Calibration of the Volatility Surface

Erik Nilsson
erinil@kth.se
840428-0292

June 12, 2008

Abstract

This thesis consists of two parts, one concerning implied volatility and one concerning local volatility. The SABR model and SVI model are investigated to model implied volatility. The performance of the two models were tested on the Eurcap market in March 2008. Two ways of extracting local volatility are reviewed by a test performed on data from European options based on the S&P 500 index. The first method is a way of solving regularized Dupire's equation and the other one is based on finding the "most likely path".

Acknowledgements

I would like to thank Robert Thorén and his coworkers at Algorithmica Research AB for guiding me during my work. I am also grateful to Professor Anders Forsgren at KTH, who has been my tutor.

Contents

1 Introduction

1.1 Purpose and background

The purpose of this thesis is to give a theoretical background of which methods that are used to model volatility. The thesis can be divided into two sections, one about modeling implied volatility, and one modeling local volatility. The first goal is to find an implied volatility method which is robust, stable and fast on the option interest rate market. The second goal is to investigate whether there is a method which can recover a plausible local volatility surface from a market implied volatility surface. For the first section, Quantlab has been the tool for implementation. It is a powerful numerical computing environment, developed by *Algorithmica Research AB*, specialized for finance. The second part is implemented in Matlab. The reader is expected to have basic knowledge in optimization and stochastic calculus.

1.2 Volatility

Volatility is the standard deviation of the time series for a stock, interest rate etc. Volatility has traditionally been used in finance as a measure of risk. There are three different types of volatility which will be treated in this thesis: Implied volatility, stochastic volatility and local volatility. It is important to understand the differences between these. The implied volatility is the volatility used in Black-Scholes formula to generate a given option price. Stochastic volatility is an extension to the Black-Scholes model where the volatility itself is a stochastic process. There are a lot of different stochastic volatility models which will be covered in a later section. Local volatility is the volatility used in the general Black-Scholes model and it is a deterministic function of expiration time and price of the underlying.

1.3 Introduction to arbitrage-free option pricing

An option is a financial contract dependent on an underlying. The underlying is usually a stock, an interest rate or an equity. An easy example of an option is the European call option which gives the buyer the right but not the obligation to buy a stock for the strike price K at maturity T. A European put option is the same, but this time the buyer has the right but not the obligation to sell the underlying. Mathematically this means that at the maturity T, the value is $\max\{0, S_T - K\}$ for a European call option. The contracts can be a lot more complex, for example exotic options where the payoff at maturity not only depends on the value of the underlying at maturity but its value several times during the contract's life or it could depend on more than one underlying.

A product which will be treated in this thesis is a so called cap. An interest rate cap is a derivative in which the buyer receives payments at the end of each period in which the interest rate exceeds the agreed strike rate. The payment is compensating the excess rate, hence the buyer never has to pay more interest rate than the strike rate. An example of a cap would be an agreement to receive a payment for each 6 months the LIBOR rate exceeds 2.5 %. A cap can be treated as a sum of caplets. A caplet is a european option with an interest rate as an underlying and having the expiration time equal to the period of payments in the cap contract, 6 months in the case mentioned above. A one year cap with a 6 month payment period is the sum of one 6 month caplet and another 6 month caplet starting in 6 months. Later on, methods for extracting the implied volatility surface on the cap market will be tested.

The difficulty with options is to determine a fair price. The more complex a contract is, the more difficult it is to actually know the "right" price for it, especially for ill-liquid markets, where just a few contracts are settled. Therefore a mathematical theory was needed to find some kind of consensus in option pricing. This theory is called Arbitrage-Free Option pricing and was among others developed by Fischer Black and Myron Scholes, see [2]. They believed that a price is determined relative to other prices quoted in the market in such a manner as to preclude any arbitrage opportunities. Arbitrage opportunity means that there is a financial strategy which guarantees no loss and has positive probability of a profit. There is a famous expression often used in financial literature, "There is no such thing as a free lunch", which gives a better understanding for the meaning of arbitrage-free.

Black and Scholes postulated that the price S of an underlying behaves like a diffusion process, i.e.

$$\frac{dS}{S} = rdt + \sigma dW_t \tag{1.1}$$

where r is the risk-free rate, σ is the volatility and W is a Brownian motion. To derive the remainder of the Black-Scholes formula, the readers have to be familiar with martingale theory, risk neutral pricing and more in the theory of stochastic calculus. This thesis cannot cover all of this, so therefore a huge step is made to obtain the following option price formula

$$C_{BS} = e^{-rT} E^Q[\max\{0, S - K\}|\Omega], \tag{1.2}$$

i.e. the option price is given by the discounted conditional expectation of the payoff function $\max\{0, S - K\}$ under the risk neutral measure Q given the information today Ω. This expression has a closed form formula called the Black-Scholes formula. For interested readers, we give the explicit form

$$C_{BS}(S, K, r, T, \sigma_{BS}) = S\phi(d_1) - Ke^{-rT}\phi(d_2), \tag{1.3}$$

where

$$d_1 = \frac{ln(S/K) + (r + \sigma_{BS}^2/2)T}{\sigma_{BS}\sqrt{(T)}}, \tag{1.4}$$

$$d_2 = \frac{ln(S/K) + (r + \sigma_{BS}^2/2)T}{\sigma_{BS}\sqrt{(T)}}, \tag{1.5}$$

and ϕ is the standard normal cumulative distribution function, see [2]. Henceforth just $C_{BS}(S, K, r, T, \sigma_{BS})$ will be used as the option price given by the Black-Scholes formula.

The Black-Scholes formula worked well before the huge crashes 1987 and 1989, but since then, it has been observed that implied volatility has a skewness and smile structure, which contradicts an assumption of constant volatility. This is due to the fact, that the model assumes log-normally distributed returns for stocks, which not always agrees with observed returns. These observations usually give a skewed distribution with fatter tails. A lot of research has been done to handle the weaknesses of the Black-Scholes model. Since Black-Scholes is widely used across the world by banks, it has been more popular to just model implied volatility and then keeping the Black-Scholes analytical formula as a tool to quote option prices from implied volatilities. Riccardo Rebonato's view of implied volatility is "the wrong number put in the wrong formula to obtain the right price", see [12].

