(\tau, t_p)(R_s(\tau) - K)^+}{L(\tau)} \right] \tag{6.4}$$

The ratio $D(\tau, t_p)/L(\tau)$ is a martingale[8] under the Q^L measure and therefore, its time t expectation is equal to its time 0 value

$$\mathbb{E}_0^{Q^L}[D(\tau, t_p)/L(\tau)] = D(0, t_p)/L(0) \tag{6.5}$$

Using this result, we can rewrite (6.4)[9] and write today's value as

$$V_{\text{caplet}}(0) = D(0, t_p) \, \mathbb{E}_0^{Q^L} \left[(R_s(\tau) - K)^+ \frac{D(\tau, t_p)/L(\tau)}{D(0, t_p)/L(0)} \right] \tag{6.6}$$

which, in turn, we can split up further writing

$$V_{\text{caplet}}(0) = D(0, t_p) \, \mathbb{E}_0^{Q^L} \left[(R_s(\tau) - K)^+ \right] +$$
$$D(0, t_p) \, \mathbb{E}_0^{Q^L} \left[(R_s(\tau) - K)^+ \left(\frac{D(\tau, t_p)/L(\tau)}{D(0, t_p)/L(0)} - 1 \right) \right] \tag{6.7}$$

---

[7]See (2.13) on page 14.

[8]Cf. Hagan (2003).

[9]Dividing by $D(0, t_p)/L(0)$ inside the [ ]'s and multiplying by $D(0, t_p)/L(0)$ outside the [ ]'s.

Comparing the first term in (6.7) to the expression in (6.3) we recognize it as being a standard European payer swaption written on a notional of $D(0, t_p)/L(0)$. We have already established the required framework to price contracts of this kind, and therefore we will leave this term for now and instead we will focus on the second term of (6.7) which is the term that yields the *convexity correction*.

As it is presented in (6.7), the convexity correction is a function of several representations of the yield curve[10]—swap rates, zero coupon prices and discount factors.[11] As a first step on our path to evaluating the convexity correction, we will express the entire term as a function of only one representation of the yield curve: the par swap rate $R_s(t)$.

First, we take a look at $L(t)$. Originally, we defined $L(t)$ in (6.1) as the sum of discounted coverages. We can rearrange this into

$$L(t) = D(t, s_0) \sum_{j=1}^{n} \delta_j^s \frac{D(t, s_j)}{D(t, s_0)} \tag{6.8}$$

After some thought we recognize this as a discount factor times a sum of discount factors between t and $s_1, \ldots, s_n$ discounted back to s_0 weighted by their respective coverages—a series of forward discount factors. The time t value of the stream of discounts must be approximately equal to continuously discounting by the par swap rate between s_0 and s_n, $R_s(t)$ in each period. Therefore, assuming further that all periods are of equal length, we can approximate (6.8) by

$$L(t) \approx D(t, s_0) \sum_{j=1}^{n} \frac{1}{q} \frac{1}{(1 + R_s(t)/q)^j} \tag{6.9}$$

where q represents the number of periods in a year for the reference swap. Using the result for the sum of a finite series

Theorem 6.1. (THE SUM OF A FINITE SERIES)[12]
Let a and k be constants. Then

$$b_n = \sum_{j=0}^{n} a\, k^j = a \frac{1 - k^{n+1}}{1 - k}$$

Setting $a = \frac{1}{q}$ and $k = \frac{1}{(1 + R_s(t)/q)}$ we are able to eliminate the summation term

---

[10]While in the theoretical realm we will assume that there is only *one* yield curve even though we have seen earlier, that indeed there are several.

[11]We distinguish between discount factors ($D(0, T)$) and zero coupon prices ($D(t_i, T)$) since the former are deterministic today (at time 0) and the latter are stochastic until we are in fact at time t_i.

[12]Cf. Fuglede et al. (1999).

and rewrite (6.9). Note that Theorem 6.1 gives the sum from 0 to n and not the sum from 1 to n which we need. Therefore we must subtract $a\,k^0 = a$ which is the term added for the case of $j = 0$. We find

$$L(t) \approx D(t, s_0) \left(\frac{1}{q} \frac{1 - \left(\frac{1}{(1+R_s(t)/q)} \right)^{n+1}}{1 - \frac{1}{(1+R_s(t)/q)}} - \underbrace{\frac{1}{q}}_{-a} \right) \qquad (6.10)$$

This expression is not very appealing, but it turns out that it can be simplified quite a bit into

$$L(t) \approx \frac{D(t, s_0)}{R_s(t)} \left(1 - \frac{1}{(1 + R_s(t)/q)^n} \right) \qquad (6.11)$$

The task of rewriting (6.10) into (6.11) is not very complex, but it does require a few clever turns and it does take a few steps. We will not present it here, but for the sake of completeness it is included in Appendix A.

Second, looking at the discount factor up to the time of payment t_p, we will approximate this quantity as the discount factor up to time s_0 (the start date of the underlying swap) *further* discounted by the par swap rate for the remaining time up to the time of payment t_p. This means that we adopt the following approximation

$$D(t, t_p) \approx \frac{D(t, s_0)}{(1 + R_s(t)/q)^\gamma} \qquad (6.12)$$

where $\gamma = (t_p - s_0)\, q = \delta_p^t\, q$ represents the fraction of a (q) period between s_0 and t_p.

