FACULTY OF ECONOMICS AND BUSINESS

Final Year Dissertation

The SABR Model in Financial Practice

By: Iker Caballero Bragagnini, Marc Colom Cervera & Iván Moreno Calixto

BSc in Economics & BSc in Business Administration

Tutored by Professor Elisa Alòs

Academic Year 2021-2022

Code: EME02

ABSTRACT

The Black-Scholes model fails to reproduce some empirical properties of volatility, from real market data, which means that the model may not be adequate to price certain options. Hence, the financial industry and academia focused on developing alternative models and methods for pricing. One of these models is the SABR model, a stochastic volatility model commonly used in nowadays financial industry. In this project, we aim to analyze and discuss the SABR model from both the theoretical and the practical point of view by motivating the usage of this model and showing how it can be applied in different aspects of the financial practice such as calibration, simulation and numerical methods to obtain implied volatilities. Throughout this analysis and discussion, we also aim to give some insights regarding what practitioners do, to give some sense of the applicability of the model and the methods in the real world.

KEYWORDS:

Option Pricing, Stochastic Volatility, SABR, Calibration, Monte Carlo, Brent Method, Surface

ACKNOWLEDGEMENTS

We want to especially thank our tutor and professor Elisa Alòs, which has been there whenever we needed help and has guided us through every aspect of this dissertation while increasing our interest in both mathematics and finance through the "Advanced Option Pricing and Modelling" module. Moreover, we would also like to thank Raúl Merino, our former "Financial Derivatives and Risk Management" module teacher, who recommended us to attend Professor Alòs advanced-level modules and introduced us to the financial mathematics world with kindness and professionally-valuable help and advice.

Last but not least, we would also like to thank Caixa Bank and Caixa Bank professionals for sharing real market data, from which we have calibrated our model and has been used throughout the discussions on this project.

Iker Caballero Bragagnini, Marc Colom Cervera & Iván Moreno Calixto,
Barcelona, June 16th, 2022

Table of Contents

1. Introduction

The Black-Scholes model was a breakthrough for the option pricing problem, but empirical evidence shows that it fails to reproduce some empirical properties of volatility, an important factor that affects option prices, from real market data, which means that the model may not be adequate to price certain options. Hence, the financial industry and academia focused on creating alternative models and methods for pricing that could improve this situation. One of these models is the SABR model, a stochastic volatility model commonly used in nowadays financial industry.

In this project, we aim to analyze and discuss the SABR model from both the theoretical and the practical point of view by motivating the usage of this model and showing how it can be applied in different aspects of the financial practice such as calibration, simulation and numerical methods to obtain implied volatilities. Throughout this analysis and discussion, we also aim to give some insights regarding what practitioners do, to give some sense of the applicability of the model and the methods in the real world.

We structure the project into four main sections. The first one is dedicated to motivating the usage of stochastic volatility models like the SABR by discussing first what a valid model is, the properties of the Black-Scholes model and the implied volatility surface. We then continue by delving into the SABR model and the famous Hagan's formula, which will be used to calibrate the model to real market data and compare the results obtained with professional calibration results.

The third section covers common simulation methods used to obtain option prices in the financial industry: the standard and the conditional Monte Carlo methods. We discuss theoretical results that back the methods and we use them to simulate prices under the SABR model, to compare both methods and get some valuable insights that can justify their usage in financial practice. Finally, we combine the calibration and the simulation methods to try to get the real option prices for the dataset and apply numerical methods to obtain the implied volatilities, while also getting practical insights when comparing and discussing the results to the original surface.

2. The Option Pricing Problem

Even though one may think a model's validity is derived from its accuracy when describing historical dynamics and its assumptions, practitioners consider a model valid when it can price derivatives consistently (not allowing for arbitrage opportunities) and make the least assumptions about future conditions (Bergomi, 2016). The main goal of a model is to be able to decompose the different risks coming from different contributions in the profits and losses (P&L) of a derivative position and secure a consistent valuation of these risks that allows to control them individually. This is a crucial principle because of its implications for option pricing models and the role of volatility modelling in the financial industry. However, an opportune question is which are the principal models and their practical implications.

In this section, we will take a closer look at the most influential model, the Black-Scholes model, its implications for volatility, and other kinds of models called "stochastic volatility models", where the pillar of our project, the SABR model, belongs.

2.1. The Black-Scholes Paradigm

The Black-Scholes model, developed by Black and Scholes (1973), is the most influential option pricing model, having such an impact on the way traders price and hedge derivatives that the authors received the Nobel prize for economics (Hull, 2021). By assuming a diffusion process for asset prices with constant volatility

$$dS_t = rS_t dt + \sigma S_t dW_t \tag{2.1}$$

where r is the risk-free interest rate, S_t is W_t is a geometric Brownian motion, or equivalently, by assuming the option price satisfies the following partial differential equation

$$\frac{\partial P_t}{\partial t} + \frac{\partial P_t}{\partial S} rS_t + \frac{1}{2}\frac{\partial^2 P_t}{\partial F^2}\sigma^2 S_t^2 = rP_t \tag{2.2}$$

where P_t is the Black-Scholes price of a European option at time t, the model allows us to obtain the price of European options through simple pricing equations

$$C_{BS} = S_0 N(d_1) - Ke^{-r(T-t)}N(d_2) \qquad P_{BS} = Ke^{-r(T-t)}N(-d_2) - S_0 N(-d_1) \tag{2.3}$$

$$for\ \ d_1 = \frac{\ln\left(\frac{S_0}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}} \quad and \quad d_2 = d_1 - \frac{\sigma\sqrt{T-t}}{2}$$

where K is the strike price, T is the option's maturity and $N(\cdot)$ is a normal probability distribution function. Even though the literature on this model has proven it is not consistent with asset prices' dynamics and makes unrealistic assumptions, the pricing equations derived from the model are used in daily financial practice, as they allow for decomposing of the P&L of derivative positions such as a short-hedged option position.

However, the one parameter in the formulae which cannot be observed in the market is volatility, as it refers to the volatility that asset prices will have during the life of the option. Even though it can be estimated with the historical volatility, in financial practice, traders usually work with what is called implied volatilities (Hull, 2021). Thus, the next section is dedicated to this crucial concept in option pricing.

2.2. The Implied Volatility and Its Surface

The implied volatility is the volatility parameter σ_{BS} that allows obtaining the market price of an option through the Black-Scholes model pricing equations. Moreover, Alòs and García (2021) show how the implied volatility can be also understood as a weighted mean of the future values of the volatility process, which supports Gatheral's (2006) reasoning. Hence, it is also an estimate of the option's volatility until maturity is reached, as mentioned before. However, we will restrict ourselves to the first definition for practical purposes. Considering a European call option with strike K and maturity T, we can see that, based on the previous definition, its market price V_t would equal its Black-Scholes price for the adequate volatility

$$V_t = C_{BS}(K, T-t, S_t, \sigma_{BS}) \tag{2.4}$$

so the implied volatility σ_{BS} can be defined as follows:

$$\sigma_{BS} = C_{BS}^{-1}(K, T-t, S_t, V_t) \tag{2.5}$$

Hence, under the Black-Scholes model, we can obtain the implied volatility of a vanilla option just from its market price. Because there is no analytical expression for C_{BS}^{-1}, numerical methods are usually applied in financial practice to obtain the implied volatility, such as the bisection, the Newton-Raphson, or Brent's method. Because we use the latter method, we will discuss it in the following sections.

Nevertheless, empirical evidence using real market data shows how the implied volatility is not the same for all strikes and maturities (Dumas et al., 1998), which is a basic assumption of the Black-Scholes model. In contrast, the implied volatility of a European option depends on its strike and time to maturity, which means we can obtain a set of implied volatilities for different values

of K and T, which conforms the so-called implied volatility surface. Moreover, an empirical analysis, as explained by Alòs and García (2021), shows how the implied volatility tends to exhibit a U-shaped pattern across different strikes for a given time to maturity. We call this phenomenon the "smile effect" if this pattern is symmetric, while we use the term "skew effect" when it is not. These smiles and skews change with time to maturity, being very pronounced for short maturities and flattening as the time to maturity increases, and the empirical skews seem to follow a power law as a function of time to maturity. These properties can be spotted in Figure 1, which is the implied volatility surface using real market data provided by Caixa Bank.

Figure 1. *Implied volatility surface from real market data*

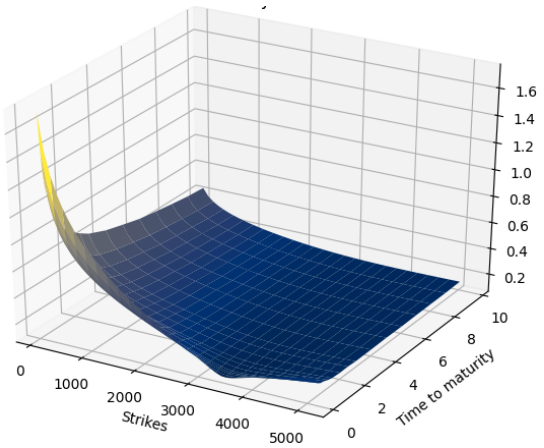

Source: *Own elaboration with real market data provided by Caixa Bank*

Thus, the Black-Scholes model is not able to reproduce volatility dynamics correctly. Moreover, there are other empirical properties apart from being non-constant that the model cannot reproduce, which motivates the use of alternative models. From a practitioner's point of view, it is necessary to use a model that can accurately reproduce the properties and behaviour of the implied volatility surface due to its effect on pricing financial derivatives and impact on the P&L of derivative positions. And this is the reason why many have shifted their focus to a kind of models that can help for this task: the stochastic volatility models, presented in the next section.