2 Stochastic volatility

2.1 Theory

A stochastic volatility model is a model where the volatility itself is a stochastic process. This is an extension to the dynamics of the Black and Scholes model. One popular model is the Heston model, where the price of the underlying is a geometric brownian motion and the volatility is a geometric brownian motion with mean reversion. Mean reversion means that the process strives to a long term mean value. The dynamic model from [8] has the following mathematical representation

$$dS_t = \mu S_t dt + \sqrt{v_t} S_t dW_1, \tag{2.1}$$
$$dv_t = \kappa(\bar{v} - v_t)dt + \eta\sqrt{v_t} dW_2, \tag{2.2}$$
$$E(dW_1 dW_2) = \rho dt, \tag{2.3}$$

where S_t is the price of the underlying, μ is the constant drift, v_t is the variance of the underlying, κ speed of reversion, $\bar{v}$ long term mean, η volatility of volatility and ρ the correlation between the two brownian motions W_1 and W_2.

Another popular model is the Bates model which is an extension to the Heston model. The difference lies within the price process where a Poisson process is added. This serves as a better explanation for discontinuous jumps in the market. The mathematical representation from [1] is

$$dS_t = (r - \lambda\mu_J)S_t dt + \sqrt{v_t} S_t dW_1 + J_t S_t dX_t, \tag{2.4}$$
$$dv_t = \kappa(\bar{v} - v_t)dt + \eta\sqrt{v_t} dW_2, \tag{2.5}$$
$$E(dW_1 dW_2) = \rho dt, \tag{2.6}$$
$$X_t \sim Po(\lambda t), \tag{2.7}$$

where r is the risk-free rate and the random variable J_t determines the jump size and follows the distribution

$$\log(1 + J_t) \sim N\Big(\log(1 + \mu_J) - \frac{\sigma_J^2}{2}, \sigma_J^2\Big). \tag{2.8}$$

Three new parameters λ, μ_J and σ_J are added to this model. A further extended model is the SVJJ, which has simultaneous jumps in the underlying and the volatility. However the more complex these models are, the more difficult they are to calibrate. The Bates model has nine parameters and SVJJ even more. Despite the fact that Stochastic volatility models are difficult to calibrate in general, there is one stochastic model that differs, the SABR model given by

$$dF_t = \sigma_t F_t^{\beta} dW_1, \quad F_t(0) = f, \tag{2.9}$$
$$d\alpha_t = \nu\alpha_t dW_2, \quad \alpha_t(0) = \alpha, \tag{2.10}$$
$$E(dW_1 dW_2) = \rho dt, \tag{2.11}$$

where F_t is the forward value, α is the volatility, ν is the volatility of volatility and W_1 and W_2 are Brownian motions, see [6]. The forward value and volatility are under the forward measure and the two processes are correlated with ρ. The forward value is $F_t = Se^{rt}$ and r is the rate. This model was created by Paul Hagan et al. The dynamics of this model is similar to the ones shown above but do not have the volatility mean reversion property and is therefore only good for short expirations theoretically. However as the expiration time $\tau \to 0$, an exact expression for the implied volatility can be obtained by using singular perturbation techniques. Due to the analytical formula, the model is easy to calibrate.

Another easy calibrated model is the SVI (Stochastic Volatility Inspired) model found by Jim Gatheral [5]. It is a clever parameterization of the implied volatility surface. From empirical observations, the implied variance is always linear in the wings and curved in the middle when plotted against the logarithmic moneyness. Therefore, the author suggested the following parameterization

$$\text{var}(k; a, b, \sigma, \rho, m) = a + b\Big\{\rho(k-m) + \sqrt{(k-m)^2 + \sigma^2}\Big\} \tag{2.12}$$

where

- a gives the overall level of variance,
- b gives the angle between the left and right asymptotes,
- σ determines how smooth the vertex is,
- ρ determines the orientation of the graph,
- m translates the graph.

The variance has the left and right asymptotes

$$\text{var}_L(k; a, b, \sigma, \rho, m) = a - b(1-\rho)(k-m) \quad k \to -\infty, \tag{2.13}$$

$$\text{var}_R(k; a, b, \sigma, \rho, m) = a + b(1+\rho)(k-m) \quad k \to \infty, \tag{2.14}$$

which agrees with the assumption of linear wings.

2.2 Calibration

In this section, the calibration of the SABR and SVI model is described. Calibration of the other stochastic volatility models are beyond the scope of this thesis.

The SABR model has four parameters α, β, ρ, ν and has the analytical closed form formula

$$\sigma_{BS}(K,f) = \frac{\alpha}{(fK)^{(1-\beta)/2}\Big\{1 + \frac{(1-\beta)^2}{24}\log^2\frac{f}{K} + \frac{(1-\beta)^4}{1920}\log^4\frac{f}{K} + O(\log^6\frac{f}{K})\Big\}}\Big(\frac{z}{x(z)}\Big)\cdot$$
$$\cdot\Big\{1 + \Big[\frac{(1-\beta)^2}{24}\frac{\alpha^2}{(fK)^{1-\beta}} + \frac{1}{4}\frac{\rho\beta\nu\alpha}{(fK)^{(1-\beta)/2}} + \frac{2-3\rho^2}{24}\nu^2\Big]t_{ex} + O(t_{ex}^2)\Big\}. \tag{2.15}$$

