CBS COPENHAGEN BUSINESS SCHOOL
HANDELSHØJSKOLEN

The SABR Model:
Modelling the Volatility Smile

Annika Abildgaard Lund (127007)

Master's Thesis
MSc in Economics and Business Administration, Finance and Investments

Department of Finance
Copenhagen Business School

15th of May 2023

Supervisor: Anders Bjerre Trolle

163,200 characters *72 normal pages*

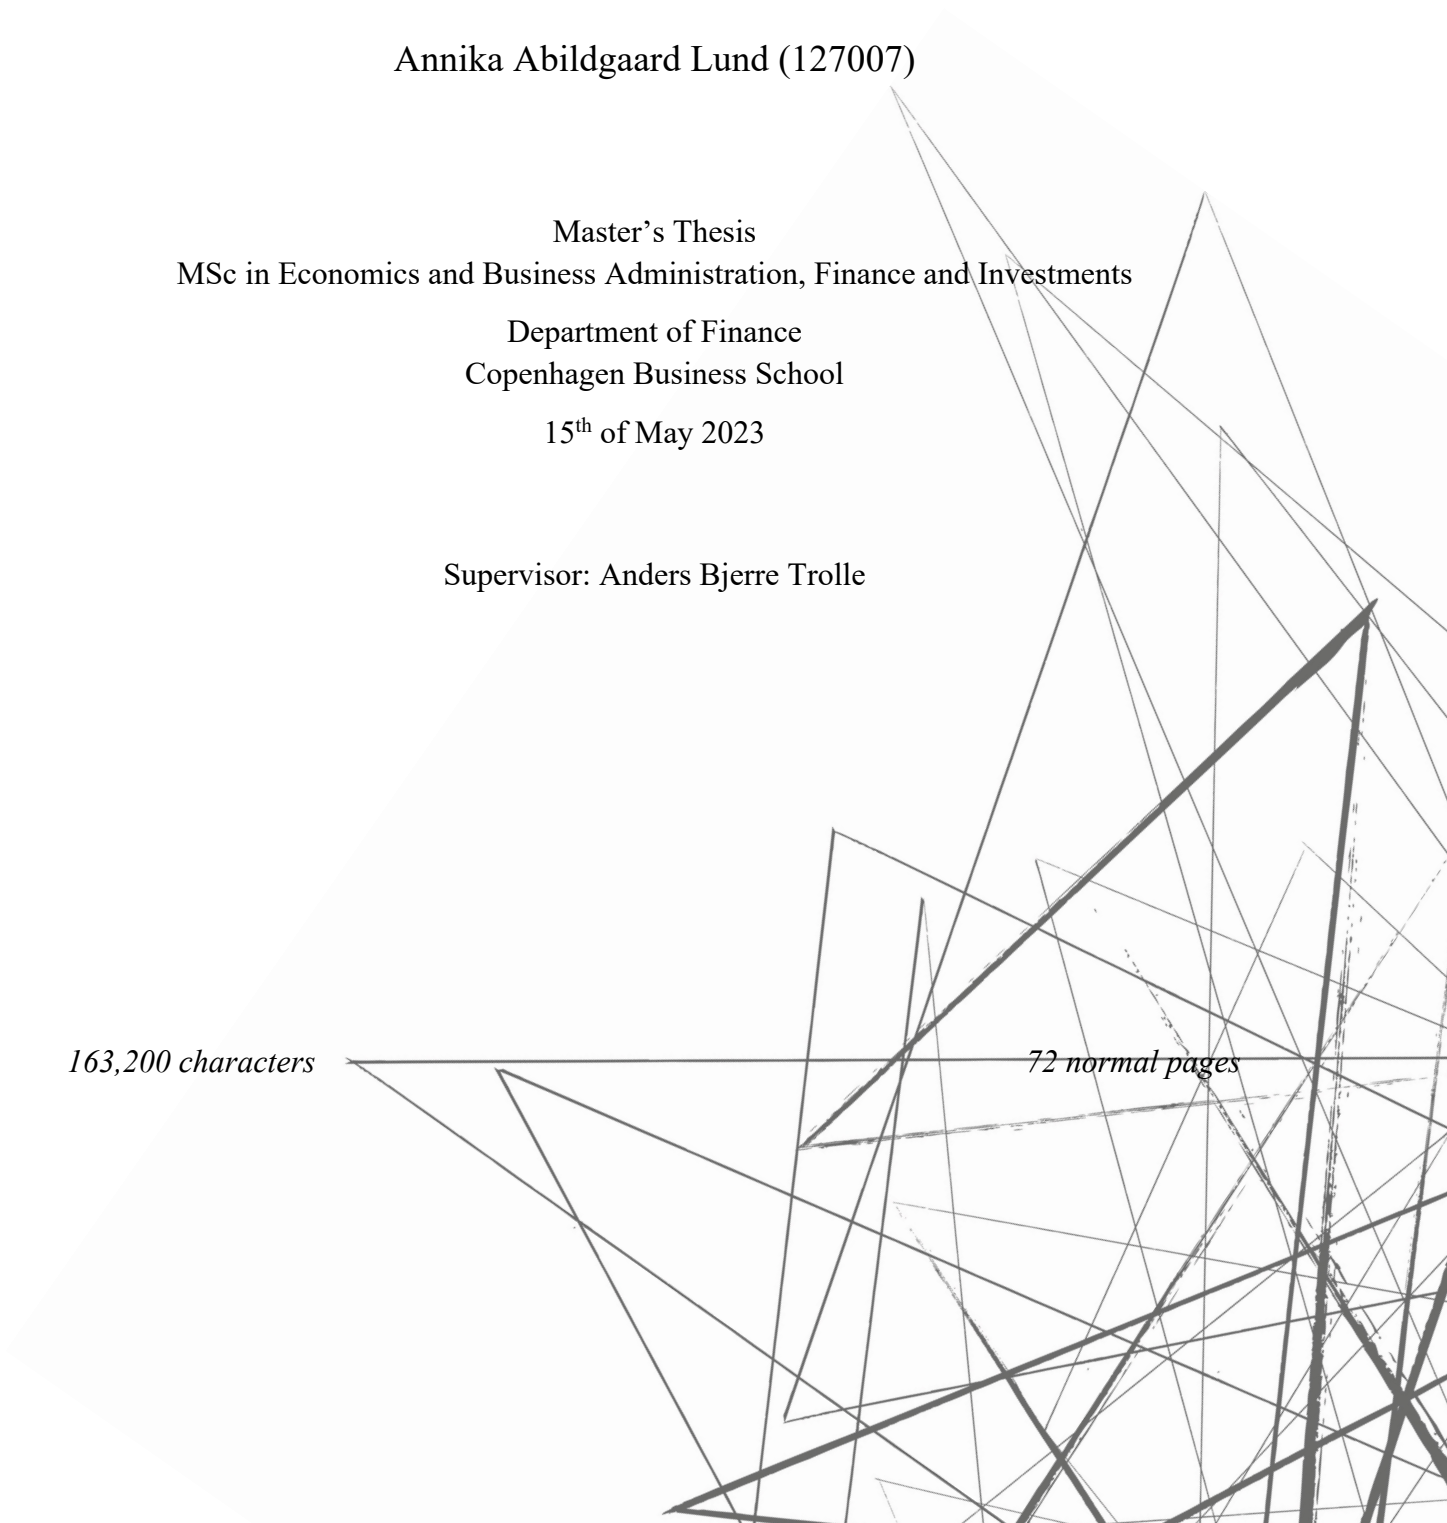

ABSTRACT

This thesis conducts an empirical analysis of the ability of the shifted SABR model, the normal SABR model, and free-boundary SABR model to manage risks in interest rate markets across different market conditions. We place a heavy focus on model calibration, parameter stability, and risk sensitivities to ensure the models are practically viable, and generally take a practitioner approach to the analysis. We calibrate each model using both Antonov's exact solution and one of Hagans approximation formulas. We do this across a wide range of model parameterisations to determine the optimal fit to our data. The Antonov normal SABR was quickly discarded due to calibration issues, while the free-boundary SABR was discarded due to a stickiness at zero making it unable to accurately model negative rates. In general, calibrating the model to prices is found to be impractical due to a very long calibration time relative to the fast and relatively accurate Hagan approximation. The normal SABR model calibrated in a normal calibration space using the Hagan approximation generally gives the most stable parameters. Still, the overarching conclusion is that a shifted SABR with $s = 5\%$ and $\beta = 1$ calibrated in a shifted Black calibration space using the Hagan approximation has the best performance both in-sample and out-of-sample, while risk measures also conveniently can be calculated. The main limitation of this model is a relatively higher pricing error for deep OTM and short-term options. In addition, we observe a slight decrease in performance for the model when rates become negative, but overall, the we conclude that it is the best suited model for empirically modelling the volatility smile across different interest rate environments.

CONTENTS

1 INTRODUCTION

The purpose of this thesis is to examine how the SABR model can be used to manage risks in interest rate markets across different market conditions. The original Black framework for option pricing assumes a constant volatility across strikes. Empirically, volatility varies along different strikes giving rise to a *volatility smile*, and modelling the dynamics of the smiles is crucial for pricing and risk management purposes in interest rate derivative markets. The SABR model introduced by Hagan et al. (2002) is a stochastic volatility model which has long been a popular tool among practitioners in the financial industry. The model is primarily used for volatilities interpolation and calculating of risks. The dynamics model can be expressed in an asymptotic closed-form formula allowing fast calibration, while producing volatility estimates which conveniently can be used in conjunction with existing pricing models. However, under certain conditions, the asymptotic solution is known to break down, and this warrants a careful understanding of the model limitations. The original SABR model has also been challenged by negative interest rates, as its usual specification is unable to model these.

We conduct an empirical analysis of the ability of certain SABR model extensions to capture volatility smiles under different market conditions such as negative interest rates and stress environments. The models considered for this purpose are the shifted SABR model, the normal SABR model, and free-boundary SABR model, all of which theoretically should be able to model negative and positive rates. Throughout the analysis, we focus on model calibration, parameter stability, and risk sensitivities to ensure the models are practically viable. We further opt to test a wide range of different parameterisations and calibration procedures for each of the SABR extensions. The research question is stated as:

> *Which specification of the SABR model is best suited for empirically modelling the volatility smile across different interest rate environments, and what are the practical limitations of the model?*

Our evaluation of model suitability for modelling our data is based on considerations of (i) the quality of fit across different market conditions, (ii) the complexity of implementing and calibrating the model, (iii) whether the parameters are consistent, stable, and intuitive, and finally (iv) the risk management capabilities under the model. The specific context of the empirical analysis is to fit the SABR models to a broad sample of European Swaptions written on 6M Euribor swaps across March 2019 to March 2023. This period includes observations of negative rates, stress scenarios, and rapidly increasing rates, and thus meets the criteria of covering different market conditions. Our exclusive focus on the EUR market and swaptions is a delimitation of the thesis. We remark that our aim is to provide an empirical perspective on already established models, which means derivation of the models is not considered.

1.1 Structure of the Thesis

The thesis is structured into four distinct parts. Part I establishes the theoretical background needed to understand the examined models of the thesis. This includes introducing basic mathematical framework (section 2), accounting for interest rate concepts (section 3), and providing context to swaps and swaption and their valuation under the classical pricing models (section 4). Part II concerns volatility models, and firstly motivates their use (section 5) and provides the framework of the original SABR model (section 6). The extensions of the SABR model which accommodate negative interest rates in the form of the shifted SABR model, the normal SABR model, and the free-boundary SABR model are introduced (section 7) in addition to an overview of risk management under SABR dynamics (section 8). The main matter of the thesis is the empirical analysis in Part III, which firstly involves data description (section 9) and specifying the empirical methodology (section 10). The analysis itself covers the assessment of model parameterisations and calibration methods, parameter stability, in-sample and out-of-sample performance, and empirical risk sensitivities (section 11). The results are wrapped up in Part IV with a discussion of the findings, their limitations, and how the analysis could be improved and extended (section 12). Finally, the conclusion of the research question is presented (section 13).

PART I

THEORETICAL BACKGROUND

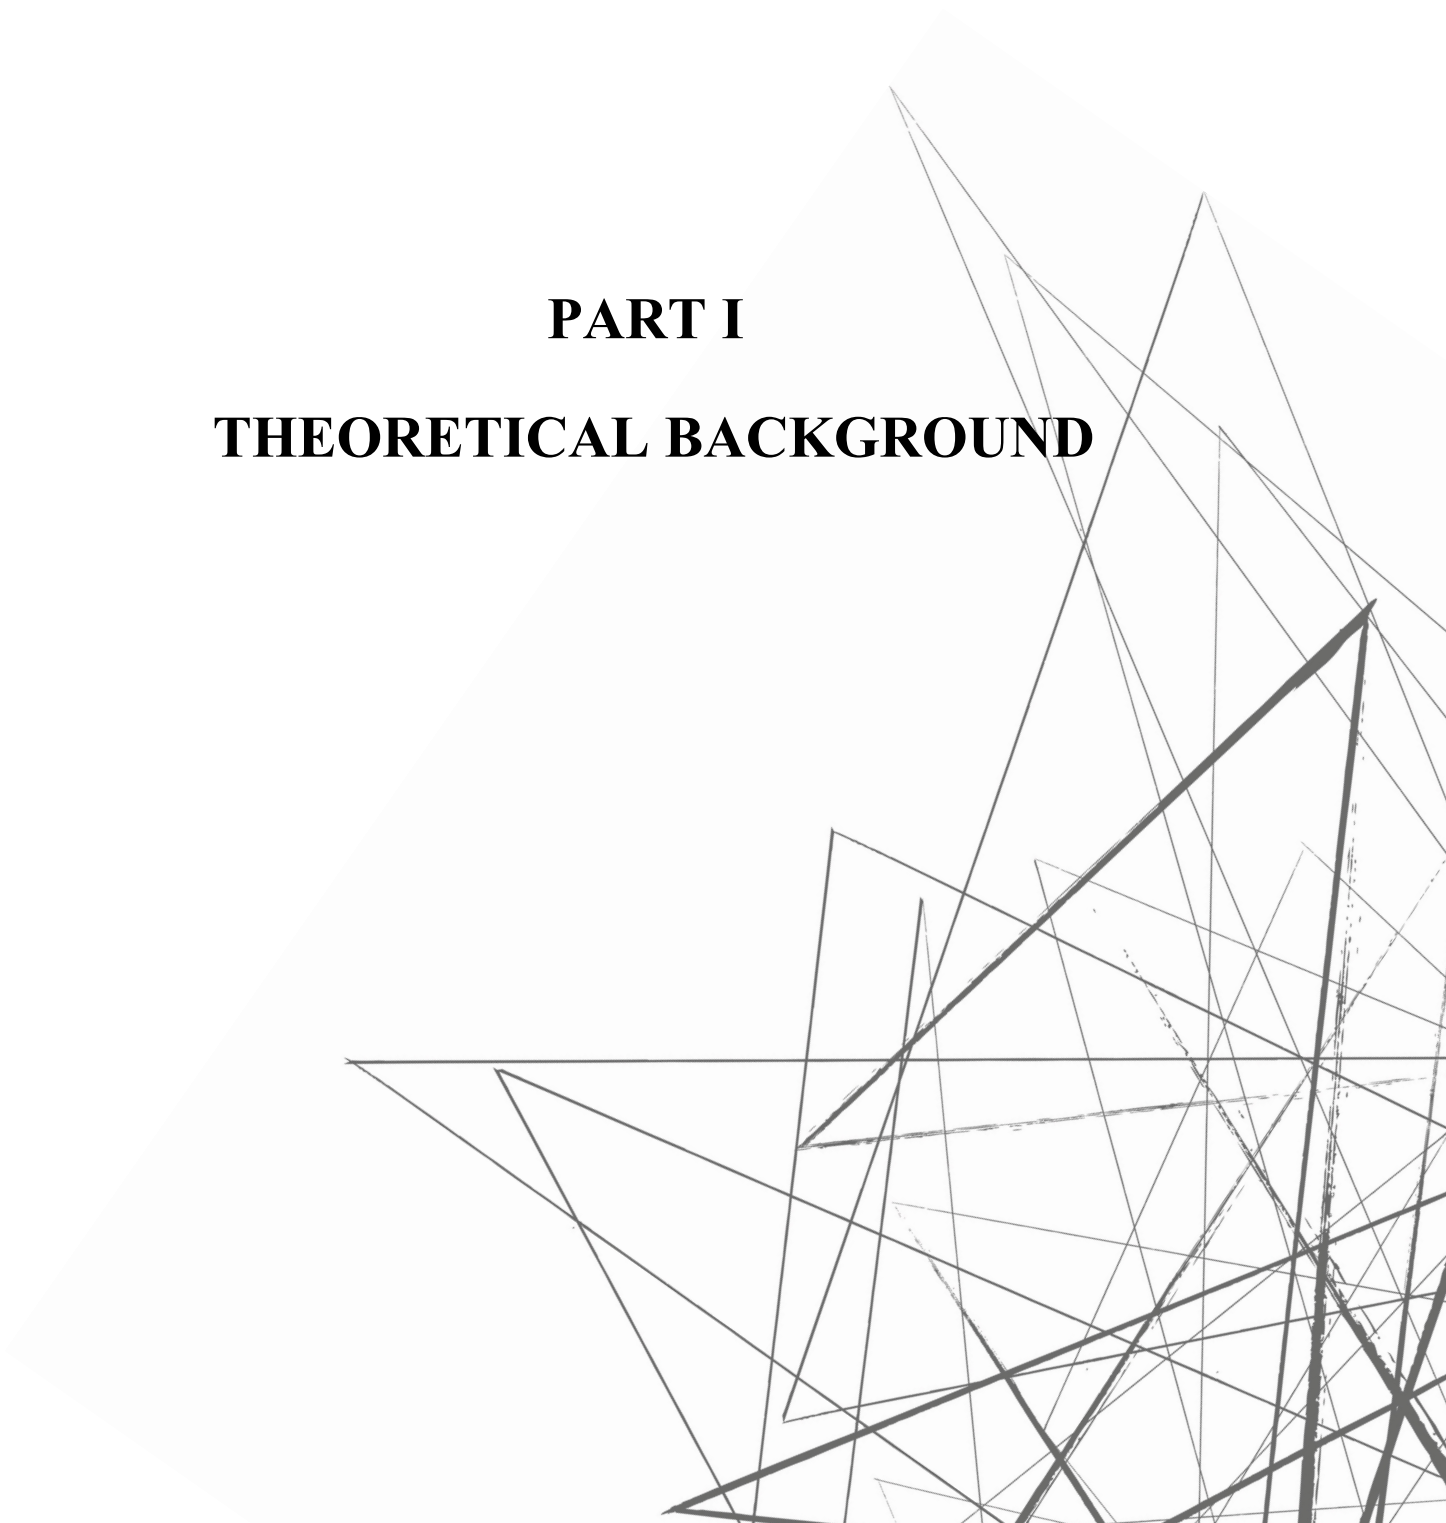

2 MATHEMATICAL FRAMEWORK

This section briefly introduces the mathematical concepts which forms the basis of valuing financial instruments and will be referenced throughout the thesis. No-arbitrage pricing is the most important notion and understanding this requires an introduction of martingales and probability measures.

Arbitrage. An arbitrage is a strategy which yields a risk-free profit, i.e., a portfolio you can buy a zero cost and later receive a positive payoff with a positive probability.

Money-market account, $M(t)$. Let $M(t)$ be the balance of a money-market account at time $t \geq 0$. The interest accrues on a continuous basis, so the time-t value of investing a unit amount at time 0 is:

$$M(t) = \exp\left(\int_0^t r_s \, ds\right) \tag{2.1}$$

Discount factor, $D(t,T)$. The discount factor $D(t,T)$ is the time-t value of a one unit of currency cash flow at time T. With stochastic interest rates, the discount factor itself is a stochastic process decreasing continuously at the instantaneous rate r_t (Brigo & Mercurio, 2006):

$$D(t,T) = \frac{M(t)}{M(T)} = \exp\left(-\int_t^T r_s \, ds\right) \tag{2.2}$$

Zero-coupon bond, $P(t,T)$. A zero-coupon bond (ZCB) is a contract paying no coupons and guaranteeing a payment of one unit of currency at maturity T. Let $P(t,T)$ be the time-t price for a ZCB maturing at time T. As the bond is risk-free, we must have $P(T,T) = 1$ (Brigo & Mercurio, 2006).

Wiener process, $W(t)$. A Wiener process, or a standard Brownian motion, is a continuous-time stochastic process commonly used to describe random movements in financial models, e.g., of interest rates. Mathematically, a Wiener process $W(t)$ satisfies that for any $0 < s < t < u$ and any $h > 0$:

- *Independent increments:* $W(u) - W(t)$ is independent of $W(t) - W(s)$
- *Stationary increments:* $W(t+h) - W(s+h) \sim W(t) - W(s)$
- *Gaussian increments:* $W(t) - W(s) \sim \mathcal{N}(0, t-s)$

This implies $W(t)$ follows a Markov process, so its conditional probability distribution of future states depends only on the present (Oosterlee & Grzelak, 2019). Further, $W(t)$ fulfils the martingale property.

Martingale property. A martingale is a zero-drift stochastic process with the property that its expected value is equal to its current value. Formally, a continuous process $V(t)$ is a martingale w.r.t. the filtration $\mathcal{F}_t$ (all information available at time t) under a probability measure $\mathbb{Q}$, if for all $t < \infty$ and $t < T$:

$$\mathbb{E}[|V(t)|] < \infty \qquad \text{and} \qquad \mathbb{E}[\, V(T) \mid \mathcal{F}_t \,] = V(t) \tag{2.3}$$

where $\mathbb{E}$ is the conditional expectation under measure $\mathbb{Q}$ (Oosterlee & Grzelak, 2019).

Fundamental theorem of arbitrage pricing. A probability measure $\mathbb{Q}$ under which the martingale property holds is an equivalent martingale measure. The first fundamental theorem of arbitrage pricing states that a financial model is free from arbitrage if there exists a martingale measure. In the absence of arbitrage, the present value of a claim is equal to the conditional expectation of its discounted future value under the $\mathbb{Q}$-measure with respect to the filtration (Björk, 2020).

Change of numeraire. Under stochastic interest rates, a change of numeraire in the no-arbitrage pricing theory is helpful. A numeraire is a tradeable asset with price process $N(t)$, so for a martingale measure $\mathbb{Q}^N$ associated with the numeraire, the arbitrage-free price $V(t)$ of a security is given:

$$\frac{V(t)}{N(t)} = \mathbb{E}^{\mathbb{Q}^N}\left\{\frac{V(T)}{N(T)} \,\middle|\, \mathcal{F}_t\right\} \tag{2.4}$$

This means that the relative (to the numeraire) price process of any asset is a martingale under the measure associated with that numeraire. This gives us the general pricing formula above, as the time-t risk-neutral price is not impacted by a change of numeraire (Brigo & Mercurio, 2006).

Probability measures. We can find the value $V(t)$ of a security under different measures, where measure refers to the unit in which we value the security (the numeraire) (Oosterlee & Grzelak, 2019).

- The **risk-neutral measure** (spot measure) $\mathbb{Q}$ has the money market account $M(t)$ as numeraire. With $M(t) = 1$:

$$V(t) = M(t)\mathbb{E}^{\mathbb{Q}}\left[\frac{V(T)}{M(T)} \,\middle|\, \mathcal{F}_t\right] = \mathbb{E}^{\mathbb{Q}}\left[\frac{V(T)}{M(T)} \,\middle|\, \mathcal{F}_t\right] \tag{2.5}$$

- The T-**forward measure** $\mathbb{Q}^T$ has the ZCB price $P(t,T)$ as numeraire. With $P(T,T) = 1$, we have:

$$V(t) = P(t,T)\mathbb{E}^{\mathbb{Q}^T}\left[\frac{V(T)}{P(T,T)} \,\middle|\, \mathcal{F}_t\right] = P(t,T)\mathbb{E}^{\mathbb{Q}^T}[\,V(T) \mid \mathcal{F}_t\,] \tag{2.6}$$

- The **swap measure** $\mathbb{Q}^A$ has the annuity factor $A_{\alpha,\beta}(t)$ defined in equation (4.2) as numeraire:

$$V(t) = A_{\alpha,\beta}(t)\mathbb{E}^{\mathbb{Q}^A}\left[\frac{V(T)}{A_{\alpha,\beta}(T)} \,\middle|\, \mathcal{F}_t\right] \tag{2.7}$$

Martingale representation theorem. A martingale can be written as a stochastic integral with respect to the underlying Wiener process $W(t)$, i.e.:

$$dF(t) = C(\cdot)\, dW(t) \tag{2.8}$$

where $F(t)$ is a martingale process under $\mathcal{F}_t$. The uncertainty of the martingale process $F(t)$ arises from the Wiener process multiplied with a (stochastic) process $C(\cdot)$. The process $C(\cdot)$ must be defined by a model in way so the process is a martingale (Björk, 2020).

3 INTEREST RATES

At its core, an interest rate represents the cost of money. The following sections define common reference rates used in derivative contracts, provide context on negative interest rates, and finally introduce basic interest rate concepts needed for swap and swaption valuation.

3.1 IBOR and IBOR Transition

The historically most common reference rates are *interbank offered rates* (IBORs), which is a series of interest rate benchmarks intended to illustrate the average rate at which major banks can borrow unsecured funds in the interbank market. The rates are published daily in an array of different currencies and maturities by various organisations, and the calculation is based on a truncated average of interest rates submissions from a panel of banks (Hull, 2018). As an example, the EURIBOR reflects the average rate at which European banks can borrow in the interbank market denominated in EUR. Other common IBORs are the GBP and USD LIBOR, CIBOR (DKK), and TIBOR (JPY).

IBOR transition. In recent years, IBORs have been undergoing change. The interbank market is less liquid since the financial crises, especially in tenors longer than overnight. In addition, there has been concerns regarding certain IBORs, where most notably manipulation of LIBOR submissions was revealed in 2012 (Nordea, n.d.). Combined this has tiggered a global reform to replace IBOR with alternative overnight risk-free reference rates showing the rate paid on unsecured overnight borrowings. The rates include SOFR (USD), ESTR (EUR), SONIA (GBP), and TONA (JPY) (Bloomberg Markets, 2020). At the time of writing, there are no plans to out phase the CIBOR or EURIBOR. From here on, when we refer to an IBOR, it may refer to any interbank rate used as a reference in financial derivatives.

3.2 Implications of Negative Interest Rates

Conventional monetary policy involves central banks fixing their short-term interest rate to influence economic activity to achieve certain goals, e.g., inflation level. Increasing the rate makes borrowing more expensive which dampens the economy, while decreasing rates achieve the opposite. In the wake of the financial crisis in 2008, conventional easing measures were rendered exhausted, and in 2014 the ECB introduced the first negative policy rate in newer times in a European context (ECON, 2021). The policy rate defines the interest rate charged on overnight deposits with the ECB, and the negative rate policy eventually carried over to cause reference rates and yield curves to be negative. In the context of derivative pricing, rates were generally assumed to have a zero lower bound before this. The empirical observation of negative rates violated the assumptions of traditional pricing models rendering them useless and revealing arbitrage in prices. Since then, ways of handling negative rates have been

conceptualised, and the zero lower bound assumption has been replaced by an *effective lower bound* assumption. The ECB returned to a positive deposit facility rate in September 2022 as a response to the increasing inflation after the Covid-crisis. Thus, for now negative rates are a feature of the past in the EUR market, but the pricing capabilities under negative rates remain an important consideration when choosing a pricing model and assessing its robustness.

3.3 Interest Rate Concepts

We now concern ourselves with basic interest rate concepts, which lie at the core of valuing swaps and swaptions. Firstly, a *day count convention* determines how interest accrues over time, which determine how we calculate the year fraction, or *coverage*, between two payment dates. Let $\tau(t,T)$ be the coverage between dates t and T, and assume we have two dates $D_1 = (d_1, m_1, y_1)$ and $D_2 = (d_2, m_2, y_2)$ between which we want to calculate the coverage. The most common day count conventions for doing this are introduced here (Brigo & Mercurio, 2006):

- *Actual/365*. A year is assumed to have 365 days, so the coverage becomes $(D_2 - D_1)/365$
- *Actual/360*. A year is assumed to be 360 days, so the coverage becomes $(D_2 - D_1)/360$
- *30/360*. A month is assumed to be 30 days and a year to be 360 days, so the coverage becomes:

$$\frac{360 \times (y_2 - y_1) + 30 \times (m_2 - m_1) + (d_2 - d_1)}{360}$$

The next concept is *compounding*, which defines how an interest rate is measured, that is, how interest is earned on interest. The most common ways this is done is using (i) simple compounding, where interest is accrued proportionally to the time of investment, (ii) annual compounding, where interest is reinvested once a year, and (iii) continuous compounding, where interest is continuously reinvested (Brigo & Mercurio, 2006). IBORs are typically simply-compounded rates, so we focus on this method.

The *spot rate* is the constant rate at which an investment of $P(t,T)$ units of currency at time t should be made to produce an amount of one unit of currency at maturity T. Let $L(t,T)$ be the simple spot rate prevailing at time t for maturity T, then:

$$L(t,T) = \frac{1 - P(t,T)}{\tau(t,T)\,P(t,T)} \tag{3.1}$$

Given spot rates $L(t,T)$ for different maturities, we can construct the *zero curve* (term structure of interest rates), effectively mapping maturities into rates at time t. Rearranging equation (3.1), we have:

$$P(t,T) = \frac{1}{1 + L(t,T)\tau(t,T)} \tag{3.2}$$

Plotting the ZCB bond price $P(t,T)$ for different maturities gives us the *discount curve*, which is a (usually) a decreasing curve starting from $P(t,t) = 1$ (Brigo & Mercurio, 2006). Spot rates can be

contrasted to *forward rates*, which are interest rates that can be locked in today for a future period. Let $F(t;S,T)$ be the simple forward rate prevailing at time t for a future period from time S to T. Then:

$$F(t;S,T) = \frac{1}{\tau(S,T)}\left(\frac{P(t,S)}{P(t,T)} - 1\right) \tag{3.3}$$

It can be shown that the simple forward rate $F(t;S,T)$ is a martingale under the T-forward measure $\mathbb{Q}^T$, which means it can be presented via the martingale representation theorem. Under the $\mathbb{Q}^T$ measure the forward rate $F(t;S,T)$ is the expectation of the future simple spot rate $L(t,T)$, so for $0 \leq t < S < T$:

$$\mathbb{E}^{\mathbb{Q}^T}[\, L(S,T) \mid \mathcal{F}_t \,] = F(t;S,T) \tag{3.4}$$

As we will see in the following, this result is crucial for valuing interest rate swaps and swaptions. In fact, by assuming forward rates are realised we can utilise the fundamental theorem of arbitrage pricing.

4 INTEREST RATE DERIVATIVES

An interest rate derivative is a financial instrument whose payoff depends on an underlying interest rate. Their typical purpose is to manage and hedge interest rate risk, but they may also be used for speculative purposes (Hull, 2018). In 2022, the BIS estimated the gross market value of the global over-the-counter (OTC) interest rate derivative market to be approximately $11.8 trillion, making it by far the largest derivative markets. The most traded interest rate derivative are interest rate swaps followed by interest rate options, which include caps, floors, and swaptions (BIS, 2022). Finding robust procedures for pricing interest rate derivatives is difficult, as the behaviour of rate is complicated to model, and the valuation of many products requires a model describing the behaviour of the entire zero-coupon curve (Hull, 2018). In the following, we consider the traditional framework for valuing interest swaps and swaptions.