Having expressed both $L(t)$ and $D(t, t_p)$ in terms of the swap rate $R_s(t)$ in (6.11) and (6.12) respectively, we now define the function $G(R_s(t))$

$$G(R_s(t)) \equiv \frac{D(t, t_p)}{L(t)} \approx \frac{R_s(t)}{(1 + R_s(t)/q)^\gamma} \frac{1}{1 - (1 + R_s(t)/q)^{-n}} \qquad (6.13)$$

Indeed, more complex models can be applied in order to increasingly improve the accuracy of the replication. However, for our purposes we will settle for (6.13). For more elaborate models we refer to Hagan (2003). With our choice of the $G(\cdot)$ function, we return to the convexity correction term (the second term) of the CMS caplet in (6.7). We rewrite the convexity correction using the $G(\cdot)$ function from (6.13) to get

$$CC(0) = D(0, t_p)\, \mathbb{E}_0^{Q^L} \left[(R_s(\tau) - K)^+ \left(\frac{G(R_s(\tau))}{G(R_s(0))} - 1 \right) \right] \qquad (6.14)$$

In order to evaluate the expectation, we apply a general result from Carr (2005),

69

stating that for any twice differentiable function h and any scalar $\kappa \geq 0$ the following holds

$$
\begin{aligned}
h(S) = & h(\kappa) + h'(\kappa)\,(S - \kappa) + \\
& \int_\kappa^\infty h''(x)\,(S - x)^+ dx + \int_0^\kappa h''(x)\,(x - S)^+ dx
\end{aligned}
\tag{6.15}
$$

The first line represents the tangent approximation while the second line is the continuous correction to the tangent as one moves from the initial point κ to the end point S. As a special case of (6.15) we can write a function of $R_s(t)$

$$
\begin{aligned}
h(R_s(t)) = & h'(K)\,(R_s(t) - K)^+ + \\
& \int_K^\infty h''(x)(R_s(t) - x)^+ dx \quad \text{for } R_s(t) > K
\end{aligned}
\tag{6.16}
$$

where we choose $h(x) = (x - K)\left(\frac{G(x)}{G(R_s(0))} - 1\right)$ making $h(K) = 0$. Substituting $h(R_s(\tau))$ into (6.14) and rearranging leaves us with

$$
\begin{aligned}
CC(0) = D(0, t_p)\{\, & h'(K)\,\mathbb{E}_0^{Q^L}[(R_s(\tau) - K)^+] + \\
& \int_K^\infty h''(x)\,\mathbb{E}_0^{Q^L}[(R_s(\tau) - x)^+]\,dx \,\}
\end{aligned}
\tag{6.17}
$$

Combining the convexity correction with the first term from the original expression for the value of the CMS caplet in (6.7) and letting $C(K) = L(0)\,E_0^{Q^L}[(R_s(\tau) - K)^+]$ denote today's value of a payer swaption with strike K we are now able to write

$$
V_{\text{caplet}}(0) = \frac{D(0, t_p)}{L(0)}\left\{ (1 + h'(K))\,C(K) + \int_K^\infty h''(x)\,C(x)\,dx \right\}
\tag{6.18}
$$

Thus, at time 0 we can replicate the payoff of (and thereby price) a CMS caplet through a portfolio of plain vanilla European payer swaptions struck on different notionals at different strikes. This is where the SABR model comes into play. The SABR model provides us with the exact implied volatility interpolation setup we need in order to price swaptions at a wide range of strike prices. Further, by applying the SABR model to the pricing of the swaptions, the Δ and $\mathcal{V}$ risks of the resulting swaption portfolio are easily aggregated in a self-consistent manner using the SABR risk framework presented in section 5.6.3.

The CMS floorlet

The argumentation and derivation of the replication formula for the CMS floorlet are similar to those of the CMS caplet. The main difference being that the floorlet is replicated through the use of European plain vanilla receiver swaptions where the CMS caplet is replicated through payer swaptions. The value of the CMS floorlet is

$$V_{\text{floorlet}}(0) = \frac{D(0, t_p)}{L(0)} \left\{ (1 + h'(K)) \, P(K) - \int_K^\infty h''(x) \, P(x) \, dx \right\} \qquad (6.19)$$

where $P(K) = L(0) \, \mathbb{E}_0^{Q^L}[(K - R_s(\tau))^+]$ is today's value of a plain vanilla European receiver swaption, and $h(x)$ is identical to that used for the caplet case.

6.4.2 The CMS swaplet

At the time of the payout t_p, having bought a caplet and having sold a floorlet— both struck at K—will result in a payoff of

$$(R_s(\tau) - K)^+ - (K - R_s(\tau))^+ = R_s(\tau) - K \qquad (6.20)$$

Thus, being long a caplet and short a floorlet yields the a payoff equivalent to that of being long a swaplet and having sold K zero coupon bonds with maturity t_p. Therefore, we can value the CMS swaplet as

$$V_{\text{swaplet}}(0) = V_{\text{caplet}}(0) - V_{\text{floorlet}}(0) + D(0, t_p) \, K \qquad (6.21)$$

Using the expressions for the values of the CMS caplet and CMS floorlet, given in (6.18) and (6.19) respectively, we can write today's value of the CMS swaplet as

$$V_{\text{swaplet}}(0) = \frac{D(0, t_p)}{L(0)} \left\{ \int_{R_s(0)}^\infty h''_{ATM}(x) \, C(x) \, dx + \right.$$
$$\left. \int_{-\infty}^{R_s(0)} h''_{ATM}(x) \, P(x) \, dx \right\} + D(0, t_p) \, R_s(0) \qquad (6.22)$$

where $h_{ATM}(x) = (x - R_s(0)) \left(\frac{G(x)}{G(R_s(0))} - 1 \right)$.[13]

We have now successfully derived an expression to price a CMS swaplet paying out $R_s(\tau)$ at time t_p. The task of pricing a CMS swap leg paying out $R_s(\tau_p)$ at each $t_p \in \{t_1, \ldots, t_m\}$ now becomes an arbitrary matter of summing a range of

---

[13]For the complete derivation of (6.22) see Appendix B.

CMS swaplets. However, before we move on to an actual pricing example, we remind ourselves that rather than paying out the actual rate $R_s(\tau)$ at time t_p, a CMS swaplet pays out according to the money market convention making the actual time t_p payoff of the CMS swaplet $\delta_p^t R_s(\tau)$, where δ_p^t is the coverage for the p'th period of the CMS swaplet—as opposed to the coverage of the reference swap that has coverages δ_p^s.[14] However, since the coverages are merely constants, these can simply be multiplied by the already established pricing expression in (6.22) to obtain the prices of CMS swaplets paying off $\delta_p^t R_s(\tau)$ rather than $R_s(\tau)$.

6.5   A pricing example using the SABR model

Previously in this chapter we have shown how the SABR model's ability to inter- and extrapolate known volatility smile points can be used in pricing CMS products. We have established formulae to replicate the payoffs of different CMS products using ranges of our well-known European plain vanilla payer and receiver swaptions. In this section we apply the theoretical pricing setup to actual market data and finally we compare our theoretical price to the price found in the market.