Notice that this formula is expressed in strike K and the forward value, i.e. $f = Se^{rt_{ex}}$, where S is the value today, r is the rate and t_{ex} is the expiration time. Only two parameters ρ and ν have to be calibrated. The scalar β is determined either from a loglog-plot with historical data, or an assumption. Since β controls the distribution function, an a priori view of the distribution function could be a way to set the β. $\beta = 0$ gives raise to a normally distributed change of the underlying, and $\beta = 1$ a lognormally distributed change of the underlying, $\beta = \frac{1}{2}$ can also be used. This decision is made on the basis of market experience. A special case of the formula above is used when the strike and forward value are equal,

$$\sigma_{ATM} =$$
$$\sigma_{BS}(f,f) = \frac{\alpha}{f^{(1-\beta)}}\left\{1 + \left[\frac{(1-\beta)^2}{24}\frac{\alpha^2}{f^{2-2\beta}} + \frac{1}{4}\frac{\rho\beta\alpha\nu}{f^{1-\beta}} + \frac{2-3\rho^2}{24}\nu^2\right]t_{ex} + O(t_{ex}^2)\right\}, \quad (2.16)$$

from this equation, α is calculated whenever needed on the fly by inserting values of ρ and ν. It is a third order polynomial equation, which has an algebraic solution. The idea is to choose α so that σ_{ATM} is fixed at the level given from the market. When σ_{ATM} is not given, a linear interpolation from the nearest neighbors is used. Another method is to relax α completely and then receive a least square fit to the SABR model.

The calibration procedure leads to a non-convex and non-linear optimization problem. The method used to solve the problem is Levenberg-Marquardt (LMA). It is a useful algorithm for non-linear least square fitting problems. The LMA interpolates between the Gauss-Newton algorithm and the gradient descent method, see, e.g., [10], [11]. The LMA is more robust than the Gauss-Newton, which means that in many cases it finds a solution even if it starts very far off the final minimum. For well behaved functions and with good initial guesses, LMA tends to be a bit slower than Gauss-Newton.

$$\min_{\rho,\nu} \sum_{(K,T)\in\Omega} ||\sigma_{SABR}(K,T) - \sigma^*(K,T)||_2 \quad (2.17)$$

where $||\cdot||_2$ is the L_2 norm. The L_2 norm of a vector $\mathbf{a}$, with elements a_i is $\sqrt{\sum_i a_i^2}$. The SVI model has five parameters and all of them have to be calibrated. Same optimization procedure as above is used to calibrate the SVI model.

$$\min_{a,b,\sigma,\rho,m} \sum_{(K,T)\in\Omega} ||\sigma_{SVI}(K,T) - \sigma^*(K,T)||_2 \quad (2.18)$$

There are many local minima of this function, therefore it is important to try different starting points.

When extracting implied volatility surfaces from the CAP-market, a certain preprocess is needed. It is called Caplet-stripping. As explained in section 1.3, a Cap is a sum of caplets and it is the caplet volatilities that

are needed to generate the correct surface. The caplet-stripper uses the Cap-values to implicitly calculate the corresponding caplet volatilities. For example, if the values of the one year cap $\text{CAP}_{0\to1}$ and the two years cap, $\text{CAP}_{0\to2}$ are given, then the 6 month caplets in-between can be extracted in the following way. $\text{CAPLET}_{0\to\frac{1}{2}} + \text{CAPLET}_{\frac{1}{2}\to1} = \text{CAP}_{0\to1}$. The sum of the two remaining caplets is $\text{CAP}_{0\to2}$ - $\text{CAP}_{0\to1}$. Consider this like two sections, i.e. 0 to 1 and 1 to 2. Within each section there are two 6 month caplets. Create a function dependent on the last implied volatility in the section. This function draws a straight line through the last and first volatility on that section and calculates the sum of the caplet values which are evaluated with BS-formula. Newton's method is used to find the exact volatility which will be consistent with the CAP-value. For the first section, a flat line is calculated, since there is no start-value in that section. For the second section a linear interpolation is made according to the procedure. In this way, the linearly interpolated caplet volatilities are extracted. The LMA algorithm and the caplet stripping procedure were already implemented in Quantlab, therefore it was favorable to implement SABR and SVI in Quantlab.

3 Local volatility

3.1 Theory

Local volatility is needed when pricing exotic options, when an option depends on an underlying several times during the life time of the contract. There is no closed form formula for these contracts, therefore Monte Carlo simulation has to be used, and the preferable volatility is the local volatility.

One of the developers of local volatility theory was Bruno Dupire, who extended the volatility to be a state-dependent function of the price of the underlying and the time to expiration. The new dynamics becomes

$$\frac{dS}{S} = rdt + \sigma_L(S,t)dW_t, \tag{3.1}$$

whis is also called the general Black-Scholes dynamics. The main idea of finding the local volatility is to derive the risk-neutral density from market prices of European options. The great breakthrough was when Dupire showed that under risk-neutrality there was a unique diffusion process consistent with these risk-neutral distributions. The corresponding unique state-dependent diffusion coefficient $\sigma_L(S,t)$, consistent with the given European option prices, is the local volatility function. The unique local volatility function is the solution to Dupire's equation

$$\frac{\partial C}{\delta T} = \frac{\sigma_L^2 K^2}{2}\frac{\delta^2 C}{\delta K^2} + (r_t - D_t)\left(C - K\frac{\delta C}{\delta K}\right) \tag{3.2}$$

where D_t is the dividend yield and C is the European option price $C(S_0, K, T)$.