4.1 Interest Rate Swaps

A swap is an OTC agreement between two parties to exchange cash flows in the future. In an *interest rate swap* (IRS), two parties agree to pay each other cash flows equal to the interest on a notional amount at predetermined future dates for a predetermined number of years. The most common type of swap is a plain vanilla IRS, where a series of fixed-rate payments are exchanged for a series of floating-rate payments, with the floating rate commonly being linked to an IBOR (Oosterlee & Grzelak, 2019). This is the type of swap we consider henceforth. We distinguish between a position in a *payer swap* (PS), where one receives a floating rate and pays a fixed rate, and a *receiver swap* (RS), where one pays a floating rate and receives a fixed rate. A position in a PS may be used to transform a floating-rate loan into a fixed-rate loan, and vice versa for a RS, illustrating the hedging capabilities of an IRS.

The mechanics of the swap are as follows. At each payment date T_i, the fixed leg has a cash flow of $N\tau_i K$, where K is the fixed rate, N is the notional of the swap, and τ_i is the time between payments. The floating leg has a cash flow of $N\tau_i L(T_{i-1}, T_i)$, where $L(T_{i-1}, T_i)$ is the floating rate for the period $[T_{i-1}, T_i]$, which resets at time T_{i-1} (Oosterlee & Grzelak, 2019). In other words, the floating rate is set before and paid at the end of each accrual period. The legs of the swap may have different payment frequencies depending on currency market conventions and the terms of the specific contract. In addition, the legs may have different day count conventions, further complicating the payment schedule.

4.1.1 Valuation of Interest Rate Swaps

Interest rate swaps may be valued using the result that in absence of arbitrage, the price of a claim is equal to the expectation of its discounted future value under the $\mathbb{Q}$-measure. In the following, we consider a plain vanilla payer swap (PS), but the corresponding values for a receiver swap may be found

by taking the additive inverse. We assume payments occur at $T_{\alpha+1}, \dots, T_\beta$, and the floating rate resets at $T_\alpha, \dots, T_{\beta-1}$. The payment frequencies of the legs may differ, so we define a set of coverages for the floating and fixed leg as $\tau_{\alpha+1}^{\text{float}}, \dots, \tau_\beta^{\text{float}}$ and $\tau_{\alpha+1}^{\text{fix}}, \dots, \tau_\beta^{\text{fix}}$, respectively. As before, let K be the fixed rate and $L(T_{i-1}, T_i)$ be the floating spot rate. For ease of notation, we assume the notional is $N = 1$, and omit it this term. The **payoff** of a PS is the floating leg payments subtracted the fixed leg payments:

$$V^{PS}\big(T_{\alpha+1}, \dots, T_\beta\big) = \sum_{i=\alpha+1}^{\beta} \tau_i^{\text{float}} L(T_{i-1}, T_i) - \sum_{i=\alpha+1}^{\beta} \tau_i^{\text{fix}} K \tag{4.1}$$

Using the money-market account $M(t)$ as numeraire, the expectation of the discounted future cash flows under the equivalent martingale measure $\mathbb{Q}$ is given from equation (2.5). Assuming $M(t) = 1$, the value of the swap at time t under this measure is:

$$V^{\text{PS}}(t) = \sum_{i=\alpha+1}^{\beta} \tau_i^{\text{float}}\, \mathbb{E}^{\mathbb{Q}}\left[\frac{L(T_{i-1}, T_i)}{M(T_i)} \,\middle|\, \mathcal{F}_t \right] - \sum_{i=\alpha+1}^{\beta} \tau_i^{\text{fix}}\, \mathbb{E}^{\mathbb{Q}}\left[\frac{K}{M(T_i)} \,\middle|\, \mathcal{F}_t \right]$$

We can change the numeraire to the ZCB price $P(t,T)$, effectively changing from the $\mathbb{Q}$-measure to the T-forward measure $\mathbb{Q}^T$. According to equation (3.4), the forward rate $F(t; T_{i-1}, T_i)$ is the expectation of the future spot rate $L(T_{i-1}, T_i)$ under this measure. Thus, applying equation (2.6) we obtain following:

$$\begin{aligned} V^{\text{PS}}(t) &= \sum_{i=\alpha+1}^{\beta} P(t,T_i)\tau_i^{\text{float}}\, \mathbb{E}^{\mathbb{Q}^T}[\, L(T_{i-1}, T_i) \mid \mathcal{F}_t \,] - \sum_{i=\alpha+1}^{\beta} P(t,T_i)\tau_i^{\text{fix}}\, \mathbb{E}^{\mathbb{Q}^T}[\, K \mid \mathcal{F}_t \,] \\ &= \sum_{i=\alpha+1}^{\beta} P(t,T_i)\tau_i^{\text{float}} F(t; T_{i-1}, T_i) - \sum_{i=\alpha+1}^{\beta} P(t,T_i)\tau_i^{\text{fix}} K \end{aligned}$$

To simplify the notation, we define an annuity factor as a linear combination of ZCBs:

$$A_{\alpha,\beta}(t) = \sum_{i=\alpha+1}^{\beta} \tau_i^{\text{fix}} P(t,T_i) \tag{4.2}$$

Given the annuity factor, the expression for the forward rate in equation (3.3), and the notion of telescopic summation (see Oosterlee & Grzelak (2019) for explanation), the **price** of a payer swap is:

$$V^{\text{PS}}(t) = \Big(P(t,T_\alpha) - P\big(t,T_\beta\big)\Big) - A_{\alpha,\beta}(t)K \tag{4.3}$$

The fixed rate is set, so the IRS has a present value of zero at inception. The rate fulfilling this is the swap rate $S_{\alpha,\beta}(t)$, which is found by setting equation (4.3) equal to zero and solving:

$$S_{\alpha,\beta}(t) = \frac{\sum_{i=\alpha+1}^{\beta} \tau_i^{\text{float}} F(t; T_{i-1}, T_i) P(t,T_i)}{\sum_{i=\alpha+1}^{\beta} \tau_i^{\text{fix}} P(t,T_i)} = \frac{P(t,T_\alpha) - P\big(t,T_\beta\big)}{A_{\alpha,\beta}(t)} \tag{4.4}$$

It can be shown that the swap rate is a martingale under the swap measure. With this defined, we can express the present value of a payer swap in a more convenient way:

$$V^{\mathrm{PS}}(t) = A_{\alpha,\beta}(t)\,(S_{\alpha,\beta}(t) - K) \tag{4.5}$$

Thus, the swap value is proportional to the difference between the swap rate and the fixed rate. The concise formula also makes it easy to infer the value: if you pay a fixed rate higher than the time-t swap rate $S_{\alpha,\beta}(t)$, the position has a negative value, and vice versa.

4.2 Swaptions

A swaption is an option written on an interest rate swap. This holder of a swaption has right, but not the obligation, to enter an interest rate swap upon the option maturity at a pre-specified strike K. The option maturity is called the *term* of the swaption. We assume the maturity of the swaption coincides with the first reset date of the underlying interest swap, i.e., T_α. The length of the underlying IRS $T_\beta - T_\alpha$ is the *tenor* of the swaption. The strike determines the fixed rate at which the underlying swap can be entered, so the value of the option is naturally driven by the difference between the strike and the swap rate $S_{\alpha,\beta}(t)$. We also distinguish between a payer and receiver swaption. Logically, a payer swaption gives the holder the right to enter a payer swap (paying the fixed rate), while a receiver swap gives the right to enter a receiver swap (received the fixed rate). We can express swaptions in terms of their *moneyness*. If the swaption is stuck with a strike equal to the forward swap rate, it is at-the-money (ATM). A payer swaption is in-the-money (ITM) if the strike is less than the swap rate and out-of-the-money (OTM) if the strike is above the swap rate, and vice versa for a receiver swaption (Brigo & Mercurio, 2006).

An important caveat in valuing swaptions is their settlement type. With *physical* settlement, a swap is entered at option expiration, and actual exchanges of cash flows take place. This can be contrasted with *cash* settlement, where the swaption is settled through an exchange of cash corresponding to the value of the swap at the time of exercise. The cash settlement is usually calculated by discounting all future cash flows using the swap rate, but this complicates the valuation procedure (Crispoldi et al., 2015). In the EUR market, the most swaptions have cash settlement, yet we henceforth only consider physical swaption settlement in the interest of keeping things simple.

4.2.1 Valuation of Swaptions

In the following, we present only formulas for payer swaptions, but analogous arguments can be made for receiver swaptions. The swaption payoff at its maturity T_α is found by considering the value of the underlying payer IRS at its first reset date. Thus, its payoff written in terms of the swap rate is:

$$\left(A_{\alpha,\beta}(T_\alpha)\,(S_{\alpha,\beta}(T_\alpha) - K)\right)^+ \tag{4.6}$$

The value of a swaption at time-t is the expectation of its discounted future payoff, which under the swap measure $\mathbb{Q}^A$ associated with the annuity factor as numeraire is given as:

$$PS\big(t, T_\alpha, T_\beta, K\big) = A_{\alpha,\beta}(t) \cdot \mathbb{E}^{\mathbb{Q}^A}\left[\big(S_{\alpha,\beta}(T_\alpha) - K\big)^+ \,\middle|\, \mathcal{F}(t_0) \right] \tag{4.7}$$

To avoid arbitrage, the swap rate must be a martingale under the swap measure (Oosterlee & Grzelak, 2019). The martingale representation theorem implies that the swap rate evolves according to:

$$dS_{\alpha,\beta}(t) = C(\cdot)\, dW(t) \tag{4.8}$$

for a coefficient $C(\cdot)$, and where dW is a Brownian motion. To be able to price the swaption, a model for $C(\cdot)$ must be postulated. In the following, we assess the classical ways of modelling $C(\cdot)$ to arrive at a closed-form solution for the price of a payer swaption.

4.2.1.1 Normal Model

The Bachelier model, or *normal* model, assumes that under the swap measure $\mathbb{Q}^A$, the swap rate $S_{\alpha,\beta}$ follows a normal diffusion process with the following stochastic differential equation (SDE):

$$dS_{\alpha,\beta}(t) = \sigma_N\, dW(t) \tag{4.9}$$

where σ_N is the instantaneous normal volatility of the swap rate, and the Weiner process W is defined under the swap measure. Given these dynamics, the time-t value of a payer swaption maturing at T_α is:

$$\mathrm{PS}^{\mathcal{N}}\big(t, T_\alpha, T_\beta, K, \sigma_N\big) = A_{\alpha,\beta}(t)\mathcal{N}\big(S_{\alpha,\beta}(t), K, t, T_\alpha, \sigma_N\big) \tag{4.10}$$

with the normal pricing formula given as (Crispoldi et al., 2015):

$$\mathcal{N}(F_t, K, t, T, \sigma_N) = \sigma_N\sqrt{T-t}\big(d_N\Phi(d_N) + \phi(d_N)\big), \qquad \text{where } d_N = \frac{F_t - K}{\sigma_N\sqrt{T-t}} \tag{4.11}$$

The terms $\Phi(\cdot)$ and $\phi(\cdot)$ are the normal cumulative distribution and probability density function, respectively. The normal pricing model assumes that the swap rate follows a normal distribution. As this distribution is symmetrical around zero, the normal model can price options with negative rates without the need for adjustments. This feature means that the model is widely used in fixed income markets, which has been characterised by negative rate in recent years, and swaptions are typically quoted in terms of implied normal volatility (Choi et al., 2022). *Implied normal volatility* is the volatility which will yield the market price of the option when put into the normal model, while *implied Black volatility* is the equivalent version for matching the market volatility using the Black model.

4.2.1.2 Black Model

The Black model, or *lognormal* model, assumes that under the swap measure $\mathbb{Q}^A$, the swap rate follows a lognormal diffusion process with the following SDE:

$$dS_{\alpha,\beta}(t) = \sigma_B S_{\alpha,\beta}(t)\, dW(t) \tag{4.12}$$

where σ_B is the instantaneous lognormal volatility of the swap rate. Given these dynamics, the time-t value of a payer swaption maturing at T_α is:

$$\mathrm{PS}^{\mathcal{B}}\big(t, T_\alpha, T_\beta, K, \sigma_B\big) = A_{\alpha,\beta}(t)\mathcal{B}\big(S_{\alpha,\beta}(t), K, t, T_\alpha, \sigma_B\big) \tag{4.13}$$

with the Black formula given by (Crispoldi et al., 2015):

$$\mathcal{B}(F_t, K, t, T, \sigma_B) = F_t\Phi(d_B^+) - K\Phi(d_B^-), \qquad \text{where } d_{1,2} = \frac{\ln\left(\frac{F_t}{K}\right) \pm \sigma_B^2/2(T-t)}{\sigma_B\sqrt{T-t}} \tag{4.14}$$

The Black model assumes that the swap rate follows a lognormal distribution, which cannot take negative values. This means that swaptions prices are not defined for negative rates (Choi et al., 2022).

4.2.1.3 Shifted Black Model

The shifted Black, or displaced diffusion model, solves the non-negative issue of the Black model, by adding a positive shift to the swap rate and the strike in the Black model, so the dynamics become:

$$dS_{\alpha,\beta}(t) = \sigma_D\big[S_{\alpha,\beta}(t) + s\big]\, dW(t) \tag{4.15}$$

where s is the shift parameter and σ_D is the instantaneous shifted Black volatility of the swap rate. The shift parameter defines the lower bound of the swap rate, which must be chosen a priori to avoid negative values. The shortcoming of the model is thus that we must anticipate how low rates can go. If we define $\tilde{F}_t = F_t + s$ and $\widetilde{K} = K + s$, the swaption price under the shifted Black model is calculated using equation (4.13) and (4.14) with $\mathcal{B}\big(\tilde{F}_t, \widetilde{K}, t, T, \sigma_D\big)$ (Russo & Fabozzi, 2017). Because of the presence of the shift parameter, the volatility of the shifted Black and Black models is not identical.

4.2.2 Risk Management

In option markets, the sensitivity of the price of the option with respect to a particular variable is quantified under the Greeks. We introduce the two measures which will be relevant later. The delta of a swaption measures the sensitivity of the swaption price with respect to the swap rate:

$$\Delta^{\mathcal{N}} = \frac{\partial\mathcal{N}(\cdot)}{\partial F_t} = A_{\alpha,\beta}(t)\cdot\Phi(d_N), \qquad \Delta^{\mathcal{B}} = \frac{\partial\mathcal{B}(\cdot)}{\partial F_t} = A_{\alpha,\beta}(t)\cdot\Phi(d_1) \tag{4.16}$$

The vega of a swaption measures the sensitivity of the swaption price with respect to volatility:

$$\Lambda^{\mathcal{N}} = \frac{\partial\mathcal{N}(\cdot)}{\partial\sigma_N} = A_{\alpha,\beta}(t)\cdot\phi(d_N)\sqrt{T-t}, \qquad \Lambda^{\mathcal{B}} = \frac{\partial\mathcal{B}(\cdot)}{\partial\sigma_B} = A_{\alpha,\beta}(t)\cdot F_t\phi(d_1)\sqrt{T-t} \tag{4.17}$$

PART II

VOLATILITY MODELS

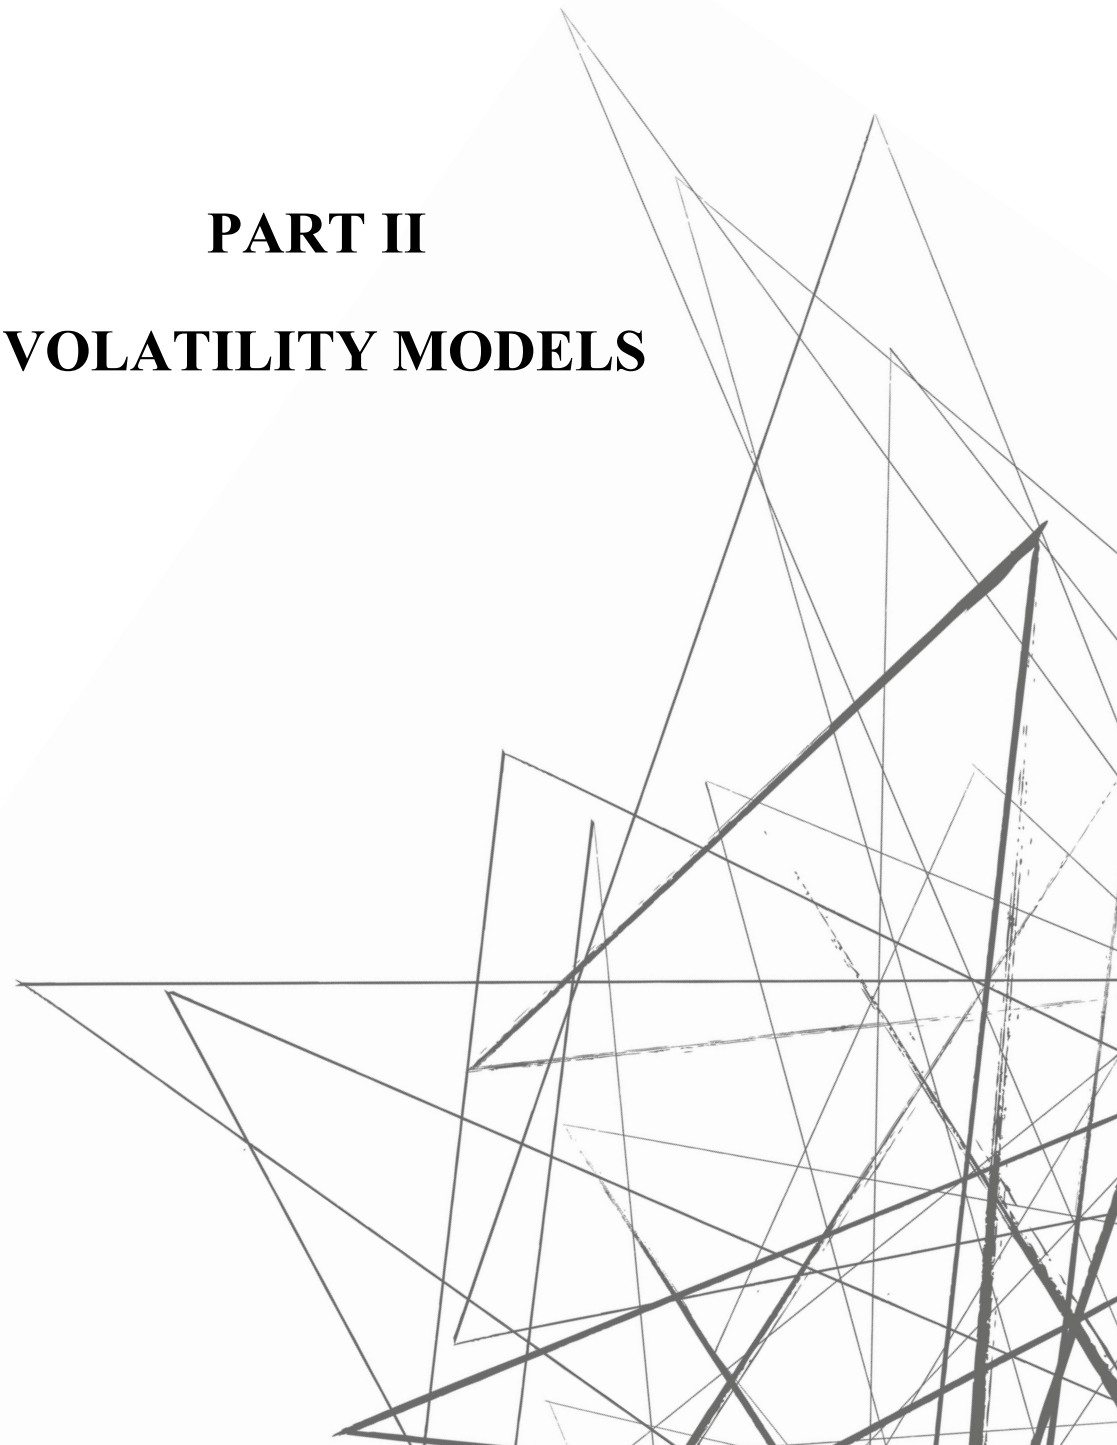

5 MOTIVATING VOLATILITY MODELS

The classical option pricing models presented in section 4.2.1 assume constant volatility of the forward rate process across strikes and maturities, which enables closed-form pricing formulas and convenient calculation of risk sensitivities. In the following, we see how the constant volatility assumption breaks down in the swaption market, which empirically exhibits a non-flat structure. We use this to motivate the need for a *stochastic volatility model*, where the idea is that the volatility of the forward rate process itself is a stochastic process which should be modelled as its own component.

5.1 Volatility Smiles

Illustrating the implied volatility across strike levels while holding maturity fixed shows what is known as a *volatility smile*. Adding the swaption term (maturity) and the tenor of the underlying swap to the strike dimension yields a three-dimensional *volatility cube* with implied volatility varying across all dimensions. We illustrate this phenomenon in Figure 5.1, which shows volatility smiles on different dates for swaptions written on fixed-for-floating 6M Euribor swaps constructed using market data from Bloomberg. Note that moneyness refers to the strike relative to the par swap rate. Looking at each subplot individually, the smiles vary in level and curvature, which is clear proof that the implied volatility varies across strike, term, and tenor. The implied volatility tends to empirically be lower for ATM strikes compared to OTM or ITM strikes, and the smiles tend to become flatter when the term and tenor

Figure 5.1: Market volatility smiles for selected swaptions observed on the first Wednesday of March 2020-2023.

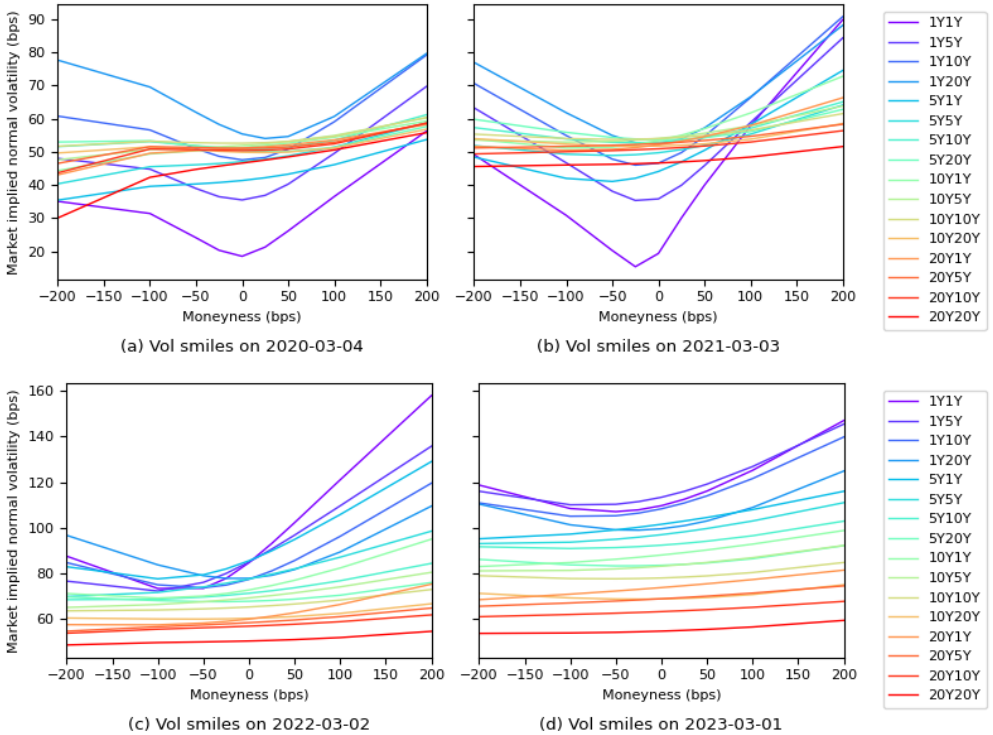

of the swaption increases. Comparing the different dates, the shape of the smiles also changes across time. For example, the smiles observed on March 1, 2023, appear flatter and are shifted upwards compared to those observed on March 4, 2019.

The volatility smile occurs because the probability distribution of forward rates assumed in the normal Bachelier and lognormal Black models does not match the empirical distribution. In fact, the empirical distribution has fatter tails (leptokurtosis), implying that the classical models underestimate the probability of extreme outcomes of forward rates. This translates into a higher volatility of deep ITM or OTM options relative to ATM options. We may also explain the volatility smile with reference to demand pressures. Options protect against adverse movements, and depending on market exposures and expectations, there will be a larger demand pressure for away-from-the-money options compared to ATM options (Hull, 2018). This pressure naturally drives up their price which translates into a larger implied volatility. Finally, liquidity effects should also be considered, where a lower market liquidity at far strikes or maturities will result in higher prices and thus larger implied volatilities (Deuskar et al., 2004).

5.2 Importance of Smile Modelling

There is clear empirical evidence that volatility smiles are a feature of the interest-rate market, and we now motivate why modelling them is important from a risk management perspective. If our pricing model is unable to capture the smile dynamics, we cannot consistently price all options in the market which has consequences for the calculation of risk sensitivities. As an example, consider the Black model where a different volatility is needed for each strike to fit the market smile. This effectively means that a different model is fitted for each option, and it becomes unclear whether risk sensitivities are consistent across strikes (Hagan et al., 2002). Further, since volatility depends on the strike, it likely also depends on the forward rate, meaning that a proportion of the vega risk in the Black model should instead be hedged as a delta risk. The overarching conclusion is that if our pricing model cannot capture the smile dynamics, our sensitivity measures and thus hedging become inaccurate.

Local volatility models are the simplest solution for incorporating non-constant volatility in option pricing, which are based on the idea of calibrating a local volatility function to the market prices of liquid options across strikes and tenors at each point in time. However, this approach may yield highly unstable parameters and the predicted smile dynamics contradict the empirical facts making the model unsuited for hedging purposes (Brigo & Mercurio, 2006). A more promising approach are *stochastic volatility models*, which are designed to both capture the stochastic behaviour of volatility and to accommodate market smiles. For the remainder of this thesis, we focus on the SABR model, which is a stochastic volatility model that long has been the standard for smile interpolation in swaption markets.

6 THE SABR MODEL

The *Stochastic Alpha Beta Rho* (SABR) model is a stochastic volatility model introduced by Hagan et al. (2002) to capture the dynamics of the implied volatility smile and improve the pricing and hedging capabilities of the normal and Black models. It is possible to derive closed-form approximations for the implied volatility of an option depending on the current forward rate f, the strike K, time to maturity T, and four intuitive model parameters (α, β, ρ, ν) which determine the dynamics of the produced smile. These approximations enable fast calibration of the model parameters to the market data, while the model still manages to capture the market-observed volatility smile comparatively well (Deloitte, 2016). For these reasons, practitioners use the SABR model for volatility modelling and smile-interpolation, especially in the context of interest rate derivatives such as swaptions. In the forthcoming sections, we introduce the features of the classic SABR model.

6.1 Model Specification

The original SABR model is specified as a two-factor model which models a single forward rate F_t and its volatility α_t as two correlated stochastic processes. The model dynamics are given by:

$$\begin{aligned} dF_t &= \alpha_t C(f) dW_t^2 \\ d\alpha_t &= \nu \alpha_t \, dW_t^1 \\ dW_t^1 dW_t^2 &= \rho \, dt \\ F(0) = f > 0, \alpha(0) &= \alpha > 0 \end{aligned} \tag{6.1}$$

where W_t^1 and W_t^2 are two standard Wiener processes with a constant correlation of $\rho \in [-1, 1]$. The forward process uses the current forward rate f as initial condition, which is assumed strictly positive. The volatility of α_t (vol-vol) is assumed constant $\nu > 0$. The model dynamics are characterised by a function $C(f)$ of the forward rate, which should be defined for $F_t > 0$ and is assumed to be positive, smooth, and integrable around zero (Hagan et al., 2015). The original SABR model is specified as $C(f) = F_t^{\beta}$ with $\beta \in [0, 1]$. In section **Error! Reference source not found.**, we introduce the SABR approximations for the implied volatility of an option and use this as the foundation for conducting an in-depth analysis of how each parameter impacts the shape of the smile in section **Error! Reference source not found.**.

Before proceeding, we note that the SABR model is only intended to fit volatility smiles for a single exercise date, which still makes it appropriate for vanilla swaptions. Under the forward measure, the forward rate and volatility processes are martingales only if we consider a single forward rate at a time

(Crispoldi et al., 2015). This means that in its original form, the SABR model cannot fit volatility surfaces or produce volatility estimates for options with multiple exercise dates.

6.2 Solving the Model

The stochastic processes in equation (6.1) with $C(f) = F_t^{\beta}$ can be solved, e.g., using Monte Carlo techniques, but these methods are quite time-consuming and are not viable for calibration purposes (Crispoldi et al., 2015). It turns out that it is possible to derive an asymptotic approximation formula for the implied volatility of a vanilla option under the SABR dynamics. This volatility can then conveniently be placed into either the normal or Black model to obtain an option price. In the following, let $\sigma_N(K)$ and $\sigma_B(K)$ be the implied normal and Black volatility, respectively, contingent on the strike of the option K. We start by looking at the original approximations of Hagan et al. (2002), and then briefly introduce an alternative approximation, which claim a higher degree of precision.