6.5.1   Choice of product

For our pricing example we choose to price a 5Y CMS swap swapping the 10Y par swap rate against a floating EURIBOR3M payment. Hence, our product is similar to that depicted in Figure 6.1 on page 64, with the N year swap rate being the 10Y swap rate and EURIBORxM being EURIBOR3M. We note that the 10Y swap rate is the swap rate against EURIBOR**6M**,[15] while the rate it is being swapped for is floating EURIBOR**3M**. The CMS swap is set in advance, paid in arrears and has quarterly payments. We price the product assuming that today is June 1st 2010. On this date we observe the swap curves shown in Figure 6.3. The two swap curves show the swap rate for maturities up to 30 years for swaps against EURIBOR3M and EURIBOR6M. Further, we observe the four volatility smiles presented in Table 6.2. Since payments in our CMS swap are quarterly, the volatility smiles in Table 6.2 will not suffice. In order to price the individual CMS swaplets we require a volatility smile for each payment date t_p. In our case we will assume that accrual periods are all of length 0.25, implying that we need a volatility smile for each $t_p \in \{0.25, 0.50, \ldots, 4.75, 5.00\}$ where time 0 is today (June 1st 2010). We address this issue by applying simple linear interpolation between the observed smiles.

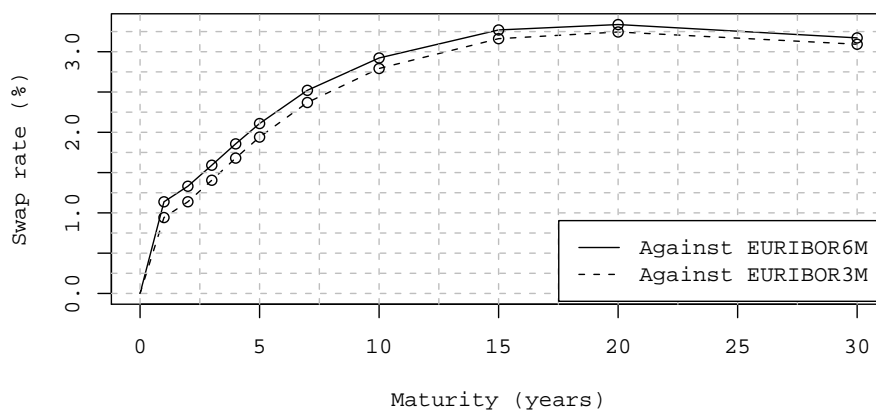---

[14]See Figure 6.2 on page 66.
[15]Standard for EUR swap rates.

Figure 6.3: *Swap curves per June 1st 2010.*

Expiry	−200 bps	−100 bps	−50 bps	ATM	+50 bps	+100 bps	+200 bps
0.25	55.71%	37.21%	31.75%	28.08%	25.51%	24.36%	24.23%
1.00	46.78%	33.60%	29.57%	26.95%	24.97%	24.05%	23.68%
5.00	32.16%	24.01%	21.20%	19.37%	18.20%	17.87%	18.30%
10.00	29.93%	22.11%	19.44%	17.66%	16.59%	16.25%	16.58%

Table 6.2: *Volatility smiles for different swaption expiries per June 1st 2010. The underlying swap is the 10Y swap against EURIBOR6M.*

6.5.2 The pricing procedure

Having established the characteristics of our CMS product we turn to the pricing procedure itself. We have seen the theoretical pricing setup, and now we will go through the steps involved in an actual pricing using market data. The pricing of the CMS swap will follow the approach presented in section 6.4.2 where we show how to price a CMS swaplet. Once we have priced all the CMS swaplets we simply add up their values to obtain the value of the CMS swap leg. Since the CMS swap valuation is merely a matter of adding CMS swaplet values we will focus on the pricing procedure of the individual CMS swaplet.

In the previous section we discussed the data at hand and how we interpolated the observed volatility smiles to end up with a sufficient set of volatility smiles. Therefore, in this section we will simply assume, that swap curves and volatility smiles are given (observed in the market). From here on we will present the pricing procedure chronologically according to our own implementation. In addition to an explanation of each step in the proces, results will be reported for the calculation of the CMS swaplet expiring at time 5 (the final swaplet).

For every pay date $t_p \in \{0.25, 0.50, \ldots, 4.75, 5.00\}$ perform the following steps

i. Lay out the schedule for the reference swap. We lay out the schedule for the 10Y swap against EURIBOR6M (semi-annual payments) starting at $s_0 = t_{p-1}$ and ending at $s_n = s_0 + 10Y$. As with the schedule for the CMS payments, we assume equal coverages $\delta_p^s = \delta^s = 0.5$ for all periods p in the reference swap. **Result**: for the final swaplet the schedule for the reference swap is (from s_0 to s_n)

$$s = \{4.75, 5.25, \ldots, 14.25, 14.75\}$$

ii. Calculate the annuity for the reference swap according to (6.1) on page 66:

$$L(0) = \sum_{i=1}^{n} \delta^s D(0, s_i) \tag{6.23}$$

Result: $L(0) = 7.4991$

iii. Calculate the relevant (forward) par swap rate. We calculate the 10Y par swap rate between s_0 and s_n as

$$R_s(0) = \frac{D(0, s_0) - D(0, s_{20})}{L(0)} \tag{6.24}$$

where the discount factors are calculated using the swap curve against EURIBOR6M.
Result: $R_s(0) = 3.8547\%$

iv. Fit the SABR model to the volatility smile for the relevant swaption. In our case, the relevant swaption is the 10Y swaption with expiry date[16] t_p—or using our standard notation: the t_pY10Y swaption. We fit the SABR model by fixing the β and then calibrating the α, ρ and ν freely. **Result**: setting $\beta = 0.5$ we obtain

$$\{\alpha, \rho, \nu\} = \{11.37\%, -41.47\%, 52.50\%\}$$

v. Define a function $I(x; \cdot)$ that is the function to be integrated over in (6.22) on

---

[16]Expiry date for the *option* element of the swaption.