Proof from [4].
Suppose the stock price diffuses with a risk-neutral drift $\mu_t = r_t - D_t$ and local volatility $\sigma_L(S,t)$ according to the equation

$$\frac{dS}{S} = \mu_t dt + \sigma_L(S_t, t)dW. \tag{3.3}$$

The undiscounted risk-neutral value $C(S_0, K, T)$ of a European option with strike K and expiration T is given by

$$C(S_0, K, T) = \int_K^\infty dS_T \varphi(S_T, T; S_0)(S_T - K), \tag{3.4}$$

where $\varphi(S_T, T; S_0)$ is the probability density of the final spot at time T. It evolves according to the Fokker-Planck equation

$$\frac{1}{2}\frac{\partial^2}{\partial S_T^2}\left(\sigma_L^2 S_T^2 \varphi\right) - \frac{\varphi}{\partial S_T}(\mu S_T \varphi) = \frac{\partial \varphi}{\partial T}. \tag{3.5}$$

Differentiating (3.4) with respect to K gives

$$\frac{\partial C}{\partial K} = -\int_K^\infty dS_T \varphi(S_T, T; S_0), \tag{3.6}$$

$$\frac{\partial^2 C}{\partial K^2} = \varphi(K, T; S_0). \tag{3.7}$$

Now, differentiating (3.4) with respect to time T gives

$$\frac{\partial C}{\delta T} = \int_K^\infty dS_T \left\{ \frac{\partial}{\partial T} \varphi(S_T, T; S_0) \right\} (S_T - K). \tag{3.8}$$

By using (3.5) in (3.8)

$$\frac{\partial C}{\delta T} = \int_K^\infty dS_T \left\{ \frac{1}{2} \frac{\partial^2}{\partial S_T^2} \left(\sigma_L^2 S_T^2 \varphi \right) - \frac{\partial}{\partial S_T} (\mu S_T \varphi) \right\} (S_T - K). \tag{3.9}$$

Integrating the first term of (3.9) by parts, $\int_a^b fg'\ dx = [fg]_a^b - \int_a^b f'g\ dx$, gives

$$\int_K^\infty dS_T \left\{ \frac{1}{2} \frac{\partial^2}{\partial S_T^2} (\sigma_L^2 S_T^2 \varphi) \right\} (S_T - K) = \left[\frac{1}{2} \frac{\partial}{\partial S_T} (\sigma_L^2 S_T \varphi)(S_T - K) \right]_K^\infty - \int_K^\infty dS_T \frac{1}{2} \frac{\partial}{\partial S_T} (\sigma_L^2 S_T^2 \varphi). \tag{3.10}$$

Using that $\lim_{K \to \infty} S_T = 0$ the term in brackets vanishes and the second term becomes $\frac{\sigma_L^2 K^2}{2} \varphi$. Integrating the second term of (3.9) by parts yields

$$-\int_K^\infty dS_T \left\{ \frac{\partial}{\partial S_T} (\mu S_T \varphi) \right\} (S_T - K) = -\left[\mu S_T \varphi (S_T - K) \right]_K^\infty + \int_K^\infty dS_T (\mu S_T \varphi). \tag{3.11}$$

In the same way as before, the term in brackets vanishes. This leads to the following equation

$$\frac{\partial C}{\partial T} = \frac{\sigma_L^2 K^2}{2} \varphi + \int_K^\infty dS_T \mu S_T \varphi. \tag{3.12}$$

The second term of (3.12) may be written as

$$\int_K^\infty dS_T \mu S_T \varphi = \mu \left[\int_K^\infty dS_T \varphi (S_T - K) + K \int_K^\infty dS_T \varphi \right]. \tag{3.13}$$

Notice that the first term of the right hand side above is exactly the undiscounted option value from (3.4). By using (3.6) to the second term of (3.13) and (3.7) to the first term of (3.12) finally gives

$$\frac{\partial C}{\partial T} = \frac{\sigma_L^2 K^2}{2} \frac{\partial^2 C}{\partial K^2} + \mu(T) \left(C - K \frac{\partial C}{\partial K} \right). \tag{3.14}$$

□

3.2 Calibration

3.2.1 Local volatility by solving Dupire's equation

Though the theory ensures a unique local volatility it is a non-trivial problem to recover it from real option data. This is due to the fact, that the theory assumes a well defined European option price space, which is not the case on real markets. As a matter of fact there are only a few dozens of option prices available, which of course makes the problem severely underdetermined. This results in an ill-posed optimization problem

$$\min_{\sigma_L} \sum_{(S,t)\in\Omega_\mathrm{m}} ||C_\mathrm{m}(S,t) - C(S,t)||_2^2 \tag{3.15}$$

subject to

$$\frac{\partial C}{\partial T} = \frac{\sigma_L^2 K^2}{2}\frac{\partial^2 C}{\partial K^2} + \mu(T)\left(C - K\frac{\partial C}{\partial K}\right) \tag{3.16}$$

$$C(S_0, K, 0) = \max(K - S_0, 0), \tag{3.17}$$

where C_m is the given market prices, Ω_m is the set of pairs (S,t) for which market prices are given. This is a famous inverse problem in computational finance. It is known for being very sensitive to noisy input data, i.e. the solution changes dramatically to small changes in the data. Ill-posed problems like this one need to be reformulated for numerical treatment. This is done by introducing some additional information of the solution, such as an assumption of the smoothness or a bound on the norm. This process is called regularization. For the inverse problem above, Tikhonov regularization and entropy regularization are used by different authors. Tikhonov regularization penalizes either the first or the second derivatives of the surface, for example the Frobenius norm of the hessian. This reduces the noise and generates the smoothest solution to the problem above. One shortcoming of this procedure is that one penalty factor has to be determined a priori. The entropy regularization generates the solution, whose corresponding probability distribution has the shortest entropy distance to an a priori probability distribution. The a priori guess is of course crucial for the calibration and needs to be chosen carefully on the basis of market information and gut feeling. There is no consensus of which solution is the best, the smoothest, the one closest to some a priori guess or some other. That is why calibration is considered as an art by some authors [9].