6.2.1 Hagan et al. (2002) Approximations

The original approximation formulas by Hagan et al. (2002) calculates the implied volatility as a function of the current forward rate f, strike K, maturity T, and a set of model parameters (α, β, ρ, ν) calibrated to the market-observed volatility smile. Considering the SABR model under a *Black calibration space*, we can approximate the implied Black volatility as:

$$\begin{aligned}\sigma_B(K) = {} & \frac{\alpha}{(fK)^{(1-\beta)/2}\left\{1 + \frac{(1-\beta)^2}{24}\ln^2(f/K) + \frac{(1-\beta)^4}{1920}\ln^4(f/K)\right\}} \cdot \left(\frac{\zeta}{\xi(\zeta)}\right) \\ & \cdot \left\{1 + \left[\frac{(1-\beta)^2}{24}\frac{\alpha^2}{(fK)^{1-\beta}} + \frac{1}{4}\frac{\rho\beta\nu\alpha}{(fK)^{(1-\beta)/2}} + \frac{2-3\rho^2}{24}\nu^2\right]T + \cdots\right\}\end{aligned} \tag{6.2}$$

where '$+\cdots$' are higher order expansions omitted from the formula due to their negligible effect on the implied volatility, andζ and $\xi(\zeta)$ can be expressed as:

$$\zeta = \frac{\nu}{\alpha}(fK)^{(1-\beta)/2}\ln(f/K) \qquad \text{and} \qquad \xi(\zeta) = \ln\left(\frac{\sqrt{1-2\rho\zeta+\zeta^2}+\zeta-\rho}{1-\rho}\right) \tag{6.3}$$

For ATM options where $K = f$, the formula for the implied Black volatility simplifies to:

$$\sigma_B^{\text{ATM}} = \frac{\alpha}{f^{1-\beta}} \cdot \left\{1 + \left[\frac{(1-\beta)^2}{24}\frac{\alpha^2}{f^{2-2\beta}} + \frac{1}{4}\frac{\rho\beta\alpha\nu}{f^{1-\beta}} + \frac{2-3\rho^2}{24}\nu^2\right]T + \cdots\right\} \tag{6.4}$$

Alternatively, we can operate under a *Normal calibration space*, where the implied normal volatility is:

$$\sigma_N(K) = \alpha(fK)^{\beta/2} \cdot \frac{1 + \frac{1}{24}\ln^2(f/K) + \frac{1}{1920}\ln^4(f/K) + \cdots}{1 + \frac{(1-\beta)^2}{24}\ln^2(f/K) + \frac{(1-\beta)^4}{1920}\ln^4(f/K) + \cdots} \cdot \left(\frac{\zeta}{\xi(\zeta)}\right) \cdot \left\{1 + \left[\frac{-\beta(2-\beta)}{24}\frac{\alpha^2}{(fK)^{1-\beta}} + \frac{1}{4}\frac{\rho\beta\nu\alpha}{(fK)^{(1-\beta)/2}} + \frac{2-3\rho^2}{24}\nu^2\right]T + \cdots\right\} \tag{6.5}$$

where ζ and $\xi(\zeta)$ are defined as in equation (6.3). For the ATM case, the formula in (6.5) simplifies to:

$$\sigma_N^{\text{ATM}} = \alpha f^{\beta} \cdot \left\{1 + \left[\frac{-\beta(2-\beta)\alpha^2}{24f^{2-2\beta}} + \frac{\rho\nu\alpha\beta}{4f^{1-\beta}} + \frac{2-3\rho^2}{24}\nu^2\right]T + \cdots\right\} \tag{6.6}$$

6.2.2 Antonov et al. (2013) Solution

The approximations derived by Hagan et al. (2002) have some major faults. Firstly, the risk-neutral probability density function (PDF) implied by the approximations can take negative values when considering long maturities, which implies that option prices no longer are convex in the strike dimension and are thus subject to arbitrage (Crispoldi et al., 2015). Further, the approximations are derived under the condition $\nu^2 T \ll 1$, and empirical violations of this causes a loss in accuracy and makes the model parameters unreliable. It also means that the approximation becomes increasingly inaccurate when strikes approach zero, especially for long maturities. Finally, high β values have been known to cause explosive behaviours in the approximations, vastly overestimating the implied volatility for high strikes.

There have been many attempts at improving the accuracy of the Hagan approximations, but we focus on the work of Antonov et al. (2013), who provide an alternative solution for the option price when $\beta \in (0, 1)$. Let $\mathcal{C}(T, K)$ be the undiscounted (forward) value of a European call option. For the zero-correlation case, $\rho = 0$, this can be calculated as:

$$\mathcal{C}_S(T,K) = (f-K)^+ + \frac{2}{\pi}\sqrt{Kf}\left\{\int_{s_-}^{s_+} \frac{\sin\big(\eta\phi(s)\big)}{\sinh(s)} G(\nu^2 T, s)\ ds + \sin(\eta\pi)\int_{s_+}^{\infty} \frac{e^{-\eta\psi(s)}}{\sinh(s)} G(\nu^2 T, s)\ ds\right\} \tag{6.7}$$

where the following definitions hold

$$G(t,s) = 2\sqrt{2}\frac{e^{-t/8}}{t\sqrt{2\pi t}}\int_s^{\infty} u e^{-\frac{u^2}{2t}}\sqrt{\cosh(u) - \cosh(s)}\, du, \tag{6.8}$$

$$\phi(s) = 2\arctan\left(\sqrt{\frac{\sinh^2(s) - \sinh^2(s_-)}{\sinh^2(s_+) - \sinh^2(s)}}\right), \tag{6.9}$$

(6.10)

$$\psi(s) = 2\arctan\left(\sqrt{\frac{\sinh^2(s) - \sinh^2(s_+)}{\sinh^2(s) - \sinh^2(s_-)}}\right), \tag{6.11}$$

$$s_- = \operatorname{arcsinh}\left(\frac{\nu|q - q_0|}{\alpha}\right), \quad s_+ = \operatorname{arcsinh}\left(\frac{\nu(q + q_0)}{\alpha}\right), \tag{6.12}$$

$$q = \frac{K^{1-\beta}}{1-\beta}, \quad q_0 = \frac{f^{1-\beta}}{1-\beta} \tag{6.13}$$

$$\eta = \left|\frac{1}{2(\beta - 1)}\right|$$

To ease the computational burden, the kernel function $G(t, s)$ can be approximated by:

$$G(t,s) \approx \sqrt{\frac{\sinh(s)}{s}}\, e^{-\frac{s^2}{2t} - \frac{t}{8}}\big(R(t,s) + \delta R(t,s)\big) \tag{6.14}$$

$$R(t,s) = 1 + \frac{3tg(s)}{8s^2} - \frac{5t^2\big(-8s^2 + 3g^2(s) + 24g(s)\big)}{128s^4} + \frac{35t^3\big(-40s^2 + 3g^3(s) + 24g^2(s) + 120g(s)\big)}{1024s^6} \tag{6.15}$$

$$g(s) = s\coth(s) - 1 \tag{6.16}$$

$$\delta R(t,s) = e^{\frac{t}{8}} - \frac{3072 + 384t + 24t^2 + t^3}{3072} \tag{6.17}$$

We can further improve the computation time by replacing the function $R(t, s)$ by its fourth-order expansion for small s (Antonov & Spector, 2012). Thus, we obtain the following:

$$R(t,s) \approx 1 + (1/8)t + (1/128)t^2 + (1/3072)t^3 + (-(1/120)t - (1/4032)t^2 - (1/15360)t^3)s^2 + \big((1/1260)t + (1/56320)t^3 - (1/40320)t^2\big)s^4 \tag{6.18}$$

$$G(t,s) \approx \sqrt{1 + s^2/6 + s^4/120}\, e^{-s^2/(2t)}\big(R(t,s) + \delta R(t,s)\big) \tag{6.19}$$

The zero-correlation approximation can be used as a mimicking model to map the solution to the general-correlation case, $\rho \neq 0$, but we do not consider that here. Note that the approximation imposes an absorbing boundary at zero, so negative rates are not allowed. The approximation should correct the tail behaviour of the distribution, but when $|\rho|$ is close to one and β is close to zero, the option price may not be convex for small strikes (Antonov et al., 2013).

6.3 Understanding the SABR Parameters

The Hagan et al. (2002) approximations enables us to readily explore how each of the parameters impact shape of the volatility smile produced by the model. We do this based on a naïve sensitivity analysis where each parameter is shifted away from its initial value for us to assess the ceteris paribus effect on the shape of the fitted volatility smile. The analysis is conducted using the Hagan et al. (2002) Black approximation in equation (6.2). We start off by arbitrarily choosing initial parameter values of $\alpha = 0.03$, $\beta = 0.60$, $\rho = -0.40$, and $\nu = 0.30$, and note that we always maintain $T = 1$ and $f = 0.03$.

Initial volatility, α. This effect of α is primarily seen on the smile level. Figure 6.1 shows that increasing α shifts the smile upwards, while the slope also slightly increases. This effect is intuitive, as a larger initial volatility should result in larger implied volatilities. The parameter is a volatility measure, $\alpha > 0$

Figure 6.2: Volatility smiles for different values of β (own creation)

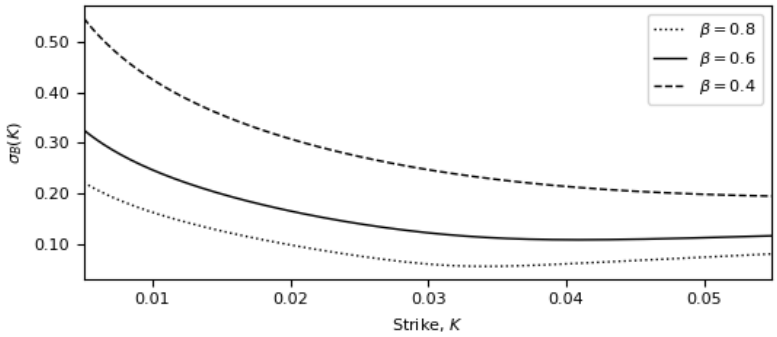

Figure 6.1: Volatility smiles for different values of α (own creation)

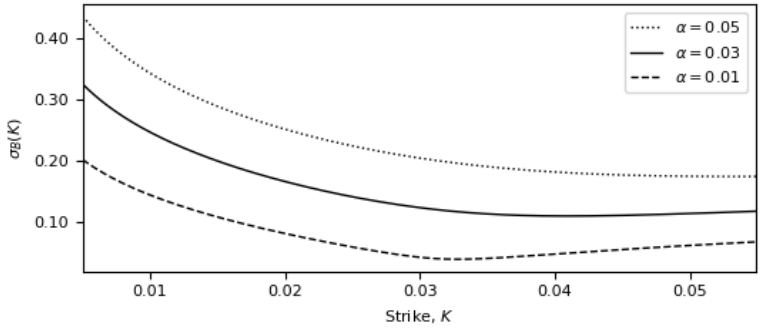

Power parameter, $\boldsymbol{\beta}$. The effect of β is primarily seen on the level of the smile and secondly on its slope. Figure 6.2 shows that increasing β shifts the smile down and makes it flatter. The slope of the smile gets more pronounced as β decreases. The β value is restricted to $\beta \in [0, 1]$, while the martingale property is only fulfilled if $0 \leq \beta < 1$ or if $\beta = 1$ and $\rho \leq 0$ (Crispoldi et al., 2015).

Correlation, $\boldsymbol{\rho}$. The correlation parameter $\rho \in [-1, 1]$ impacts the steepness and rotation of the smile. Figure 6.3 illustrates how the smile rotates around the ATM as ρ increases, effectively increasing the curvature of the smile. The smile also gets steeper as we move from 1 to -1. Empirically, we typically observe a negative correlation, where volatility increases as rates go down.

Volatility-of-volatility, $\boldsymbol{\nu}$. The effect of the ν parameter is primarily seen on the curvature of the smile. Figure 6.4 shows how the curvature of the smile becomes more pronounced as ν increases, with the curvature changing around the ATM point. Thus, increasing ν makes the smile effect more pronounced with higher implied volatility for away-from-the-money options. The parameter is a volatility measure, which means it is restricted to be positive, i.e., $\nu > 0$.

Figure 6.3: Volatility smiles for different values of ρ (own creation)

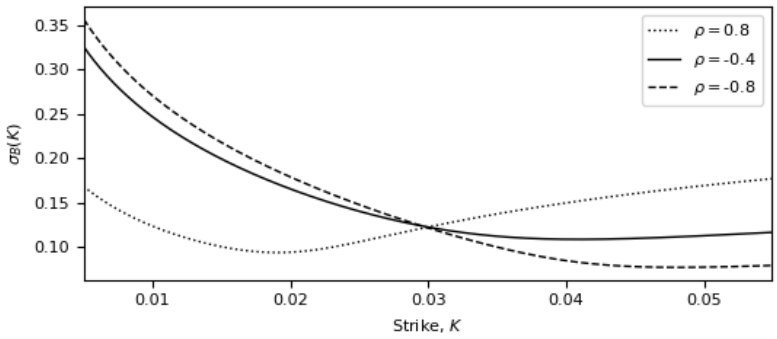

6.4 Calibrating the SABR Model

To implement the SABR model in practice, the parameters must be calibrated to market data. As the model dynamics are only valid for modelling a single forward at a time, we must calibrate a set of

Figure 6.4: Volatility smiles for different values of ν (own creation)

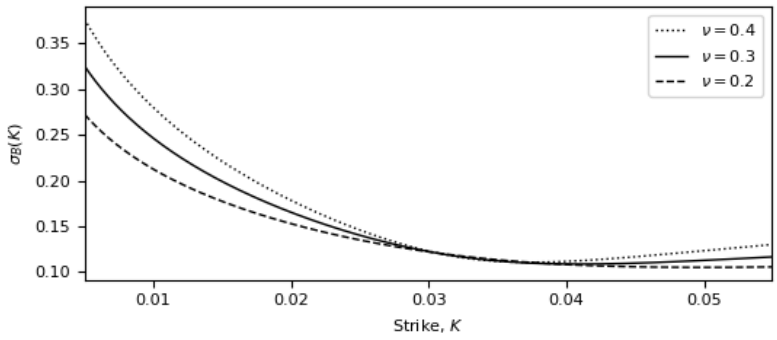

parameters for each option maturity. The calibration is done using an iterative procedure to identify the set of parameters which minimise the model sum of squared errors across strike levels observed in the market. In section **Error! Reference source not found.**, we found that some of the model parameters impact the shape of the smile in a similar way, which indicates that the model is overparameterized. This issue becomes more pronounced as the number of market quotes on a given date for swaptions with the same term and tenor likely is quite limited, so with up to four free parameters, there is a risk of overfitting the data. Practitioners have different ways of addressing this issue in the calibration procedure, which most commonly involves fixing β ahead of the other parameters. We present the ways of choosing the β parameter in section 6.4.1. Further, the SABR parameters can either be calibrated to (a) implied option volatilities in the market using the Hagan et al. (2002) approximation formula, or (b) to option market prices using the exact solution of Antonov et al. (2013) for the option value. Both methods are presented below.

6.4.1 Determining the β Parameter

There are generally three approaches to determining the β parameter. We cover each of them below.

Free β. A straightforward approach is to let β be estimated freely in the calibration along with the remaining parameters, although this is not considered standard practice. Hagan et al. (2002) make the point that β and ρ impact the volatility smile in a similar way and find that the model can fit market smiles reasonably well for very different values of β. This means that allowing β to be free in the calibration risks fitting market noise, which could cause poor performance.

Fixing β. The most common approach is to fix the value of β ahead of the other parameters based on prior beliefs about the distribution of the forward rate. This will usually result in $\beta = \{0, 0.5, 1\}$, although any value within $\beta \in [0, 1]$ is theoretically viable. It should be noted that the choice of β can affect the risk sensitivities of the model. Setting $\beta = 0.5$ is quite common in interest rate markets, and this specification yield the stochastic CIR model where $C(f) = \sqrt{F_t}$.

Fitting β. It is hypothesised that β can be determined from empirical observations of the *backbone,* which is the curve traced by the ATM implied volatilities as they vary due to changes in the forward rate. Specifically, Hagan et al. (2002) suggests inferring the value of β by regressing historical values of the log-rates against ATM log-volatilities. Under a Black calibration space, we do a log-transformation on equation (6.4) to obtain the following expression for fitting β:

$$\sigma_B^{\mathrm{ATM}} = \frac{\alpha}{f^{(1-\beta)}} + \cdots \qquad \Leftrightarrow \qquad \ln\left(\sigma_B^{\mathrm{ATM}}\right) = \ln(\alpha) + (\beta - 1)\ln(f) + \cdots \tag{6.20}$$

The second term of the formula is omitted when fitting β, which should not impact the results as the fitting procedure is already quite noisy due to the stochastic nature of α and f. With an empirical

estimate of the slope of the regression, the fitted β is found as the slope plus 1. Under a normal calibration space, we instead perform the log-transformation on equation (6.6), which gives us:

$$\sigma_N^{\mathrm{ATM}} = \alpha f^{\beta} + \cdots \qquad \Leftrightarrow \qquad \ln\left(\sigma_N^{\mathrm{ATM}}\right) = \ln(\alpha) + \beta \ln(f) + \cdots \tag{6.21}$$

The fitted β is then simply the regression slope. The fitted β parameter is generally found to be time-sensitive and deteriorates towards zero with maturity (West, 2005). This would suggest fitting a β for each standardised option maturity. The implication of $\beta = 0$ is a downward sloping backbone, where the ATM volatility is decreasing in the forward rate, while $\beta = 1$ implies a flat backbone, where changes in the forward rate only causes a parallel shift in the curve.

6.4.2 Calibrating to Option Volatilities

Calibrating the model parameters using an approximation formula for the implied volatility is the more traditional and perhaps sensible choice, as the procedure does not involve complicated operations. We assume β has been fixed beforehand c.f. section 6.4.1. This leaves us with the task of estimating α, ρ, and ν. Practitioners have two different methods for handling this task which are presented below.

6.4.2.1 Calibration Method 1: Free α

The intuitive solution for estimating α, ρ, and ν estimating them using an iterative optimization algorithm to find the set of parameters which yields the smallest sum of squared errors between the market implied volatility and the fitted SABR implied volatility. Given market observations of implied volatility quotes σ_i^{mkt} with corresponding strikes K_i, we should solve the constrained problem:

$$\begin{aligned}(\hat{\alpha}, \hat{\rho}, \hat{\nu}) &= \underset{\alpha,\rho,\nu}{\arg\min} \sum_{i}^{n} \left(\sigma_i^{\mathrm{mkt}} - \sigma^{\mathrm{SABR}}(K_i, f, \alpha, \beta, \rho, \nu)\right)^2 \\ &= \text{s.t.} \quad \alpha \leq 0, \quad -1 < \rho < 1, \quad \nu \leq 0\end{aligned} \tag{6.22}$$

where σ_i^{SABR} represents the model implied volatility for strike K_i under either a Black or normal calibration space, where the calibration space crucially should match the volatility type of the market quotes. We note that the current forward rate f conveniently can be inferred from bootstrapped forward or swap rate curves, depending on the type of option considered.

6.4.2.2 Calibration Method 2: Implied α

The second estimation method is to estimate α from the ATM volatility by inverting the approximation formula and determining α as the root of the cubic function. This means that α no longer is considered a model parameter, as the smile dynamics instead are described by σ^{ATM}, β, ρ, and ν. Empirically, the ATM volatilities usually needs to be updated frequently (e.g., daily), while the smiles need much less frequent updates (e.g., bi-weekly), and using this parameterisation allows us to update σ^{ATM} on the fly

while ρ and ν can be updated as needed (Hagan et al., 2002). Under a Black calibration space, we can imply α by inverting equation (6.4) and solving for the root:

$$\begin{aligned} \sigma_B^{\mathrm{ATM}} &= \frac{\alpha}{f^{1-\beta}}\left\{1+\left[\frac{(1-\beta)^2}{24}\frac{\alpha^2}{f^{2-2\beta}}+\frac{1}{4}\frac{\rho\beta\alpha\nu}{f^{1-\beta}}+\frac{2-3\rho^2}{24}\nu^2\right]T+\cdots\right\} \quad \Leftrightarrow \\ 0 &= \left[\frac{(1-\beta)^2T}{24f^{2-2\beta}}\right]\alpha^3+\left[\frac{\rho\beta\nu T}{4f^{1-\beta}}\right]\alpha^2+\left[1+\frac{2-3\rho^2}{24}\nu^2T\right]\alpha-\sigma_B^{\mathrm{ATM}}f^{1-\beta} \end{aligned} \tag{6.23}$$

If we are under a normal calibration space, we instead invert equation (6.6) and solve for the root:

$$\begin{aligned} \sigma_N^{\mathrm{ATM}} &= \alpha f^{\beta}\cdot\left\{1+\left[\frac{-\beta(2-\beta)\alpha^2}{24f^{2-2\beta}}+\frac{\rho\nu\alpha\beta}{4f^{1-\beta}}+\frac{2-3\rho^2}{24}\nu^2\right]T+\cdots\right\} \quad \Leftrightarrow \\ 0 &= \left[\frac{-\beta(2-\beta)T}{24f^{2-2\beta}}\right]\alpha^3+\left[\frac{\rho\beta\nu T}{4f^{1-\beta}}\right]\alpha^2+\left[1+\frac{2-3\rho^2}{24}\nu^2T\right]\alpha-\frac{\sigma_N^{\mathrm{ATM}}}{f^{\beta}} \end{aligned} \tag{6.24}$$

In either case, the cubic expression may have up to three real roots, but we should always opt to estimate α as the smallest possible root for an optimal result.

Given market observations of implied volatility quotes σ_i^{mkt} with corresponding strikes K_i, we can formalise the combined minimisation problem under this parameterisation as:

$$\begin{aligned} (\hat{\rho},\hat{\nu}) &= \arg\min_{\rho,\nu}\sum_i\left(\sigma_i^{\mathrm{mkt}}-\sigma^{\mathrm{SABR}}\left(K_i,f,\alpha(\rho,\nu,\sigma^{\mathrm{ATM}}),\beta,\rho,\nu\right)\right)^2 \\ &\text{s.t.} \quad -1<\rho<1, \quad \nu\leq 0 \end{aligned} \tag{6.25}$$

It is again crucial that the calibration space of the model volatilities σ_i^{SABR} matches the volatility type of the fitted market quotes. In practice, the procedure is implemented by implying out α as the smallest possible root of the cubic in each iteration of the optimisation algorithm depending on the values of ρ and ν in the specific iteration and given σ^{ATM} and β.

6.4.3 Calibrating to Option Prices

The model parameters can be calibrated directly to option prices using the exact solution of Antonov et al. (2013), although doing this is computationally a lot heavier than calibrating to option volatilities as it involves numerically solving complicated integrals. We thus expect a slower convergence of the optimal solution. The β parameter can be fixed beforehand using a priori expectations c.f. section 6.4.1, however, it seems reasonable to let β be estimated freely in the zero-correlation case, as a part of the argument for fixing its value was it having similar effect on the smile as ρ. This means we opt to freely estimate α, β, and ν using an iterative optimization algorithm to minimise the sum of squared errors between the market option prices and the fitted SABR option prices. Formally, we express:

$$\left(\hat{\alpha},\hat{\beta},\hat{\nu}\right) = \underset{\alpha,\rho,\nu}{\arg\min} \sum_{i}^{n} \left(\mathcal{C}_i^{\mathrm{mkt}} - \mathcal{C}^{\mathrm{SABR}}(K_i, f, \alpha, \beta, \nu)\right)^2 \\ = \text{s.t.} \quad \alpha \leq 0, \quad 0 < \beta < 1, \quad \nu \leq 0 \tag{6.26}$$

where $\mathcal{C}_i^{\mathrm{mkt}}$ is the market price of an option with strike K_i and $\mathcal{C}^{\mathrm{SABR}}$ is estimated model price of the same option. It is less relevant whether we consider the minimisation in terms of discounted or undiscounted option prices if we are consistent in our choice across the market and the model prices.

7 NEGATIVE RATES AND THE SABR MODEL

The original specification of the SABR model does generally not allow negative rates or strikes, as for any choice of $\beta > 0$, the model relies on a lognormal assumption of the distribution of forward rates, which is not defined for zero or negative values. Furthermore, the Hagen et al. (2002) approximations are known to lose accuracy when rates approach zero. The model was introduced under conditions where negative rate where a rarity, but the frequent and consistent observation of negative rates in, e.g., the EUR market in the past years seem has driven the development of a class of SABR models, which can model negative interest rates. In the following, we introduce the most common models, which will also be the foundation of our empirical analysis.

7.1 Shifted SABR

The simplest extension of the SABR model to accommodate negative interest rates is the *shifted SABR model*, also known as the displaced SABR model. The underlying idea is based on the same concept as the shifted Black model: There exists a lower boundary to how negative interest rates can become, but the boundary is below zero. We move the lower boundary of the model dynamics by applying a positive shift s to the forward process, where we define $\tilde{F}_t = F_t + s$ for some $s > 0$. This then gives us the following system of stochastic differential equations describing the shifted SABR model dynamics:

$$\begin{aligned} d\tilde{F}_t &= \alpha_t \tilde{F}_t^{\beta} dW_t^2 \\ d\alpha_t &= \nu \alpha_t \, dW_t^1 \\ dW_t^1 dW_t^2 &= \rho \, dt \\ \tilde{F}(0) = f + s > 0, &\alpha(0) = \alpha > 0 \end{aligned} \tag{7.1}$$

The shifted SABR model inherits the approximations of the classic SABR if the shift is consistently applied to both the forward rate and the strike, where we define $\tilde{F}_t = F_t + s$ and $\tilde{K} = K + s$ (Hagan et al., 2018). With a shift of $s = 0.05$, the model can handle rates larger than -5%. We note that using the shifted SABR model in a Black calibration space makes the approximation produce shifted Black volatilities. In the Bachelier model, adding a shift to both the forward rate and the strike cancels out and does thus not impact the calculated price or volatility. Still, the shift must be applied in the SABR approximation under a normal calibration space to produce correct normal volatilities.

Drawback of the shift parameter. The shifted SABR model requires an a priori determination of the shift parameter, which can be done using historical data. Practitioners seem to apply a shift between 1% and 5%, but there is no clear consensus on what the optimal value is. The issue that arises is that we are not certain how low the rates can go. If rates move below our lower bound of $-s$, the shift must be

readjusted, which will cause undesirable jumps in the estimated values and sensitivities. We finally note that the shift feature also applies to the PDF, which means the arbitrage issues around zero explained in section **Error! Reference source not found.** are mitigated, although the function still degenerates near the lower boundary. The implication of this is that hedges should not be done near the boundary (Antonov et al., 2015b).

7.2 Normal SABR

Setting $\beta = 0$ in the original SABR model gives us the *normal SABR model*, which can model negative rates without imposing a shift or free boundary condition. Under this model, the dynamics are given by:

$$\begin{aligned} dF_t &= \alpha_t \, dW_t^2 \\ d\alpha_t &= \nu \alpha_t \, dW_t^1 \\ dW_t^1 dW_t^2 &= \rho \, dt \\ F(0) &= f, \alpha(0) = \alpha > 0 \end{aligned} \tag{7.2}$$

The model dynamics in equation (7.2) imply $C(f) = 1$ so the increments of the forward process do not depend on the current level of the forward rate. This is inconsistent with our expectation based on empirical data and raises the concern whether specifying $\beta = 0$ is appropriate. The normal SABR can be seen as a generalisation of the Bachelier model with a stochastic volatility process, although the distribution of the forward rate is only truly normal if α_t is taken as constant (Crispoldi et al., 2015). We now present two methods of approximating option values under the normal SABR, namely the original Hagan et al. (2002) implied volatility approximation and the explicit solution by Antonov et al. (2015b). We also describe how each method can be used to calibrate the model parameters.

7.2.1 Hagan et al. (2002) Approximation

We consider the normal SABR model only in the context of a normal calibration space, as the model dynamics allow the forward rate to become arbitrarily negative making a lognormal distribution of forward rates impossible and thus using Black volatilities is inappropriate. The implied normal volatility can be found using the original Hagan et al. (2002) approximation seen in equation (6.5). Setting $\beta = 0$ in this formula, simplifies the expression to:

$$\sigma_N(K) = \alpha \cdot \left(\frac{\zeta}{\xi(\zeta)}\right) \cdot \left[1 + \frac{2 - 3\rho^2}{24} \nu^2 T + \cdots\right] \tag{7.3}$$

where ζ and $\xi(\zeta)$ can be expressed as:

$$\zeta = \frac{\nu}{\alpha}(f - K) \qquad \text{and} \qquad \xi(\zeta) = \ln\left(\frac{\sqrt{1 - 2\rho\zeta + \zeta^2} + \zeta - \rho}{1 - \rho}\right) \tag{7.4}$$

For the ATM case, the formula in (7.3) simplifies to:

$$\sigma_N^{\text{ATM}}(K) \approx \alpha \left[1 + \nu^2 \frac{2 - 3\rho^2}{24} T\right] \tag{7.5}$$

When calibrating the normal SABR model to the market normal implied volatilities, we are faced with the same parameterisation choices for α, ρ, and ν as earlier. Thus, we must choose between freely estimating α, ρ, and ν via the minimisation problem in equation (6.22), or freely estimating ρ and ν while inferring α form the ATM volatility as in the minimisation problem in equation (6.25).

7.2.2 Antonov et al. (2015b) Solution

Under the normal SABR with a free boundary, a closed-from solution for the undiscounted (forward) value of a call option in the general-correlation case, $\rho \neq 0$, can be found as (Antonov et al., 2015b):

$$\mathcal{C}_N(T, K) = (f - K)^+ + \frac{V_0}{\pi} \int_{s_0}^{\infty} \frac{G(\nu^2 T, s)}{\sinh s} \sqrt{\sinh^2 s - (k - \rho \cosh s)^2}\, ds \tag{7.6}$$

where the following holds

$$\cosh s_0 = \frac{-\rho k + \sqrt{k^2 + \bar{\rho}^2}}{\bar{\rho}^2} \tag{7.7}$$

$$k = \frac{K - f}{V_0} + \rho \tag{7.8}$$

$$V_0 = \frac{\alpha}{\nu} \tag{7.9}$$

$$\bar{\rho} = \sqrt{1 - \rho^2} \tag{7.10}$$

The kernel function $G(t, s)$ is identical to under the original SABR model and can be found in equation (6.8). In computations, we still approximate this function using equation (6.14)-(6.17). An important feature of this solution is that the behaviour of the forward rate is not sticky around zero. We can use this solution to calibrate the SABR parameters to the market-observed prices. In particular, we calibrate the parameters $\alpha, \rho,$ and ν freely using the optimisation problem presented in equation (6.22).

7.3 Free-Boundary SABR

The final extension to the SABR framework we consider is the *free-boundary SABR* model originally proposed by Antonov et al. (2015a). By setting $C(f) = |F_t|^\beta$ and assuming a free boundary (in contrast to the absorbing boundary at zero in the original SABR model), the dynamics are given by:

$$dF_t = \alpha_t |F_t|^\beta \, dW_t^2 \tag{7.11}$$

$$F(0) = f,\ \ \alpha(0) = \alpha > 0$$

with $\beta \in [0, 0.5)$. The specification of $C(f)$ ensures that negative interest rates can be modelled without the need for specifying an arbitrary shift parameter. The PDF of the model has a singularity at 0, which will mean that in Monte Carlo simulations, the forward process will primarily take values of the same sign as the initial forward rate (Antonov et al., 2019). This will impact the stability of the parameters and the correctness of risk sensitivities. Overall, the extreme stickiness at zero is very undesirable in the context on modelling negative rates, and the model is likely not used much in practice. Still, we include the model in our empirical analysis to compare the performance. For the zero-correlation case, $\rho = 0$, the time value of the option can be solved exactly, and we present the solution in the next section. For the general correlation case, $\rho \neq 0$, no exact solution exists but it is possible to use a mapping technique to approximate it (Antonov et al., 2015a). We will not consider this further in our context.