2.3. The Stochastic Volatility Models

Bergomi (2016) shows how directly modelling the implied volatility dynamics (by assuming a diffusion process for S_t and another for I_t) is impractical given that the dynamics of the underlying determine the dynamics of the implied volatility. Additionally, Bergomi (2016) defends that more sophisticated models would allow to precisely characterize the P&L and the conditions when it vanishes, which makes it easier to spot the contribution of each effect rather than grouping them by omitting market microstructure-related parameters, as simpler models do.

Consequently, many researchers and practitioners have tried to construct models which allow specifying implied volatility dynamics in a less restricted way.

One type of models developed are the stochastic volatility models. Based on Alòs and García (2021), these are models where the volatility σ_t is equal to a function f of a diffusion process Y_t which satisfies a stochastic differential equation of the form

$$dY_t = a(t, Y_t)dt + b(t, Y_t)dB_t \tag{2.6}$$

where a and b are deterministic functions and B_t is a Brownian motion. Hence, these models consider that the volatility follows a diffusion process which is driven by a different Brownian motion from the one which drives the underlying prices. Literature on these models show how some empirical properties of the implied volatility surface are reproduced, while others are not, such as the dependence on time to maturity (Renault and Touzi, 1996; Lee, 2002; Medvedev and Scaillet, 2007). This last fact is important because it determines how models are calibrated and used in the financial industry and our approach in the following section.

Even though these models do not perfectly reproduce the implied volatility surface, they are commonly used. One of the most popular models used nowadays is the SABR model, proposed by Hagan et al. (2002). We will discuss the model and its calibration in the following sections.

3. The SABR Model

The stochastic alpha-beta-rho model or SABR model, proposed by Hagan et al. (2002) is a stochastic volatility model which is being widely used in the financial industry due to its ability to reproduce some important volatility surface properties through just three parameters and because of the existence of a closed-form solution from which one can derive a formula for the corresponding implied volatility, called Hagan's formula. However, there are some problems with this model which limits its potential and applicability. Furthermore, it is worth it to delve into the usage of the model in financial practice, so we calibrate the model with real market data and compare it to a professional calibration to get valuable insights about methodology and application.

In this section, we introduce and discuss some properties of the SABR model and Hagan's formula, and we shift to the discussion of our own calibration of the model with real market data and the results obtained with respect to professional ones.

3.1. The Model

The SABR model is a two-factor model for the forward prices of the underlying, whose second factor (the spot volatility σ_t) is a stochastic process. Based on the original model, it can also be mathematically defined as follows:

$$\begin{cases} dF_t = \sigma_t F_t^{\beta} dW_t \\ d\sigma_t = \alpha \sigma_t dB_t \end{cases} \quad and \quad \rho = \frac{d\langle W,B \rangle_t}{dt} \quad for \ t \in [0,T] \tag{3.1}$$

where F_t is the forward price of the underlying, W_t and B_t are Brownian motions and ρ is the correlation between them. In this model, α is a parameter $\alpha > 0$ which can be understood as the volatility of volatility, while β is a parameter $\beta \in [0,1]$ that determines how at-the-money volatility changes when forward price changes.

The selection of β could be done through historical data of forwards and implied volatility pairs, but it is normally done through *a priori* considerations (Hagan et al., 2022). The three principal exponents considered in the industry are $\beta = 1$ (which makes the model stochastic lognormal), $\beta = 0$ (which makes the model stochastic normal), and $\beta = 1/2$ (which makes the model stochastic Cox-Ingersoll-Ross). Because proponents consider lognormal models are "more natural" and it is difficult to distinguish the effect of β from that of ρ because both cause a downward sloping skew as the strike varies (Hagan et al., 2002), practitioners usually choose $\beta = 1$ for calibration and analysis purposes. Moreover, $\beta = 1$ means that at-the-money volatility

should remain constant for changes in forward prices (as $K = F_t$). This is relevant as we will use this exponent when calibrating the model to real market data.

The effect of these parameters on the implied volatility and the implications for modelling can be better understood by looking at Hagan's formula, which is derived from the model and is discussed in the next section.

3.2. Hagan's Formula

Hagan's formula allows us to express the implied volatility of an option in terms of its strike price, time to maturity and the parameters previously mentioned. This formula is derived from the Black's formula for European option prices (Black, 1976)

$$C_B = e^{-rT}[F_0 N(d_1) - KN(d_2)] \qquad P_B = V_0^{call} + e^{-rT}[K - F_0] \tag{3.2}$$

$$for\ \ d_1 = \frac{\ln\left(\frac{F_0}{K}\right) + \frac{1}{2}\sigma_B^2 T}{\sigma_B\sqrt{T}} \quad and\ \ d_2 = d_1 - \sigma_B\sqrt{T}$$

where σ_B is the constant volatility parameter assumed by Black's model (which will be discussed later). Given these formulae, Hagan et al. (2002) derive the following analytical expression for the implied volatility σ_B from the SABR model:

$$\sigma_B(F_0, K) = \frac{\sigma_0}{(F_0K)^{\frac{1-\beta}{2}}\left[1+\frac{(1-\beta)^2}{24}\log^2\left(\frac{F_0}{K}\right)+\frac{(1-\beta)^4}{1920}\log^4\left(\frac{F_0}{K}\right)+\cdots\right]} \cdot \left(\frac{z}{x(z)}\right)$$
$$\cdot \left[1 + \left(\frac{(1-\beta)^2}{24}\frac{\sigma_0^2}{{F_0}^{2(1-\beta)}} + \frac{1}{4}\frac{\rho\beta\alpha\sigma_0}{{F_0}^{1-\beta}} + \frac{2-3\rho^2}{24}\alpha^2\right)T\right] + \cdots \tag{3.3}$$

where σ_0 is the initial volatility and the z and $x(z)$ terms are given by these expressions:

$$z = \frac{\alpha}{\sigma_0}\left(\frac{F_0}{K}\right)^{\frac{1-\beta}{2}}\log\left(\frac{F_0}{K}\right) \qquad x(z) = \log\left(\frac{\sqrt{1-2\rho z+z^2}+z-\rho}{1-\rho}\right) \tag{3.4}$$

Moreover, the formula can be further simplified in the case of at-the-money options (where $K = F_0$) because of some logarithms vanishing:

$$\sigma_B(F_0, K) = \frac{\sigma_0}{(F_0)^{1-\beta}} \cdot \left(\frac{z}{x(z)}\right) \cdot \left[1 + \left(\frac{(1-\beta)^2}{24}\frac{\sigma_0^2}{{F_0}^{2(1-\beta)}} + \frac{1}{4}\frac{\rho\beta\alpha\sigma_0}{{F_0}^{1-\beta}} + \frac{2-3\rho^2}{24}\alpha^2\right)T\right] + \cdots \tag{3.5}$$

As Hagan et al. (2002) indicate, the terms that account for "..." can be omitted because of their small size, even though this omission can lead to relative errors when pricing options. Alòs and

García (2021) provide an approximation for both formulae (3.3) and (3.5) when ignoring the small-sized terms and when $\beta = 1$

$$\sigma_B(F_0, K) \approx \sigma_0 \left[1 + \left(\frac{\rho\alpha\sigma_0}{4} + \frac{2-3\rho^2}{24}\alpha^2\right)T\right]\left(\frac{z}{x(z)}\right) \tag{3.6}$$

$$\sigma_B(F_0, K) \approx \sigma_0 \left[1 + \left(\frac{\rho\alpha\sigma_0}{4} + \frac{2-3\rho^2}{24}\alpha^2\right)T\right] \tag{3.7}$$

The complete formulae and the approximations can show how the implied volatility depends on the main parameters of the model. Both α and ρ affect on the implied volatilities skew, as high values of α make the curvature increase, while high values of $|\rho|$ make the skew more pronounced. One important result related to the skew of the implied volatility surface of the SABR model and the effect of the parameters is presented by Alòs and García (2021). Through the use of Malliavin calculus, the authors obtain the short-time limit of the skew slope for the implied volatility for the SABR model

$$\lim_{T \to t} \frac{\partial \sigma_{B,t}}{\partial k}(k_t^*) = \frac{\alpha\rho}{2} \tag{3.8}$$

where $\sigma_{B,t}$ is the implied volatility in time t, k is the log-strike and k_t^* is the at-the-money log-strike in time t. What this result means is that the skew for short maturities tends to a finite value which depends on the product $\alpha\rho$. Nevertheless, real market data shows the existence of a blow-up in the short-end of the implied volatility surface, which means that the limit of the skew tends to infinity. Consequently, the SABR model is not able to reproduce this short-end blow-up. Furthermore, classical stochastic volatility models which assume a diffusion process cannot reproduce it either because the limit of the short-time skew slope is also finite, as shown by Alòs et al. (2007).