One way to solve the Dupire equation is a method described in [7]. It uses natural cubic splines and Tikhonov regularization. In this case the regularization is based on penalizing the discretized gradient. The idea is to interpolate the given option prices by a natural cubic spline per time slice. The natural cubic spline is defined as the minimizer $\hat{g}$ of the optimization

problem below

$$\min_g \sum_{i=1}^{n} ||y_i^* - g(u_i)||_2^2 + \lambda \int_{u_1}^{u_n} g''(v)dv, \tag{3.18}$$

where y_i^* are the given option prices, u_i the strike values and λ is a penalty parameter. The optimization problem (3.18) may be written as a quadratic program

$$\min_x -y^T x + \frac{1}{2} x^T \mathbf{B} x \tag{3.19}$$

$$\text{subject to} \qquad \mathbf{A}^T x = 0 \tag{3.20}$$

where

$$x = \begin{pmatrix} g \\ \xi \end{pmatrix}, \quad y = \begin{pmatrix} y^* \\ \mathbf{0} \end{pmatrix}, \tag{3.21}$$

$$\mathbf{A} = \begin{pmatrix} \mathbf{Q} \\ -\mathbf{R}^T \end{pmatrix}, \tag{3.22}$$

$$\mathbf{B} = \begin{pmatrix} \mathbf{I}_n & 0 \\ 0 & \lambda \mathbf{R} \end{pmatrix}, \tag{3.23}$$

$\mathbf{Q}$ is a $n \times (n-2)$-matrix and $h_i = u_{i+1} - u_i$, then $\mathbf{Q}$ is defined in the following way

$$q_{j-1,j} = h_{j-1}^{-1}, \quad q_{i,j} = -h_{j-1}^{-1} - h_j^{-1}, \quad q_{j,j+1} = h_j^{-1} \quad \forall\ j = 2, \ldots, n-1$$

$$q_{i,j} = 0 \quad \forall\ (i,j) \in \{(i,j) |\ |i-j| \geq 2\}$$

$\mathbf{R}$ is $(n-2) \times (n-2)$ defined by its elements $r_{i,j}\ \forall\ i,j = 2, \ldots, n-1$

$$r_{i,j} = \frac{1}{3}(h_{i-1} + h_i) \quad \forall\ i = 2, \ldots, n-1$$

$$r_{i,i+1} = r_{i+1,i} = \frac{1}{6} h_i \quad \forall\ i = 2, \ldots, n-2$$

$$r_{i,j} = 0 \quad \forall\ (i,j) \in \{(i,j) |\ |i-j| \geq 2\}$$

the (n-2)-vector ξ is the corresponding second derivatives for the interior nodes of the natural cubic spline. Note that the second derivatives at the boundary nodes are zero by definition. The matrix $\mathbf{I}_n$ is the $n \times n$ identity matrix.

It is convenient to use natural cubic splines, since they are known to be smooth, continuous and twice differentiable everywhere. Therefore C_K,

C_{KK} can be calculated by just differentiating $\hat{g}$. C_T is obtained using a derivative approximation.

$$C_T(K_i, T_j) = \frac{1}{\tau_j + \tau_{j+1}}\Big(\frac{\tau_j}{\tau_{j+1}} g_{j+1}(K_i) + (\frac{\tau_{j+1}}{\tau_j} - \frac{\tau_j}{\tau_{j+1}}) g_j(K_i) - \frac{\tau_{j+1}}{\tau_j} g_{j-1}(K_i)\Big) \quad (3.24)$$

with $\tau_j = T_j - T_{j-1},\ j = 2, \ldots, m-1$. Euler forward and Euler backwards are used at the end-points. For discounted option prices, and with zero dividend, which will be the case in the test later on, the Dupire equation can be written like

$$\sigma_L^2 = \left(\frac{2(C_T + rKC_K)}{K^2 C_{KK}}\right) \quad (3.25)$$

Having all these quantities, there is only one unknown, the local variance σ_L^2. Therefore the local volatility surface can be obtained by solving a linear equation system, hence this method is a lot faster than other methods. Unfortunately there is no guarantee that the solution is positive, which could yield complex volatilities. Therefore the problem needs to be regularized. The first key approach was to solve

$$\mathbf{D}z \approx \mathbf{b} \quad (3.26)$$

where $\mathbf{D}$ is a diagonal matrix with the denominator of (3.25) as its elements, z is σ_L^2 and$\mathbf{b}$ is the corresponding vector of the enumerator of (3.25). By applying the regularization, the problem becomes

$$\min_z ||\mathbf{D}z - \mathbf{b}||_2^2 + \gamma ||\mathbf{L}z||_2^2 \quad (3.27)$$

where $\gamma > 0$ is a regularization parameter and $\mathbf{L}$ is a discrete gradient operator. The optimization problem above can be written as a linear equation system

$$(\mathbf{D}^2 + \gamma \mathbf{L}^T \mathbf{L})z = \mathbf{D}\mathbf{b} \quad (3.28)$$

$\mathbf{L}$ is defined as

$$\mathbf{L} = \begin{bmatrix} \mathbf{I}_{m+1} \otimes \mathbf{L}_K \\ \mathbf{L}_T \otimes \mathbf{I}_n \end{bmatrix} \quad (3.29)$$

and $\mathbf{L}_K \in R^{n-1,n}$, $\mathbf{L}_T \in R^{m,m+1}$

$$\mathbf{L}_K = \begin{bmatrix} -1 & 1 & & & \\ & -1 & 1 & & \\ & & \ddots & \ddots & \\ & & & -1 & 1 \end{bmatrix}, \mathbf{L}_T = \begin{bmatrix} -\frac{1}{\tau_1} & \frac{1}{\tau_1} & & & \\ & -\frac{1}{\tau_2} & \frac{1}{\tau_2} & & \\ & & \ddots & \ddots & \\ & & & -\frac{1}{\tau_m} & \frac{1}{\tau_m} \end{bmatrix}$$

by increasing γ sufficiently, a positive solution is guaranteed. A priori guesses has to be used when choosing the parameters λ for the splines and γ for the regularization.