7.3.1 Antonov et al. (2015b) Solution

We here present the exact solution for the undiscounted value of a call option in the zero-correlation case as seen in Antonov et al. (2015b). We have:

$$\mathcal{C}_{FB}(T, K) = (f - K)^+ + \frac{1}{\pi}\sqrt{|Kf|}\{\mathbf{1}_{K\geq 0}A_1 + \sin(|\gamma|\pi)\, A_2\} \tag{7.12}$$

with integrals

$$A_1 = \int_0^{\pi} \frac{\sin\phi \sin(|\gamma|\phi)}{b - \cos\phi} \frac{G\left(\nu^2 T, s(\phi)\right)}{\cosh s(\phi)} d\phi \tag{7.13}$$

$$A_2 = \int_0^{\infty} \frac{\sinh\psi\, (\mathbf{1}_{K\geq 0}\cosh(|\gamma|\psi) + \mathbf{1}_{K<0}\sinh(|\gamma|\psi))}{b + \cosh\psi} \frac{G\left(\nu^2 T, s(\psi)\right)}{\cosh s(\psi)} d\psi \tag{7.14}$$

The following definitions apply

$$\gamma = -\frac{1}{2(1-\beta)} \tag{7.15}$$

$$b = \frac{|f|^{2(1-\beta)} + |K|^{2(1-\beta)}}{2|fK|^{1-\beta}} \tag{7.16}$$

$$\sinh s(\phi) = \frac{\nu}{\alpha}\sqrt{2\frac{|fK|^{1-\beta}}{(1-\beta)^2}(b - \cos\phi)} \tag{7.17}$$

$$\sinh s(\psi) = \frac{\nu}{\alpha}\sqrt{2\frac{|fK|^{1-\beta}}{(1-\beta)^2}(b + \cosh\psi)} \tag{7.18}$$

The kernel function $G(t, s)$ is given as in equation (6.8), while we use the approximations in equations (6.14)-(6.17) in computations. We can use this solution to calibrate α, ρ, and ν to the market-observed

prices via the optimisation problem presented in in equation (6.22), but the complexity of the integrals makes the calibration time consuming and convergence if often poor.

7.3.2 Approximating the Solution à la Hagan

It may be possible to approximate the implied normal volatility under the free-boundary model (as Black volatilities become inappropriate under a free boundary assumption) using the same idea as the original Hagan et al. (2002) approximations (Kienitz, 2015). The approximation formula for the implied normal volatility in the form presented in Hagan et al. (2014) gives us the expression:

$$\sigma_N(T,K) = \frac{\alpha(f-K)}{\int_K^f \frac{dx}{C(x)}} \cdot \left(\frac{\zeta}{\xi(\zeta)}\right) \cdot \left\{1 + \left[g\alpha^2 + \frac{\rho\nu\alpha}{4}\frac{C(f)-C(K)}{f-K} + \frac{2-3\rho^2}{24}\nu^2\right]T\right\} \tag{7.19}$$

where ζ, $\xi(\zeta)$, and the geometric function g are given as:

$$\zeta = \frac{\nu}{\alpha}\int_K^f \frac{dx}{|x|^\beta}, \qquad \xi(\zeta) = \ln\left(\frac{\sqrt{1-2\rho\zeta+\zeta^2}-\rho+\zeta}{1-\rho}\right) \tag{7.20}$$

$$g = \ln\left(\frac{1}{f-K}\int_K^f \frac{\sqrt{C(f)C(K)}}{C(x)}dx\right) \Big/ \left(\int_K^f \frac{dx}{C(x)}\right)^2 \tag{7.21}$$

Letting $C(f) = |F_t|^\beta$, the integral in the formulas above can be solved analytically as:

$$\int_K^f \frac{dx}{|x|^\beta} = \frac{\frac{f}{|f|^\beta} - \frac{K}{|K|^\beta}}{1-\beta} \tag{7.22}$$

Kienitz (2015) claims that this approximation formula can easily be calibrated to market-observed implied normal volatilities which yields almost identical prices to those found using the heavy procedure exact solution. Recall that the method of the Hagan approximation assumes that the function $C(f)$ is integrable around zero. As this does not hold when $C(f) = |F_t|^\beta$, the performance of this approximation formula may be subpar. The singularity also makes the solution unreliable when the forward and the strike take on opposite signs. We observe an issue when considering the ATM volatility, as the integral evaluates to zero, causing numerical issues in the geometric function. This could be addressed using a Taylor expansion as in Hagan et al. (2002), but as we already expect the approximation to perform poorly, we instead opt to calculate the ATM volatility by interpolating between two fictive strikes placed immediately above and below the forward rate f. However poor the suspected performance is, we can use this approximation to calibrate α, ρ, and ν in a normal calibration space via the minimisation problem in equation (6.22).

8 RISK MANAGEMENT UNDER SABR DYNAMICS

A part of the motivation for modelling the volatility smile with a stochastic volatility model was to improve the accuracy of risk management measures such as the *Greeks,* which trading desks use to hedge their open positions. In section 8.1, we present the formulas for calculating delta and vega risks under SABR dynamics to highlight how they differ from the constant volatility models. To give a fuller picture, we show how parameter sensitivities can be analysed as a complement the risk management capability of the models. Both topics are relevant under the empirical analysis.

We require the hedges implied by our model to be consistent under all market conditions which renders the classic SABR and the free-boundary SABR models useless. Thus, we only consider risk management under the shifted SABR and the normal SABR. For the shifted SABR, the formulas should be adjusted to accommodate the shift by using $\tilde{f} = f + s$ and $\tilde{K} = K + s$. We focus only on risk management measures under the Hagan approximations, and present the formulas only in the context of the Black calibration space, where $\sigma_B \equiv \sigma_B(K, f;\ \alpha, \beta, \rho, \nu)$ is found using the Hagan Black approximation and $\mathcal{B} = \mathcal{B}(K, f, t, T, \sigma_B)$ is the option time value found by inserting the SABR approximation into the Black model. The normal SABR model and the shifted SABR model under the normal calibration space requires that we use the Bachelier model for calculating risk measures. By replacing the SABR Black volatility and Black option time value with their normal equivalents, $\sigma_N \equiv \sigma_N(K, f;\ \alpha, \beta, \rho, \nu)$ and $\mathcal{N} = \mathcal{N}(K, f, t, T, \sigma_N)$, the formulas presented remain valid. Finally, we draw attention to that partial derivatives empirically can be calculated using finite difference, which is explained in section 10.4.

8.1 SABR Greeks

We limit our treatment of the Greeks under SABR dynamics to only delta and vega risks, as these risk measures differ most substantially from what is defined under the constant volatility models. We note that considering only these two Greeks means that we cannot comment on the full suite of hedging capabilities under the SABR dynamics, and instead keep the analysis on a conceptual level.

8.1.1 Delta

Delta risk measures the change in the option price due to a one-unit change in the underlying forward rate, which naturally is a central risk for option traders to hedge against. Under the SABR dynamics, changes in f impact both the option price directly and the volatility, so a correction term must be added to the classic delta under the Black model. For the case where α is explicitly calibrated to the market data (i.e., *not* inferred from the ATM volatility), Hagan et al. (2002) define the delta as:

$$\Delta^{\text{Hagan}} = A_{\alpha,\beta}(t) \cdot \left(\frac{\partial \mathcal{B}}{\partial f} + \frac{\partial \mathcal{B}}{\partial \sigma_B} \cdot \frac{\partial \sigma_B(K, f; \alpha, \beta, \rho, \nu)}{\partial f} \right)$$

This way of expressing delta neglects the effect of having correlated forward and volatility processes, which implies that when f changes, the volatility α also changes, and vice versa. By rewriting the model as two independent Wiener processes, a more consistent expression for the delta risk is (Bartlett, 2006):

$$\Delta^{\text{Bartlett}} = A_{\alpha,\beta}(t) \cdot \left(\frac{\partial \mathcal{B}}{\partial f} + \frac{\partial \mathcal{B}}{\partial \sigma_B} \cdot \left(\frac{\partial \sigma_B(K, f; \alpha, \beta, \rho, \nu)}{\partial f} + \frac{\partial \sigma_B(K, f; \alpha, \beta, \rho, \nu)}{\partial \alpha} \cdot \frac{\rho \nu}{f^{\beta}} \right) \right) \tag{8.1}$$

With this specification, the SABR delta is composed of the Black delta, $\Delta^{\text{Black}} = \partial \mathcal{B} / \partial f$, in addition to a correction for (a) volatility movements caused by a shift in the forward rate and (b) volatility movements caused by a shift in the initial volatility α through its correlation with the forward process.

8.1.2 Vega

Vega risk measures the change in the option price due to changes in the volatility of the forward rate. Under the SABR dynamics, this is usually defined in terms of the initial volatility α. Hagan et al. (2002) provide the following expression:

$$\Lambda^{\text{Hagan}} = A_{\alpha,\beta}(t) \cdot \frac{\partial \mathcal{B}}{\partial \sigma_B} \cdot \frac{\partial \sigma_B(K, f; \alpha, \beta, \rho, \nu)}{\partial \alpha}$$

As before, this expression neglects the effect of the correlation, and Bartlett (2006) instead adjusts it to:

$$\Lambda^{\text{Bartlett}} = A_{\alpha,\beta}(t) \cdot \frac{\partial \mathcal{B}}{\partial \sigma_B} \left(\frac{\partial \sigma_B}{\partial \alpha} + \frac{\partial \sigma_B}{\partial f} \cdot \frac{\rho f^{\beta}}{\nu} \right) \tag{8.2}$$

With this specification, the SABR vega is composed of the Black vega, $\Lambda^{\text{Black}} = \partial \mathcal{B} / \partial \sigma_B$, in addition to a correction for (a) volatility movements caused by a shift in the initial volatility α and (b) volatility movements caused by to a shift in the forward rate through its correlation with the forward process.

8.2 SABR Parameter Sensitivities

An analysis of the option price sensitivity with respect to each of the parameters in the SABR model can help us quantify model misspecification risk, which encompasses the risk of incorrect parameter values, e.g., due to a faulty calibration procedures or inability to update the parameter estimates frequently enough. The ρ and ν parameters move with the market and thus requires frequent updating, which means the risk of them changing is extremely relevant. Assessing parameter sensitivities is viewed as an important supplement to the Greeks, as understanding the inherent model risks will improve the robustness of hedges. A report by Deloitte (2016) states that banks using a SABR model will opt to hedge the SABR parameters rather than the traditional Greeks.

The price sensitivity to a parameter is defined as the first-order partial derivative with respect to it, which in practice is calculated using a finite difference method (see section 10.4). Sensitivity to α is capture by the SABR vega so we omit it from the sensitivity analysis. The price sensitivity $\mathcal{S}$ with respect to each of the remaining parameters ρ, β, and ν can be calculated as using equation **Error! Reference source not found.-Error! Reference source not found.**.

$$\mathcal{S}(\rho) = \frac{\partial V}{\partial \rho} = A_{\alpha,\beta}(t) \cdot \frac{\partial \mathcal{B}}{\partial \sigma_B} \cdot \frac{\partial \sigma_B(K, f; \alpha, \beta, \rho, \nu)}{\partial \rho} \tag{8.3}$$

$$\mathcal{S}(\nu) = \frac{\partial V}{\partial \nu} = A_{\alpha,\beta}(t) \cdot \frac{\partial \mathcal{B}}{\partial \sigma_B} \cdot \frac{\partial \sigma_B(K, f; \alpha, \beta, \rho, \nu)}{\partial \nu} \tag{8.4}$$

$$\mathcal{S}(\beta) = \frac{\partial V}{\partial \beta} = A_{\alpha,\beta}(t) \cdot \frac{\partial \mathcal{B}}{\partial \sigma_B} \cdot \frac{\partial \sigma_B(K, f; \alpha, \beta, \rho, \nu)}{\partial \beta} \tag{8.5}$$

The first sensitivity $\mathcal{S}(\rho)$ measures the risk of changes in ρ, which is referred to as the *vanna* risk. Vanna risk is usually taken as the second-order partial derivative with respect to the forward rate and volatility, which is exactly measured by the correlation parameter ρ in a SABR context. It thus represents the risk of the skew increasing. $\mathcal{S}(\nu)$ is the sensitivity to changes in ν, which can be referred to as the *volga* risk. Volga is traditionally the second-order partial derivative with respect to volatility, i.e., sensitivity to volatility-of-volatility. In the SABR context, this is exactly measured by the vol-vol parameter ν and represents the risk of the volatility smile becoming more pronounced (Hagan et al., 2002). Finally, $\mathcal{S}(\beta)$ is the sensitivity to changes in β, which we found mainly impacts the level and slope of the smile. Based on Hagan et al. (2002), we expect a low sensitivity to the β parameter.

A final sensitivity can be computed for the shift parameter s under the shifted SABR model. While the shift is assumed constant, we previously found a risk of having to update its value if rates drop too low. Updating this parameter causes shifts in prices and hedges, so understanding the swaption price sensitivity to it is extremely relevant. We define the sensitivity of the shift parameter as:

$$\mathcal{S}(s) = \frac{\partial V}{\partial s} = A_{\alpha,\beta}(t) \cdot \frac{\partial \mathcal{B}}{\partial \tilde{\sigma}_B} \cdot \frac{\partial \tilde{\sigma}_B\left(\tilde{K}, \tilde{f}; \alpha, \beta, \rho, \nu\right)}{\partial s} \tag{8.6}$$

PART III

EMPIRICAL APPLICATION

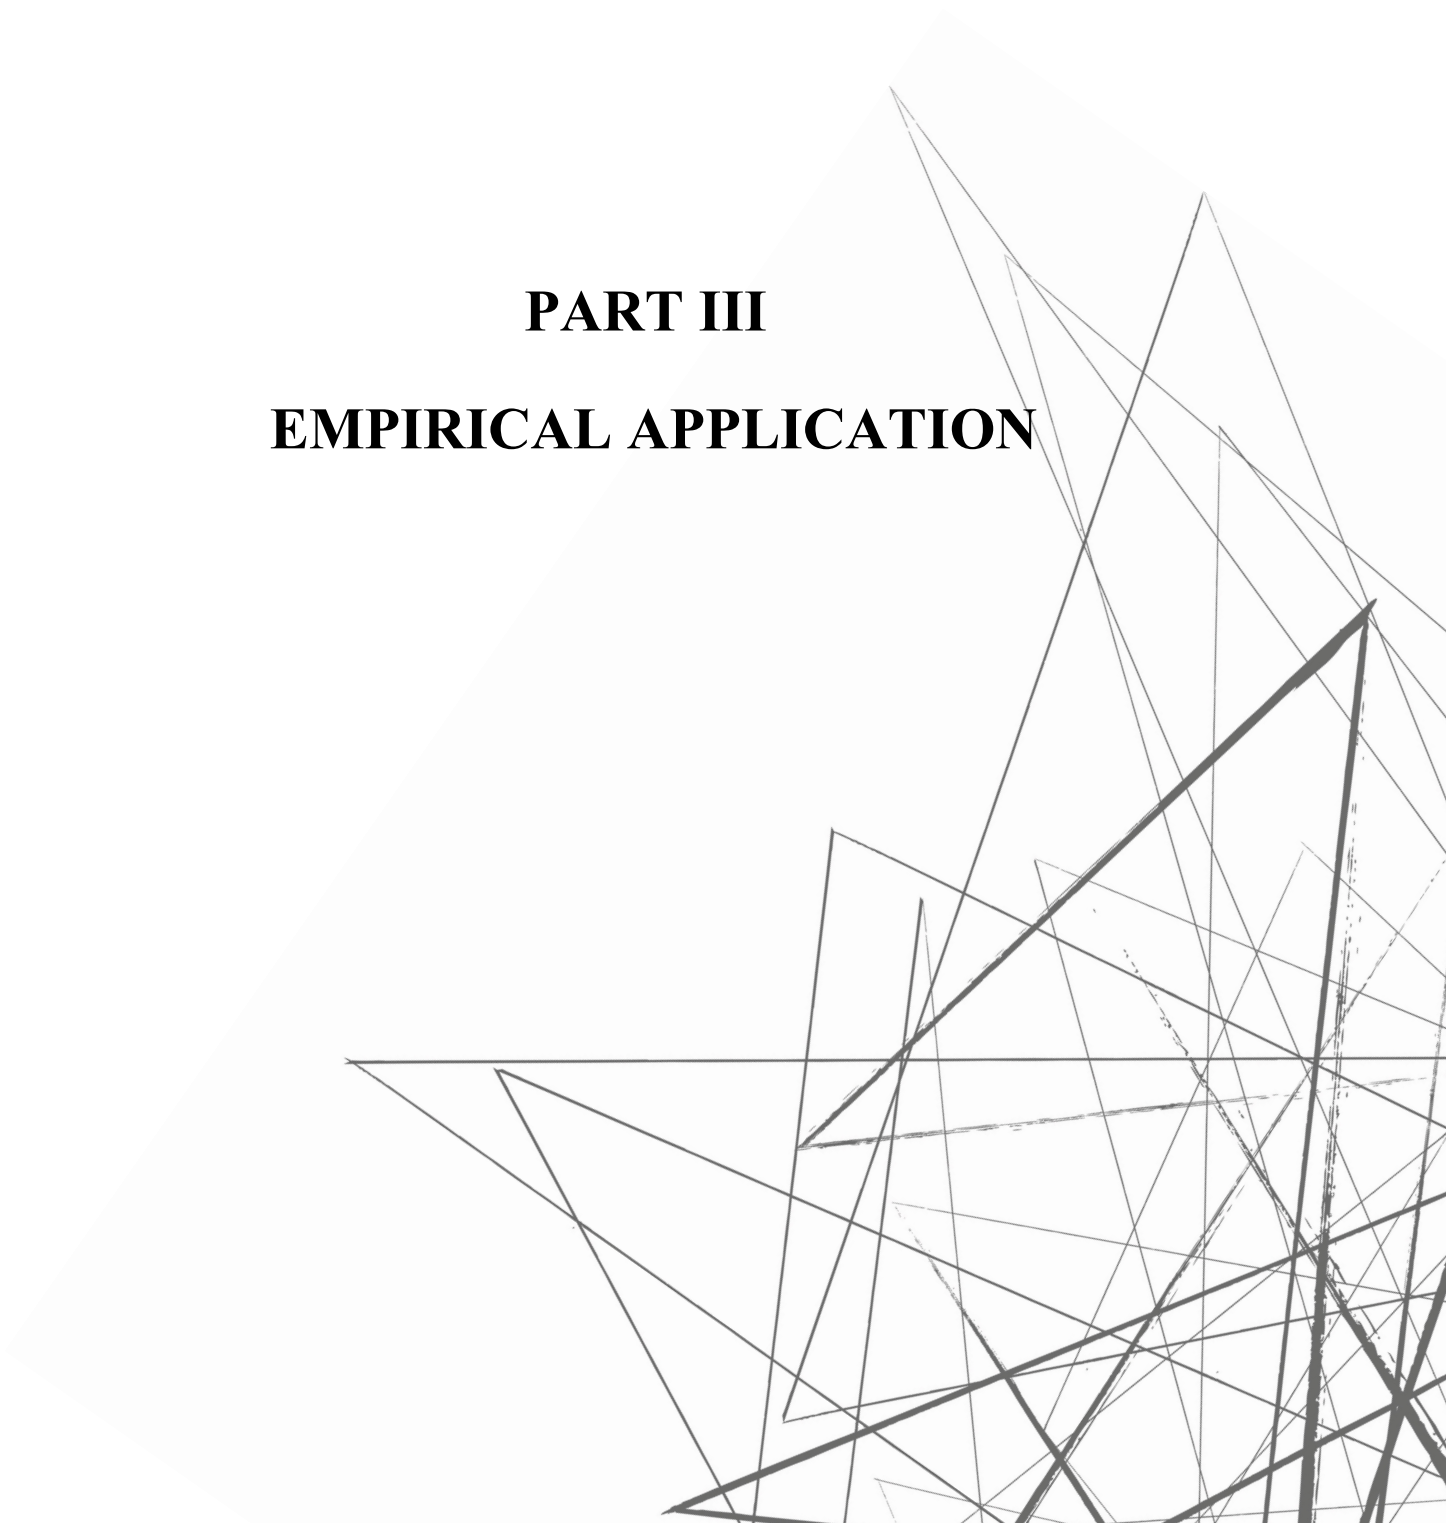

9 DATA DESCRIPTION

The empirical analysis focuses on evaluating the ability of the SABR models to model volatility smiles across different market conditions using real market data to enable a comparison of the robustness of the models and identify any potential shortcomings. This section accounts for the data used in the analysis by providing key summary statistics of the sample. Section 9.1 accounts for the data used to bootstrap the zero curve, where we also include considerations of the multi-curve framework. The main data sample, consisting of swaption implied volatility quotes, is described in section 9.2

9.1 Swap Data

To be able to price swaptions, we need a forward curve and a discount curve. We utilise a *single-curve framework*, meaning the same bootstrapped zero curve is used both for calculating forward rates and discounting. The curve is bootstrapped using par swap rate quotes for fixed-for-floating plain vanilla interest rate swaps, where the floating rate is the 6M Euribor. The data is collected with daily frequency from 27 March 2019 to 1 March 2023 using Bloomberg (ticker *EUSA*) with tenors 1-10Y, 15Y, 20Y, 30Y, 35Y, 40Y, 45Y, 50Y, and 60Y to enable curve bootstrapping to price up to a 30Y30Y swaption without extrapolating. The fixed leg of the swap has annual payments and uses 30/360 day count convention, while the floating leg has semi-annual payments and uses an Act/360 day count. The swaps have a two business days spot lag, and payments due on a non-business day are adjusted to occur on the following business day. We account for these conventions when calculating coverage. Table 9.1 shows par swap rates for selected tenors and dates. On the first two dates shown, the rates are increasing in the swap tenor up to 30Y, after which they slightly decrease. This implies a generally upward sloping normal zero curve. A different scenario is seen on 1 March 2023, where the rates are decreasing in the swap tenor, implying a downward sloping *inverted zero curve*. Inverted zero curves are an uncommon market phenomenon, which implies a market expectation of declining rates in the future.

Table 9.1: Par swap rate quotes for 6M Euribor swaps across selected tenors and dates

	1Y	3Y	5Y	7Y	10Y	15Y	20Y	30Y	40Y	50Y	60Y
2019-30-27	-0.23%	-0.16%	-0.01%	0.16%	0.44%	0.77%	0.93%	1.01%	1.00%	0.97%	0.96%
2021-03-12	-0.51%	-0.44%	-0.32%	-0.17%	0.04%	0.29%	0.41%	0.43%	0.39%	0.36%	0.33%
2023-03-01	3.77%	3.67%	3.45%	3.34%	3.29%	3.25%	3.10%	2.74%	2.53%	2.35%	2.24%

A few remarks are made on the curve construction. In practice, zero curves are bootstrapped using the most liquid interest rate instruments available for a given maturity. This usually means that the short end of the curve is bootstrapped using IBOR rates, the middle part using FRAs, and the far end using swap rates (Oosterlee & Grzelak, 2019). To ease the data collection, we choose to bootstrap the entire curve using swap rates. The bootstrapping procedure is implemented using the Quantlib python

package, which offers functionalities for handling day count conventions, payment adjustments, and has a wide range of interpolation methods. We found that using log-cubic interpolation of discount factors yields the smoothest bootstrapped curves, and we thus use this method.

9.1.1 A Note on the Multi-Curve Framework

Since the financial crisis of 2008, it has become standard practice to use a *multi-curve framework* when valuing derivatives. This means that a forward curve is constructed for each tenor structure (usually 1M, 3M, 6M, and 12M), while discounting is done with a single unique discounting curve. The risk-free rate used for discounting should be based on an instrument with the smallest possible credit risk, i.e., the curve with the shortest possible tenor. In the eurozone, the market standard is using €STR, which reflects the rate for unsecured overnight borrowing for euro area banks. The implication of a multi-curve framework is that the no-arbitrage relations valid under a single curve no longer hold, which complicates changing from the spot measure to the forward measure (Oosterlee & Grzelak, 2019). There are practical ways to circumvent this issue, but we note that curve construction is a complex topic which is not our focus area. Thus, we choose to implement the simpler single-curve framework with the implication that discounting is not done using a truly risk-free rate.

9.2 Swaption Data

The swaption data consists of market quotes of implied normal volatility for European swaptions written on fixed-for-floating plain vanilla interest rate swaps, where the floating rate is the 6M Euribor. We collect the volatility quotes every Wednesday from 27 March 2019 to 1 March 2023 using Bloomberg. The swaptions have terms and tenors representing all possible combinations of 1-5Y, 7Y, 10Y, 15Y, 20Y, and 30Y, and strikes (moneyness) of ± 200bps, ± 100bps, ± 50bps, ± 25bps, and ATM. Thus, each week, we collect quotes for 100 different swaption contracts each with 9 different strike levels. With 206 dates in the sample period, we obtain a dataset of 185,400 implied normal volatility quotes. As we only collect volatility quotes, we must infer the par swap rates for underlying swaps using the bootstrapped forward and discounting curves. In these calculations, we utilise that the floating leg of the underlying swaps has semi-annual payments with an Act/360 day count, while the fixed leg has annual payments with a 30/360 day count. We assume a spot lag of two business days, so a quote on a Wednesday is for a swaption contract starting on Friday (assuming no holidays). Payments occurring on non-business days are adjusted to occur on the following business day, following market conventions.

Figure 9.1 shows the average inferred swap rate and the observed range per date for the 100 different swaption contracts in the sample. For the majority of the sample dates, at least some of the quoted swaptions have a negative swap rate, which means that ability to accommodate negative rates is required from the models used to fit the dataset. From 16 March 2022 and onwards, all swaptions have strictly

positive swap rates. Another thing to note is that the average swap rate per day remains relatively stable at around 0% to 0.5% in the years 2020-2022, after which it increases steeply to around 2%. This observation is in line with the recent interest rate increases seen in the euro area.

Figure 9.1: Average swap rate and observed range over time in the sample for the 100 different swaption term and tenor combinations with the rate inferred from the bootstrapped forward and discounting curves.

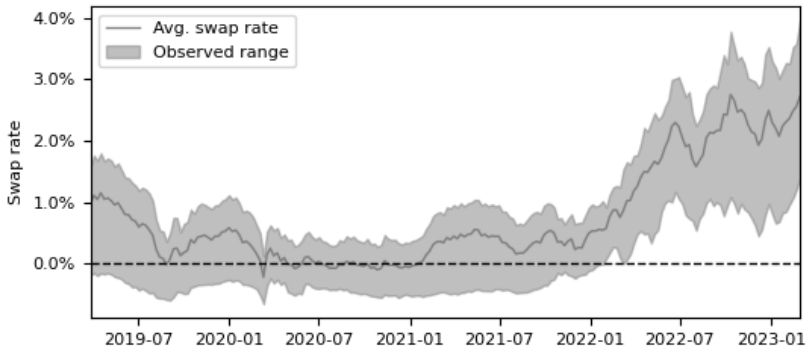

The sample distribution of the swap rates, prices, and normal volatilitises is seen in **Figure 9.2:** Distribution of swap rates, prices, and volatilities in the swaption datasetFigure 9.2. Considering first the distribution of the swap rates, it has a positive skew with a fat right tail and most observations placed around 0% to 0.50%. The average swap rate is 0.74% with a standard deviation of 0.91%, and the observations range from -0.66% to 3.96%. Looking at the distribution of the swaption price, it has a strong positive skew with most observations being near the lower bound of zero. The price is measured as the number of cents paid per € of swap notional. Across the sample, the price ranges from ≈0 to 72.29 cents with an average of 6.58 cents and standard deviation of 8.89 cents. Finally, the distribution of the normal volatility also has a positive skew with a fat right tail. The average volatility is 62.08 bps with

Figure 9.2: Distribution of swap rates, prices, and volatilities in the swaption dataset

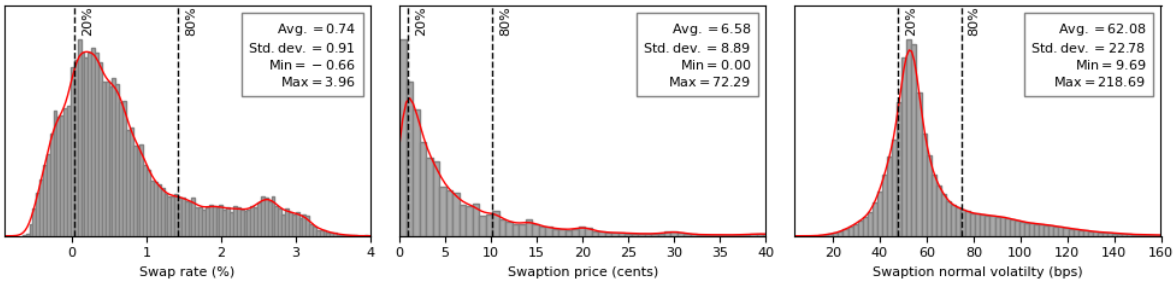

a standard deviation of 22.78 bps, while the observations range from 9.69 to 218.69 bps.

In appendix A, we summarise the average implied volatility and price by swaption term and tenor. 1Y1Y swaptions has the largest average implied normal volatility and lowest average price, while the inverse is true for 30Y30Y swaptions. The average implied volatility decreases monotonically in both the term and tenor dimension, while conversely the average swaption price increases monotonically in both the term and tenor dimension. Finally, we note that Figure 5.1 provides examples of the market implied volatility smiles based on this data across different dates and swaption contracts.

10 EMPIRICAL METHODOLOGY

The tools needed to understand the practical aspects of the analysis are provided in this section. In section 7, we introduced the theory of three different SABR extensions which all share the feature of being able to model negative interest rates, namely the shifted SABR, the normal SABR, and the free-boundary SABR. The models are viewed as potential contenders for our purpose of identifying the most robust SABR model across market conditions. We therefore opt to implement all three in the empirical analysis. We initially need a method of converting the market quotes of implied normal volatility into implied Black volatilities which is described in section 10.1. The all-important part of the empirical analysis is the implementation of the calibration procedures for each of the SABR models, which is accounted for in section 10.2. The analysis of the implemented models is, among other things, done using an in-sample and out-of-sample analysis, and we provide the context of this in 10.3. Finally, we want to empirically test the risk management measures of the SABR model, and we explain how these quantities are numerically calculated in section 10.4. The empirical analysis is implemented in python, and the full implementation is provided in an external attachment.

10.1 Converting Volatilities

Under the normal and Black models, the volatility parameter has a different meaning. The Black volatility is measured in terms relative to the rate, so an increase in from 1.0% to 1.2% gives a Black volatility of 20%. A consequence of this is that small changes in near-zero rates will make the values abnormally large. The normal volatility is measured in absolute terms, so the normal volatility of the example would be 0.2%. The swaption data is quotes in terms of the implied normal volatility, so to be able to fit the SABR models under a Black calibration space, we need a way to convert between the volatilities. Given a normal volatility σ_N, we can price the swaption under the normal model (where we derive the par swap rate using the bootstrapped zero curve). Under a no-arbitrage assumption, this price must be equal to the price under the Black model for a certain implied Black volatility σ_B^*. Thus:

$$\mathcal{N}(F_t, K, t, T, \sigma_N) = \mathcal{B}(F_t, K, t, T, \sigma_B^*) \tag{10.1}$$

We can infer out σ_B^* using an iterative procedure. Root finding algorithms are usually sensitive to the initial guess, so having a good initial guess means a more accurate volatility conversion. Choi et al. (2022) provide the following approximation for converting normal volatility to Black volatility:

$$\sigma_B(K) \approx \frac{\sigma_N}{F_0\sqrt{k}}\left(1 + \frac{\sigma_N^2}{24\,k\,F_0^2}T\right) \Big/ \left(1 + \frac{\ln^2 k}{24}\right), \qquad \text{for } k = \frac{K}{F_0} \tag{10.2}$$

When solving for Black volatilities, we thus use the approximation in equation (10.2) as the initial guess and then apply the fsolve root finder algorithm from the scipy.optimize package in python. The shifted Black volatility is found in the same way except for a shift is added to swap rate and the strike.