Another problem with the model arises from the fact that Hagan's formula presented is not arbitrage-free, as the approximate density function generated by the model leads to negative values for low strike prices (Alòs and García, 2021). This will likely yield incorrect results and create arbitrage opportunities, which is a key issue for its financial application and has led to different methods for solving it (Grzelak and Oosterlee, 2017; Le Floc'h and Kennedy, 2016). On the other hand, these simple formulae allow practitioners to easily calibrate the model, so some of these structural problems with the model are offset by its practicality.

In the following section, we delve more into the model and Hagan's formula usage to compare the results of our own calibration with a professional calibration and to give valuable insight into how it is applied in the financial world.

3.3. SABR Model Calibration

To give some sense of how the so-called quants and traders use the SABR model in daily financial practice and to apply both the model and Hagan's formula, we calibrate the SABR model using real market data provided by Caixa Bank. This dataset contains the implied volatility for many strike prices and maturities for a European call option, which has been represented by the surface in Figure 1, and the underlying's forward price for each maturity. The goal of calibrating a model is to obtain parameters that allow reproducing market volatility surface for all strike prices and maturities.

However, as mentioned before, stochastic volatility models such as the SABR are not able to reproduce the dependence of the implied volatility with respect to time to maturity, so trying to obtain a set of parameters for all maturities would not be effective. A common practice in the financial industry is to calibrate the model obtaining a set of parameters for each fixed maturity, so we will use this approach for calibrating. Hence, we need to minimize the sum of squared errors to obtain a triple of α, ρ and σ_0 for each time to maturity

$$(\widehat{\alpha}, \widehat{\rho}, \widehat{\sigma_0})_T = \underset{\alpha,\rho,\sigma_0}{\operatorname{argmin}} \sum_{K_T} \left(\sigma_{K,T}^{mkt} - \sigma^{SABR}(F_0, T, K, \alpha, \rho, \sigma_0)\right)^2 \tag{3.9}$$

where $\sigma_{K,T}^{mkt}$ is the market implied volatility and σ^{SABR} is the SABR implied volatility for a given strike K and maturity T. In this context, the implied volatility of the SABR model is the implied volatility obtained through Hagan's formula, and we use approximation (3.6) in the minimization problem as the gains of precision for including other terms are not significant (Hagan et al., 2002). In the Python code used for calibration (available through a link in the Appendix), we use the "*curve_fit"* function from the "*Scipy"* package to obtain the parameters for each maturity.

The resulting parameters are summarized in Table B1 (in the Appendix) and Figure 2, which show the calibrated parameters for each time to maturity. The results show how the alpha parameter changes a lot depending on time to maturity, having higher values for short maturities and lower values for longer ones. However, the other parameters seem to be more or less constant across different maturities, even though we can observe that they increase for longer times to maturity. By comparing it to Figure 3, which shows the professionally-calibrated parameters by Caixa Bank, we can see that our results are very similar, especially the alpha parameter.

Figures 2 (left) & 3 (right). *Graphs of alpha (blue), rho (yellow) and initial sigma (green) for different times to maturity*

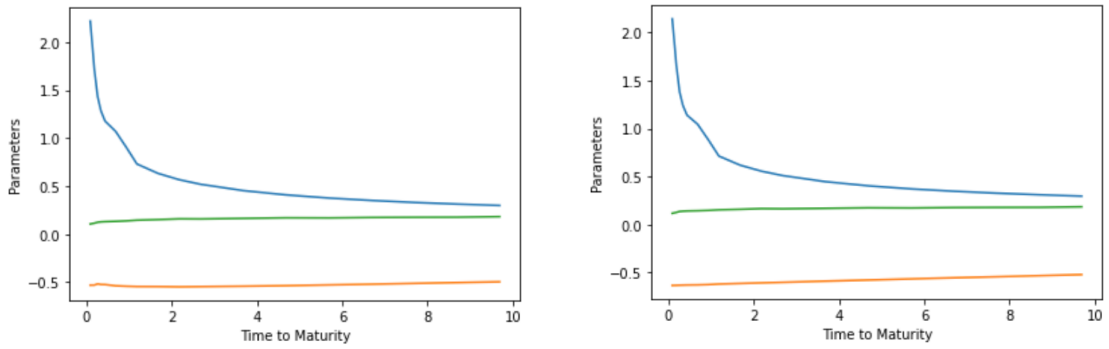

Source: *Own elaboration with real market data provided by Caixa Bank*

We can show that rho and the initial sigma are increasing in time to maturity in a clearer way by looking at Figures 4 and 5, which plot the parameters separately. However, comparing the results to Caixa Bank's calibration with Figures 6 and 7, we can see that there are some differences. When it comes to the values of rho, we can check that the values obtained are greater for shorter maturities while they are lower for longer ones. We can also see that in our calibration there is a range of values for which rho decreases, while the calibrations of Caixa Bank show that it always increases with maturity. Looking now at the initial sigma, we can look that the values are almost equal to Caixa Bank's calibrated values, which indicates an accurate calibration for the initial sigma.

Figures 4 (left) & 5 (right). *Graph of rho (right) and initial sigma (left) for different times to maturity*

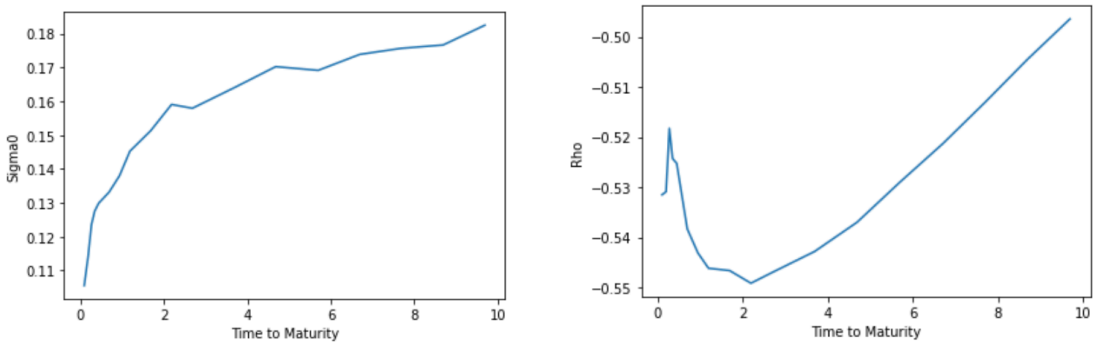

Source: *Own elaboration with real market data provided by Caixa Bank*

Figures 6 (left) & 7 (right). *Graph of rho (right) and initial sigma (left) for different times to maturity calibrated by Caixa Bank*

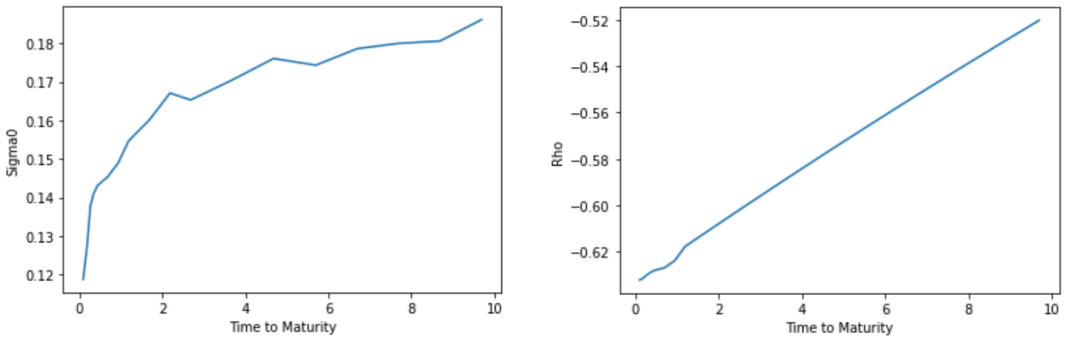

Source: *Own elaboration with real market data provided by Caixa Bank*

These differences are due to the methodology used for calibration: in our case, we just used an optimization function from a Python package (code available through a link in the Appendix), but Caixa Bank relies on professional and more accurate methods for this optimization problem, using some restrictions that allow obtaining better results. Hence, it is expected to see some discrepancies with professional calibrations.

As shown earlier, both the alpha α and the rho ρ affect the skew of the implied volatility and it also determines the short-time limit of the slope of the skew in the SABR model. So, by multiplying both parameters for each of the maturities, we can represent the skew as a function of time to maturity. When doing so for the parameters calibrated by us and the ones calibrated by Caixa Bank, we can observe that a power law is followed, graphically representing this result in Figures 8 and 9 below.