3.2.2 Local volatility by most likely path calibration

Due to the numerical difficulties of solving Dupire's equation, there is a more robust way of finding the local volatility, according to [9]. It is based on a concept called most likely path first found by Jim Gatheral [4]. He shows that the implied variance is well approximated by the time integral from $t = 0$ to $t = T$ of the local variance along the most likely price path $E[S_t|S_T = K]$. This calculation is complicated when the normal dynamics $\frac{dS_t}{S_t} = \mu_t dt + \sigma(t, S_t)dW_t$ is used. Therefore a simpler version is introduced in [9].

$$\frac{dS_t}{S_t} = \mu_t dt + \sigma(E[S_t|S_T = K])dW_t = \mu_t dt + \sigma_*(t)dW_t \tag{3.30}$$

There is a closed form formula to calculate the conditional expected value.

$$E[S_t|S_T = K] = F_t F_T^{-\alpha_{t,T}} K^{\alpha_{t,T}} e^{\frac{1}{2}\alpha_{t,T}\int_t^T \sigma^2(s)ds} \tag{3.31}$$

where $F_t = S_0 e^{\int_0^t \mu_s ds}$ is the forward value and

$$\alpha_{t,T} = \frac{\int_0^t \sigma^2(s)ds}{\int_0^T \sigma^2(s)ds} \tag{3.32}$$

is a regression coefficient between 0 and 1. The calibration is an iterative process, where each iteration consists of two steps:

1) Find the most likely path,
2) Calculate the new local volatility.

The procedure to find the most likely path is itself an iterative process, where the initial guess is set to the forward price path. This guess is used in the equation above to obtain the next guess of the most likely path. The procedure is repeated until it converges, then the final path is the most likely path. When having the most likely path, the $i+1^{th}$ implied volatility surface can be calculated by

$$\sigma_{BS}^{i+1} = \frac{1}{T}\int_0^T \sigma_L^i(\bar{S}_t, t)dt, \tag{3.33}$$

where $\bar{S}_t$ is the most likely path. Now the $i + 1^{th}$ local volatility surface is obtained according to the following point-wise adjustment,

$$\sigma_L^{i+1}(S, t) = \frac{\sigma_{\text{market}}(S, t)}{\sigma_{BS}^{i+1}(S, t)}\sigma_L^i(S, t). \tag{3.34}$$

The procedure continues until the change between two iterations is sufficiently small.

As will be seen in the test section 5, the generated surface tends to be very noisy. Since that is not expected, an additional smoothing process is used. Tikhonov regularization with a discrete Hessian is used to eliminate small noise. The smooth surface is the solution to the following optimization problem.

$$\min_{\tilde{\sigma}_L} \sum_{(S,t)\in\Omega} ||\tilde{\sigma}_L(S,t) - \sigma_L(S,t)||_2^2 + \zeta||H(\tilde{\sigma}_L(S,t)||_F^2, \tag{3.35}$$

where σ_L is the surface obtained from the most likely path calibration, Ω is the set of all states and times. $||\cdot||_F$ is the Frobenius norm. The Frobenius norm of a matrix A with elements a_{ij} is $\sqrt{\sum_{i,j} a_{ij}^2}$. H is an operator calculating an approximation of the Hessian in (S,t) of $\tilde{\sigma}_L$. The operator H gives

$$H(\tilde{\sigma}_L(S,t) = \begin{bmatrix} \partial_S^2\tilde{\sigma}_L & \partial_{tS}\tilde{\sigma}_L \\ \partial_{St}\tilde{\sigma}_L & \partial_t^2\tilde{\sigma}_L \end{bmatrix}, \tag{3.36}$$

$$\partial_S^2\tilde{\sigma}_L = \frac{\tilde{\sigma}_L(S_{i+1},t_j) + \tilde{\sigma}_L(S_{i-1},t_j) - 2\tilde{\sigma}_L(S_i,t_j)}{(\Delta S)^2}, \tag{3.37}$$

$$\partial_{St}\tilde{\sigma}_L = \frac{\tilde{\sigma}_L(S_i+1,t_j) - \tilde{\sigma}_L(S_{i-1,t_j}) - \tilde{\sigma}_L(S_i,t_{j+1}) + \tilde{\sigma}_L(S_i,t_{j-1})}{\Delta S \Delta t}. \tag{3.38}$$

Of course these approximations are slightly altered at the boundary points. To speed up the optimization, the gradient of the objective function is calculated. Solving the first necessary optimality equation for a quadratic program $\nabla f = 0$ is equivalent to solving a linear equation system, which is solved quickly and yields the global optimal solution.

4 Test of the implied volatility models

4.1 Procedure

Two methods of the ones mentioned in section 2 are tested to model the implied volatility on the Eurcap market in March 2008. It is one stochastic volatility model, SABR and one parametric model SVI. The test is to investigate how stable they are. The goal is to find a method which can be running without any supervision. The calibration is done for all caps each day in March 2008.

4.2 Results

As can be seen, both SVI and SABR perform very well. Both of them are almost identical and coincide almost perfectly with each data point. Performance-wise the SABR model is calibrated almost instantly, the SVI takes between 1 to 20 seconds to calibrate, since each calibration uses 10 randomly chosen initial points.