10.2 Model Calibration

When calibrating the models, we fit a set of model parameters for each date and swaption contract in the sample, which means we fit a total of 20,600 smile. Expect for in the out-of-sample analysis, we fit each smile based on 9 market quotes and use an optimisation algorithm to minimise the sum of squared errors (SSE) between the market values and the model values. Models calibrated to option volatilities fit under the constrained optimisation problems formalised in section 6.4.2. Likewise, models calibrated to option prices fit under the constrained optimisation problem formalised in section 6.4.3. We apply a shift of $s = 5\%$ to models requiring a shift parameter, i.e., the shifted SABR model and the Black model. This appears to be a larger shift than what is otherwise applied in the literature, but we maintain it as the lowest strike in the sample is -2.5%, and we wish to apply a buffer to mitigate the risk of the bound being hit in the future as a practitioner would do in real applications.

The largest challenge of the calibration is the choice of an **optimisation algorithm**, and we conduct extensive testing on a subsample of the dataset to ensure we arrive at the optimal method. For any algorithm requiring an *initial guess*, we choose it arbitrarily for the first calibration in the dataset after which we continuously update it to be equal to the optimal parameters of the previous calibration. We do this to aid convergence of the optimisers, but we also risk not reaching a global solution. There is a clear and very expected trade-off between computational efficiency and accuracy. This combined with the varying levels of complexity in the optimisation problems mean that we deem it appropriate to apply different procedures to different models. In all cases, the algorithms are from scipy.optimize.

- **Shifted/Normal SABR (Hagan et al.):** Calibrating the shifted SABR and normal SABR models to option volatilities using the approximation formulas of Hagan et al. (2002) is relatively straight forward to implement. We use the SLSQP optimisation method, where we use the method of continuously updating the guess. The convergence seems quite good, and there are only a few data points in the sample with more questionable solutions. Overall, it takes around 10 minutes to fully calibrate a shifted or normal SABR model across the entire dataset.
- **Free-boundary SABR (Hagan et al.):** Calibrating the free-boundary SABR to option volatilities using the approximation formulas of Hagan et al. (2002) proves more of a challenge. We attempt to force a better convergence and apply the differential evolution algorithm, which works as a global optimiser and does not require an initial guess. The calibration time of the free-boundary SABR across the entire dataset is around 1 hour, i.e., ten-fold the other models under the Hagan models.

- **Shifted/Normal/Free-boundary SABR (Antonov et al.):** The remaining models are calibrated to option prices using the Antonov solutions. We encounter major problems in terms of the convergence of the optimisation algorithm, and the calibration time becomes explosive, even when using the approximation for the kernel function. The long calibration time hinders the possibility of using a global optimiser, while most local optimisers prove too sensitive to the initial guess. The only feasible choice is to settle on using SLSQP with a very small tolerance and update the initial guess after each smile calibration. To further aid the convergence, the infinite upper bounds of the integrals are changed to a bound of 100 out of necessity. In the end, calibrating each model to the dataset takes between 10 and 32 hours depending on the model and the number of free parameters.

Different steps could have been taken as a part of aiding the calibration convergence, e.g., approximating the initial guess using the Hagan et al. (2002) approximation formula or attempt to identify a superior setting for the optimiser. As we define the ease of implementation as a consideration when assessing the models, we opt to not take any further steps in improving the calibration to option prices.

10.3 In-Sample and Out-of-sample Testing

The in-sample and out-of-sample testing is our primary way of uncovering the shortfalls of our models and testing the quality of the model fit across different market conditions, which is an important perspective in connection with the analysis of the parameter stabilities and risk sensitives. The **in-sample analysis** is done by taking the calibrated model parameters and using them to price all swaptions in the dataset. For models calibrated to option volatilities, we use the calibrated parameters for each observation in the sample to calculate the SABR implied volatility, and then use this in either the Bachelier or Black model to price each swaption. Models calibrated to option prices readily provide the swaption price without any extra steps. We define the performance relative to the market in different ways to mitigate the risk of overstating or understanding errors. In particular:

$$\text{MPE} = \frac{1}{n}\sum_{i=1}^{n}\frac{(\hat{x}_i - x_i)}{\hat{x}_i} \tag{10.3}$$

$$\text{MAPE} = \frac{1}{n}\sum_{i=1}^{n}\left|\frac{(\hat{x}_i - x_i)}{\hat{x}_i}\right| \tag{10.4}$$

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n}|\hat{x}_i - x_i| \tag{10.5}$$

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(\hat{x}_i - x_i)^2 \tag{10.6}$$

where we let $\hat{x}_i$ be the market value and x_i be our estimated value over n observations. The magnitude of errors is measured by MAE (in absolute terms) and MAPE (in percentage terms). MPE is used to indicate the direction of the errors, with a negative value indicating our model estimates the value to be too

largest. MSE can be thought of as a measure for the volatility of errors. We apply the biggest emphasis on the MAPE, as it provides a nicely scaled relative comparison. These functions area alo used outside the in-sample analysis when assessing the fit between market and model volatilities.

The **out-of-sample analysis** is done to test how accurate our models can reproduce volatilities which it has not seen before. This is done by excluding a certain level of moneyness from our dataset, recalibrating the models to the remaining data points, and using these new parameters to estimate the volatility for all data points including those removed. We split this analysis into two distinct parts. The first part concerns the ability of the models to price near-the-money swaptions by removing ± 25bps observations, respectively, from the calibration. The second part assess the ability of the models to price deep ITM and OTM swaptions by removing ± 200bps observations. In all cases, we compare the pricing capabilities of the models by looking at the MAPE of the price relative to the market price.

10.4 Computation of Risk Sensitivities

The final perspective on the SABR model is its risk sensitivities, which are given as partial derivatives and the swaption value. We numerically approximate these measures in the analysis by using finite difference methods, and the most accurate way to do this is using the method of central difference (Crispoldi et al., 2015). This, first-order partial derivative can be approximated as:

$$\frac{\partial V}{\partial u} \approx \frac{V(u + \Delta u) - V(u - \Delta \mathrm{u})}{2\Delta u} \tag{10.7}$$

The second-order partial derivative can be found in the same fashion as:

$$\frac{\partial^2 V}{\partial u^2} \approx \frac{V(u + \Delta u) - 2V(u) + V(u - \Delta \mathrm{u})}{(\Delta u)^2} \tag{10.8}$$

In these formulas, $V(u)$ is the function we want to differentiate with respect to u, while Δu represents an infinitesimal bumping amount which is applied to the variable we are differentiating with respect to. We set $\Delta u = 1e - 6$ in the empirical analysis, which seems to give satisfactory results.

11 EMPIRICAL ANALYSIS

The main matter of thesis is found in this section, where we conduct a thorough empirical analysis of the shifted SABR, normal SABR, and free-boundary SABR models with the purpose of determining which of the models is best suited for modelling the volatility smiles of our dataset. The analysis consists of six parts. Section 11.1 concerns the parameterisation and calibration of the models. Here we test the fit of a broad range of model parameterisations to determine the optimal choices. We also compare the performance of calibrating to market volatilities using the Hagan formulas against calibrating to market prices using the Antonov solutions. We refer to each of these calibration spaces with a prefix H or A in the name of the model. Section 11.2 is a visual analysis on the fit of the market smiles to uncover any structural issues. Section 11.3 concerns extensive testing of the stability of the parameters. Section 11.4 deconstructs the in-sample performance to uncover under which conditions the models perform poorly. Section 11.5 performs an out-of-sample analysis by re-fitting our models while leaving out certain quotes to test if the models can accurately interpolate the smiles. Section 11.6 finally calculates parameter sensitivities and empirical Greeks to understand the hedging capabilities.

11.1 Choosing a Model Parameterisation

The SABR models can be specified in a wide range of ways, and a thorough assessment of each of the methods is required to identify which model parameterisation yields the best fit to the observed market data. We separate this analysis into two parts, with the parameterisation of the models under the Hagan approximations and Antonov solutions considered in section 11.1.1 and 11.1.2, respectively.

11.1.1 Calibrating to Swaption Implied Volatilities

We analyse the optimal parameterisation of the H-shifted, H-normal, and H-free-boundary SABR models when calibrating their parameters to market quotes of implied volatility via the Hagan approximation formulas. The practical implementation of the optimisation procedure is outlined in section 10.2.

11.1.1.1 H-Shifted SABR

The H-shifted SABR model offers the broadest range of ways to specify the model. Firstly, the model can be fitted in either a normal or Black calibration space. Secondly, the model parameterisation can be done in different ways with varying consequences. We maintain the standard practice of freely estimating ρ and ν, leaving us with the α and β parameters to specify. In the following, we compare the fit to market data across 24 shifted SABR models each with a different parameterisation or calibration space. For all models, we apply the shift $s = 5\%$ to the rate and strike of the swaptions.

Calibration space. We first consider the implication of calibrating the model to normal or shifted Black volatilities, which means we empirically test the performance of the Hagan et al. (2002) normal approximation against the equivalent shifted Black approximation. We fit 12 different model parameterisations under each calibration space and summarise the average MPE and MAPE across them in Table 11.1. We find that calibrating the models in a normal calibration space on average yields a *slightly* lower MPE and MAPE compared to under a shifted Black calibration space. This result is however not strong enough for us to draw any conclusions on accuracy without conducting a more in-depth analysis of the risk-neutral PDF under each approximation. For now, we consider the differences in fit negligible, making the choice of calibration space rather arbitrary.

Specifying α. The initial volatility parameter α can either be implied from the ATM volatility or estimated freely. For each α specification, we fit 12 different models and summarise the average MPE and MAPE across them in Table 11.1. We find that implying α from the ATM volatility on average gives a lower MAPE than estimating α freely. We also conclude that under a free α, the models tend to on average overestimate the market volatility as seen by the negative MPE, while the opposite is true under an implied α. Overall, there seems to be evidence in our data that implying α improves the fit of the models. This is not an unexpected finding, as implying α reduces the number of free parameters and thus aids the alleged over-parameterisation of the SABR model.

Table 11.1: Average MPE and MAPE from calibrating the H-shifted SABR model to the market implied volatilities

Volatility	**No. models**	**Avg. MPE**	**Avg. MAPE**
Normal	12	0.015%	0.723%*
Shifted Black	12	0.054%	0.736%

$\boldsymbol{\alpha}$	**No. models**	**Avg. MPE**	**Avg. MAPE**
Free	12	-0.026%	0.756%
Implied	12	0.095%	0.703%*

$\boldsymbol{\beta}$	**No. models**	**Avg. MPE**	**Avg. MAPE**
0.25	4	0.068%	0.979%
0.50	4	0.047%	0.846%
0.75	4	0.031%	0.702%
1.00	4	0.012%	0.643%
Fitted	4	0.024%	0.632%
Free	4	0.026%	0.573%*

Determining β. When it comes to the parameterisation of β, we generally have three options: (i) fixing it at an a priori value, (ii) fitting it based on historical observations of the backbone, or (iii) estimating it freely. We first consider the implications of *fitting* β. Following section 6.4.1, we regress historical observations of the log-rates against ATM log-volatilities and fit a β estimate for each swaption contract to account for the findings of West (2005). The regression results for a subset of swaptions fitted using a shifted Black calibration space can be seen in appendix B, and we summarise the results via Figure 11.1. Graph (a) shows the log-log plot for a 5Y5Y swaption fitted using weekly data from 27 March 2019 to 1 March 2023, where the regression slope implies an estimate of $\beta = 1.76 \notin [0, 1]$. Graph (b) shows the same plot for a 15Y15Y swaption, where the slope implies a permitted value of $\beta = 0.54$. Clipping the implied β values to the parameter bounds, the backbone of most swaptions in the sample imply that a $\beta = 1$ is optimal. Overall, we find a tendency for the β implied by the backbone to become smaller as the swaption maturity increases, which is in line with West (2005). Our analysis is inherently

biased with data leakage as we fit the regressions across the whole sample, and we find that the β estimates are very sensitive to how much data is included in the regression. Our dataset includes periods of (i) negative rates, (ii) high levels of uncertainty around the Covid-crisis, and (iii) rapidly increasing rates and inverted yield curves. These very different market conditions make it unsurprising that the fitted β is sensitive to the data provided, and we could argue that it should change over our sample time.

Figure 11.1: Log-log plots of shifted swap rate against ATM shifted Black volatilities for 5Y5Y and 15Y15Y swaptions.

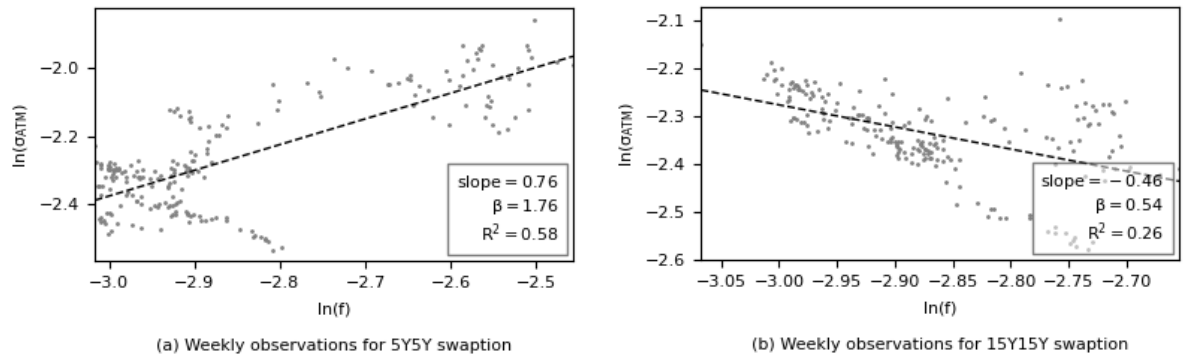

(a) Weekly observations for 5Y5Y swaption

(b) Weekly observations for 15Y15Y swaption

We can now compare how the three different β specifications fit our data, where we specify the apriori fixing as $\beta \in \{0.25, 0.50, 0.75, 1.00\}$ to encapsulate a large part of the permitted values. We fit 4 different models under each β specification and summarise their performance in Table 11.1. We find that the β specification has a dramatic impact on the fit of the model, which is in direct contradiction to the findings of Hagan et al. (2002). In the apriori fixings, we find that a higher β generally yields a better model fit. Letting β be free yields the lowest average MAPE overall, but this parameterisation is often not considered as it increases overfitting risks and promotes unstable parameters. With reference to the very different market conditions of our data, we expect to see these adverse effects if we let β be free. Fitting β per swaption performs second best on an average, although the performance is comparable to fixing $\beta = 1$. This is expected as the most swaptions have $\beta = 1$ when fitting the backbone.

Overall model selection. Having assessed the average performance of each specification, we now turn to the individual models. Table 11.2 shows the SSE, MPE, and MAPE of each H-shifted SABR model with parameters calibrated to market implied volatilities. We still consider the differences in fit between the calibration spaces as negligible and opt to apply the H-shifted SABR model under a shifted Black calibration space to nicely align the shift specification. The model specification with the lowest MAPE across our sample uses an implied α and free β. Based on the discussion above, we opt to not allow a free β, although we return to the issue in section 11.3.3. By the process of elimination, the optimal specification of the H-shifted SABR model for our data is found under a shifted Black calibration space with α implied from the ATM volatility and $\beta = 1$, which is the model we consider in the analysis.

11.1.1.2 H-Normal SABR

The model selection process for the H-normal SABR is more straightforward as per default $\beta = 0$ and we only consider a normal calibration space. We let ρ and ν be free, so the only parameterisation to

Table 11.2: Results from in-sample calibration of the H-shifted, H-normal, and H-free-boundary SABR models to the market implied volatility using the Hagan approximation formulas with varying model parameterisations. The best performing models in terms of MAPE are marked with a star, while the models selected for further analysis have a stippled border.

Model specification	Volatility	α	β	ρ	ν	SSE	MPE	MAPE
H-Shifted SABR	Normal	Free	0.25	Free	Free	0.00125	-0.028%	0.975%
H-Shifted SABR	Normal	Free	0.50	Free	Free	0.00107	-0.027%	0.833%
H-Shifted SABR	Normal	Free	0.75	Free	Free	0.00107	-0.022%	0.720%
H-Shifted SABR	Normal	Free	1.00	Free	Free	0.00162	-0.027%	0.697%
H-Shifted SABR	Normal	Free	Fitted	Free	Free	0.00097	-0.020%	0.666%
H-Shifted SABR	Normal	Free	Free	Free	Free	0.00083	-0.025%	0.584%
H-Shifted SABR	Normal	Implied	0.25	Free	Free	0.00135	0.100%	0.960%
H-Shifted SABR	Normal	Implied	0.50	Free	Free	0.00115	0.076%	0.814%
H-Shifted SABR	Normal	Implied	0.75	Free	Free	0.00100	0.051%	0.668%
H-Shifted SABR	Normal	Implied	1.00	Free	Free	0.00089	0.029%	0.588%
H-Shifted SABR	Normal	Implied	Fitted	Free	Free	0.00091	0.037%	0.608%
H-Shifted SABR	Normal	Implied	Free	Free	Free	0.00086	0.039%	0.559%*
H-Shifted SABR	Shifted Black	Free	0.25	Free	Free	0.72761	-0.023%	1.018%
H-Shifted SABR	Shifted Black	Free	0.50	Free	Free	0.83284	-0.040%	0.918%
H-Shifted SABR	Shifted Black	Free	0.75	Free	Free	0.67564	-0.032%	0.748%
H-Shifted SABR	Shifted Black	Free	1.00	Free	Free	0.80350	-0.040%	0.690%
H-Shifted SABR	Shifted Black	Free	Fitted	Free	Free	0.42531	-0.018%	0.640%
H-Shifted SABR	Shifted Black	Free	Free	Free	Free	0.39787	-0.014%	0.583%
H-Shifted SABR	Shifted Black	Implied	0.25	Free	Free	0.67109	0.223%	0.963%
H-Shifted SABR	Shifted Black	Implied	0.50	Free	Free	0.58225	0.176%	0.819%
H-Shifted SABR	Shifted Black	Implied	0.75	Free	Free	0.50390	0.128%	0.674%
H-Shifted SABR	Shifted Black	Implied	1.00	Free	Free	0.45746	0.086%	0.596%
H-Shifted SABR	Shifted Black	Implied	Fitted	Free	Free	0.45564	0.097%	0.614%
H-Shifted SABR	Shifted Black	Implied	Free	Free	Free	0.43566	0.102%	0.566%*
H-Normal SABR	Normal	Free	0.00	Free	Free	0.00150	-0.027%	1.123%
H-Normal SABR	Normal	Implied	0.00	Free	Free	0.00163	0.142%	1.111%*
H-Free-boundary SABR	Normal	Free	0.10	Free	Free	0.12516	3.812%	9.139%
H-Free-boundary SABR	Normal	Free	0.25	Free	Free	0.20824	6.197%	12.486%
H-Free-boundary SABR	Normal	Free	0.49	Free	Free	0.37600	10.483%	19.285%
H-Free-boundary SABR	Normal	Free	Free	Free	Free	0.00519	0.125%	1.415%*

consider is the choice of implying α or letting it be free. The results of the two specifications are seen in Table 11.2, where we see their performance is quite similar. The MAPE is slightly higher with a free α, while the MPE implies that this model on average slightly overestimates the market volatility. As we expect that implying α aids the over-parameterisation issues, we consider the H-normal SABR model under a normal calibration space where α is implied from the ATM volatility and $\beta = 0$ in the analysis.

11.1.1.3 H-Free-Boundary SABR

We finally consider the model parameterisation of the H-free-boundary SABR, where we stick to the natural choice of a normal calibration space. We suspect that the mathematical tractability of the

approximation formula is poor, so we do not consider the possibility of implying α from the ATM volatility. This means that we always allow α, ρ, and ν to be free in our parameterisation. Recall that the forward process of the free-boundary SABR is a martingale only if $\beta \in [0, 0.5)$. Thus, we test a priori fixings of $\beta \in \{0.10, 0.25, 0.49\}$ and compare it to letting β be free. The results of the calibrations are seen in Table 11.2 above. The best fitting model specification across all performance measures by quite a large margin is letting all parameters be free. Based on the large difference in performance between the free β and the a priori fixings, we suspect an overfitting issue and disregard this parameterisation. In contrary to the H-shifted SABR, the fit of the H-free-boundary SABR improves as the β parameter decreases. Thus, for the forthcoming analysis, we consider the H-free-boundary SABR model calibrated under a normal calibration space with $\beta = 0.10$ and all other parameters are free.

11.1.2 Calibrating to Swaption Prices

We now analyse the optimal parameterisation of the A-shifted, A-normal, and A-free-boundary SABR models when calibrating their parameters to market prices via the Antonov solutions.

11.1.2.1 A-Shifted SABR

The exact solution of the A-shifted SABR model is only valid when $\rho = 0$, and we let α and ν be free in the estimation. This leaves us with the task of determining the β specification, where we opt to compare the performance of a free β to apriori fixings of $\beta \in \{0.25, 0.50, 0.75, 0.99\}$. Table 11.3 shows the calibration results. We find that letting β be free yields the best fit to our data. Hagan et al. (2002) argue that ρ and β have similar impacts on the volatility smile and use this when justifying the choice of fixing β. We are considering the opposite case of a fixed $\rho = 0$, so by inverting the argument we can justify allowing β to be free. Thus, given our dataset, the optimal model parameterisation of the A-shifted SABR model calibrated using Antonov's solution has $\rho = 0$ and allows α, β, and ν to be free.

11.1.2.2 A-Normal SABR

The A-normal SABR model has only one logical choice in terms of model parameterisation as the exact solution is valid under general correlation, which means we set $\beta = 0$ and allow α, ρ, and ν to be free. The calibration results are seen in Table 11.3. The model has an extremely large MAPE of 50.3% across the whole sample, which is alarming. When inspecting the calibrated data closer, we find extensive evidence of a faulty calibration procedure. For example, the model consistently estimates the ATM volatility to be very near zero when rates are negative. The ν parameter takes on explosive values throughout the sample, α remains at near-zero levels, and ρ exhibit random jumps between positive and negative correlations. At this stage, we cannot credit the issues to anything other than convergence, so the only option is to use a heavier optimiser or re-evaluate the entire procedure. The complexity of implicating and calibrating the model is a key consideration in our analysis, so we conclude that the A-normal SABR model is unsuited due to heavy implementation work and large calibration time.

Table 11.3: Results from in-sample calibration of the A-shifted, A-normal, and A-free-boundary SABR models to market prices using the Antonov solutions with varying model parameterisations. The best performing models in terms of MAPE are marked with a star, while the models selected for further analysis have a stippled border.

Model specification	α	β	ρ	ν	SSE	MPE	MAPE
A-Shifted SABR	Free	0.25	0.00	Free	0.80085	3.229%	5.322%
A-Shifted SABR	Free	0.50	0.00	Free	0.67661	1.896%	4.615%
A-Shifted SABR	Free	0.75	0.00	Free	1.01486	0.566%	4.671%
A-Shifted SABR	Free	0.99	0.00	Free	1.74727	-0.566%	5.345%
A-Shifted SABR	Free	Free	0.00	Free	0.25515	1.506%	2.885%*
A-Normal SABR	Free	0.00	Free	Free	14.67367	-43.012%	50.294%*
A-Free-boundary SABR	Free	0.10	0.00	Free	4.72736	-5.195%	16.026%
A-Free-boundary SABR	Free	Free	0.00	Free	4.35022	-4.166%	13.428%*

11.1.2.3 A-Free-Boundary SABR

The exact solution of the A-free-boundary SABR model is only valid when $\rho = 0$, and we always let α and ν be free in the estimation procedure. As was the case for the A-shifted SABR model, we are left with the task of determining the β specification. The calibration procedure of the A-free-boundary SABR take an exceptionally long time, so we do only a single a priori fixing of $\beta = 0.10$ and compare this to letting β be free. The calibration results are seen in Table 11.3, and we find that a free β performs better. Using the same logic as earlier, we justify a free β when correlation is taken constant. Thus, given our dataset, the optimal model parameterisation of the A-free-boundary SABR model calibrated using Antonov's solution has $\rho = 0$ and allows α, ρ, and ν to be free.

11.2 Fitting Market Smiles

The analysis of the optimal model parameterisation for our dataset in the preceding section concluded on five optimal model across the three SABR types, which we now limit the analysis to. We disregard the A-normal SABR due to extensive issues when calibrating to market prices.

Practitioners primarily use the SABR model to interpolate swaption volatilities across strikes, which naturally requires accurate and smooth modelling of the smiles observed in the market. We use this as a starting point for a visual analysis of the model fitted smiles compared to the market ones to uncover any structural issues. The ability of the models to fit a smile is expected to be impacted by the horizontal placement of the smile, i.e., the level of the swap rate. Based on this, we illustrate the calibrated smiles for each model on 1 March 2023 and 11 March 2020, which are the dates with the highest and lowest average swap rates, respectively, in the sample. Conveniently, 11 March 2020 coincides with the extreme market conditions seen at the start of the Covid-19 crisis and can thus be seen as a stress test scenario. In addition to the level of the swap rate, we also expect the smile fit to vary with the liquidity of the swaption due to stale quotes, so we illustrate the fit for 1Y5Y, 5Y7Y, 10Y10Y, and 20Y20Y swaptions where we expect less liquidity with longer maturity and tenor.

Shifted SABR. The volatility smiles fitted using the optimal H-shifted SABR model are seen in Figure 11.2. The smiles in plot (a) show a near-perfect fit across all swaptions, which is an initial indication of a good ability to fit the market smiles under normal market conditions. Plot (b) shows a much poorer performance with none of the fitted smiles managing to capture even a single market observation. If we look at the placement of the market implied volatilities, we notice quite abnormal shapes to the smiles, especially for 20Y20Y and 10Y10Y swaptions, with evidence of humps and even inverted smiles. We do not expect the models to be able to capture such extreme market conditions, but we will comment on this issue later. The fitted volatility smiles under the optimal A-shifted SABR model are seen in Figure 11.3. The smiles in plot (a) show a worse performance for the A-shifted SABR, with deviances from the market quotes for at the shortest maturity, although not enough to warrant any definitive conclusions. The smiles in plot (b) under the stress scenario also do not capture the abnormal market smiles. The fitted smiles of the A-shifted SABR generally tends to be placed higher than in the H-shifted SABR, which might imply a structural error in the model.

Figure 11.2: Sample of fitted volatility smiles under the H-shifted SABR with $s = 5\%$ and $\beta = 1$

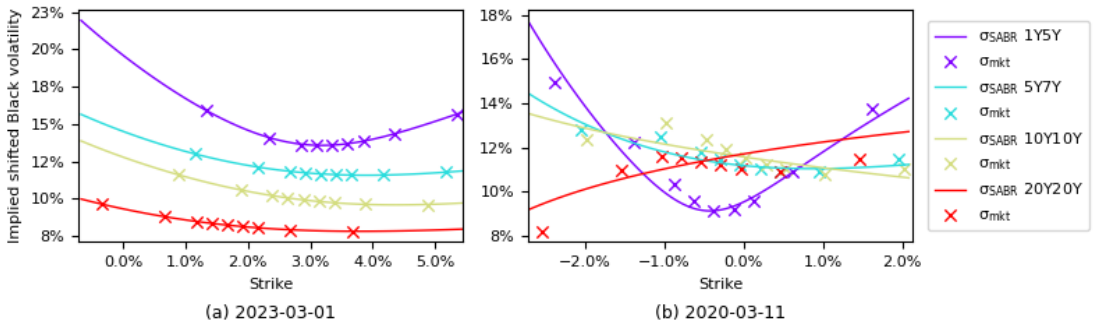

Figure 11.3: Sample of fitted volatility smiles under the A-shifted SABR with $s = 5\%$, $\rho = 0$, and a free β.

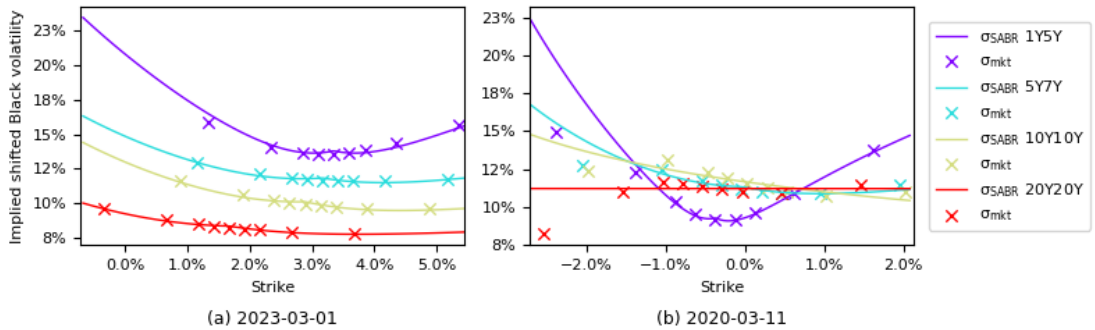

Normal SABR. Turning to the smiles fitted with the H-normal SABR model in Figure 11.4, the performance is quite similar the shifted SABR ones. Plot (a) again has a near-perfect fit with no immediately visible errors, while plot (b) shows the same inability to fit the market smiles under the stress conditions. Based purely on visual considerations, it may appear as if the normal SABR fits the stress scenario better than the shifted SABR, although the smiles are quoted in two different volatility terms which makes the comparison hard. We assess the comparable performance more in section 11.4.

Figure 11.6: Fitted volatility smiles for the H-normal SABR model with $\beta = 0$.

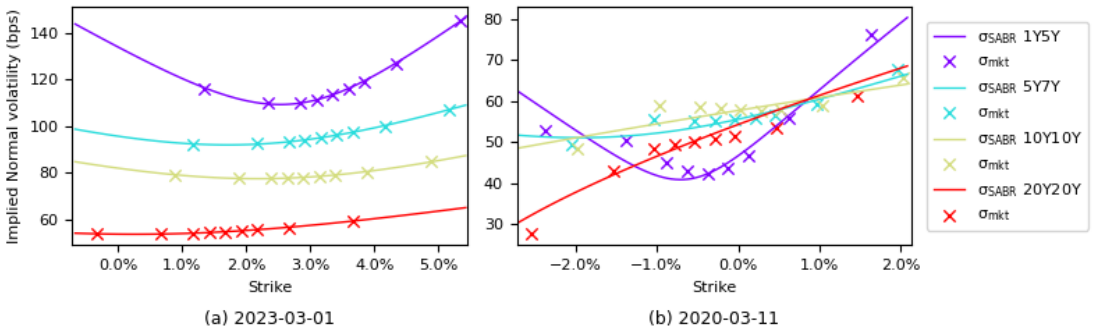

Free-boundary SABR. The smiles fitted using the H-free-boundary SABR are seen in Figure 11.5, where we immediately notice structural issues. One the high-rate date in (a), the fit to the market is all-around good, but the continuously fitted curve (along which interpolation would happen) shows issues at strikes near zero. This is explained by the function $C(f) = |F_t|^\beta$ having a singularity at zero, which also violates the assumptions of the Hagan approximation. The issue becomes more apparent in (b) where a large amount of market observations placed in the around zero. This singularity issue makes the current specification H-free-boundary SABR undesirable, as it will only produce robust estimates when all quotes are placed far away from zero. Figure 11.6 shows the smiles fitted using the A-free-boundary SABR, where structural errors are also seen. The high-strike date in (a) show kinks along the 1Y5Y curve around zero, while we see evidence of calibration issues in the stress scenario in (b), where the model yields odd and discontinuous smiles. These observations along with the abysmal calibration time of the A-free-boundary SABR, means that both free-boundary SABR models break multiple of our requirements for a good model, which causes us to omit it from the analysis going forward.