page 71 so that

$$I(x; R_s(0), \beta, \alpha, \rho, \nu, T, L, \text{type}) =$$
$$h''_{ATM}(x; R_s(0)) \times PV_{\text{swaption}}(x; R_s(0), \beta, \alpha, \rho, \nu, T, L, \textit{type})$$

where $h_{ATM}(\cdot)$ is the function from section 6.4.2 on page 71,[17] x is the strike price, $R_s(0)$ is the par swap rate, $\{\beta, \alpha, \rho, \nu\}$ are SABR parameters from (iv), T is the swaption's expiration date,[18] L is the annuity from (ii) and $\textit{type}$ is an indicator of whether the swaption is a payer or receiver swaption.

vi. Calculate the value of the swaplet according to (6.22). However, we make a slight adjustment to the lower bound of the second intergral setting it to 0 rather than $-\infty$ since we expect interest rates to be non-negative. This leaves us with

$$PV_{\text{swaplet}} = D(0, t_p) R_s(0) +$$
$$\frac{D(0, t_p)}{L(0)} \left\{ \int_{R_s(0)}^{\infty} I(x; \cdot, \textit{pay}) \, dx + \int_0^{R_s(0)} I(x; \cdot, \textit{receive}) \, dx \right\}$$

Result: $PV_{\text{swaplet}} = 4.35\%$

vii. Multiply the present value from (vi) by the CMS swap coverage. Since we assume equal coverages for each period we multiply by $\delta^t = 0.25$.
Result: $\delta^t PV_{\text{swaplet}} = 1.09\%$

These seven points are repeated for each CMS swaplet in the CMS swap, meaning that for the 5Y CMS swap with quarterly payments they are repeated twenty times. The total value of the CMS swaplets gives us the value of the CMS swap leg. Arbitrarily setting $\beta = 0.5$ yields

$$PV_{\text{CMS leg}} = 17.0407\% \tag{6.25}$$

To compare our price with a market quote we need to get the value on the form of a CMS margin payment.[19] The standard for EUR contracts is a swap against EURIBOR3M and hence, we need to value this leg also in order to come up with a net present value of the entire CMS swap. Today's value of a quarterly payment

---

[17]Using $q = 2$ and $\gamma = 0.5$ for $G(\cdot)$ (6.13) since our reference swap is semi-annual and our CMS swap payments are quarterly.

[18]In our notation, the swaptions used to price a CMS swaplet should expire at the payment date of the CMS swaplet t_p.

[19]The margin c that is set today and yields a PV of 0 when the CMS leg is swapped against some floating xIBOR rate.

of the EURIBOR3M rate during the course of the CMS swap (from t_0 to t_m) is

$$PV_{\text{EURIBOR3M}} = \sum_{i=1}^{m} \delta^t \, D(0, t_i) \, F(0, t_{i-1}, t_i) \qquad (6.26)$$

where $F(0, t_{i-1}, t_i)$ is today's forward rate between t_{i-1} and t_i. Using the definition from (2.1) on page 5 we can simplify the expression further

$$PV_{\text{EURIBOR3M}} = D(0, t_1) - D(0, t_m) \qquad (6.27)$$

Hence, in our case the PV of the floating EURIBOR3M leg is calculated by setting $t_1 = 0.25$ and $t_m = 5$ and calculating the discount factors on the EURIBOR3M swap curve. Doing so yields

$$PV_{\text{EURIBOR3M}} = 9.2489\%$$

This leaves us with a PV of both the CMS swap leg and the floating EURIBOR3M. The difference between these values is the CMS swap's net present value

$$\begin{aligned} \text{Net } PV_{\text{CMS swap}} &= PV_{\text{CMS leg}} - PV_{\text{EURIBOR3M}} \\ &= 17.0407\% - 9.2489\% \\ &= 7.7917\% \end{aligned}$$

We will now calculate the CMS swap spread c, that sets the net PV of the CMS swap equal to 0. We can think of the net PV as the value of a CMS swap on a unit notional. To determine the size of the CMS spread c, we must determine the PV for a unit spread. The value of a single unit CMS spread c paid at every payment date the sum of the spread payments is

$$\sum_{i=1}^{20} \delta^t D(0, t_i) = 4.8056 \qquad (6.28)$$

Hence, every basis point paid as a CMS spread has a value of 4.8056 basis points over the duration of the CMS swap. Subsequently, if the CMS swap has a net PV of 7.7917% and the value of a unit CMS spread is 4.8056, then the CMS spread that leaves the CMS swap with a net PV of 0 can be determined by

$$c^* = \frac{7.7917\%}{4.8056} = 1.6214\% \text{ (or 162.14 bps)} \qquad (6.29)$$

Figure 6.4: *Historical quotes for the 5Y CMS swap against the 10Y index.*

Figure 6.4 shows a screen from the Bloomberg system with historical mid prices[20] for 5Y CMS swaps on the 10Y swap rate. We see that the Bloomberg system reports a quote for the CMS swap of 175.5 bps. This puts us roughly 13.3 bps off target—a deviation of $13.3/175.5 = 7.5\%$.

Our deviation from the market quote might arise for several reasons. First, we adopt a somewhat simplified approach of handling the coverages δ^s and δ^t. We assume these to be constant, while in the real world they are of course not constant, but subject to day count conventions and the specific layout of the business calendar for the relevant period of time. Second, the historical Bloomberg data only offers a *mid* price. It does not report anything about the magnitude of a potential bid-ask spread. While we do not have acces to historical bid-ask spreads Figure 6.5 shows the bid-ask spread for several CMS swaps as per February 22nd 2011. We note that the 5Y CMS swap on the 10Y index has a bid-ask spread of 10 bps with a bid of 144.50 and an ask of 154.50. Hence, mid price to bid price can easily span (at least) 5 bps. Third, as mentioned in section 6.5.1 we are not initially endowed with a complete volatility surface. All we have are the four volatility smiles presented in Table 6.2 on page 73 out of which we can only use the three up to and including expiry time 5. To get what we need we apply a simple—and maybe over-simplifying—linear interpolation between the volatility