Figures 8 (left) & 9 (right). *Power law behaviour of $\alpha\rho$ with our calibrated parameters (left) and Caixa Bank's calibrated parameters (right)*

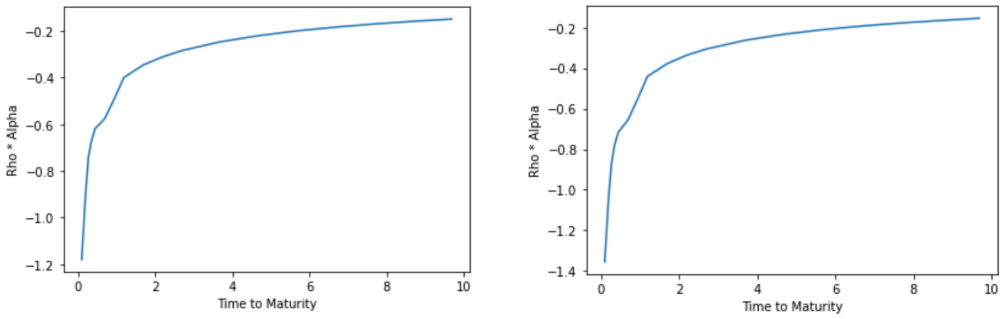

Source: *Own elaboration with real market data provided by Caixa Bank*

Recalling that one empirical property of implied volatilities is that the skews tend to follow a power law, we can see how our calibrated parameters are consistent with this property, approximately following a power law close to $-0.4388 \cdot T^{-0.43104}$.

Finally, we can compare the implied volatility surface obtained through the Hagan formula and our calibrated SABR parameters to the implied volatility surface of the real market data surfaces graphically through Figures 12 and 13. As we can observe, both surfaces seem almost equal, which indicates that the calibration was very accurate. By looking at the sections for the shortest and the longest time to maturity in Figure 12, it is even clearer that the calibrated model fits the data accurately and that some empirical properties of the implied volatility surface, such as the flattening of the skew when the time to maturity increases or the short-end blow-up.

Figures 10 (left) & 11 (right). *Implied volatility surface from the data (left) and implied SABR volatility using calibrated parameters (right)*

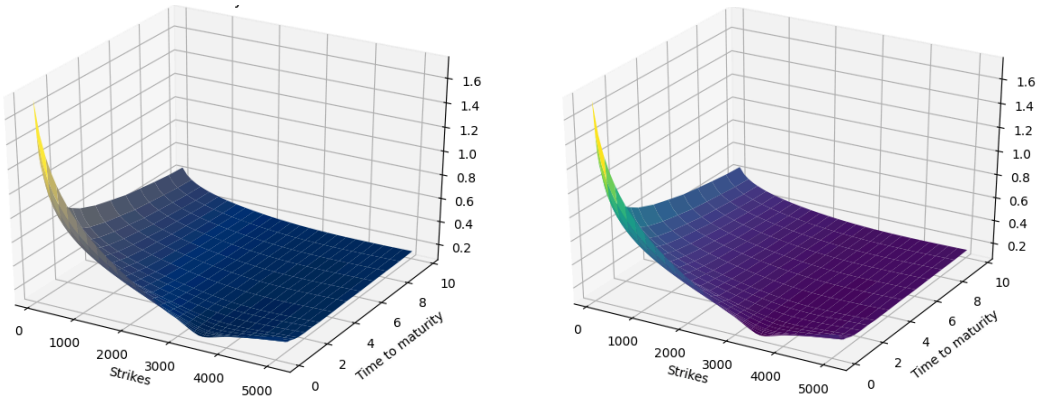

Source: *Own elaboration with real market data provided by Caixa Bank*

Figure 12. *Section of the implied volatility surface for short and long times to maturity*

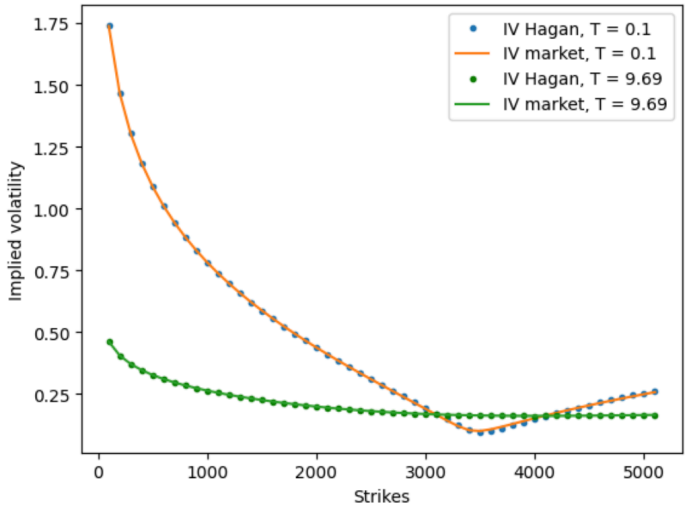

Source: *Own elaboration with real market data provided by Caixa Bank*

After discussing the obtained results from our own calibration and comparing them to Caixa Bank's calibration, we can conclude that the SABR parameters obtained through our calibration methodology are very similar to the ones obtained by Caixa Bank, yielding an almost equal implied volatility surface. Moreover, we see that the calibrated SABR model satisfies some of the empirical properties of implied volatilities. We also spot some differences regarding the calibrated parameters with the parameters calibrated by Caixa Bank, even though this might be due to the use of more sophisticated methods to obtain better results, which are not used in our calibration procedure.

4. Simulation of Option Prices

To obtain the option prices for each strike price and maturity of the real market dataset, common simulation methods used in the financial industry can be applied, such as the Monte Carlo and the Conditional Monte Carlo method. These methods are based on the principles of risk-neutral valuation, the Hull and White formula and the Black-Scholes model.

Through these simulations, we expect to show how simulation methods work, why they are used in the financial industry and some valuable insights from their comparison. Due to our dataset using forward prices, we will also discuss the Black and the Black-Scholes model usage for the prices' simulation.

4.1. The Monte Carlo Method

A Monte Carlo simulation of a stochastic process is a procedure for sampling random outcomes for the process (Hull, 2021). When applied to financial engineering and risk management, this method is often used to obtain the expected value of a random variable modelled by a stochastic process, such as the underlying asset prices, or to estimate measures such as "Value at Risk" (VaR) and the so-called "Greeks" (Chen and Hong, 2008). Furthermore, Chen and Hong (2008) defend the helpfulness of this method in the financial industry because of its easy usage, the independence between convergence and problem dimensionality and the avoidance of the solution existence verification problem (in contrast to using a PDE approach to price options).

To justify the usage of the Monte Carlo method for option pricing, we will follow the Hull (2021) approach, based on risk-neutral valuation, one of the crucial principles of derivatives pricing. In a risk-neutral world, every investor is risk-neutral, meaning that they would not ask for a risk premium in their investments, as they do not have risk preferences. Hence, they will expect to get an expected return equal to the risk-free rate (which is constant), and they will use an identical discount rate.

Knowing that the value of an option at maturity is its intrinsic value, then the expected value of the option at maturity in the risk-neutral world is $E^*[h(S_T)_+]$, where E^* denotes the expectations with risk-neutral probabilities and $h(S_T)$ is a payoff function that depends on S_T between the payoff at maturity and zero. Then, we can obtain the option price of a European option by discounting the expected value using the risk-free rate:

$$C = e^{-r(T-t)}E^*[h(S_T)_+] \qquad P = e^{-r(T-t)}E^*[h(S_T)_+] \tag{4.1}$$

Because a European option's price depends on the expected value of the option at maturity, then we can price these kinds of options by obtaining the expected value, which means that there is room for using the Monte Carlo method to compute it and obtain the option prices.

In our case, we assume a diffusion for both asset prices and the spot volatility under the SABR model, and we obtain an estimate of the option's value at maturity by simulating multiple paths, computing the option value at maturity, and then discounting its average. However, due to the use of a stochastic volatility model to price options, we need to simulate two different but correlated Brownian motions, which makes the simulation slow and with high variability (as there are two sources of randomness). A way to overcome this issue is to use an alternative method, known as the conditional Monte Carlo method.

4.2. The Black and the Black-Scholes Model

Before delving into the conditional Monte Carlo method, we first explain Black's model and its relationship with the Black-Scholes model to understand our simulation approach. As mentioned, our dataset uses forward prices of the underlying asset instead of the usual price levels, while it just gives the implied volatility for each strike and maturity. Hence, the Black-Scholes model cannot be used to model asset prices because we are not using price levels nor do we have the underlying asset price at $t = 0$, so we need to use another model.

A widely-used model for pricing forward options using forward prices is the Black model (Black, 1976), which states that forward prices follow a diffusion process of the form

$$dF_t = \sigma F_t dW_t \tag{4.2}$$

or, equivalently, that they satisfy the following partial differential equation:

$$\frac{\partial P_t}{\partial t} + \frac{1}{2}\frac{\partial^2 P_t}{\partial F^2}\sigma^2 F^2 = rP_t \tag{4.3}$$

By assuming this process, Black derives the formulae for European forward options' prices shown in (3.2), carrying the same advantages for financial practice as the ones for the Black-Scholes model do. In fact, one can notice equations in (2.3) and (3.2) are very similar, and we can show in which conditions do the equations coincide. Using the differential equation (2.2) and (4.3), we can see that if $r = 0$ then

$$F_t = e^{rT} S_t \;\Rightarrow\; F_t = S_t \tag{4.4}$$

and the Black-Scholes partial differential equation is equivalent to the Black's partial differential equation. Given this result, the Black-Scholes pricing equations and the Black pricing equations

are equivalent, as they are the solution for the equivalent partial differential equation. This is why we will use the Black-Scholes model for the conditional Monte Carlo simulations of option prices.