Figure 11.5: Fitted volatility smiles for the H-free-boundary SABR model with $\beta = 0.1$

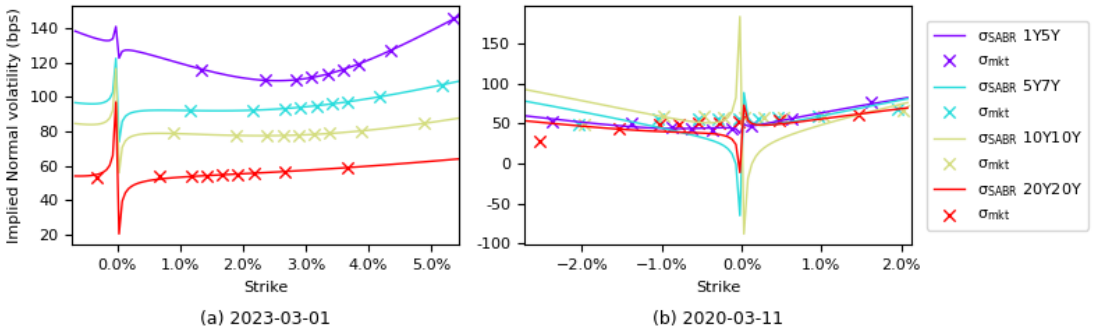

Figure 11.4: Fitted volatility smiles for the A-free-boundary SABR model with $\rho = 0$ and a free β.

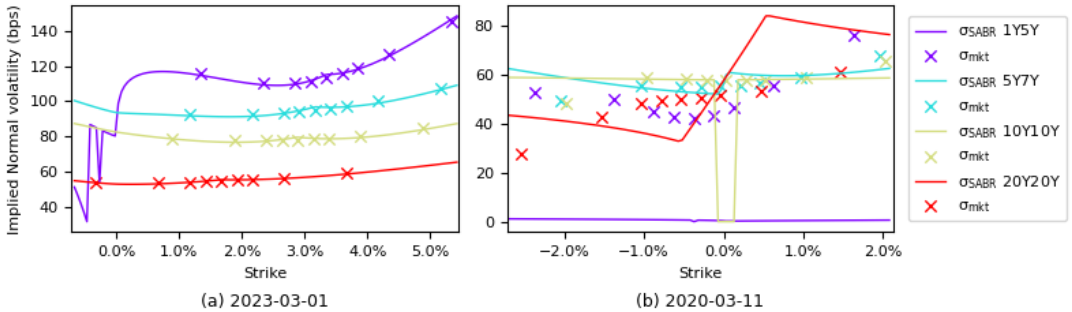

11.3 Parameter Stability Analysis

The free parameters in the SABR models are recalibrated at each point in our sample and are thus subject to change over time to accommodate different smile levels and curves. Expect for under extreme market conditions, the parameters for a given market smile should not exhibit large value jumps between calibration points, as this is an indication of either a poor implementation of the calibration or inability of the model dynamics to fit the market, i.e., a model misspecification. For real-life implementations, assessing the rate of change in the parameter values is important to ensure the parameters are updated frequently enough to capture changes in the volatility smiles. Parameter stability is also important in terms of risk sensitivities and Greeks, as these measures will be impacted by large parameter changes. We use this as motivation for dedicating a part of analysis to assessing the parameter stability of the three optimal model specifications identified for our dataset.

Table 11.4: Summary statistics for the values of α, β, ρ, and ν across the entire sample for the optimal specifications of the H-shifted SABR, H-normal SABR, and A-shifted SABR models.

	α				β			
	Avg.	Std.	CV	Observed range	Avg.	Std.	CV	Observed range
H-Shifted SABR	0.099	0.022	0.226	[0.031, 0.215]	1.000	0.000	-	-
H-Normal SABR	0.006	0.002	0.370	[0.002, 0.017]	0.000	0.000	-	-
A-Shifted SABR	0.046	0.038	0.831	[0.002, 0.200]	0.629	0.331	0.526	[0.001, 0.999]

	ρ				ν			
	Avg.	Std.	CV	Observed range	Avg.	Std.	CV	Observed range
H-Shifted SABR	-0.094	0.340	3.597	[-0.999, 0.999]	0.298	0.184	0.616	[0.001, 1.352]
H-Normal SABR	0.399	0.244	0.612	[-0.999, 0.999]	0.300	0.196	0.654	[0.001, 1.529]
A-Shifted SABR	0.000	0.000	-	-	0.341	0.249	0.731	[0.001, 1.749]

Table 11.4 show summary statistics for the parameters of each model across the entire sample, where we use the coefficient of variation ($\text{CV} = \sigma/|\bar{x}|$) to compare the variability of the parameters. The **initial volatility** α is relatively stable under the H-SABR models, with the least variability being found in H-shifted SABR. Both models are calibrated by implying α, which could have an impact on the stability. The A-shifted SABR has a much less stable α with a CV almost four times the size of the H-shifted SABR. The **vol-vol parameter** ν exhibit similar properties across all three models, although it is slightly more stable under the H-shifted SABR. Looking at the observed range of ν, each model has values ranging from the minimum of 0.01 to values above 1. This is theoretically viable but could also indicates an explosive volatility of volatility at certain data points, which is indicative of a poorly specified model. The **power parameter** β is free under the A-shifted SABR, where it takes on an average value of 0.63 with a CV on 0.52, which seems like a reasonable amount of variability compared to the ranges of the other parameters. The most concerning observation is seen for the **correlation parameter** ρ with a high degree of variability under the H-shifted SABR model.

With an initial concept of how the parameters of each of the models vary, we extend the analysis to a more intuitive dimension by looking at the parameter values over time for a 10Y10Y swaption contract in section 11.3.1. We then check if there are any specific terms or tenors in our option, which the models exhibit large variation on in section 11.3.2. Finally, we add a perspective to the analysis in section 11.3.3 by assessing the stability of the H-shifted SABR model with a free β.

11.3.1 Parameter Values over Time

To illustrate the parameter stability over time, we plot the parameter values and their relative change over each date in the sample using the example of a 10Y10Y swaption. Figure 11.5 shows this for the **H-shifted SABR model**, where we have capped the relative changes at 100%. For this contract, we generally find relatively stable parameters, although there are some extreme one-off changes in the start of the sample. Looking first at α, its value is stable, and the development is smooth over time. Its relative changes are also quite small, and we calculate an average change of 0.13% with a standard deviation of 3.36%. The correlation ρ has a larger variability and a rougher development over time, which is especially bad towards the start of the sample. The average change in ρ is -1.48% with a standard deviation of 33.20%. The vol-vol ν has a slightly rough development, but the changes in absolute terms are undramatic. In relative terms, the variability becomes more pronounced where the relative change has a mean of 1.53% and a standard deviation of 20.64%. We find that it is the parameters which influence the curve of the volatility smile, ρ and ν, that have the largest instability. Another important point is that the drop in ρ and ν occurs on the same date as the ill-fitting volatility smiles in section 11.2, and it can thus be attributed to extreme market conditions rather than a model error. All around, the parameter stabilities of the H-shifted SABR do not give cause for concern based on this particular contract.

Figure 11.7: Parameter values (left) and their relative change (right) for a 10Y10Y swaption under the H-shifted SABR.

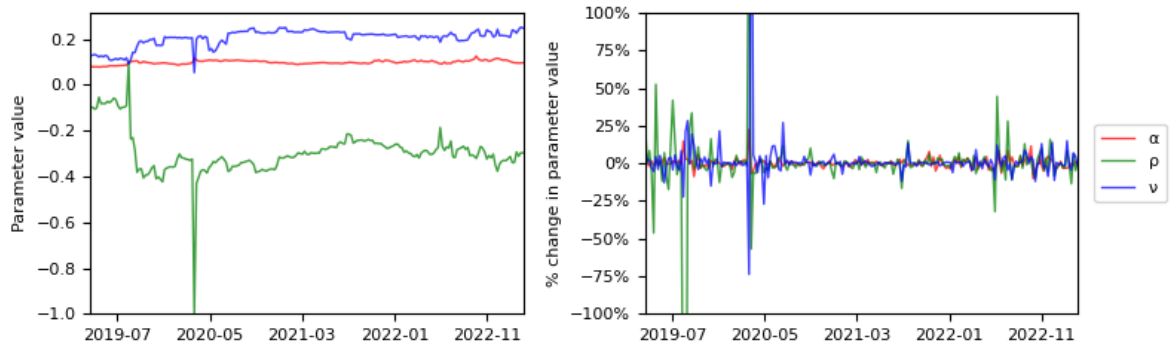

The plot for the same 10Y10Y swaption under the **H-normal SABR model** is seen in Figure 11.6. We again find that α is the most stable parameter, with a mean the relative change of 0.19% and a standard deviation of 2.54%. The actual value is diminishingly small and visually there are no changes to be seen. We notice see the same one-off deviations in ν and ρ as under the H-shifted SABR, although the direction of ρ is inverse due to the difference in β specification. The development of ν is all-around

comparable to the H-shifted SABR. ρ appears to be the most unstable parameter with large jumps both in absolute and relative terms. It has a mean change of 1.85% with a standard deviation of 29.86%. The observation that ρ is especially unstable at the end of the sample where interest rates are increasing could imply that $\beta = 0$ does not match the dynamics of the backbone. The conclusion is that the parameters of the H-normal SABR are less stable for a 10Y10Y swaption than the H-shifted SABR.

Figure 11.8: Parameter values (left) and their relative change (right) for a 10Y10Y swaption under the H-normal SABR.

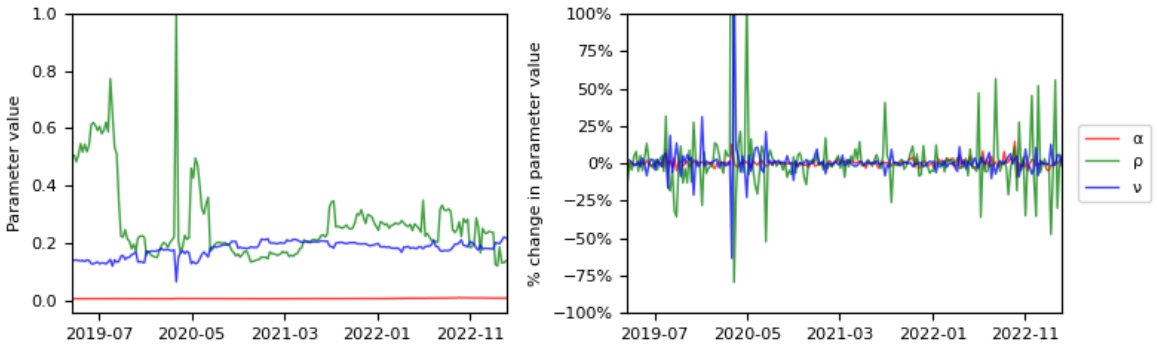

We finally consider the **A-shifted SABR** model in Figure 11.7, which seem to have the least stable parameters compared to the H-SABR models. The movements of α are larger than under the H-models with an average change of 0.68% and a standard deviation of 17.09%. The development of ν is very similar for all three models, taking on the same shape and level. Finally, the β parameter is seen to vary a lot throughout the sample, even more so than ρ does in H-SABR model. While A-shifted SABR model is least stable for this swaption, the values are not too concerning when the explained the market stress scenarios are considered, and especially the relative changes are comparable to H-normal SABR.

Figure 11.9: Parameter values (left) and their relative change (right) for a 10Y10Y swaption under the A-shifted SABR.

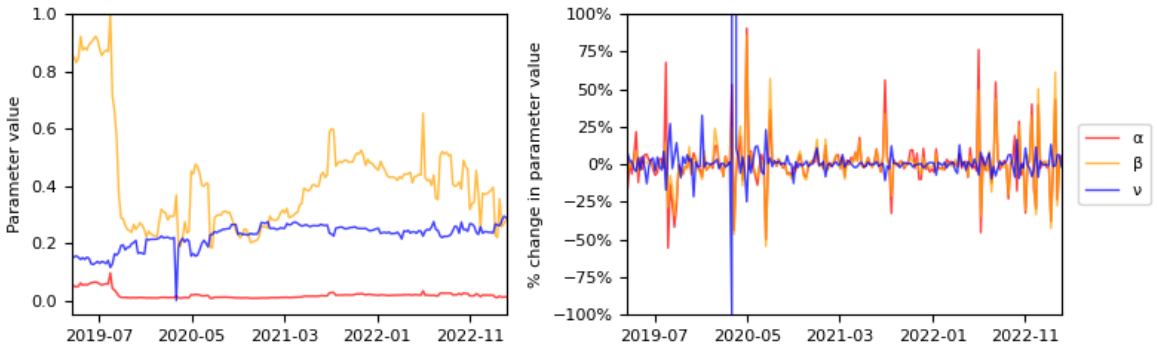

11.3.2 Parameters over Terms and Tenors

A relevant perspective on the parameter stability is how it changes parameter changes across different terms and tenors for each model. We use the standard deviation of the weekly relative change in parameter values to quantify the differences across terms and tenors. In the following, note that we cap the scale of the standard deviation at 100% although many observations lie above this. We assess the relative change rather than the absolute change as a large relative change reveals value spikes.

Figure 11.11: Standard deviation of the relative change in parameter values across term and tenor for the H-shifted SABR.

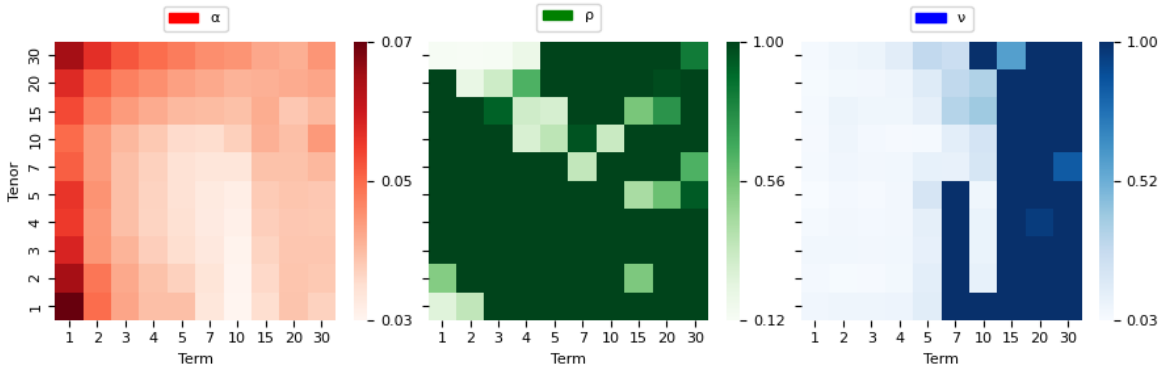

The standard deviation of the relative changes in each parameter of the **H-shifted SABR** model is seen in Figure 11.8. As expected, α has the least amount of variability across all terms and tenors with a maximum standard deviation of 6.6% for 1Y1Y swaptions. The change in the correlation parameter ρ has quite extreme standard deviations with the largest value found for 7Y15Y swaptions as 1,424.4%. We note that these extreme values are found when the parameter estimate is moving around near zero, where a small absolute change translates into a very large relative change. Nevertheless, when combining this heatmap with knowledge from plots similar to those in section 11.3.1 for different contracts, we maintain the conclusion that ρ generally has a tendency to make large movements and jumps. The vol-vol parameter ν is quite stable for short maturities across all tenors, but as the swaption maturity increases, the variability of the relative changes increases dramatically.

Figure 11.10: Standard deviation of the relative change in parameter values across term and tenor for the H-normal SABR.

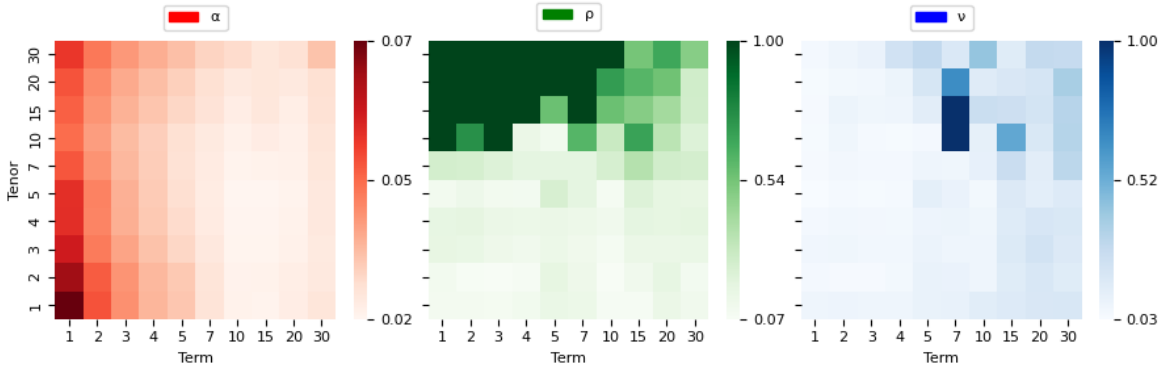

When comparing the heatmaps for the H-shifted SABR with those for the **H-normal SABR** in Figure 11.9, the immediate conclusion seems to be that the parameters of the normal SABR appear much more stable. The variability of α has a very similar level as in the H-shifted SABR, but ρ and ν have significantly lower standard deviations of the relative change for most contracts. Most of the variability in ρ is seen for short-term swaptions written on long swaps, which is in exact contrast to the H-shifted SABR. In terms of ν, the standard deviations are comparable for short maturities, but for longer maturities, the N-normal SABR exhibits a larger degree of stability.

Figure 11.10 shows the heatmaps for the **A-shifted SABR**, which reveal some stability issues. The variability in α is much worse than the H-shifted SABR model, and especially for short terms and tenors. The muddy structure in α could suggest that a non-optimal value is returned in the calibration. Both ν have a high degree of variation across short terms and long terms, more so than the H-shifted SABR model. Finally, the β parameter has the highest variation for short terms and medium-to-long tenors, which alludes to a similar pattern as what is seen for ρ under the H-normal model.

Figure 11.12: Standard deviation of the relative change in parameter values across term and tenor for the A-shifted SABR

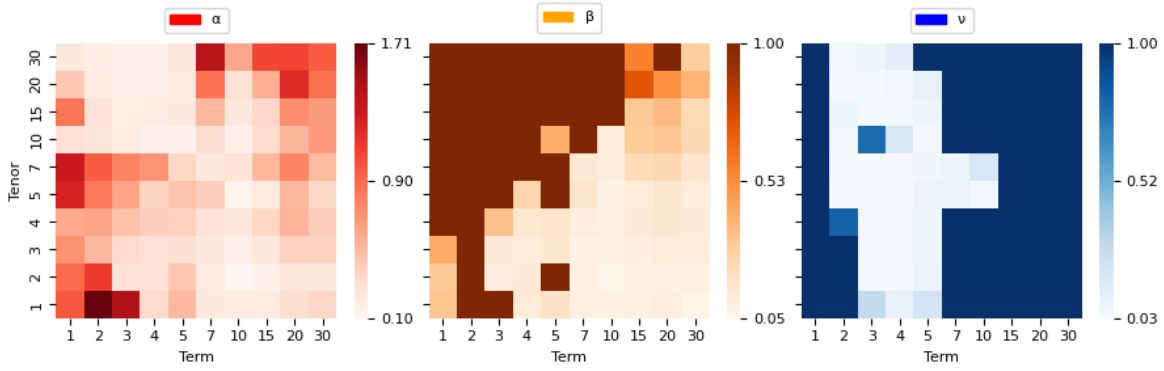

Overall, the analysis shows that the H-normal SABR model have the most stable parameter estimates, i.e., the fewest jumps in parameter values. The H-shifted SABR model has a high degree of variability in the correlation parameter controlling the curvature of the smile, while the vol-vol parameter might have some issues for longer-dated options. The A-shifted SABR model has the worst stability, with especially the vol-vol ν parameter being concerning. Hagan et al. (2002) argues that α inferred from the ATM volatility should ideally be updated several times per day, while the remaining parameters should be updated at least once per week. Our analysis here seems to imply that the parameters should be updated more frequently to avoid jumps, but an analysis with daily or inter-day data would have aided the strength of this conclusion.

11.3.3 Implication of Free Parameters

In the model calibration, we mentioned that letting β be estimated freely generally was not a preferred choice due to concerns of overparameterization and overfitting. To investigate this statement, we dedicate our attention to comparing the parameter values and their relative changes for a 10Y10Y swaption under the H-shifted SABR model with an implied α and $\beta = 1$ seen in section 11.3.1 to those under the H-shifted SABR model with an implied α and a free β. The plot for the H-shifted SABR with a free β is seen in Figure 11.13. For most of the sample, the optimal β is freely estimated to be equal to 1, so the impact on the parameter stability for this time period is negligible. From March 2022 and onwards, the estimated beta starts to slowly decline, and across the whole period, the standard deviation of the change in β is 1.36%. The impact on the relative changes in the other parameters can be seen from the right-hand plot, where especially ρ but also α exhibit more relative jumps in the parameter values.

Overall, the average relative change in all parameters increase from their values when $\beta = 1$ (3.71%, 10.66%, 20.08%), while the standard deviation of the relative changes is comparable albeit slightly higher for α and lower for ρ and ν.

Figure 11.13: Parameter values (left) and their relative change (right) over time for the H-shifted SABR with $s = 5\%$ and a free β for a 10Y10Y swaption.

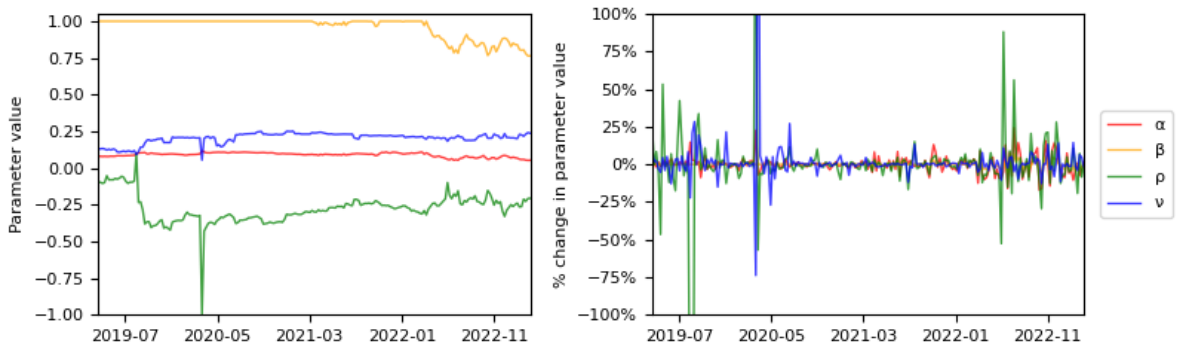

Based on this simple analysis it is hard to justify how much a free β impacts the parameter stability in the H-shifted SABR model for our dataset. While the model fit improves when letting β be free, we cannot draw conclusions without further analysis, and we maintain our optimal value $\beta = 1$.

11.4 In-sample Pricing Capabilities

We have tested the most promising models from the initial calibration for structural pitfalls and unstable parameters, and we are now left with (1) the **H-shifted SABR** calibrated to shifted Black implied volatilities using Hagan's formula with $s = 5\%$, $\beta = 1$, and α implied from the ATM volatility, (2) the **H-normal SABR** calibrated to normal implied volatilities using Hagan's formula with $\beta = 0$, and α implied from the ATM volatility, and (3) the **A-shifted SABR** calibrated to market prices using Antonov's exact solution with $s = 5\%$, $\rho = 0$, and a free β. This section conducts an analysis on the pricing capabilities of these three SABR models to test their abilities to model the smiles following section 10.3.

11.4.1 Aggregated Performance

The aggregated in-sample performance across the 185,400 swaptions in the sample is seen in Table 11.5. The **H-shifted SABR** model with an MAPE of 0.68% has superior performance compared to the other model in terms of most performance measures. The positive MPE of 0.14% indicates a slight tendency to underprice the swaptions on average. The average absolute mispricing is €24.8. The **H-normal SABR** model has a quite good performance albeit slightly worse than the H-shifted SABR model. The MAPE is 1.28% and the average absolute mispricing is €36.6. The volatility of errors measured by MSE is the lowest for the H-normal SABR. The **A-shifted SABR** model is the worst performing model, but with a MAPE of 2.88% the performance is still competitive. The average absolute mispricing is €51.3, which is almost double the H-shifted SABR. We can contrast the performance of the A-shifted

SABR to a model calibrated using the Hagen approximation with the same specification of $\rho = 0$ and a free β. The pricing errors of this model is seen in Panel B of Table 11.5. We conclude that the exact solution of Antonov appears to outperform the Hagan approximation when considering the MAPE of the models. In other words, this is an indication that calibrating the shifted SABR model to option prices improves the overall accuracy of the model. To cement this finding, we would have to engage in mapping the Antonov solutions of a general correlation case, which is beyond our scope.

Table 11.5: Total pricing errors across the whole sample for the H-shifted SABR, H-normal SABR, and A-shifted SABR assuming a notional of €100,000

Panel A. Total pricing errors across the whole sample for the three optimal models

	MSE	MAE	MPE	MAPE	Avg. est. price	Avg. mkt. price
H-Shifted SABR	12,474	24.8	0.14%	0.68%	€ 6,576.7	€ 6,578.0
H-Normal SABR	12,020	36.6	0.15%	1.28%	€ 6,575.3	€ 6,578.0
A-Shifted SABR, $\rho = 0$	13,762	51.3	1.51%	2.88%	€ 6,577.7	€ 6,578.0

Panel B. Total pricing errors across the whole sample for H-Shifted SABR with $\rho = 0$

	MSE	MAE	MPE	MAPE	Avg. est. price	Avg. mkt. price
H-Shifted SABR, $\rho = 0$	14,434	41.5	1.69%	2.97%	€ 6,576.2	€ 6,578.0

It is relevant to consider how the error of the models change throughout the sample period, as this indicates whether certain market conditions impact the pricing performance. Figure 11.14 shows the development in MAPE from 27 March 2019 to 1 March 2023 for each of the models. The following conclusions are drawn. The H-shifted SABR is consistently the best model, while the A-shifted SABR model consistently performs the worst. The shape of the H-shifted SABR and N-shifted SABR curves are very similar, and the errors move in lock-step until the H-normal SABR converges toward the H-shifted SABR towards the end of the sample. These shapes indicate that the models perform better under market conditions in the form of positive interest rates, which is seen in the last year of our sample.

Figure 11.14: Mean absolute percentage error in swaption pricing across the sample for each SABR model type.

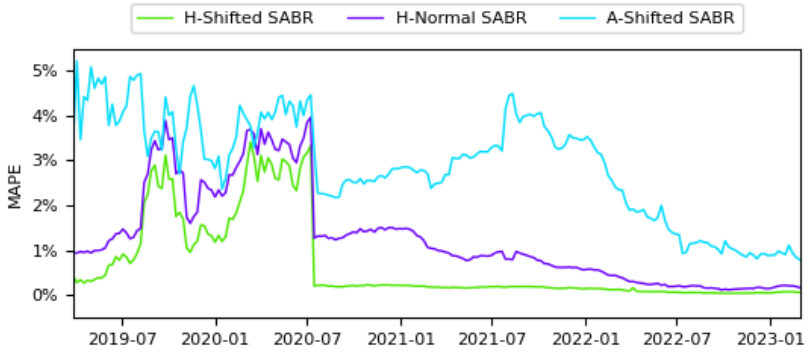

The errors of the A-shifted SABR are similar to the other two models as the start of the data sample, but during 2021 and 2022, it has a large spread to the curves of H-Shifted SABR and H-normal SABR. For all models is a noticeable drop in the MAPE from 8 July 2020 to the next observation on 15 July

2020, after which the performance of the H-models stays relatively stable and slowly converges towards 0. It is hard to explain this improvement in the models, and when inspecting the data, there are no dramatic changes in any key variables between the two dates, so the improvement remains unexplained. The analysis thus far indicates that the model that best fit our data is the H-shifted SABR model. In the coming sections we further deconstruct our dataset to determine what drives the mispricing.

11.4.2 Deconstructing the Swaptions

To test whether there are any situations where the model accuracy breaks down, we look at different cross sections of our sample to assess exactly which swaptions the models perform well or poorly on.

11.4.2.1 Moneyness

We first consider the strike level of the swaption, i.e., the moneyness. A positive moneyness means that the strike is higher than the par swap rate, and as we are pricing payer swaptions, this means that the swaption is OTM when the moneyness is positive, and vice versa. Table 11.6 shows the MAPE (panel A) and MPE (panel B) of each SABR model across different strike levels, with each level having 20,600 observations. For the H-shifted and H-normal SABR, the MAPE is smallest for ATM swaptions, while the mispricing generally increases as we move away from the money. The models perform worst on the +200bps moneyness, i.e., on deep OTM options, while deep ITM options are priced very accurately. We note that not all deep OTM options are mispriced to this extend. If we take the median absolute percentage error instead, we obtain 0.54% and 2.01% for the shifted and normal SABR, respectively. This implies that the absolute pricing errors have a positively skewed distribution, and the mean over-estimates the most common error value. For the A-shifted SABR, the most accurate pricing is seen for deep OTM swaptions, and the MAPE increase with the strike. The errors are consistently larger than the other models, and the +50bps or larger strike they become large enough to make the model unviable.

Table 11.6: Mean absolute percentage error (MAPE) and mean percentage error (MPE) across strike levels for the H-shifted SABR, N-normal SABR, and A-shifted SABR models calibrated using Hagan's formula

Panel A. Mean absolute percentage error (MAPE) when pricing swaptions across strike levels in bps

	-200	-100	-50	-25	ATM	+25	+50	+100	+200
H-Shifted SABR	0.10%	0.37%	0.37%	0.26%	0.00%	0.41%	0.74%	0.98%	2.89%
H-Normal SABR	0.16%	0.55%	0.65%	0.52%	0.00%	0.96%	1.79%	1.86%	5.01%
A-Shifted SABR	0.39%	0.78%	0.83%	0.72%	1.75%	2.11%	3.91%	6.43%	9.03%

Panel B. Mean percentage error (MPE) when pricing swaptions across strike levels in bps

	-200	-100	-50	-25	ATM	+25	+50	+100	+200
H-Shifted SABR	-0.09%	0.34%	0.32%	0.22%	0.00%	-0.31%	-0.51%	-0.34%	1.59%
H-Normal SABR	-0.16%	0.55%	0.65%	0.52%	0.00%	-0.96%	-1.79%	-1.85%	4.41%
A-Shifted SABR	-0.28%	0.16%	0.01%	-0.42%	-1.71%	1.02%	3.09%	5.10%	6.59%

The direction of the percentage pricing errors is measured by the MPE in panel B of Table 11.6. Again, the H-SABR models have a comparable performance, though the H-normal SABR generally performs

worse, and also has a more consistent under-pricing at high strikes then the H-shifted SABR. The A-shifted SABR overprices both deep ITM and ATM options, while the remaining observations are under-priced. The dip in value at the ATM point could imply a calibration issue.