---

[20]The mid price is the price right between the asking price and the bid price.

```
 20                                                    Message
200<Go> to view in Launchpad
 8:27 EUR  CMS                                         PAGE  1  /  2

  Object     Ask        Bid      Time        Object      Ask       Bid      Time
  2 Year Index                               10 Year Index
1)  5Y Swap  66.3000   60.3000    7:00    9)   5Y Swap 154.5000 144.5000    7:00
2) 10Y Swap  56.1000   50.1000    7:00   10)  10Y Swap 129.3000 119.3000    7:00
3) 15Y Swap  50.0000   44.0000    7:00   11)  15Y Swap 110.1000 100.1000    7:00
4) 20Y Swap  44.9000   38.9000    7:00   12)  20Y Swap 103.3000  83.3000    7:00


  5 Year Index                               20 Year Index
5)  5Y Swap 112.5000  103.5000    7:00   13)   5Y Swap 173.8000 153.8000    7:00
6) 10Y Swap  92.9000   83.9000    7:00   14)  10Y Swap 129.1000 109.1000    7:00
7) 15Y Swap  79.4000   70.4000    7:00   15)  15Y Swap  98.6000  78.6000    7:00
8) 20Y Swap  69.9000   60.9000    7:00   16)  20Y Swap  87.7000  67.7000    7:00




  EUR CMS quoted A/360 vs 3M Euribor


Australia 61 2 9777 8600 Brazil 5511 3048 4500 Europe 44 20 7330 7500 Germany 49 69 9204 1210 Hong Kong 852 2977 6000
Japan 81 3 3201 8900        Singapore 65 6212 1000      U.S. 1 212 318 2000      Copyright 2011 Bloomberg Finance L.P.
                                                                SN 697401 G612-977-0 22-Feb-2011 08:27:17
```

Figure 6.5: *CMS broker screen as per February 22nd 2011.*

SABR β	0.20	0.40	0.60	0.80	0.85
CMS spread (bps)	161.44	161.8	162.7	167.53	176.87

Table 6.3: *CMS spreads for different β values.*

points over time. Fourth, as mentioned earlier we arbitrarily chose to set $\beta = 0.5$. We have seen earlier, that the SABR model fits the points of the volatility smile very nicely, seemingly regardless our choice of β. Could it be, however, that while the value of β has little to no effect on the observed smile, it *does* indeed have an effect on the extrapolation of the smile? We remind ourselves that the procedure of pricing a CMS swap through a series of swaptions require us to integrate over an infinite range of strikes, and therefore extrapolating the volatility smile correctly becomes an important task. Effectively, the swaption prices tend to be close to 0 at strikes $\approx 15\%$, but this still leaves a lot of "unknown territory" between our last observed volatility—which is ATM + 200 bps—and the point when the swaption prices approach 0. Table 6.3 shows CMS spreads resulting from different choices of SABR β. We observe that while the CMS spread estimates are fairly consistent when using a $\beta < 0.60$, choosing a $\beta > 0.60$ causes significantly higher CMS spreads. While this inconsistency of CMS spreads might at first glance seem to be an inconvenience, this actually allows for a way to calibrate a SABR model

to market data using CMS quotes.[21] Inverting our pricing process into a SABR-calibrating process our CMS pricing model tells us to use a $\beta \in [0.82, 0.86]$ to price the 5Y CMS swap within the quoted mid price ± 5 bps.

Concluding on our pricing example, we have shown that the SABR model, together with our fairly simple framework for plain vanilla European interest rate swaption pricing, can price more complicated financial derivatives. However, we have also seen, that when expanding the volatility smile boundaries beyond what is market-observable, the choice of β starts to have an effect on our results. We found that in our pricing procedure a β around 0.82 would put us on the market mid price.

---

[21]Cf. Mercurio and Pallavicini (2006).

Chapter 7

Conclusion

We have provided evidence supporting our claim that empirically the return volatility of a financial asset is hardly constant. Rather, we have seen that volatility seems to fluctuate over time. We have also shown that the Black-Scholes assumption of constant volatility is inconsistent with the market's pricing of European plain vanilla options. The market clearly shows a tendency to a smile-like volatility curve when plotting volatility against strike prices. Thus, seen in the context of the Black-Scholes setup, volatility is not only varying in time but also in strike price. We presented two different approaches in our efforts to model the behavior of volatility over time as well as over strike prices.

The first approach was the local volatility model. Initially, we considered the local volatility model in the discretized case of binomial tree option pricing. We showed that from a given volatility smile—corresponding to known European plain vanilla option prices—we were able to reverse the usual binomial pricing scheme and back our way into a grid (tree) of implied volatilities. Further, the reversal of the binomial pricing yielded a tree showing how the value of the underlying asset evolves over time (under the local volatility model) along with transition probabilities. Hence, the method enables us to back our way into the several grids, containing numerous pieces of information, from observable European plain vanilla option implied volatilities. Following, these grids can be used for pricing various more exotic derivatives in a market-consistent manner. After showing, in the discretized case, how the local volatility model could be used to obtain the information needed to price consistent with market observables, we moved on to the analysis of the dynamics implied by the local volatility model when the forward price of the underlying asset shifts away from its initial state. We expected the volatility smile to shift in the same direction as the forward price, but instead we

saw that the model shifts the volatility smile in the opposite direction. Besides being counterintuitive this is also inconsistent with what is observed in the market. Finally, we discussed the implications this has in a risk management context, and concluded that while the local volatility model might be a reasonable choice for purposes of pricing, when it comes to risk management the model is inadequate.

As a second approach to modeling the volatility smile we presented the SABR model of Hagan et al. (2002). We showed how the implied Black-Scholes volatility, predicted by the SABR model, can be approximated in closed form, and how the various parameters affect the shape of the resulting volatility smile. We saw that some of the parameters had somewhat similar effects, which we, in turn, utilized when estimating the SABR parameters. We presented various estimation methods, but eventually we found that the SABR model is very capable of fitting a volatility smile regardless of the estimation method. Further, the fit of the model also did not seem to be very dependent on the choice of β.[1] However, by shifting the β we saw that while the fit of the model is not significantly affected, the parameters of the SABR model are indeed. We demonstrated a potential problem arising from the instability of the SABR parameters by calculating Δ risk for different choices of β, and saw that different βs yielded different SABR parameters, which eventually led to very different Δ risk profiles. We remedied this discrepancy by stepping back and reconsidering the model, taking into account the correlation between the forward price level and the instantaneous volatility that is inherent in the SABR model and incorporating this into our risk measures. By doing so, we showed that the SABR model does predict very similar Δ risk profiles regardless of the β used.