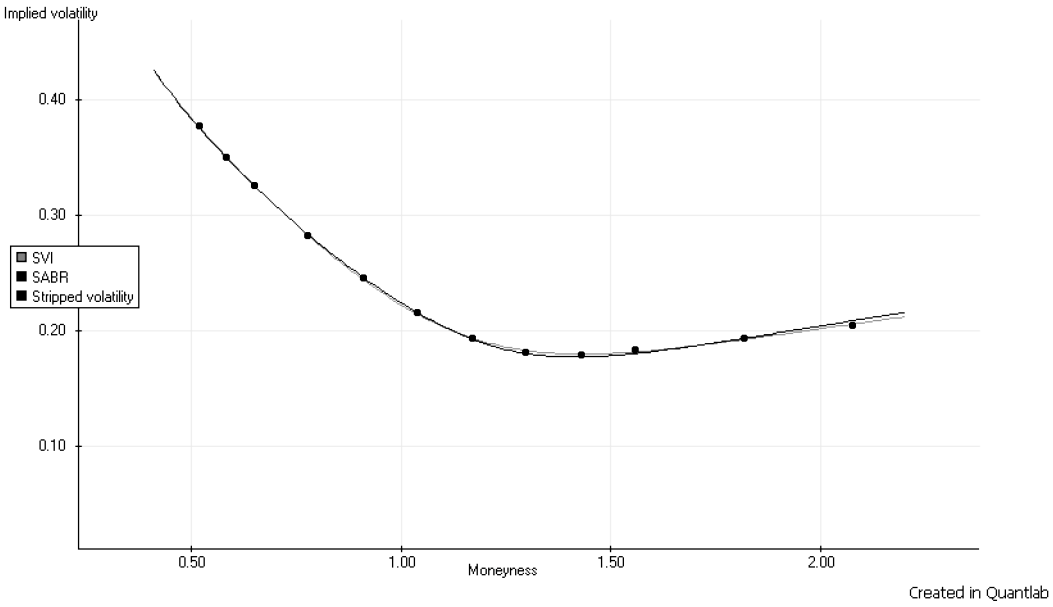

Figure 1: SVI and SABR fits to the 1.5 year volatility slice, 24/3-2008. SABR is calibrated according to the article. The dots are the stripped volatilities

The figure 2 shows a case were there is a noticeable difference between the data points and the SABR fit. Compare to the fit in figure 4 with a relaxed α a few pictures below.

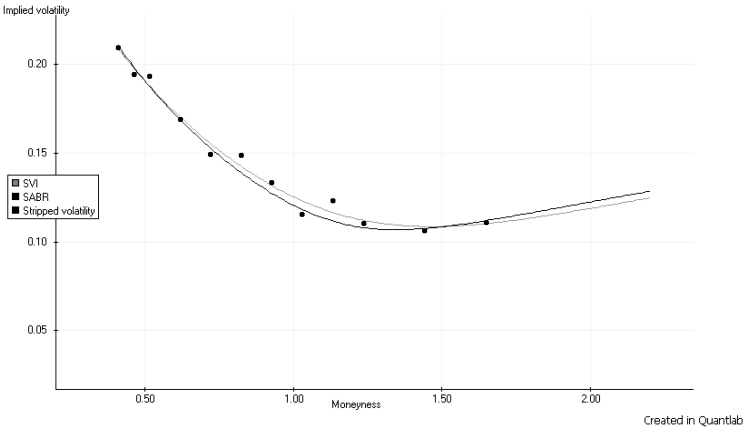

Figure 2: SVI and SABR fits to the 8 year volatility slice, 24/3-2008. SABR is calibrated according to the article. The dots are the stripped volatilities

Figure 3 shows an example where SABR and SVI differs from each other. Both curves are pretty close to the data points but when extrapolating these curves, SABR and SVI will go in different directions. It is far from clear which model to choose and luckily it is not the job of the author to choose, but it is interesting to notice that they may differ.

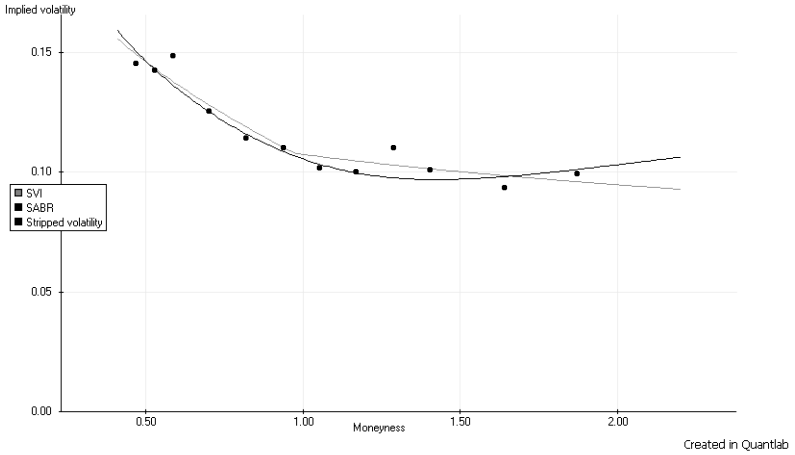

Figure 3: SVI and SABR fits to the 25 year volatility slice, 24/3-2008. SABR is calibrated according to the article. The dots are the stripped volatilities

Figure 4 shows the calibration on the same data as Figure 1, but now the SABR model is calibrated with a relaxed α. The fit is very good.

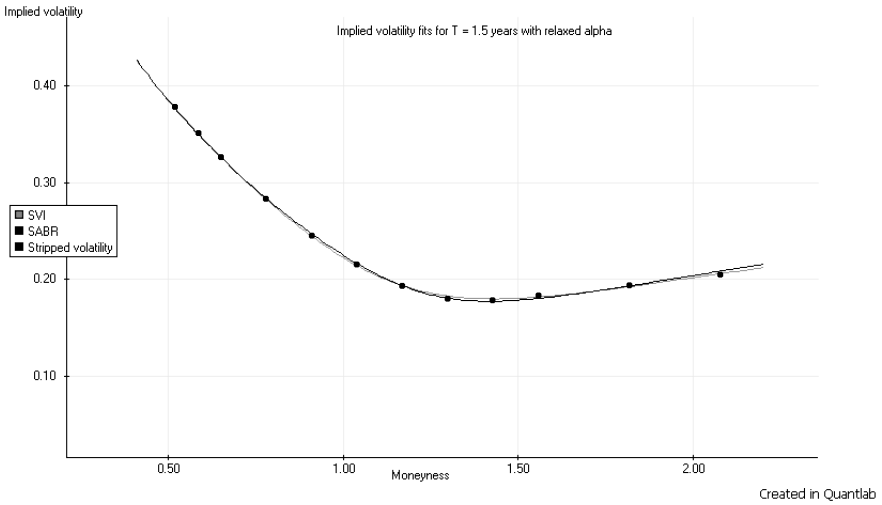

Figure 4: SVI and SABR fits to the 1.5 year volatility slice, 24/3-2008. SABR is calibrated with relaxed α. The dots are the stripped volatilities

Figure 5 shows another fit which also is very good. The fit is closer to the data than when α is chosen according to the article.