11.4.2.2 Term and Tenor

A key dimension in swaptions is the term and tenor structure and how it impacts the model performance. Let us first consider the swaption term. Panel A of Table 11.7 shows the MAPE of each model across the tenors of the sample. The H-shifted SABR model does not have particularly strong term effects, although short-dated options of 1Y and long-dated options of 30Y have the largest pricing errors. We expect that the models perform poorly on longer options as the approximations may break down, but it is surprising that the model performs worse on short-term options. The H-normal SABR model exhibits similar term effects, with the lowest MAPE found for medium-term options with a tenor of 7Y-15Y. Again, the model performs worst on the short-dated options.

Table 11.7: Mean absolute percentage error (MAPE) and mean percentage error (MPE) across tenors for the H-shifted SABR, H-normal SABR, and A-shifted SABR models calibrated using Hagan's formula

Panel A. Mean absolute percentage error (MAPE) when pricing swaptions across option expiries

	1Y	2Y	3Y	4Y	5Y	7Y	10Y	15Y	20Y	30Y
H-Shifted SABR	1.46%	0.93%	0.67%	0.53%	0.43%	0.33%	0.27%	0.37%	0.68%	1.13%
H-Normal SABR	3.22%	2.11%	1.57%	1.23%	0.93%	0.72%	0.54%	0.56%	0.78%	1.12%
A-Shifted SABR	11.67%	4.42%	3.06%	2.34%	1.74%	1.34%	0.95%	0.92%	1.04%	1.37%

Panel B. Mean percentage error (MPE) when pricing swaptions across option expiries

	1Y	2Y	3Y	4Y	5Y	7Y	10Y	15Y	20Y	30Y
H-Shifted SABR	0.74%	0.42%	0.28%	0.20%	0.15%	0.09%	0.07%	-0.01%	-0.18%	-0.39%
H-Normal SABR	0.42%	0.45%	0.35%	0.28%	0.21%	0.12%	0.09%	0.00%	-0.12%	-0.27%
A-Shifted SABR	7.81%	2.34%	1.59%	1.18%	0.85%	0.52%	0.28%	0.20%	0.15%	0.14%

The A-shifted SABR model has clearer term effects, with the largest errors seen for short maturities and the error decreasing in maturity. In terms of the mispricing direction seen in panel B, both the H-shifted and H-normal SABR tend to overprice long-dated options and underprice short-dated options. This contrasts with the free-boundary SABR, which consistently underprices and to an increasingly extreme degree as the maturity increases.

Table 11.8: Mean absolute percentage error (MAPE) and mean percentage error (MPE) across swap tenors for the H-shifted SABR, H-normal SABR, and A-shifted SABR models calibrated using Hagan's formula

Panel A. Mean absolute percentage error (MAPE) when pricing swaptions across swap tenors

	1Y	2Y	3Y	4Y	5Y	7Y	10Y	15Y	20Y	30Y
H-Shifted SABR	0.79%	0.80%	0.72%	0.67%	0.65%	0.61%	0.60%	0.60%	0.65%	0.70%
H-Normal SABR	1.39%	1.42%	1.33%	1.26%	1.23%	1.17%	1.17%	1.18%	1.27%	1.35%
A-Shifted SABR	6.80%	4.23%	3.33%	3.04%	2.76%	2.24%	1.64%	1.46%	1.62%	1.72%

The other time dimension of the swaption is the tenor of the underlying swap. Table 11.8 shows the MAPE across the tenors of the underlying swaps in the sample. There is no clear structure in the mis-pricing of the H-SABR models when it comes to the aggregate performance across tenors, while the A-shifted SABR model preforms much on shorter tenors and the performance increasing in the tenor.

Having assessed the isolated performance of the models across moneyness, term, and tenor, it seems relevant to combine the dimensions to uncover any inherent structures.

11.4.2.3 Cube Mispricing

We wish to investigate how the models perform when looking both across swaption term and tenor and term and moneyness. We first focus our attention on the **H-shifted SABR model**, where heatmaps of errors across the whole sample period is seen in Figure 11.15. The left-hand plot is the MAPE across swaption term and tenors. As expected from the previous section, the pricing errors are largest for short maturities of 1Y and long maturities of 30Y, but we now also see how the errors vary with the tenor of the swap. The model is particularly poor at pricing short-term options on short-term swaps, with an absolute percentage error of up to 2.0%. There also seems to be some performance deterioration for long-term 30Y options on medium-to-long term swaps of 7Y-20Y. The model performs best on medium term options written on medium term swaps, which may be explained by these being the most liquid options found in the market. Looking instead at the right-hand plot, we see that the model performs quite well across all terms and levels of moneyness expect deep OTM options. Especially deep OTM 1YZY swaptions have a high average percentage error of 8.0%. As alluded to previously, we can compare this to the median APE of 1.95% and conclude that the distribution of absolute percentage errors is right skewed. When further inspecting the data, we find that the large errors for deep OTM short-term options are especially apparent when the swap rate is negative or near zero.

Figure 11.15: Mean absolute percentage error in pricing for the H-shifted SABR model with $s = 5\%$ and $\beta = 1$ over the whole sample period across swaption term and tenor (left) and term and strike level (right).

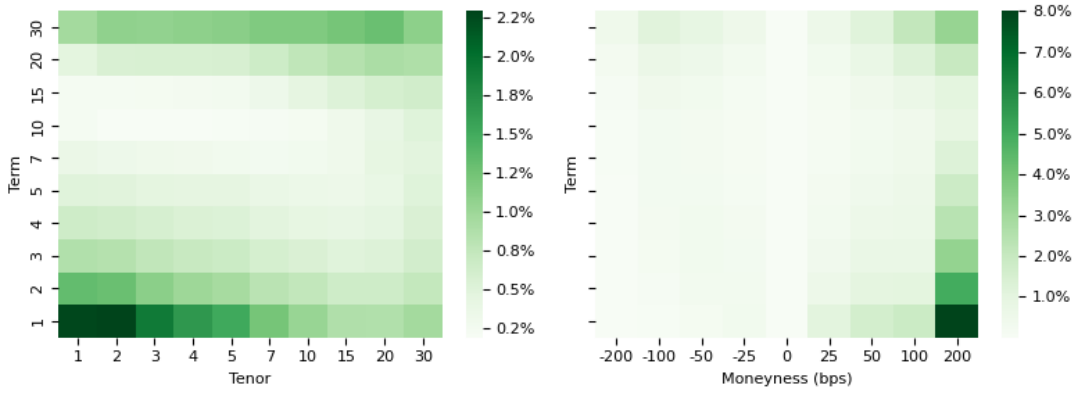

The same plot is seen for the H-normal SABR model in Figure 11.17. The conclusions of the H-shifted SABR generally also applies here, with the model performing worst on short-dated options written on short-dated swaps, and especially on deep OTM short-dated options. The poor performance for short-dated options is here more consistent across tenor. The model has the same issue as the shifted SABR

when it comes to short-dated OTM options which is mainly driven by observations with negative swap rates. We also note how the range of the errors is almost double that of the shifted SABR, implying that when the normal SABR performs poorly, it performs much worse than the shifted SABR. Still, for certain term and tenor combinations and moneyness levels, the performance is quite comparable.

Figure 11.16: Mean absolute percentage error in pricing for the H-normal SABR model with $\beta = 0$ model over the whole sample period across swaption term and tenor (left) and term and strike level (right).

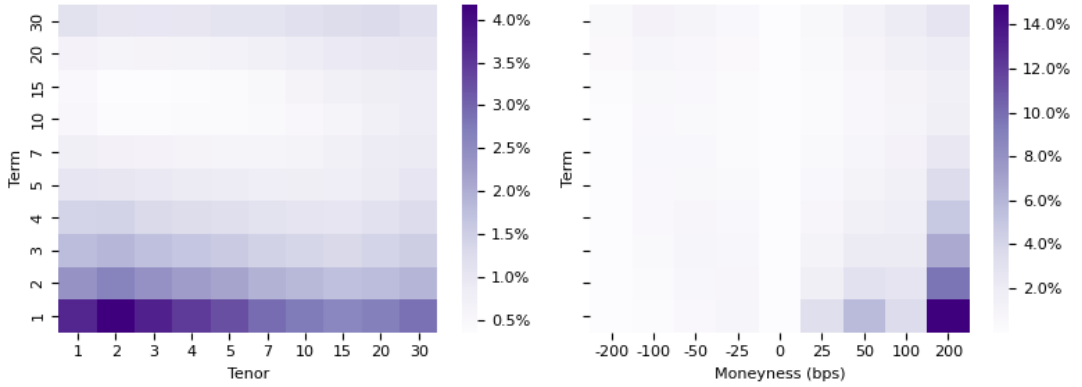

Finally, we have the plot for the A-shifted SABR in Figure 11.16. In terms of the term and tenor errors in the left-hand plot, we almost see the same errors as the H-SABR models had with the worst performance found on short-dated options written on short-dated swaps. The error magnitude is, however, a lot larger under the A-shifted SABR. The right-hand plot shows that the model has the largest errors for deep OTM options with a short term with the errors being in the magnitude of 35%.

Figure 11.17: Mean absolute percentage error in pricing for the A-shifted SABR model with $\rho = 0$ and a free β over the whole sample period across swaption term and tenor (left) and term and strike level (right).

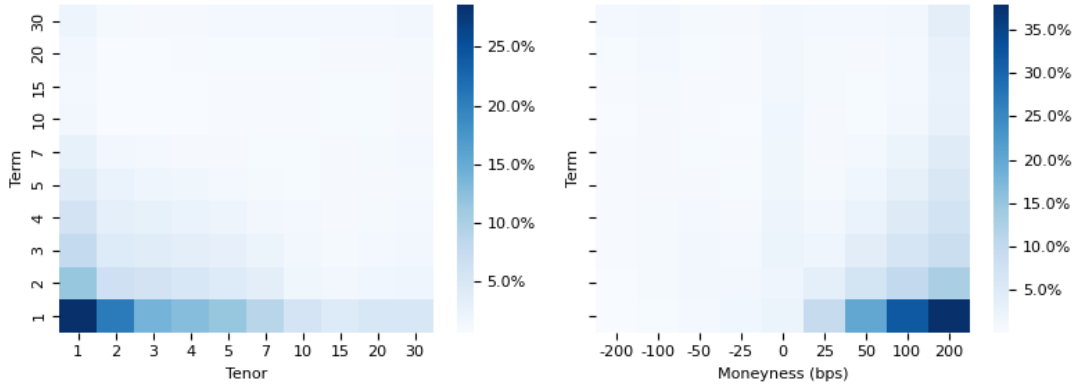

11.4.3 Deconstructing the Market Conditions

Given the evolution in the MAPE of the models over time and the volatility smiles from earlier, we see clear indications of differences in performance depending on the market conditions. To understand whether systematic issues consist in any of our models, we break up our data into different buckets. The most obvious way to do this is to consider the level of the strike and swap rate of each observation both independently and together, that is, assess whether an observation has a negative rate, a near-zero rate,

or a positive rate. We thus create buckets for the swap rate and strike in the following manner: If the rate is between -0.05% and +0.05% we classify the observation a near-zero rate, while values outside these bounds are considered negative and positive rates, respectively. We also combine the strike and swap rate buckets, so if either is classified as near zero, the combined classification is near zero; if one is negative the combined classification is negative; and only if both are positive do we classify the observation as positive rate. The price MAPE for each model across these buckets are seen in Table 11.9, where a darker colour means a higher value for the particular model and aggregation relative to the other observations in the aggregation. All three models have the same pattern. The MAPE is increasing in the strike, which is logical as the option moves more OTM, prices become smaller, and the market quotes may become staler and more illiquid. In terms of the swap rate, the MAPE decreases as the swap rate increases, which tells us that the models are better at pricing high-swap rate options. Finally, the combined buckets of the swap rate and strike reveals that all three models perform worst at when at least one of the strike or the swap rate is in the near neighbourhood of zero. Note that out of the 10,908 observations in this bucket, 1,212 of them are ATM swaptions with both rates placed near zero. While the worst performance is for near-zero rates, the difference between near-zero rates and negative rates is not too different for the H-SABR models. We also note that out of the 60,216 observations, 17,146 have both rates as negative. The models perform best for positive rates.

Table 11.9: Mean absolute percentage error (MAPE) when grouping the data by the value of the swap rate, the strike, and the combination of the two relative to 0 for the H-shifted SABR, H-normal SABR, and A-shifted SABR

	Combined rate buckets			**Strike buckets**			**Swap rate buckets**		
	Negative	Near zero	Positive	Negative	Near zero	Positive	Negative	Near zero	Positive
Observations	60,216	10,908	114,276	53,167	5,602	126,631	29,043	10,908	145,449
H-Shifted SABR	1.06%	1.23%	0.43%	0.39%	0.45%	0.81%	1.88%	1.23%	0.40%
H-Normal SABR	1.81%	2.11%	0.92%	0.66%	1.06%	1.55%	3.22%	2.11%	0.83%
A-Shifted SABR	4.12%	3.51%	2.17%	1.18%	3.17%	3.59%	7.79%	3.51%	1.86%

This analysis is slightly biased by not considering moneyness or maturity effects, and the definition of near-zero is also rather vague. To further the analysis of the market conditions impact on the MAPE, we could consider different levels of prices and volatility. We here provide the summary finding. If we section the price into percentiles of 20%, we do generally find larger MAPE errors at lower prices. A low price will generally imply an OTM option, so there is a liquidity effect to consider here. In addition, the bottom 20% of the sample have a market price of less than 1 cent per € notional of the swaption, so a small absolute mispricing might be exaggerated in relative terms. Doing the same with the swaption volatility we find less of a pattern except for a tendency to have a slightly lower MAPE for higher levels of volatility. The effect of different volatility levels is quite large for the free-boundary model, with a MAPE of only 1.8% in the 20% of the sample when sorting volatility, which should be compared to 9.8% to 13.6% for the other percentiles. We may be able to explain this observation with reference to

that this chunk of high-vol observations also have few observations with near-zero rates and generally a higher rate level.

11.5 Out-of-Sample Pricing Capabilities

The in-sample performance is only of value to us if the models are also able to price out-of-sample. This can be thought of as a volatility interpolation test. The calibration time of the A-shifted SABR means that we do not consider it for the out-of-sample analysis. Instead, we limit the focus to the H-shifted SABR and H-normal SABR models. We expect a difference in the ability of the models to price far-away-from-the money and near-the-money options due to liquidity effect and the fact that the tail properties of the model are known to break down. The implementation of the out-of-sample procedure is outlined in section 10.3,

Table 11.10: Total mean absolute percentage error (MAPE) for the full in-sample calibration and each of the out-of-sample calibrations for the H-shifted SABR and H-normal SABR models across all dates, contracts, and strikes

	Fit to all	Fit excl. -200bps	Fit excl. +200bps	Fit excl. -25bps	Fit excl. +25bps
H-Shifted SABR	0.68%	0.24%*	1.27%	1.01%	0.69%
H-Normal SABR	1.28%	0.71%*	2.24%	1.28%	1.28%

Table 11.10 shows the price MAPE calculated across the full sample for the models fitted to the full sample from the preceding sections and each of the out-of-sample calibrations. The total performance of the H-shifted SABR and H-normal SABR model dramatically improves when excluding deep ITM swaptions with strikes of -200bps from the model calibration. Both models had a quite high average in-sample precision when it came to pricing this strike level, so this improvement implies that the full calibration may have overfitted this strike level at the expense of generalised performance. Both models seem to be able to price the whole sample to a satisfactory level regardless of which strike is excluded from the calibration. Excluding near-the-money strikes only has little impact on the performance of the models compares to the in-sample analysis. The model performance is worst when excluding deep ITM swaptions, which is expected given the in-sample analysis. We are, however, not able to base the full analysis purely on this aggregated view. In the following sections we consider the price MAPE across strike levels for each calibration and finally look at the resulting smiles.

11.5.1 Near and Far Moneyness Interpolation

The price MAPE across all dates and contracts broken down into strike levels for each SABR model is seen in Table 11.11 with the excluded strike marked in bold and the cell colours indicating the highest (dark) and lowest (light) MAPE strike level for each calibrated model. Panel A shows the performance of the H-shifted SABR. The pricing performance for strikes from -100bps to +100bps is quite similar across all calibrations, but they do differ in their abilities to accurately price deep ITM and OTM swaptions. The out-of-sample model excluding -200bps strikes from the calibration performs best out of any

model for deep OTM swaptions with a MAPE of 0.61%. This model consistently has the lowest MAPE for all strikes expect deep ITM of -200bps. The deep ITM performance is however still better than the deep OTM performance. Overall, the H-shifted SABR model seems to be able to interpolate volatilities and thus prices to a satisfactory level.

Table 11.11: Mean absolute percentage error (MAPE) for the full in-sample calibration and each of the out-of-sample calibrations for the H-shifted SABR and H-normal SABR models per strike level

Panel A. Mean absolute percentage error (MAPE) comparison for H-shifted SABR with $s = 5\%$ and $\beta = 1$									
	-200	-100	-50	-25	ATM	+25	+50	+100	+200
Avg. market price	€17,979	€10,734	€7,631	€6,298	€5,158	€4,224	€3,479	€2,415	€1,284
H-Shifted SABR	0.10%	0.37%	0.37%	0.26%	0.00%	0.41%	0.74%	0.98%	2.89%
H-Shifted SABR (excl. -200)	**0.47%**	0.05%	0.06%	0.06%	0.00%	0.16%	0.32%	0.42%	0.61%
H-Shifted SABR (excl. +200)	0.07%	0.35%	0.31%	0.20%	0.00%	0.20%	0.46%	2.06%	**7.78%**
H-Shifted SABR (excl. -25)	0.10%	0.37%	0.38%	**0.28%**	0.01%	0.48%	0.93%	1.47%	5.05%
H-Shifted SABR (excl. +25)	0.10%	0.37%	0.37%	0.26%	0.00%	**0.42%**	0.78%	1.02%	2.85%

Panel B. Mean absolute percentage error (MAPE) comparison for H-normal SABR with $\beta = 0$									
	-200	-100	-50	-25	ATM	+25	+50	+100	+200
Avg. market price	€17,979	€10,734	€7,631	€6,298	€5,158	€4,224	€3,479	€2,415	€1,284
H-Normal SABR	0.16%	0.55%	0.65%	0.52%	0.00%	0.96%	1.79%	1.86%	5.01%
H-Normal SABR (excl. -200)	**0.81%**	0.13%	0.14%	0.18%	0.00%	0.55%	1.14%	1.56%	1.92%
H-Normal SABR (excl. +200)	0.11%	0.47%	0.49%	0.34%	0.00%	0.38%	0.47%	3.24%	**14.65%**
H-Normal SABR (excl. -25)	0.16%	0.56%	0.66%	**0.52%**	0.00%	0.99%	1.85%	1.98%	4.84%
H-Normal SABR (excl. +25)	0.16%	0.55%	0.66%	0.52%	0.00%	**0.98%**	1.83%	1.95%	4.86%

The performance of the normal SABR model seen in panel B of Table 11.11 gives rise to almost identical conclusions as for the shifted SABR. The tendency for all models to perform the same except for when it comes to deep ITM or OTM swaptions is seen again, and excluding -200bps strikes from the calibration data makes the model fit even better across all remaining strikes than the in-sample calibration. We do note that for the out-of-sample calibration we fit the data to only 8 data points, and with 2 free parameters, there is a risk of overfitting the datapoints to extrapolation to even more ITM or OTM swaptions could become unreliable. Overall, the normal SABR also works reasonably well for interpolating volatilities. It works best near the money and the further away from the money the larger the error. While satisfactory, the shifted SABR model still seems to be preferred due to its smaller errors.

11.5.2 Recalibrated Volatility Smiles

As a final perspective on the interpolation capabilities of the SABR models, we compare the volatility smiles implied by the out-of-sample models with the in-sample model and the market for the same dates as in section 11.2, but we only look at a 1Y5Y swaption as an example. Figure 11.18 shows the smiles for the shifted SABR model. The differences in fit between the in-sample and out-of-sample models is

only small on the high-rate date in panel (a), with only slight shifts in the tail of the smiles. The low-rate date in panel (b) has a virtually no differences in fit between excluding near-the-money strikes and the in-sample fit, although neither manage to capture the smile adequately. In this case, we see that when we calibrate the model excluding deep ITM strikes, the resulting smile is able to capture most of the market implied smile, although the error for the -200bps strike also becomes much larger.

Figure 11.18: Fitted volatility smiles for the H-shifted SABR model with $s = 1$ and $\beta = 1$ in-sample and out-of-sample for a 1Y5Y swaption on selected dates. The parenthesis in the legend implies which strike is excluded from the calibration.

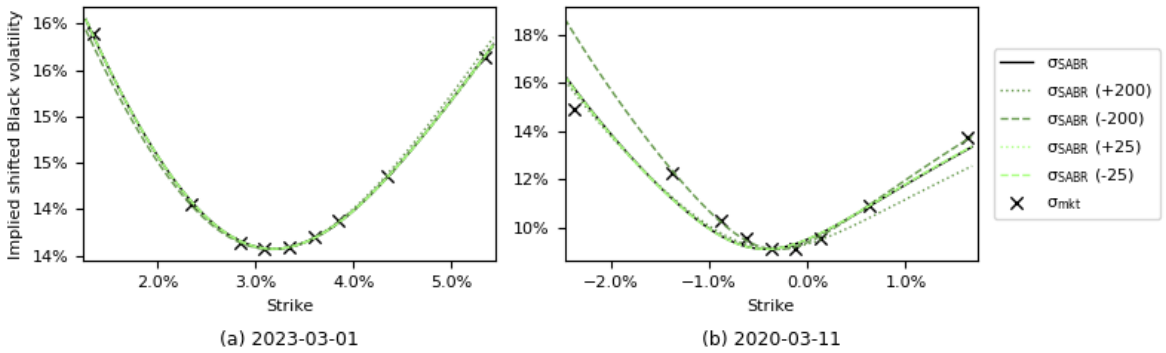

A similar pattern is seen for the H-normal SABR in Figure 11.19. The high-rate date in (a) shows that extrapolating away from the money strikes in the out-of-sample analysis yields different smile tails and thus a larger error, while the near-the-money fit is visually on par with the in-sample calibration. The low-rate date in (b) has more distinct visually differences in fit in terms of the shape of the smile, but the excluding the -200bps strike from the calibration manages to fit the remaining points better, while extrapolating is less accurate.

Figure 11.19: Fitted volatility smiles for the normal SABR model with $\beta = 0$ in-sample and out-of-sample for a 1Y5Y swaption on selected dates. The parenthesis in the legend implies which strike is excluded from the calibration.

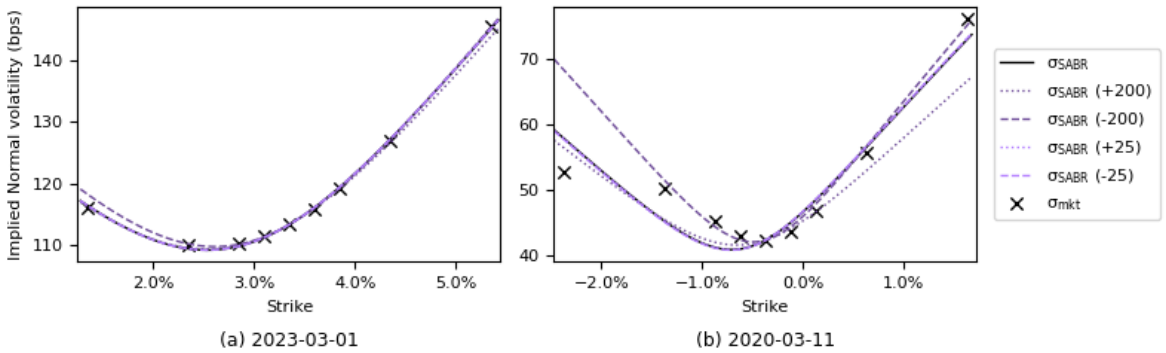

Overall, the conclusion of the out-of-sample analysis is that particularly the shifted SABR model but also to a high degree the normal SABR model are able to successfully interpolate volatilities. The interpolation capabilities are best for near-the-money strikes, and the accuracy deteriorate the when the extrapolating volatilities for strikes on the smile wings. We notice a clear trade-off when it comes to fitting the model as well. It seems that if you fit the smiles to all quotes including deep OTM or ITM ones which may be stale and illiquid, the overall performance of the model actually decreases as the parameters pay too much attention to these outer points. A way to mitigate this in a practical application

would be to apply weights in the optimisation procedure depending on the purpose of the model. We expects derivative trading desks would mostly deal with near-the-money pricing and interpolation, so applying larger weights to these observations and a smaller weight/excluding any far away from the money observations would likely yield a better model.

11.6 Empirical Risk Sensitivities

This final part of our empirical analysis of the capabilities of the SABR model will focus on risk sensitivities to complete the perspective on the model. Based on our findings thus far, we conclude that the H-shifted SABR model with $\beta = 1$ is superior to the other specifications considered, so we focus our attention only on this model. The risk sensitivity analysis will follow the intuitions presented in section 8.2, where we first analyse the empirical delta and vega Greeks under the SABR dynamics relative to the Black model in section **Error! Reference source not found.**. The risk framework is then further extended in section 11.6.2 by calculating the parameter sensitivities. To limit the data quantity, we use an illustrative calculation case of XY10Y swaptions (i.e., the term is fixed at 10Y) and calibrate the shifted SABR model to market quotes on 1 March 2023 across 10 different terms, thus producing 10 different smiles. We discount the risk measures using annuity factors inferred from the discounting curve based on the term and tenor.

11.6.1 Delta and Vega Risk

We first compare delta risk under the shifted SABR model to delta risk under the shifted Black model. The shifted Black delta is calculated using the analytical solution defined in equation (4.16), where the constant volatility assumption means that we must use the ATM volatility when calculating delta values away-from-the-money. The delta of the Shifted SABR model is calculated following the definition of Bartlett's delta in equation (8.1). This is done numerically using central differences, where we also use the calibrated SABR parameter to account for the movements in volatility through the forward rate f and its correlation with α. The calibrated parameters are also used to interpolate volatilities when calculating delta values away-from-the-money.

Figure 11.20: Delta across different strike levels under the Shifted SABR model (left) and the shifted Black model (right) for XY10Y payer swaptions using market quotes of March 1, 2023

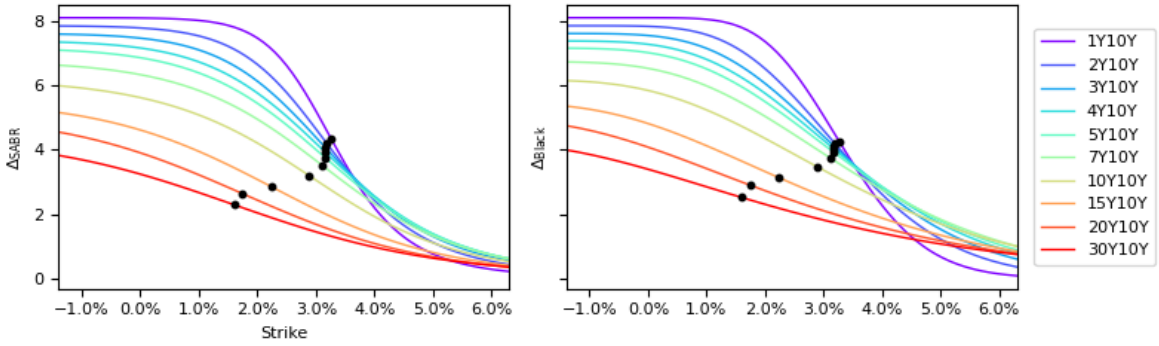

The calculated delta curves across different strike levels and terms are see in Figure 11.19, where we have marked the ATM strike for each curve with a black bullet. We first comment on the general shape of the curves. First note that the par swap rate is decreasing in maturity which explains the horizontal placement of the delta curves. The delta is consistently estimated to be positive, which implies that increases in the par swap rate will increase the swaption price. As we are considering payer swaptions, this finding is intuitive. As we move deep OTM at high strike levels, the delta sensitivity become increasing smaller as the likelihood of moves ITM decreases, and conversely for deep ITM options at low strikes. The curves as the steepest slope around the ATM point, which implies the highest delta sensitivity here (gamma). The delta curves become flatter as the time to maturity increases, which is intuitive as longer maturities have more time for the rate to move. We now compare the delta curves of the shifted SABR model (left) to the delta curves of the shifted Black model in the (right). The most noticeable difference in the shape of the curves is seen for high strikes (deep OTM), where the shifted SABR model imply a lower delta. We can explain this by considering values of the calibrated parameters. In general, increasing f has the effect of shifting the volatility smile to the right and making it

Figure 11.21: Delta across different strike levels under the Shifted SABR model (left) and the shifted Black model (right) for XY10Y payer swaptions using market quotes of March 1, 2023

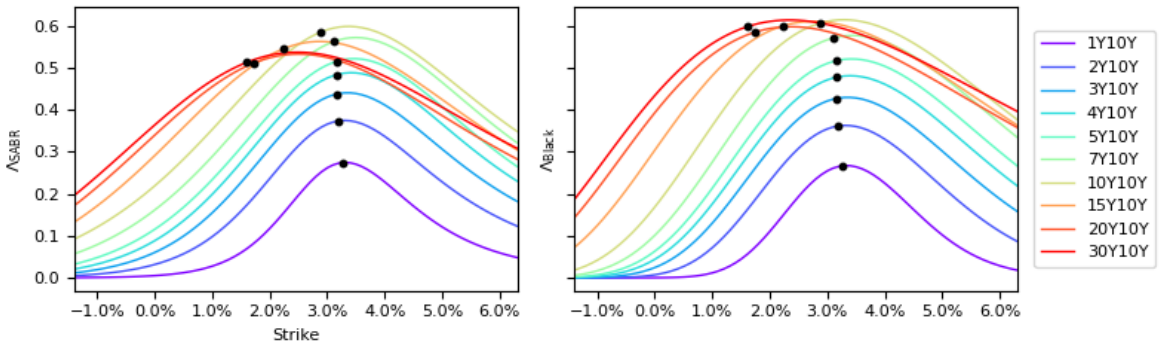

slightly flatter, which means a lower volatility for high strikes. We observe that the vol-vol parameter ν is decreasing in maturity, while the correlation ρ is negative and decreasing in maturity. The net effect of the SABR adjustment to the delta becomes negative for high strikes, which explains the downward shift seen. The 1Y10Y curve instead shifts up at high strikes under the shifted SABR, which is explained by a small negative correlation and high value of ν making the net effect positive for this contract.

The other Greek to consider is vega, which we calculate under the shifted SABR model and the shifted Black model using the same procedure as for delta with the respective formulas found in equation (8.2) and (4.17). Figure 11.20 shows the resulting vega curves. Commenting first on the shape of the curves, we see vega is positive across all terms and strikes, which implies that an increase in volatility will increase the swaption price. This observation is in line with typical market behaviour. Vega is increasing with maturity, as there is more time for the swaption to end up ITM. The abnormal placement of the longest maturities can be explained by an inverted yield curve which causes longer-term swaptions to have a smaller annuity factor than shorter-term ones. The vega is largest as the ATM point as volatility

is a determining factor in the swaption ending up ITM. Comparing the vega curves of the shifted SABR model (left) to the vega curves of the shifted Black model in the (right), we see the most noticeable difference for long maturity options at ATM strikes, where the shifted SABR implies a lower vega. Increasing α under the SABR model shifts the smile level up which implies a higher volatility and thus a higher price. Bartlett's correction term offsets this effect for most maturities, but as ρ becomes more negative and ν decreases with maturity in our data, the net effect for large maturities become negative.