4.3. The Conditional Monte Carlo Method

The conditional Monte Carlo method is a simulation method very similar to the standard Monte Carlo Method, which extracts samples from the conditional distribution of a random variable given a value for some concrete function (Lindqvist et al., 2020). At first, this method had limited use in financial practice because of the presumption that one can just use it when the factors were instantaneously uncorrelated (Willard, 1997), but once this presumption was proved false, it became one important research topic in mathematical finance and has many applications in the financial world.

When applied to option pricing, we consider a Monte Carlo simulation for the conditional expected value given all the available information of the brownian motions at t, $\mathcal{F}_t$, which will allow us to only simulate one of the two Brownian motions present in stochastic volatility models (one for each diffusion process). Thus, the variability and the computation time of the results should be reduced due to one randomness source not being computed. That is why this method is often classified as a variance reduction method and carries many practical advantages.

As we did with the standard Monte Carlo method, we justify this method through the Hull and White formula (Hull and White, 1990) extension for correlated Brownian motions found by Romano and Touzi (1997) and Willard (1997), which states that

$$C_t = E^* \left[C_{BS} \left(T - t, X_t \xi_t, \sqrt{1-\rho^2} v_t \right) \mid \mathcal{F}_t \right] \tag{4.5}$$

$$for\ \ \xi_t = \rho \int_t^T \sigma_u dW_u - \frac{\rho^2}{2} v_t^2 (T-t) \ \ and\ \ v_t = \sqrt{\frac{1}{T-t} \int_t^T \sigma_u^2 du}$$

where X_t is the log-price. This result follows from the diffusion processes assumed and the existent correlation between them, and its derivation is available in the Appendix. Nevertheless, the results show that we just need to simulate the brownian motion W_t to obtain the price of the European option.

Notice that to apply this method we need to use Black-Scholes pricing equations, but the dataset uses forward prices. Thanks to the results for Black's and Black-Scholes model presented above, there will be no problem using this method as long as $r = 0$, because the two models would be equivalent. Consequently, we will assume $r = 0$ when applying the simulation methods.

In the following section, we discuss the results of using this method compared to the standard Monte Carlo method.

4.4. Comparing Methods

To exemplify the usage of these methods and support the claims made on them, we simulate option prices assuming the SABR model dynamics, using a standard and a conditional Monte Carlo method for different number of simulation trials.

The first thing we noticed is that the standard Monte Carlo method is very easy to program, while the conditional Monte Carlo one involves more programming due to the presence of integrals. Yet, one should not be misled by this fact, as it is expected to be faster than the standard method. The Python code used for the simulations with both methods is available through a link in the Appendix.

Using the code, we simulated prices for a European call option with strike $K = 80$, underlying initial price $S_0 = 100$, risk-free interest rate $r = 0$, initial volatility $\sigma_0 = 0.15$ and maturity $T = 1$. Moreover, we assumed a SABR model with alpha $\alpha = 0.4$ and $\rho = 0.6$, and we always use 1000 time steps, so the only aspect we are changing is the number of simulated option prices and the number of simulated paths for the underlying. The results of these simulations are shown in Tables 2, 3 and 4 below.

Table 2. *Numerical results with different nº of simulated option prices and n=100*

Method	*Time Steps*	*Nº of Simulated Paths*	*Nº of Simulated Option Prices*	*Time (seconds)*	*Mean Price*	*Std. Dev.*
MC	1000	100	100	0.6962	19.6487	3.1845
CMC	1000	100	100	0.6493	19.5615	2.7946
MC	1000	100	500	0.7731	19.7439	2.4736
CMC	1000	100	500	0.7318	19.7602	1.9861
MC	1000	100	1000	0.7154	20.2455	1.7986
CMC	1000	100	1000	0.6634	19.9201	0.8921

Source: *Own elaboration*

Table 3. *Numerical results with different nº of simulated option prices and n=250*

Method	*Time Steps*	*Nº of Simulated Paths*	*Nº of Simulated Option Prices*	*Time (seconds)*	*Mean Price*	*Std. Dev.*
MC	1000	250	100	1.7622	19.9182	2.9581
CMC	1000	250	100	1.6070	19.5608	2.7944
MC	1000	250	500	1.7690	20.0389	2.2666
CMC	1000	250	500	1.6228	19.7598	1.9860
MC	1000	250	1000	1.0655	20.0667	1.9238
CMC	1000	250	1000	0.9842	19.9356	1.7860

Source: *Own elaboration*

Table 4. *Numerical results with different nº of simulated option prices and n=500*

Method	*Time Steps*	*Nº of Simulated Paths*	*Nº of Simulated Option Prices*	*Time (seconds)*	*Mean Price*	*Std. Dev.*
MC	1000	500	100	4.3294	19.5743	2.8670
CMC	1000	500	100	3.2489	19.5589	2.7941
MC	1000	500	500	4.3628	19.8743	2.1378
CMC	1000	500	500	4.3529	19.7920	2.1086
MC	1000	500	1000	4.4021	19.9685	1.3742
CMC	1000	500	1000	3.3649	19.8912	1.1504

Source: *Own elaboration*

For each table, we can see how the standard deviation and the computation time of the conditional Monte Carlo method (CMC in the tables) are less than for the standard Monte Carlo method. Moreover, we can also observe how the standard deviation decreases when the number of simulated prices increases, meaning that if we simulated a high number of prices, we would tend to the real value of the option. This latter result is a consequence of the backbone of Monte Carlo methods, the Central Limit Theorem, which states that, as we increase the number of trials (in this case, the number of simulated option prices), the sample average will tend to the mean, which would be the real option price. However, we highlight the fact that we have not simulated a large number of prices nor paths as it is done in the financial industry, so the mean value could be more precise and the standard deviation could be lower. Yet, these simulations have the purpose of exemplifying and backing some well-known results, so this should not be considered an important caveat.

To sum up, we can see how Monte Carlo methods can be applied to obtain option prices using different models such as the Black-Scholes or the SABR model. Furthermore, the more simulated paths and option prices, the more reliable will be our estimation of the price. The results of applying both methods evidence how we obtained very similar prices, but the computation time and the variability of the standard Monte Carlo method are higher than for the conditional Monte

Carlo. This fact can explain why some practitioners would rather use the conditional method when it comes to efficiency, even though the standard method can be applied easier. This is the main motive explaining why we use the conditional method for obtaining the implied volatility surface from the prices computed using this method, as we will explain in the next section.

5. Implied Volatility Surface Comparison

In this section, we combine the discussed simulation methods and the SABR model's calibration to reproduce the implied volatility surface for the options from market data using numerical methods. The rationale behind this is the following: using the calibrated parameters of the SABR model and the conditional Monte Carlo method, we can simulate prices which should tend or be very close to the real market prices if enough simulations are made. And to test this hypothesis, we can use these simulated prices as the market prices and compute the implied volatility for the Black-Scholes prices to construct the implied volatility surface, which should be the same as the one for the real market data if simulated prices are equal to the market prices.

To do so, we first compute the forward prices using the conditional Monte Carlo simulation method and the calibrated parameters of our SABR model for each maturity. Then, we apply numerical methods to obtain the implied volatility for the Black-Scholes prices using the simulated prices as market prices, and we finally compare the obtained implied volatility surface with the one from the dataset to extract some valuable conclusions from the SABR model. We start by explaining the numerical method used, and we proceed by looking at the results.

5.1. Brent's Method

As mentioned before, we will be using Brent's method to obtain the implied volatilities for each strike and time to maturity. The method created by Brent (1971) is a hybrid method which combines the bracketing method of the bisection method and the inverse quadratic interpolation to find the zero of a given a function, even though the idea of this method was first proposed by Dekker (1969), which combines the bisection with the secant method. We briefly explain this method to better understand the procedure used for obtaining the implied volatilities.

Every iteration involves a current iterate b_k, a counterpoint a_k such that $sign\big(f(b_k)\big) \neq sign\big(f(a_k)\big)$ and a previous iterate b_{k-1} (which is $b_{k-1} = a_0$ for the first iteration), and two potential values are provided for the following iterate (representing the potential iterate value by s), where the first one belongs to the inverse quadratic interpolation method and the second belongs to the bisection method.