As a final act, we demonstrated how the SABR model's ability to inter- and extrapolate a volatility smile can be applied in a pricing scenario. We used our basic setup for pricing payer and receiver swaptions, along with the SABR model to price constant maturity caplets and constant maturity floorlets, which then, in turn, we combined in order to price a constant maturity swap. We compared our results with market prices and gave a discussion of our findings.

7.1 The next step

To round off, a few comments on possible extensions of the work presented in this thesis will be provided. A natural extension of the work done in this thesis

---

[1] Cf. the presentation of the SABR model on page 36.

would be a further investigation of ways to estimate a suitable β. As we suggest ultimately in section 6.5.2 on page 79, more complex derivatives might be utilized in order to calibrate the β estimate to that inherent in the market. Further, the expansion of the SABR model from a single-maturity (volatility smile for varying strike prices) to a model that is continuous in maturity as well (volatility surface spanned by strike price and time) would also be an obvious choice for additional research—indeed a combination of the two possible extensions mentioned here might lead to interesting results.

Appendix A

Derivation of (6.11)—the approximation of $L(t)$

In this appendix we present the derivation of the approximation of the term

$$L(t) = D(t, s_0) \sum_{j=1}^{n} \delta_j^s \frac{D(t, s_j)}{D(t, s_0)}$$

On page 69 we show that we can approximate $L(t)$ by the expression

$$L(t) \approx D(t, s_0) \left(\frac{1}{q} \frac{1 - \left(\frac{1}{(1+R_s(t)/q)} \right)^{n+1}}{1 - \frac{1}{(1+R_s(t)/q)}} - \frac{1}{q} \right)$$

Further, we claim that it can be simplified into

$$L(t) \approx \frac{D(t, s_0)}{R_s(t)} \left(1 - \frac{1}{(1 + R_s(t)/q)^n} \right)$$

In the following we will verify this claim.

A.1 The derivation

The derivation takes a few steps and a few tricks, but it is not very complex. Therefore, comments will be kept at a minimum.

$$L(t) \approx D(t, s_0) \left(\frac{1}{q} \frac{1 - \left(\frac{1}{1 + R_s(t)/q} \right)^{n+1}}{1 - \frac{1}{(1 + R_s(t)/q)}} - \frac{1}{q} \right)$$

$$= D(t, s_0) \left(\frac{1}{q} \frac{1 - \frac{1}{(1 + R_s(t)/q)^{n+1}}}{\frac{1 + R_s(t)/q - 1}{1 + R_s(t)/q}} - \frac{1}{q} \right)$$

$$= D(t, s_0) \frac{1}{R_s(t)} \left(\frac{1 - \frac{1}{(1 + R_s(t)/q)^{n+1}}}{\frac{1}{1 + R_s(t)/q}} - \frac{R_s(t)}{q} \right)$$

$$= \frac{D(t, s_0)}{R_s(t)} \left((1 + R_s(t)/q) - \frac{1 + R_s(t)/q}{(1 + R_s(t)/q)^{n+1}} - \frac{R_s(t)}{q} \right)$$

$$= \frac{D(t, s_0)}{R_s(t)} \left(1 - \frac{1}{(1 + R_s(t)/q)^n} \right)$$

Q.E.D.

84

Appendix B

Derivation of the replication formula for the CMS swaplet

To derive the formula for replicating the CMS swaplet through the use of plain vanilla European payer and receiver swaptions the following entities are required:

The caplet

$$V_{\text{caplet}}(0) = \frac{D(0, t_p)}{L(0)} \left\{ (1 + h'(K)) \, C(K) + \int_K^\infty h''(x) \, C(x) \, dx \right\} \tag{B.1}$$

where $C(K) = L(0) \, \mathbb{E}_0^{Q^L}[(R_s(\tau) - K)^+]$ is the price of a plain vanilla European payer swaption with strike K.

The floorlet

$$V_{\text{floorlet}}(0) = \frac{D(0, t_p)}{L(0)} \left\{ (1 + h'(K)) \, P(K) - \int_K^\infty h''(x) \, P(x) \, dx \right\} \tag{B.2}$$

where $P(K) = L(0) \, \mathbb{E}_0^{Q^L}[(K - R_s(\tau))^+]$ is the price of a plain vanilla European receiver swaption with strike K.

The $h(\cdot)$ function

$$h(x) = (x - R_s(0)) \left(\frac{G(x)}{G(R_s(0))} - 1 \right) \tag{B.3}$$

Here $G(R_s(t)) \equiv \frac{D(t, t_p)}{L(t)} \approx \frac{R_s(t)}{(1 + R_s(t)/q)^\gamma} \frac{1}{1 - (1 + R_s(t)/q)^{-n}}$ where q denotes the number of periods in a year,[1] and γ is the fraction of a period between the start date of the reference swap and the payout date t_p.

---

[1] For the reference swap starting at s_0 and ending N years later at s_n quarterly payments would imply $q = 4$, semiannually would imply $q = 2$ and so on.

B.1 The derivation

The payoff at time t_p of being long a CMS caplet and being short a CMS floorlet both struck at K and fixed at time τ is

$$(R_s(\tau) - K)^+ - (K - R_s(\tau))^+ = R_s(\tau) - K \tag{B.4}$$

which is equal to that of being long a CMS swaplet ($+ R_s(\tau)$) and short K zero coupon bonds ($- K$) with maturity t_p. Therefore, we have

$$V_{\text{swaplet}}(0) = V_{\text{caplet}}(0) - V_{\text{floorlet}}(0) + D(0, t_p)\, K \tag{B.5}$$

Now, using (B.1) and (B.2) we can solve for today's price of the CMS swaplet. We start out by subtracting the value of the floorlet from that of the caplet[2]

$$\frac{D(0, t_p)}{L(t_0)} \left\{ (1 + h'(R_s(0)))\, C(R_s(0)) + \int_{R_s(0)}^{\infty} h''(x)\, C(x)\, dx \right\} -$$

$$\frac{D(0, t_p)}{L(0)} \left\{ (1 + h'(R_s(0)))\, P(R_s(0)) - \int_{R_s(0)}^{\infty} h''(x)\, P(x)\, dx \right\}$$

$$= \underbrace{\frac{D(0, t_p)}{L(0)} \left\{ (1 + h'(R_s(0)))\, (C(R_s(0)) - P(R_s(0))) \right\}}_{\text{I}} +$$

$$\underbrace{\frac{D(0, t_p)}{L(0)} \left\{ \int_{R_s(0)}^{\infty} h''(x)\, C(x)\, dx + \int_{-\infty}^{R_s(0)} h''(x)\, P(x)\, dx \right\}}_{\text{II}} \tag{B.6}$$