11.6.2 Sensitivity of SABR Parameters

The sensitivity analysis of swaption prices with respect to each of the SABR parameters in the SABR model can help us quantify misspecification risk. This is particularly important in our context, as we saw clear evidence of faulty calibration or model specification under the shifted SABR model with reference to, e.g., the extreme jumps in value of the correlation parameter ρ. As we did for the SABR Greeks, we calculate the parameter sensitivities defined in section 8.2 using central differences.

Sensitivity of ρ. We first consider the vanna sensitivity, i.e., the sensitivity of the swaption price relative to changes in the correlation ρ. The calculated sensitivities are illustrated in the left-hand plot of Figure 11.21. The ρ sensitivity is increasing in maturity and largest for deep OTM high strikes. Recall from section 6.3 that ρ has the impact of rotating the smile counter clockwise around the ATM point as it increases. Thus, for deep OTM options, the increasing ρ will rotate the smile up and thus increase the volatility and swaption price. Conversely, for deep ITM strikes, the smile will rotate down causing the swaption price to decrease. This fact explains the negative ρ sensitivities seen for low strikes across all maturities. We see that the sensitivity curves are steepest at the ATM point around which the smiles rotate.

Figure 11.22: Parameter risk sensitivities in terms of ρ/vanna (left) and ν/volga (right) under the Shifted SABR model for XY10Y payer swaptions using market quotes of March 1, 2023

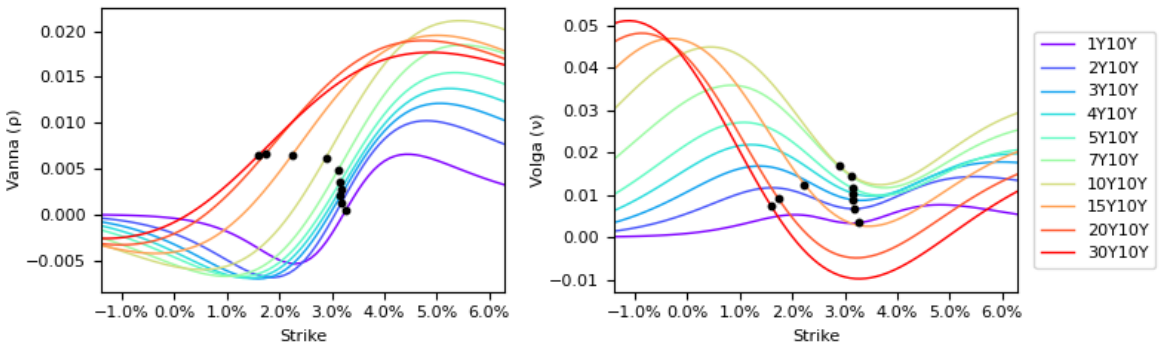

Sensitivity of ν. The Volga sensitivity is the swaption price sensitivity to changes in the vol-vol parameter ν, which we recall magnifies the smile effect. The calculated risk sensitivities are illustrated in the right-hand plot of Figure 11.21. Across most maturities and strikes, the ν sensitivity is positive, which intuitively means that overestimating the volatility-of-volatility causes higher swaption prices. The

Volga sensitivity becomes negative for long maturities with high strikes. For the 30Y10Y case, we observe a vol-vol of $\nu = 0.12$ and correlation of $\rho = -0.45$. Increasing ν makes the volatility process larger, which the negative correlation with the forward process translates into a decrease in the forward rate and thus the swaption price. The sensitivity of short-maturity swaptions to changes in the vol-vol is negligible. Finally, we a higher sensitivity is generally seen for low strike options while the lowest is seen for ATM strikes, consistent which how increasing the vol-vol parameter mostly impacts away-from-the-money strikes.

Sensitivity of $\boldsymbol{\beta}$. In our implementation of the shifted SABR model, we keep $\beta = 1$ fixed. Despite this, there is risk of changes in the market dynamics which could make it necessary to update β, and we are interested in understanding how a possible misspecification would impact the swaption prices. The calculated β sensitivities are seen in the right-hand plot of Figure 11.22. The sensitivities are negative across all maturities and strikes, while there is a tendency for longer maturities to be more sensitive to changes than shorter maturities. The sensitivity is generally largest at the ATM strike. Thus, a ceteris paribus decrease (as it is currently placed on the perimeter of its permitted bound) in β will cause all of our swaptions to decrease in value.

Figure 11.23: Delta across different strike levels under the Shifted SABR model (left) and the shifted Black model (right) for XY10Y payer swaptions using market quotes of March 1, 2023

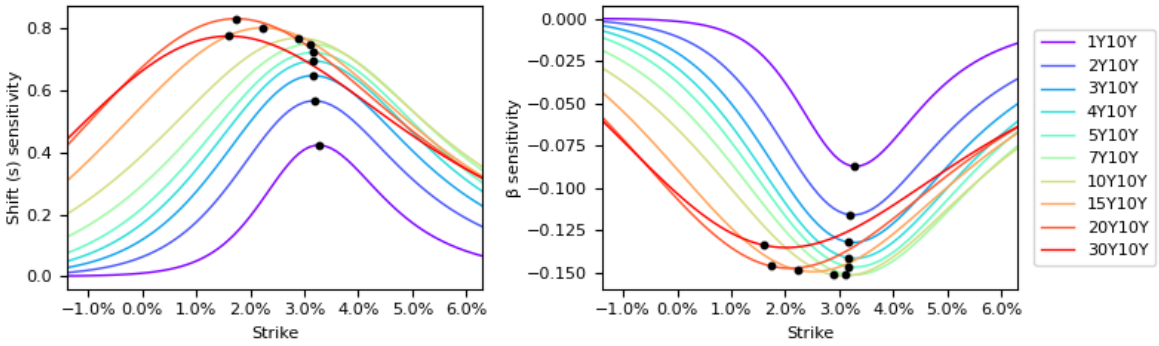

Sensitivity of shift, $\boldsymbol{s}$. The final parameter sensitivity is the shift sensitivity, whose sensitivity curves are seen in the left-hand plot of Figure 11.22. The observed sensitivity is positive across all curves, which means that an increase in the shift parameter intuitively will cause swaption prices to increase. Visually, the sensitivity curves are similar to the vega ones with the largest sensitivity around ATM strikes. We also notice how the β and shift sensitivities appear inversely related, but explaining this observation is harder.

PART IV

DISCUSSION AND CONCLUSION

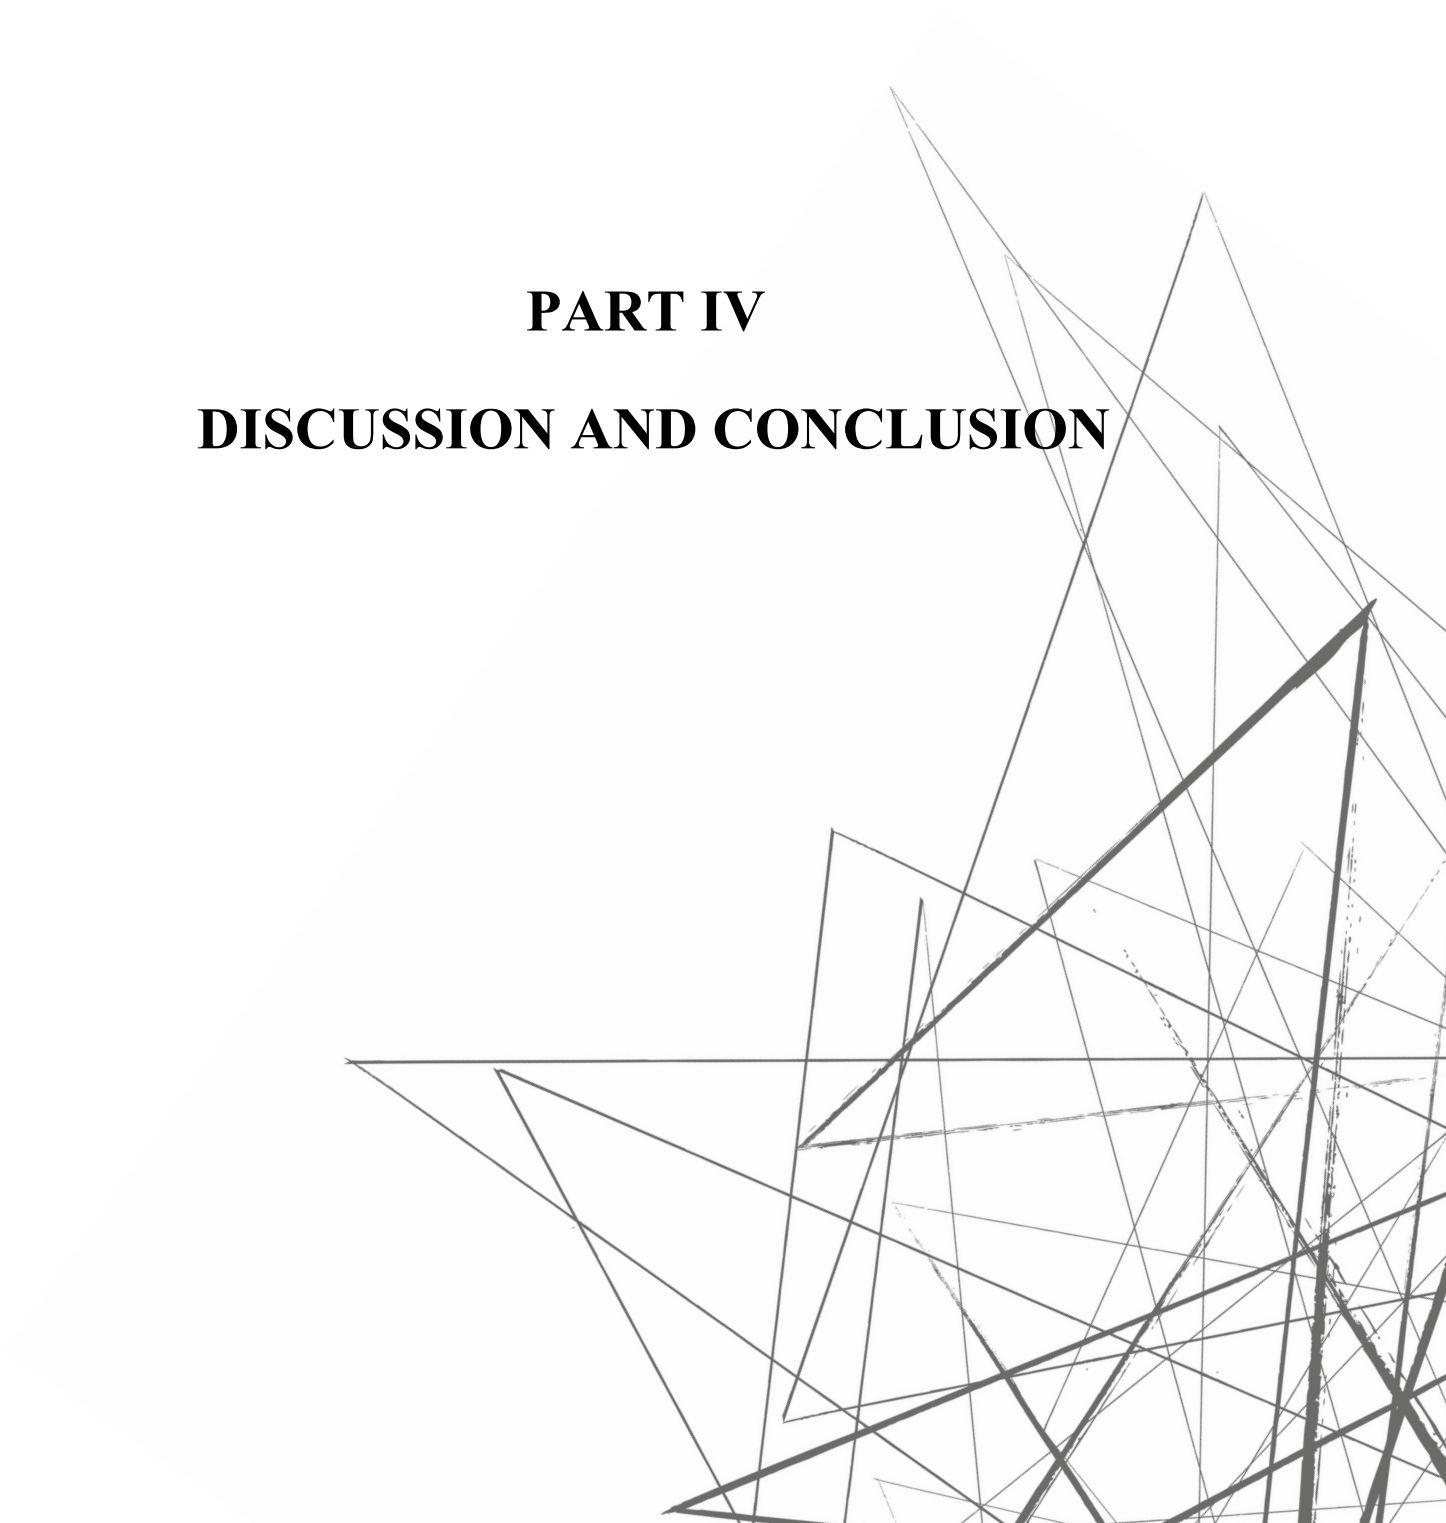

12 IMPLICATION AND LIMITATIONS OF RESULTS

Before we conclude on our results, we open a discussion of the implications and limitations of the empirical analysis. The first part was dedicated to identifying the optimal parameterisation of each SABR model type and for each method of calibrating the models. The results for the models calibrated to market prices using the exact solutions were strongly impacted by **convergence issues** in the calibrations. The Antonov normal SABR did not yield any results and was disregarded straight away. The Antonov free-boundary SABR managed to converge, but the produced data showed evidence of unexpected kinks and shifts in the smiles. The Antonov shifted SABR was successfully calibrated, but still had generally a worse performance than the Hagan approximations. The implication here is that we cannot conclude whether the models themselves are mis-specified, or if the implementation of the calibration procedure is faulty. Thus, more extensive testing of the price calibration procedure is needed. In addition, we also did not do the mapping of the zero-correlation models to the general-correlation case, so we cannot comment on whether this will help the accuracy performance.

A good compliment to the analysis would have been to include **alternative volatility approximation** formulas to assess the relative accuracy of the original Hagan ones. The analysis could also be expanded with an empirical test of the points where the approximations break down. Monte carlo simulations of the model processes and exploration of the pdf would provide good insights this aspect of the models.

The **parameters sensitivity** analysis could be expanded. For example, shifting the model parameters by one week and price the swaptions while implying alpha out from ATM would be equivalent to having a weekly recalibration and could help better understand how the parameters move. In addition, doing an outlier analysis by removing extreme market observations from the data and assessing how this impacts the change in parameter values would also add an interesting perspective.

We did numerical calculations of **risk sensitivities** under the SABR dynamics and saw evidence of differences in their properties relative to under the Black model. The natural extension of this would include a profit and loss analysis based on real market data.

The scope of this thesis was to determine which specification of the SABR model is best suited for empirically modelling the volatility smile across different interest rate environments. We put a limit on the model selection by requiring that it could model negative rates. This was necessary to fit the model to EUR swaption data as this historically has exhibited negative rates. There are, however, other models which could have been relevant to include. A forward process of the SABR is not mean reverting, which empirically is wrong as forward rates tends to revert in the long run. As higher-volatility periods are associated with a flatter volatility smile, a stochastic volatility framework with mean reverting term in the volatility process would improve the model. A such model is suggested in (Henry-Labordère, 2008).

13 CONCLUSION

This thesis conducted an empirical analysis of the ability of the Shifted SABR, Normal SABR, and Free-Boundary SABR to capture volatility smiles across different interest rate environments. The calibration of the models was implemented in two different ways. The first method calibrated the model parameters to market quotes of implied volatility via the Hagan approximations, which is considered the traditional way of implementing model, and has the clear advantage of almost instant calibration to a single smile. The second method calibrated the model parameters to market prices via the Antonov solutions, which are much more computationally heavy but should theoretically be more consistent than the Hagan approximation. For each of the two calibration spaces, we considered a wide range of model parameterisations, and eventually concluded on three which we used in our main analysis. The Free-Boundary SABR model was found to be unfeasible for our purpose due to its stickiness at zero and heavy calibration procedure. Antonov's exact solution for the Normal SABR was also disregarded due to deep-rooted implementation issues.

This left us with three models: the Hagan Shifted SABR with $s = 5\%$ and $\beta = 1$ and an implied α, the Hagan Normal SABR with an implied α, and the Antonov Shifted SABR with $s = 5\%$, $\rho = 0$, and a free β. The analysis then shifted to focusing on parameter stability, where we found that the Hagan Normal SABR generally had less jumps in the parameter values over time, making this model the optimal choice for risk management purposes. The in-sample pricing performance was, however, found to be superior under the Hagan Shifted SABR with a total price MAPE of 0.68% across the sample. Across all models, we found a tendency to have the largest pricing error for deep OTM high strikes, while the models also consistently performed worse on short swaption tenors. And these are found to be the limitation of the models. To address the concept of negative rates, we find that grouping our data by the combined placement of the swap rate and strikes, the models consistently perform worse on near-zero and negative rates relative to positive rates. The out-of-sample analysis assessed how the models fit the data if certain observations are omitted from the calibration procedure. The overarching conclusion was that the Hagan Shifted SABR was able to do this to a satisfactory level. Finally, we illustrated how to calculate Greeks and parameter sensitivities under the Hagan Shifted SABR model. Based on the initial model suitability criteria defined, the Hagan Shifted SABR model is the empirically best suited for modelling the volatility smile across different interest rate environment, with the limitation that the performance decreases slightly under negative rates, and the fit seems to perform poorly on deep OTM high strikes and short-term swaptions.

REFERENCES

Antonov, A., Konikov, M., & Spector, M. (2013). SABR spreads its wings. *Risk*, *26*, 58–63. https://www-risk-net.esc-web.lib.cbs.dk/derivatives/interest-rate-derivatives/2284866/sabr-spreads-its-wings

Antonov, A., Konikov, M., & Spector, M. (2015a, January). The Free Boundary SABR: Natural Extension to Negative Rates. *Numerix*. https://doi.org/10.2139/ssrn.2026350

Antonov, A., Konikov, M., & Spector, M. (2015b, August). Mixing SABR models for Negative Rates. *Numerix*. http://ssrn.com/abstract=2653682

Antonov, A., Konikov, M., & Spector, M. (2019). *Modern SABR analytics: Formulas and insights for quants, former physicists and mathematicians* (First). Springer International Publishing. https://doi.org/10.1007/978-3-030-10656-0

Antonov, A., & Spector, M. (2012). Advanced analytics for the SABR model. *Numerix Quantitative Research*. https://doi.org/10.2139/ssrn.2026350

Barlett, B. (2006). Hedging under SABR Model. *Willmott Magazine*. https://lesniewski.us/papers/published/HedgingUnderSABRModel.pdf

BIS. (2022, October 30). *Global OTC derivative markets: D5.1 Foreign exchange, interest rate, equity linked contracts*. https://stats.bis.org/statx/srs/table/d5.1?f=pdf

Björk, T. (2020). *Arbitrage Theory in Continuous Time* (Fourth). Oxford University Press. https://doi.org/10.1093/oso/9780198851615.001.0001

Bloomberg Markets. (2020, August 17). *In a post-LIBOR world, here are the benchmarks that will matter*. Bloomberg. https://www.bloomberg.com/professional/blog/in-a-post-libor-world-here-are-the-benchmarks-that-will-matter/

Brigo, D., & Mercurio, F. (2006). *Interest Rate Models — Theory and Practice: With Smile, Inflation and Credit* (Second). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-34604-3

Choi, J., Kwak, M., Tee, C. W., & Wang, Y. (2022). A Black-Scholes user's guide to the Bachelier model. *Journal of Futures Markets*, *42*(5), 959–980. https://doi.org/10.1002/fut.22315

Crispoldi, C., Wigger, G., & Larkin, P. (2015). SABR and SABR LIBOR Market Models in Practice. In *SABR and SABR LIBOR Market Models in Practice*. Palgrave Macmillan UK. https://doi.org/10.1057/9781137378644

Deloitte. (2016). *Risk management under the SABR model*. https://www2.deloitte.com/content/dam/Deloitte/global/Documents/Financial-Services/be-aers-fsi-sabr-sensitivities.pdf

Deuskar, P., Gupta, A., & Subrahmanyam, M. G. (2004). *Interest Rate Option Markets: The Role of Liquidity in Volatility Smiles* (S-DRP-04-03,). https://ssrn.com/abstract=1295852

ECON. (2021). *What Are the Effects of the ECB's Negative Interest Rate Policy?* Monetary Dialogue Papers. https://www.europarl.europa.eu/cmsdata/235691/02.%20BRUEGEL_formatted.pdf

Hagan, P., Lesniewski, A., & Woodward, D. (2015). Probability distribution in the SABR model of stochastic volatility. *Springer Proceedings in Mathematics and Statistics*, *110*. https://doi.org/10.1007/978-3-319-11605-1_1

Hagan, P. S., Kumar, D., Lesniewski, A. S., & Woodward, D. E. (2002). Managing Smile Risk. *Willmont*, 84–108. https://www.researchgate.net/publication/235622441

Hagan, P. S., Kumar, D., Lesniewski, A., & Woodward, D. (2014). Arbitrage-Free SABR. *Wilmott*, *2014*(69), 60–75. https://doi.org/10.1002/wilm.10290

Hagan, P. S., Lesniewski, A. S., & Woodward, D. E. (2018). Managing Vol Surfaces. *Wilmott*, *2018*(93), 24–43. https://doi.org/10.1002/wilm.10643

Henry-Labordère, P. (2008). Analysis, geometry, and modeling in finance: Advanced methods in option pricing. In *Analysis, Geometry, and Modeling in Finance: Advanced Methods in Option Pricing*. Chapman & Hall/CRC.

Hull, J. C. (2018). *Option, Futures, and Other Derivatives* (Ninth Global). Pearson Education Limited.

Kienitz, J. (2015). *Approximate and PDE Solution to the Boundary Free SABR Model - Applications to Pricing and Calibration*. https://doi.org/10.2139/ssrn.2647344

Nordea. (n.d.). *IBOR*. Retrieved 17 April 2023, from https://www.nordea.com/en/our-services/ibor

Oosterlee, C. W., & Grzelak, L. A. (2019). Mathematical Modeling and Computation in Finance. In *Mathematical Modeling and Computation in Finance*. World Scientific. https://doi.org/10.1142/q0236

Russo, V., & Fabozzi, F. J. (2017). Calibrating short interest rate models in negative rate environments. *Journal of Derivatives*, *24*(4), 80–92. https://doi.org/10.3905/jod.2017.24.4.080

West, G. (2005). Calibration of the SABR model in illiquid markets. *Applied Mathematical Finance*, *12*(4), 371–385. https://doi.org/10.1080/13504860500148672

APPENDIX

A Swaption Data Summary

Table A.1. summarises the average implied normal volatility measured in bps (panel A) and the average calculated swaption price in € assuming a €100,000 notional (panel B) of the swaption data used in the analysis grouped by swaption terms and tenors.

Table A.1: Summary statistics for the swaption dataset using 185,400 implied volatility quotes collected from Bloomberg with weekly frequency over the period 27 March 2019 to 1 March 2023.

Panel A. Average implied normal volatility (bps) across the sample by swaption term (vertical) and tenor (horizontal)

	XY1Y	XY2Y	XY3Y	XY4Y	XY5Y	XY7Y	XY10Y	XY15Y	XY20Y	XY30Y	**Total**
1YZY	61.61	62.58	64.69	66.38	67.68	70.33	72.53	74.19	75.00	75.54	**69.05**
2YZY	61.37	62.22	63.82	64.84	65.90	67.57	69.09	69.52	69.88	69.88	**66.41**
3YZY	64.47	64.50	65.13	65.52	65.99	67.05	67.79	67.37	67.20	66.60	**66.16**
4YZY	65.99	65.45	65.68	65.69	65.80	66.29	66.55	65.42	64.91	63.83	**65.56**
5YZY	66.01	65.35	65.35	65.18	65.18	65.32	65.29	63.68	62.90	61.54	**64.58**
7YZY	66.97	66.03	65.56	65.19	64.81	64.49	63.94	61.77	60.58	58.77	**63.81**
10YZY	65.91	64.99	64.48	64.01	63.55	62.86	61.90	59.33	57.75	55.53	**62.03**
15YZY	62.19	61.23	60.84	60.44	60.00	59.24	58.23	55.21	53.35	50.79	**58.15**
20YZY	58.96	58.08	57.75	57.31	57.01	56.19	54.96	51.88	49.77	46.94	**54.88**
30YZY	55.20	54.42	53.99	53.53	53.22	51.55	49.54	46.08	43.70	40.64	**50.19**
Total	62.87	62.48	62.73	62.81	62.91	63.09	62.98	61.44	60.50	59.00	62.08

Panel B. Average swaption price (€) across the sample assuming a notional of €100,000 by swaption term (vertical) and tenor (horizontal)

	XY1Y	XY2Y	XY3Y	XY4Y	XY5Y	XY7Y	XY10Y	XY15Y	XY20Y	XY30Y	**Total**
1YZY	492	990	1,489	1,991	2,488	3,484	4,943	7,278	9,498	13,763	**4,642**
2YZY	551	1,106	1,665	2,225	2,782	3,895	5,527	8,094	10,559	15,280	**5,169**
3YZY	611	1,220	1,832	2,438	3,040	4,248	6,003	8,739	11,370	16,406	**5,591**
4YZY	665	1,323	1,979	2,623	3,265	4,542	6,395	9,250	12,010	17,279	**5,933**
5YZY	711	1,407	2,097	2,777	3,452	4,789	6,720	9,671	12,530	17,994	**6,215**
7YZY	788	1,554	2,307	3,050	3,774	5,204	7,276	10,412	13,445	19,249	**6,706**
10YZY	863	1,702	2,520	3,321	4,107	5,648	7,869	11,224	14,457	20,638	**7,235**
15YZY	918	1,805	2,685	3,545	4,389	6,048	8,460	12,032	15,466	22,023	**7,737**
20YZY	952	1,874	2,797	3,697	4,591	6,333	8,853	12,587	16,121	22,870	**8,068**
30YZY	1,028	2,026	3,021	4,006	4,980	6,783	9,363	13,202	16,805	23,642	**8,486**
Total	**758**	**1,501**	**2,239**	**2,967**	**3,687**	**5,097**	**7,141**	**10,249**	**13,226**	**18,914**	**6,578**

B Fitting β to ATM Shifted Black Volatilities

We present selected results from regressing the log-rates against ATM log-volatilities in the shifted Black calibration space per swaption contract across all dates in the sample. We find the fitted β is generally decreasing in maturity.

Table B.1: Results for selected swaption contracts from regressing $\ln(f)$ against $\ln(\sigma_B^{\mathrm{ATM}})$ using weekly data over the entire data horizon 27 March 2019 to 1 March 2023.

Volatility type	Swaption	Intercept	Slope	SE	P-value	R-squared	Implied β
Shifted Black	1Y1Y	4.522	2.458	0.105	9.69E-60	0.729	1.000
Shifted Black	1Y5Y	2.787	1.767	0.066	5.51E-69	0.780	1.000
Shifted Black	1Y10Y	1.221	1.214	0.063	1.62E-47	0.643	1.000
Shifted Black	1Y15Y	0.825	1.080	0.080	4.11E-30	0.472	1.000
Shifted Black	1Y20Y	0.464	0.946	0.096	7.07E-19	0.321	1.000
Shifted Black	1Y30Y	-0.082	0.736	0.122	7.96E-09	0.151	1.000
Shifted Black	5Y1Y	0.870	1.083	0.050	2.12E-54	0.694	1.000
Shifted Black	5Y5Y	-0.107	0.757	0.046	7.96E-40	0.576	1.000
Shifted Black	5Y10Y	-0.946	0.470	0.052	9.75E-17	0.288	1.000
Shifted Black	5Y15Y	-1.151	0.404	0.066	5.23E-09	0.154	1.000
Shifted Black	5Y20Y	-1.368	0.324	0.078	4.50E-05	0.079	1.000
Shifted Black	5Y30Y	-1.852	0.149	0.090	9.93E-02	0.013	1.000
Shifted Black	10Y1Y	-1.227	0.374	0.037	1.34E-19	0.332	1.000
Shifted Black	10Y5Y	-1.723	0.206	0.042	1.66E-06	0.107	1.000
Shifted Black	10Y10Y	-2.283	0.007	0.050	8.85E-01	0.000	1.000
Shifted Black	10Y15Y	-2.656	-0.118	0.058	4.28E-02	0.020	0.882
Shifted Black	10Y20Y	-2.939	-0.214	0.062	6.50E-04	0.056	0.786
Shifted Black	10Y30Y	-3.264	-0.322	0.063	6.94E-07	0.114	0.678
Shifted Black	15Y1Y	-1.470	0.300	0.057	3.36E-07	0.120	1.000
Shifted Black	15Y5Y	-2.035	0.097	0.065	1.35E-01	0.011	1.000
Shifted Black	15Y10Y	-3.093	-0.276	0.059	5.22E-06	0.097	0.724
Shifted Black	15Y15Y	-3.665	-0.463	0.055	9.06E-15	0.256	0.537
Shifted Black	15Y20Y	-3.805	-0.502	0.055	4.21E-17	0.293	0.498
Shifted Black	15Y30Y	-4.052	-0.573	0.051	2.82E-23	0.385	0.427
Shifted Black	20Y1Y	-2.028	0.093	0.081	2.52E-01	0.006	1.000
Shifted Black	20Y5Y	-2.814	-0.181	0.073	1.40E-02	0.029	0.819
Shifted Black	20Y10Y	-3.759	-0.501	0.056	1.40E-16	0.285	0.499
Shifted Black	20Y15Y	-4.240	-0.648	0.046	8.30E-32	0.492	0.352
Shifted Black	20Y20Y	-4.358	-0.675	0.045	1.17E-34	0.523	0.325
Shifted Black	20Y30Y	-4.581	-0.732	0.041	4.27E-43	0.606	0.268
Shifted Black	30Y1Y	-2.544	-0.089	0.083	2.85E-01	0.006	0.911
Shifted Black	30Y5Y	-3.454	-0.394	0.066	1.27E-08	0.147	0.606
Shifted Black	30Y10Y	-4.297	-0.658	0.045	2.93E-33	0.508	0.342
Shifted Black	30Y15Y	-4.870	-0.829	0.035	9.12E-60	0.729	0.171
Shifted Black	30Y20Y	-5.085	-0.884	0.035	3.17E-64	0.755	0.116
Shifted Black	30Y30Y	-5.169	-0.886	0.037	9.85E-61	0.735	0.114