Then, the method decides the value of the next iteration by testing if two inequalities hold simultaneously. Given a numerical tolerance δ, if the previous step used the bisection method, then

$$\begin{cases} |\delta| < |b_k - b_{k-1}| \\ 2|s - b_k| < |b_k - b_{k-1}| \end{cases} \tag{5.1}$$

must hold to perform interpolation, otherwise, the bisection method is performed and the value is assigned to the next iteration. If the previous step used the inverse quadratic interpolation method, then

$$\begin{cases} |\delta| < |b_{k-1} - b_{k-2}| \\ 2|s - b_k| < |b_{k-1} - b_{k-2}| \end{cases} \tag{5.2}$$

must hold to perform interpolation, otherwise, the bisection method is performed and the value is assigned to the next iteration. We do not give a more detailed explanation of the workings of this method because we only intend to illustrate its mechanics and it is out of the scope of this project. We chose this method because it converges faster and avoids problems of more traditional methods such as the Newton-Raphson (which does not guarantee convergence) or the bisection method (which converges very slowly).

To apply this method, we use the function "*optimize.fminbound"* from "*Scipy"* Python's library (again, the code used is available through a link in the Appendix), which uses the Brent method but ensures that the minimum is achieved in the specified bracket. The minimization problem where we will apply this method is the following

$$\min_{\sigma_{BS}} \left(C_{BS}(S_0, K, T-t, \sigma_{BS}) - C_{SABR}\right)^2 \tag{5.3}$$

where C_{BS} is the Black-Scholes price of the corresponding European call, σ_{BS} is the implied volatility of the Black-Scholes price and C_{SABR} is the simulated price by the conditional Monte Carlo method assuming the calibrated SABR (assumed to be the observed market price). The result is the implied volatility for each strike and time to maturity available in the dataset, and all these volatilities will allow us to construct the implied volatility surface that will be compared to the original one in the following section.

5.2. The Obtained Surface vs The Original Surface

As stated, our aim for comparing the surface obtained by simulating the prices via the conditional Monte Carlo method with the calibrated parameters of the SABR model and applying the Brent method is to obtain the implied volatilities and test whether the option prices obtained are nearly the same option prices observed in the markets. Through the comparison of the surfaces, we will discuss whether this is the case and give some intuition about why we obtained these results, as it allows to point out some interesting properties of option prices.

We plot the implied volatility surface obtained through the Brent method jointly with the original implied volatility of Caixa Bank's dataset in Figure 13, while we also show the individual implied volatilities for each strike and time to maturity by plotting the points in Figure 14.

Figures 13 (left) & 14 (right). *Implied volatility surfaces (left) and implied volatilities for each strike and time to maturity (right)*

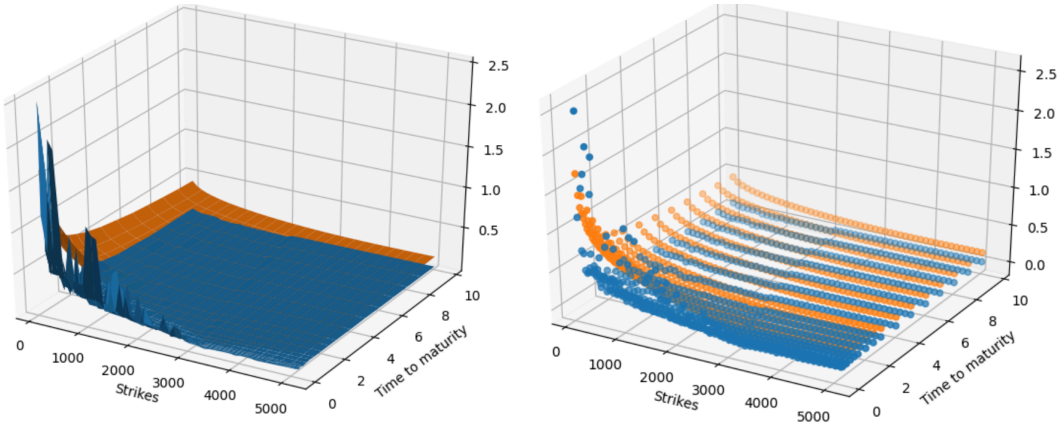

Source: *Own elaboration with real market data provided by Caixa Bank*

As we can observe, the plots are very similar for many strike prices and times to maturity, but both surfaces differ in the short-end of the surface, as the obtained implied volatilities for short times to maturity and low strike prices are larger than the original ones. At first, this fact might indicate that the simulated prices are not near enough to the real market prices, which would make the implied volatilities differ. This might be the case because we did not simulate a very high number of paths (we just simulated ten thousand paths) that will assure the convergence to the real option price. In financial practice, simulated paths could range from one thousand to one million, being higher numbers desirable. However, there is another explanation related to the methodology used to obtain the implied volatilities and the behaviour of option prices for low strikes and short times to maturity.

A well-known result is that the Vega of a European option (the sensibility of the price with respect to volatility) is at its highest when the option is at-the-money. Under the Black-Scholes model, the Vega of an option can be represented as follows

$$\mathcal{V} = \frac{\partial C_{BS}}{\partial \sigma} = S_0\sqrt{T}N'(d_1) \quad where \quad d_1 = \frac{\ln\left(\frac{S_0}{K}\right)+\frac{\sigma^2}{2}T}{\sigma\sqrt{T}} \tag{5.4}$$

where $N'(\cdot)$ is the standard probability density function, and the value of d_1 that maximizes the function is $\sigma\sqrt{T}/2$. This result can also be extended to stochastic volatility models because the Vegas are very similar (Hull and White, 1987; 1988). Hence, when an option is deep in or out-of-the-money for a given time to maturity, the Vega of an option decreases a lot, which means that the prices do not vary too much as volatility changes. If these deep out or in-the-money options have a very short time to maturity (like in the short-end of the implied volatility surface), then there is no time for the underlying asset price to change too much, and the sensibility of option prices to volatility is even lower.

Consequently, a very low sensibility of the price with respect to volatilities for the options at the short-end (which are deep in-the-money) makes the Brent method applied yield very high implied volatilities. This highlights that the method used for this project is unstable and financial practice requires more tailor-made methods.

6. Final Conclusions

Option pricing models in financial practice are not used because they capture historical dynamics or they have realistic assumptions, but because they allow to price options consistently and they have few assumptions about the future. Even if the Black-Scholes model is not empirically accurate, it is used in financial practice and allows practitioners to work with its implied volatility. However, real market data shows how the model is not able to reproduce some properties of implied volatilities, and hence this motivates the creation of models such as the stochastic volatility models, which reproduce some properties of the implied volatility surface that allow practitioners to price options more consistently.

One of the most important stochastic volatility models is the SABR model, which allows to model forward prices and volatility through diffusion processes that depend on three parameters. Moreover, this model is very popular among practitioners because a formula for implied volatility which allows obtaining implied volatilities given the parameters of the model, called Hagan's formula. Using the model and the Hagan's formula, we calibrate the SABR model to real market data and the results of this calibration show how the SABR model reproduces some important implied volatilities empirical properties and that the calibrated parameters and surface are very similar to the ones obtained by Caixa Bank professionals, even though there are some discrepancies because of more sophisticated methods usage.

When it comes to options prices obtention through the implied volatilities, we focus on common simulation methods such as the standard and the conditional Monte Carlo methods. These methods rely on mathematical and financial theory results that show how prices depend on expectations, so we can obtain prices by simulating prices under the SABR model. The results of comparing both methods show how the conditional Monte Carlo method is less time-consuming and reduces the variability of the option prices, which is why many professionals in the industry use this method to simulate.

Finally, we combine all the aspects above and numerical methods to test whether the prices obtained through the conditional Monte Carlo method with the calibrated SABR model are close enough to the real option prices. By applying Brent's method to the minimization problem posed, we get an implied volatility surface very similar to the original one except for the short-end implied volatilities, for which the values extracted from the method are higher. This could be due to not simulating enough paths for each price or to the low sensibility of the option price to volatility in the short-end, as these options are deep in-the-money, which would highlight the instability of the method used in the project and the need of using more tailor-made methods commonly used in financial practice.

This project could be extended by studying other types of option pricing models which capture more or at least different empirical properties of the implied volatility surface. Candidates could be stochastic-local volatility models, alternative stochastic volatility models like the Heston model (Heston, 1993), models that include more volatility factors like the two-factor Bergomi model (Bergomi, 2005), jump-diffusion models and even long-memory volatility models as Comte and Renault (1998) ones. Another possible extension can be to include more sophisticated methods used in financial practice for calibration of the models and to accurately represent implied volatility surfaces (for example, by interpolating or by using restrictions when calibrating).

References

Alòs, E. and García, D. (2021). *Malliavin calculus in finance theory and practice*. Boca Raton Chapman & Hall/Crc Press.

Alòs, E., Leon, J. and Vives, J. (2007). On the short-time behaviour of the implied volatility for jump-diffusion models with stochastic volatility. *Finance and Stochastics, 11,* 571-589.

Bergomi L. (2005). *Smile dynamics II*. Risk October, pp. 67-73

Bergomi, L. (2016). *Stochastic volatility modeling*. Boca Raton ; London ; New York: Crc Press, Cop.

Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, [online] 3(1-2), pp.167–179. doi:10.1016/0304-405x(76)90024-6.

Black, F. and Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, [online] 81(3), pp. 637–654. doi:10.1086/260062.