Focusing first on part I of (B.6), we note that $h'(R_s(0))$ evaluates to 0, leaving us with

$$\text{I:} \quad \frac{D(0, t_p)}{L(0)} \left\{ (C(R_s(0)) - P(R_s(0))) \right\}$$

$$= \frac{D(0, t_p)}{L(0)} \left\{ L(0)\, \mathbb{E}_0^{Q^L} [(R_s(\tau) - R_s(0))^+ - (R_s(0) - R_s(\tau))^+] \right\}$$

$$= D(0, t_p) \left\{ \mathbb{E}_0^{Q^L} [R_s(\tau) - R_s(0)] \right\} \tag{B.7}$$

The two swaptions were struck at $R_s(0)$, and thus, following (B.5), to solve for the value of the swaplet, we must add $D(0, t_p)\, R_s(0)$ to I and II from (B.6). This

---

[2]Both struck at $R_s(0)$.

leaves us with a final CMS swaplet value of

$$\underbrace{D(0, t_p) \left\{ \mathbb{E}_0^{Q^L}[R_s(\tau) - R_s(0)] \right\}}_{I} + D(0, t_p)\, R_s(0) + \text{II} \tag{B.8}$$

The $D(0, t_p)\, R_s(0)$ terms cancel out and the par swap rate is a martingale under the Q^L measure. This leaves us with our CMS swaplet replication formula

$$V_{\text{swaplet}}(0) = \frac{D(0, t_p)}{L(0)} \left\{ \int_{R_s(0)}^{\infty} h''(x)\, C(x)\, dx + \right.$$
$$\left. \int_{-\infty}^{R_s(0)} h''(x)\, P(x)\, dx \right\} + D(0, t_p)\, R_s(0) \tag{B.9}$$

showing us how to integrate over the second derivative of the $h(\cdot)$ function to find the correct notionals for the various strike prices which then in turn will yield the value of the CMS swaplet.

Appendix C

Selected code

In this appendix some selected bits of code will be presented. The code is written in the open source R software environment (R Development Core Team, 2010), which can be downloaded free of charge from *www.r-project.org*. The code bits will be organized, so that the bigger pieces are given first, and subsequently the smaller "helping functions" will be presented.

C.1   Fitting a SABR model

In this section we present two ways of calibrating a SABR model to market data. We show how to calibrate the model with freely varying parameters $\{\alpha, \rho, \nu\}$ and how to calibrate the SABR model with an α determined through ρ, *nu* and ATM volatility ($\{\alpha(\rho, \nu, \text{ATM.vol}), \rho, \nu\}$)) as shown in (5.7) on page 47.

C.1.1 Fitting a SABR model freely

```
#  Function that fits the SABR model based on:
#   beta (b), ATM forward (fATM), observed volatilities (ObsVols)
#   with corresponding strikes (strikes) and maturity of the
#   swaption (t).
#  All parameters (alpha, rho and nu) are free.
FitSABRFree <- function(beta, fATM, strikes, t, ObsVols,
  init.values = c(0.03, -0.3, 0.3),
  lower.bound = c(0.00, -1, 0.001),
  upper.bound = c(1, 1, 1), ...){

  # Defining function to be minimized:
  obj <- function(parm, beta, fATM, strikes, t, ObsVols){
    alpha <- parm[1]
    rho   <- parm[2]
    nu    <- parm[3]
    EstVols <- SABRVol(a=alpha, r=rho, v=nu,
      b=beta, f=fATM, K=strikes, t=t)

    return(sum((EstVols - ObsVols)^2))
  }

  opt <- nlminb(start=init.values,
    objective = obj,
    lower = lower.bound,
    upper = upper.bound,
  # Additional parameters for 'obj':
    beta = beta,
    fATM = fATM,
    strikes = strikes,
    t = t,
    ObsVols = ObsVols)

  # Printing information regarding optimization:
  cat("Error in convergence: ",opt$convergence, "\n")
  cat(opt$message, "\n")
  cat("Iterations: ",opt$iterations, "\n\n")

  # Storing optimization results:
  parms        <- opt$par
  names(parms) <- c("alpha", "rho", "nu")
  obj          <- opt$objective
```

```
estvols      <- SABRVol(a=parms[1], r=parms[2], v=parms[3],
                  b=beta, f=fATM, K=strikes, t=t)


# Returning results (in list form):
return(list(parms=parms, estvols=estvols, obsvols=ObsVols,
  obj=obj))
}
```

C.1.2 Fitting a SABR model with α implied by ρ, ν and ATM volatility

```
#  Function that fits the SABR model based on beta (b),
#  ATM forward (fATM), observed volatilities (ObsVols)
#   with corresponding strikes (strikes) and the
#   maturity of the swaption (t).
#  (Alpha is implied through rho, nu and ATM volatility.)
FitSABRImpA <- function(beta, fATM, volATM, strikes, t, ObsVols,
  init.values = c(-0.3, 0.3),
  lower.bound = c(-1, 0.001),
  upper.bound = c(1, 1), ...){
  obj <- function(parm, beta, fATM, volATM, strikes, t, ObsVols){
    rho    <- parm[1]
    nu     <- parm[2]
    # Determining alpha based on rho, nu and ATM vol:
    alpha <- alpha0(b=beta, r=rho, v=nu, vol=volATM, f=fATM, t=t)
    EstVols <- SABRVol(a=alpha, r=rho, v=nu, b=beta, f=fATM,
      K=strikes, t=t)

    return(sum((EstVols - ObsVols)^2))
  }

  opt <- nlminb(start=init.values,
    objective = obj,
    lower = lower.bound,
    upper = upper.bound,
    # Additional parameters for 'obj':
    beta = beta,
    fATM = fATM,
    volATM = volATM,
    strikes = strikes,
    t = t,
    ObsVols = ObsVols
  )
  cat("Error in convergence: ", opt$convergence, "\n")
  cat(opt$message, "\n")
  cat("Iterations: ", opt$iterations, "\n\n")

  parms <- c(alpha0(b=beta, r=opt$par[1], v=opt$par[2],
    vol=volATM, f=fATM, t=t),opt$par)
  names(parms) <- c("alpha", "rho", "nu")
  obj <- opt$objective
```

91

```
estvols <- SABRVol(a=parms[1], r=parms[2], v=parms[3],
  b=beta, f=fATM, K=strikes, t=t)

# Returning results (in list form):
return(list(parms=parms, estvols=estvols, obsvols=ObsVols, obj=obj))
}
```

C.2 Additional helping functions

In this section we present some additional functions that occur in the fitting of the SABR model.