Brent, R. P. (1971). An algorithm with guaranteed convergence for finding a zero of a function. *The Computer Journal, 14,* 422-425.

Chen, Nan & Hong, L. (2008). Monte Carlo simulation in financial engineering. 919-931. doi: 10.1109/WSC.2007.4419688.

Comte, F. and Renault (1998) E. Long memory in continuous-time stochastic volatility modelsMathematical Finance, Vol. 8, No. 4, pp. 291–323.

Dekker, T. J. (1969). "Finding a zero by means of successive linear interpolation". in Dejon, B.; Henrici, P. (eds.), *Constructive Aspects of the Fundamental Theorem of Algebra*, London: Wiley-Interscience

Gatheral, J. (2006). *The volatility surface: a practitioner's guide*. Hoboken, N.J.: John Wiley & Sons.

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

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The review of financial studies*, 6(2), 327- 343.

Hull, J.C. (2021). *Options, futures, and other derivatives*. Harlow (Gran Bretaña): Pearson.

Hull, J. and White, A. (1987). The Pricing of Options on Assets with Stochastic Volatilities. *The Journal of Finance*, 42(2), p.281. doi:10.2307/2328253.

Hull, J. and White, A. (1988) An Analysis of the Bias in Option Pricing Caused by a Stochastic Volatility. Advances in Futures and Options Research, 3, 27-61.

Hull, J. and White, A. (1990). Pricing Interest-Rate-Derivative Securities. *Review of Financial Studies*, 3(4), pp.573–592. doi:10.1093/rfs/3.4.573.

Lee, R.W. (2001). Implied and Local Volatilities Under Stochastic Volatility. *International Journal of Theoretical and Applied Finance*, 04(01), pp.45–89. doi:10.1142/s0219024901000870.

Lindqvist, B.H., Erlemann, R. and Taraldsen, G. (2021). Conditional Monte Carlo revisited. *Scandinavian Journal of Statistics*. doi:10.1111/sjos.12549.

Medvedev, A. and Scaillet, O. (2006). Approximation and Calibration of Short-Term Implied Volatilities Under Jump-Diffusion Stochastic Volatility. *Review of Financial Studies*, 20(2), pp.427–459. doi:10.1093/rfs/hhl013.

Renault, E. and Touzi, N. (1996). Option Hedging and Implied Volatilities in a Stochastic Volatility Model. *Mathematical Finance*, 6(3), pp.279–302. doi:10.1111/j.1467-9965.1996.tb00117.x.

Romano, M. and Touzi, N. (1997). Contingent Claims and Market Completeness in a Stochastic Volatility Model. *Mathematical Finance*, 7(4), pp.399–412. doi:10.1111/1467-9965.00038.

OTC derivatives statistics at end-December 2021. (2022). *www.bis.org*. [online] Available at: https://www.bis.org/publ/otc_hy2205.htm#:~:text=The%20gross%20market%20value%20of%20derivatives%20contracts%20%E2%80%93%20a%20measure%20of [Accessed 15 Jun. 2022].

Willard, G.A. (1997). Calculating Prices and Sensitivities for Path-Independent Derivatives Securities in Multifactor Models. *The Journal of Derivatives*, 5(1), pp.45–61. doi:10.3905/jod.1997.407982.

Appendix A. Link for the Python Code Used

In order to the take a closer look to the Python code used to do the calibration of the SABR model, the simulation of prices assuming the SABR, the obtention of implied volatility surfaces through numerical methods and the figures and tables used, we make the code available through the GitHub link below:

https://github.com/Ivanmc8/SABR_Final_Thesis.git

Appendix B. Other Figures and Tables Used

Table B1. *Calibrated parameters for each time to maturity*

Maturity	Alpha	Sigma0	Rho	Skew
0.0986301	2.21944	0.105532	-0.531501	-1.17963
0.194521	1.71631	0.114461	-0.530875	-0.911148
0.271233	1.43265	0.123497	-0.518251	-0.742472
0.347945	1.28519	0.127518	-0.524285	-0.673803
0.443836	1.17833	0.129871	-0.525228	-0.618893
0.693151	1.07016	0.133144	-0.538324	-0.576093
0.942466	0.904525	0.13799	-0.543094	-0.491242
1.19178	0.730762	0.145253	-0.546149	-0.399105
1.69041	0.632034	0.151312	-0.546623	-0.345484
2.18904	0.566245	0.159065	-0.549131	-0.310943
2.68767	0.519048	0.157943	-0.54699	-0.283914
3.68493	0.453347	0.163994	-0.542808	-0.24608
4.68219	0.408856	0.170242	-0.536986	-0.21955
5.69863	0.3757	0.169145	-0.528923	-0.198716
6.69589	0.350371	0.173873	-0.521338	-0.182662
7.69315	0.329966	0.175668	-0.513059	-0.169292
8.69041	0.313044	0.176644	-0.504486	-0.157926
9.68767	0.298635	0.182493	-0.496445	-0.148256

Source: *Own elaboration with real market data provided by Caixa Bank*

Figure B1. *Sections of the implied volatility surfaces for different maturities (calibrated and from the market data)*

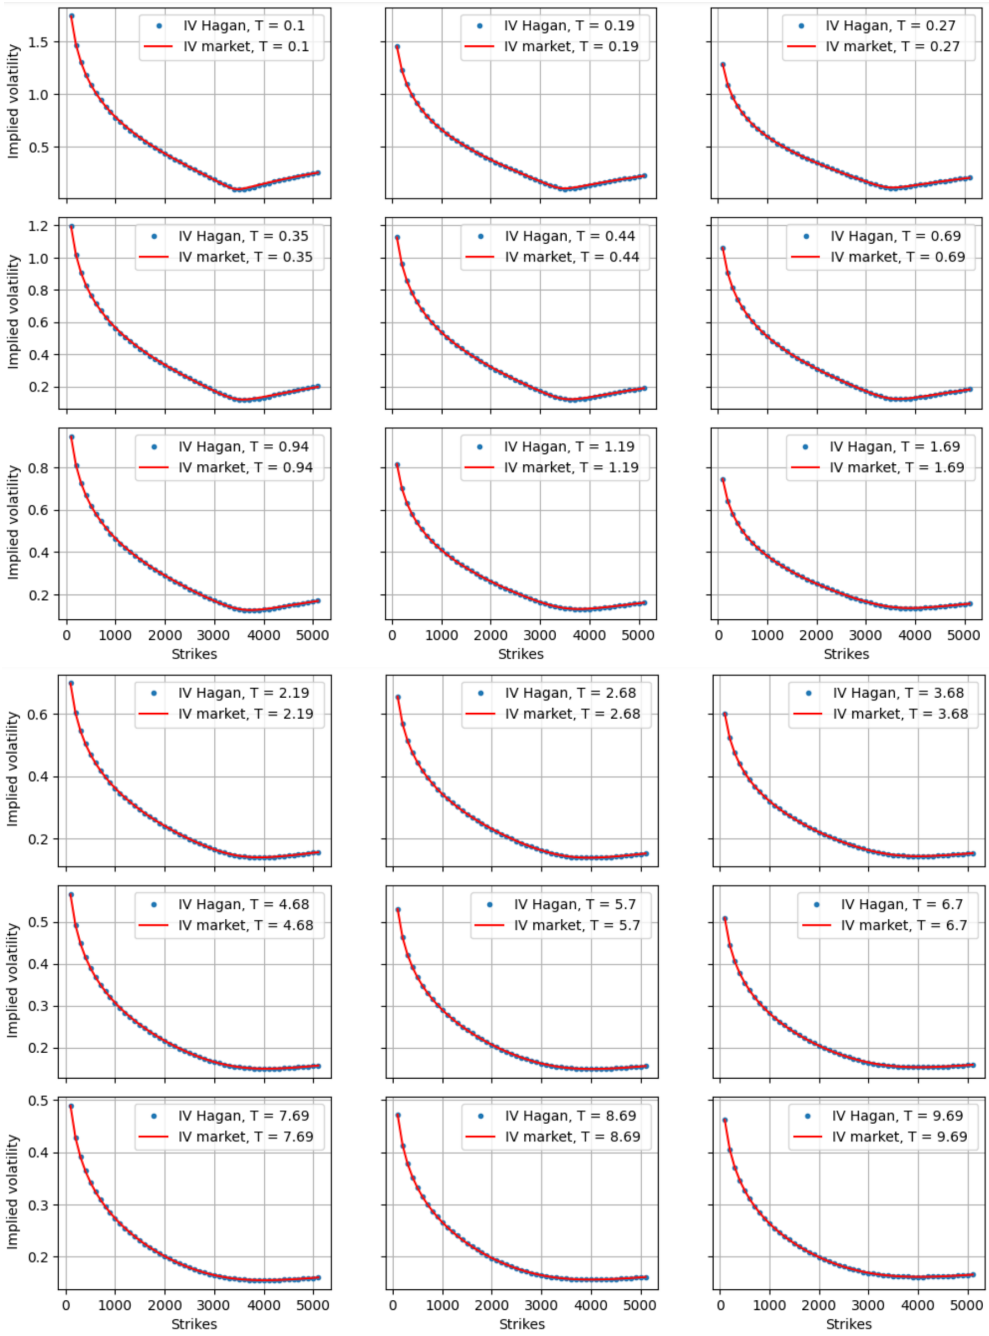

Source: *Own elaboration with real market data provided by Caixa Bank*

Figure B2 (left) & B3 (right). *Power law fitting for αρ with our calibrated parameters and (left) and Caixa Bank's calibrated parameters (right)*

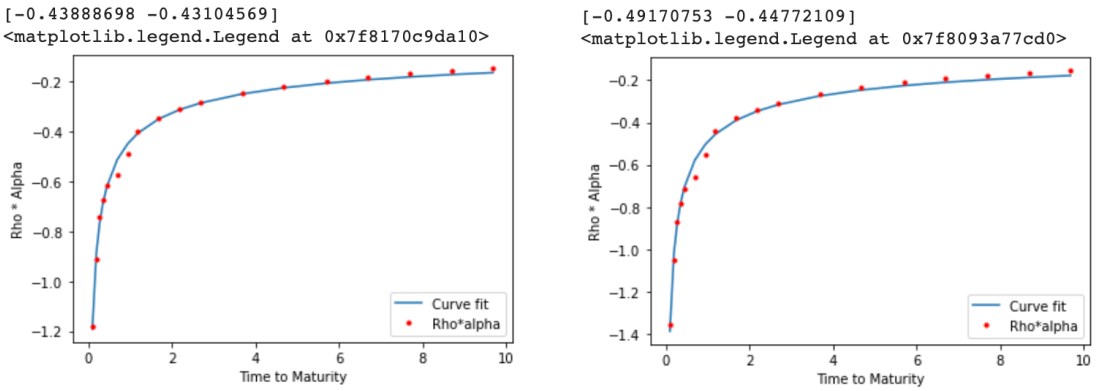

Source: *Own elaboration with real market data provided by Caixa Bank*

Figure B4 (above) & B5 (below). *Convergence plot for the standard Monte Carlo method (above) and the conditional Monte Carlo method (below)*

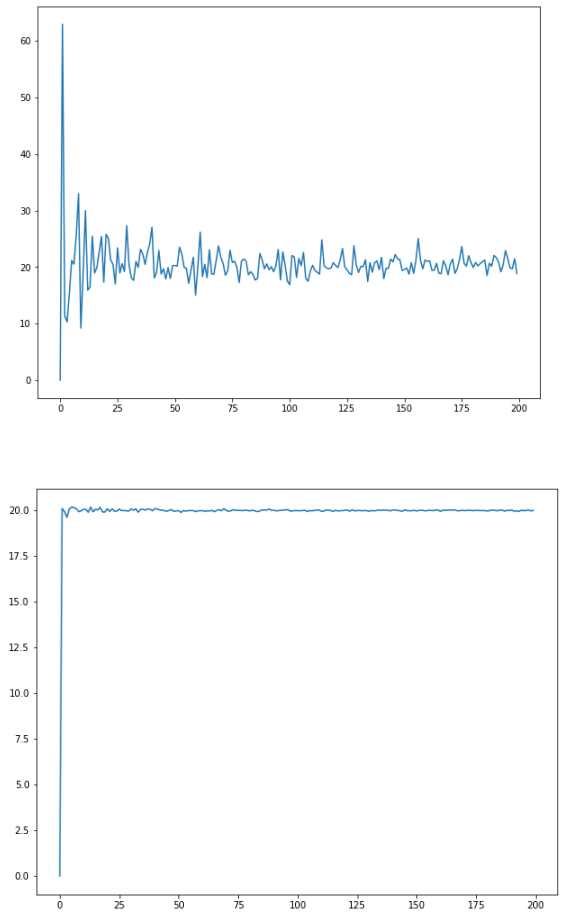

Source: *Own elaboration with real market data provided by Caixa Bank*

Appendix C. Derivation of the Extended Hull and White Formula

The formula presented in section 4.3 is a result obtained by Romano and Touzi (1997) and Willard (1997). We proceed to formally prove this result by considering a certain type of diffusion process for underlying asset prices, but it can be further generalized to other kind of processes. We assume basic knowledge of stochastic calculus, such as Itô formula.

Proof. Consider a diffusion process for the underlying asset price S_t of the form

$$dS_t = \sigma_t S_t dW_t \tag{C.1}$$

where the risk-free interest rate is $r = 0$, σ_t follows a diffusion process of the form

$$d\sigma_t = a(t, \sigma_t)dt + b(t, \sigma_t)dZ_t \tag{C.2}$$

and W_t is a Brownian motion which follows a diffusion process of the form

$$dW_t = \rho dB_t + \sqrt{1-\rho^2}dZ_t \tag{C.3}$$

where B_t and Z_t are independent Brownian motions. Applying Itô formula to the log-price X_t, which follows a diffusion of the form

$$dX_t = -\frac{\sigma_t^2}{2}dt + \sigma_t dW_t \tag{C.4}$$

we obtain the following result

$$X_T = X_t + \int_t^T \frac{1}{S_u}\sigma_u S_u dW_u - \frac{1}{2}\int_t^T \frac{1}{S_u^2}\sigma_u^2 S_u^2\, du$$

$$\Rightarrow\ X_T = X_t + \int_t^T \sigma_u dW_u - \frac{1}{2}\int_t^T \sigma_u^2\, du$$

$$\Rightarrow\ X_T = X_t + \rho\int_t^T \sigma_u dB_u + \sqrt{1-\rho^2}\int_t^T \sigma_u dZ_u - \frac{1}{2}\int_t^T \sigma_u^2\, du$$

$$\Rightarrow\ S_T = S_t \exp\left(\rho\int_t^T \sigma_u dB_u + \sqrt{1-\rho^2}\int_t^T \sigma_u dZ_u - \frac{1}{2}\int_t^T \sigma_u^2\, du\right)$$

$$\Rightarrow S_T = S_t \exp\left(\rho\int_t^T \sigma_u dB_u + \sqrt{1-\rho^2}\int_t^T \sigma_u dZ_u - \frac{\rho^2}{2}\int_t^T \sigma_u^2\, du - \frac{(1-\rho^2)}{2}\int_t^T \sigma_u^2\, du\right) \tag{C.5}$$

By defining

$$S_t' \equiv S_t \exp\left(\rho \int_t^T \sigma_u dB_u - \frac{\rho^2}{2}\int_t^T \sigma_u^2\, du\right) \tag{C.6}$$

$$v_t^2 \equiv \frac{1}{T-t}\int_t^T \sigma_u^2\, du \tag{C.7}$$

then the result obtained in (C.5) can be expressed as

$$S_T = S_t' \exp\left(\sqrt{1-\rho^2}\int_t^T \sigma_u dZ_u - \frac{(1-\rho^2)}{2} v_t^2 (T-t)\right) \tag{C.8}$$

As we can observe, S_t' is distributed log-normally with mean $-\frac{v_t^2}{2}(T-t) + \rho\int_t^T \sigma_u dB_u$ and variance $(1-\rho^2)v_t^2$. Now, based on conditional expectations properties, we can see that the price of a European option with maturity T and strike K can be expressed as

$$V_t = e^{-r(T-t)}E^*[h(X_t)_+|\,\mathcal{F}_t] = e^{-r(T-t)}E^*[E^*(h(X_t)_+|\mathcal{F}_t \vee \mathcal{F}^Z)|\,\mathcal{F}_t] \tag{C.9}$$

where E^* is the expectation using risk-neutral probability, $h(X_t)_+$ is the payoff function, $\mathcal{F}_t \equiv \mathcal{F}^B \vee \mathcal{F}^Z$ and $\mathcal{F}^B$ and $\mathcal{F}^Z$ denote the filtrations of the Brownian motions B_t and Z_t. Because the Black-Scholes price of an option is

$$V_t^{BS} = e^{-r(T-t)}E^*(h(X_t)_+|\mathcal{F}_t \vee \mathcal{F}^Z) = BS(X_t, K, T-t, \sigma_{BS}) \tag{C.10}$$

where BS is the pricing equation of the corresponding European option and σ_{BS} is the volatility parameter, then we can merge (C.9) and (C.10) to obtain

$$V_t = E^*[BS(X_t, K, T-t, \sigma_{BS})|\,\mathcal{F}_t] \tag{C.10}$$

Because the Black-Scholes model assumes a log-normal distribution of asset prices, and S_t' follows the same kind of distribution, then

$$V_t = E^*\left[BS\left(\ln S_t', K, T-t, \sqrt{1-\rho^2}v_t\right)|\,\mathcal{F}_t\right] \tag{C.11}$$

which is equivalent to the extended Hull and White formula

$$V_t = E^*\left[BS\left(X_t\xi_t, K, T-t, \sqrt{1-\rho^2}v_t\right)|\,\mathcal{F}_t\right] \tag{C.12}$$

$$when\ \ \xi_t = \rho\int_t^T \sigma_u dB_u - \frac{\rho^2}{2}\int_t^T \sigma_u^2\, du$$

Thus, we have proved the extended Hull and White formula. ■