C.2.1 SABR volatility [SABRVol()]

```
# SABR volatility (eq. 2.17a in Hagan et al. [2002])
# based on: alpha (a), rho (r), nu (v), beta (b),
# ATM forward (f), strike (K) and expiry (t).
SABRVol <- function(a, r, v, b, f, K, t){

  # Defining "building blocks" for the bigger equation:
  x <- function(z,r){
    log((sqrt(1 - 2*r*z + z^2) + z - r)/(1 - r))
  }

  z <- v/a * (f*K)^((1 - b)/2)*log(f/K)

  Denom <- (f*K)^((1 - b)/2)*
    (1 +
      (1 - b)^2/24*log(f/K)^2 +    # 0 if f==K
      (1 - b)^4/1920 *log(f/K)^4   # 0 if f==K
    )
  Term3 <- (1 + ((1 - b)^2/24*a^2/((f*K)^(1 - b)) +
            (1/4)*(r*b*v*a)/((f*K)^((1 - b)/2)) +
            (2 - 3*r^2)/24*v^2)*t
          )

  # If f==K then z/x(z,r) = 0/0 => set equal to 1.
  # Term2 <- if(f==K){
  #   1
  #   }else{
  #   z/x(z,r)
  #   }

  # Rewritten to work with vectors of strikes (K's):
  Term2 <- rowSums(cbind(is.na(z/x(z,r)),
    (1 - is.na(z/x(z,r)))*(z/x(z,r))), na.rm=TRUE)

  # Putting the pieces together and returning the result
  return( (a/Denom) * Term2 * Term3 )
}
```

C.2.2 Determining α from ρ, ν and ATM volatility [alpha0()]

```
alpha0 <- function(b, r, v, vol, f, t){
  # Find root of CubPoly in [0,1]:
  uniroot(CubPoly, interval=c(0,1),
    # Additional parameters for 'CubPoly':
    b=b, r=r, v=v, vol=vol, fwd=f, t=t)$root
}
```

C.2.3 Setting up the cubic [CubPoly()]

```
CubPoly <- function(a, b, r, v, vol, fwd, t){
  # Polynomial coefficients:
  A <- ((1 - b)^2*t)/(24*fwd^(2 - 2*b))
  B <- (r*b*v*t)/(4*fwd^(1 - b))
  C <- 1 + ((2 - 3*r^2)/24)*v^2*t

  return(A*a^3 + B*a^2 + C*a - vol*fwd^(1 - b))
}
```

Bibliography

Bartlett, B. (2006). Hedging under SABR model. *Wilmott Magazine*, July/August:2–4.

Beckers, S. (1980). The constant elasticity of variance model and its implications for option pricing. *Journal of Finance*, 35(3):661–673.

Björk, T. (2004). *Arbitrage Theory in Continuous Time Finance*. Oxford University Press, 2nd edition.

Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, 3:167–179.

Black, F. and Scholes, M. (1973). The pricing of options and corporate liablitities. *Journal of Political Economy*, 81(3):637–654.

Carr, P. (2005). Hedging with options (presentation). Presented at Columbia University, New York NY.

Cox, J., Ingersoll, J., and Ross, S. (1985). A theory of the term structure of interest rates. *Econometrica*, 53(2):385–407.

Cox, J., Ross, S. A., and Rubinstein, M. (1979). Option pricing: A simplified approach. *Journal of Financial Economics*, 7:229–263.

Derman, E. and Kani, I. (1994). Riding on a smile. *Risk Magazine*, pages 32–39.

Dumas, B., Fleming, J., and Whaley, R. (1998). Implied volatility functions: Empirical tests. *Journal of Finance*, L3(6):2059–2106.

Fuglede, B., Grubb, G., and Gutmann Madsen, T. (1999). *Analyse og optimering*. Universitetsbogladen København.

Fujii, M., Shimada, Y., and Takahashi, A. (2010). A Note on Construction of Multiple Swap Curves with and without Collateral. *SSRN eLibrary*.

Gatheral, J. (2002). Stochastic volatility and local volatility. *Case Studies in Financial Modeling Course Notes.*

Hagan, P. (2003). Convexity conundrums: Pricing CMS swaps, caps and floors. *Wilmott Magazine*, pages 38–44.

Hagan, P., Kumar, D., Lesniewski, A., and Woodward, D. (2002). Managing smile risk. *Wilmott Magazine*, pages 84–108.

Hagan, P. and Woodward, D. (1999). Equivalent Black volatilities. *Applied Mathematical Finance*, (6):147–157.

McDonald, R. (2006). *Derivatives Markets.* Pearson, 2nd edition.

Mercurio, F. and Pallavicini, A. (2006). Smiling at convexity: Bridging swaptions, skews and CMS adjustments. *Risk Magazine*, August:64–69.

R Development Core Team (2010). *R: A Language and Environment for Statistical Computing.* R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.

Schrager, D. and Pelsser, A. (2006). Pricing swaptions and coupon bond options in affine term structure models. *Mathematical Finance*, 16(4):673–694.

Sundaresan, S. (2009). *Fixed Income Markets and Their Derivatives.* Academic Press, 3rd edition.

Sydsæter, K. (2005). *Matematisk Analyse, bind 1.* Gyldendal, 7th edition.

West, G. (2005). Calibration of the SABR model in illiquid markets. *Applied Mathematical Finance*, 12(4):371–385.