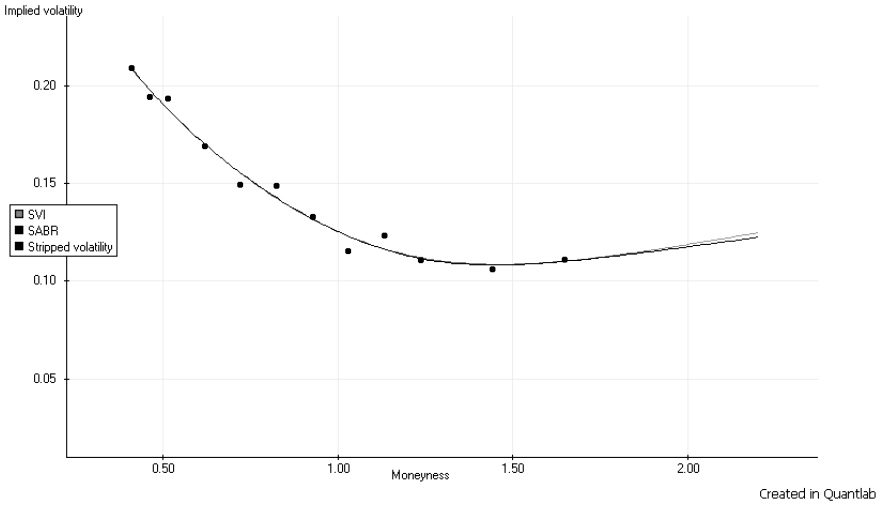

Figure 5: SVI and SABR fits to the 8 year volatility slice, 24/3-2008. SABR is calibrated with relaxed α. The dots are the stripped volatilities.

5 Test of the local volatility models

5.1 Procedure

This test is done on the European S&P 500 index European call options of October 1995, where the present value is S=\$590, and interest rate r=0.06. The market implied volatility surface is shown in figure 6, and the data is displayed in table 1. For the Dupire calibration, the splines for each time slice are evaluated for values $S \cdot [0.85 : 0.005 : 1.4]$. For the likely path calibration, the first calculation is done on a grid $S \cdot [0.85 : 0.025 : 1.4] \times [0.2 : 0.2 : 5]$. These two surfaces are then linearly interpolated to a finer grid $S \cdot [0.85 : 0.005 : 1.4] \times [0.2 : 0.05 : 5]$ using Matlab's *griddata*. The reason to this is that Monte Carlo simulation will be used to verify that the option prices can be regenerated from the local volatility surface obtained from the calibration. If the prices can be regenerated within a 95 % confidence interval, we consider that the local volatility is correct. The monte carlo simulation consists of simulating

$$S_n(t_{i+1}) = S_n(t_i) + r\Delta t_i S_n(t_i) + \sigma_L(S_n, t_i) S_n(t_i) \sqrt{\Delta t_i} N(0,1), \qquad (5.1)$$

where n is the index of simulation and i is the index of time. The scalar N, represents the number of simulations, which in this test is set to 10 000. The option price is obtained from

$$S(K_j, t_i) = \frac{e^{-rt_i}}{N} \sum_{n=1}^{N} \max\{0, S_n(t_i) - K_j\}. \qquad (5.2)$$

An approximate 95 % confidence interval is calculated as well. Below is the original implied volatility surface from the S&P500 and a table with the data. The real market prices are obtained by Black-Scholes formula from the given implied volatilities.

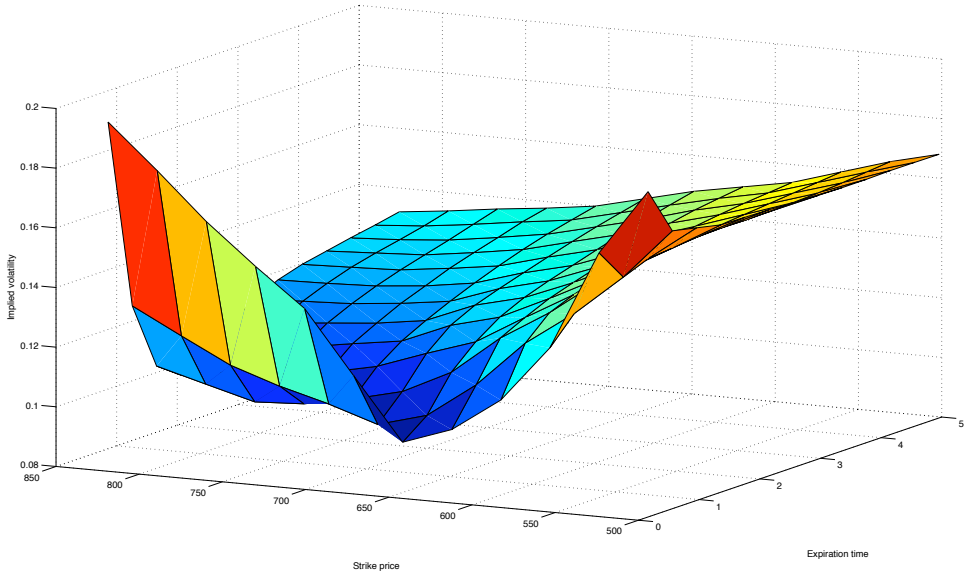

Figure 6: Implied volatility from European S&P500 index European call options of October 1995.

$T \backslash \frac{S}{S_0}$	0.850	0.900	0.950	1.000	1.050	1.100	1.150	1.200	1.300	1.400
0.175	0.190	0.168	0.133	0.113	0.102	0.097	0.120	0.142	0.169	0.200
0.425	0.177	0.155	0.138	0.125	0.109	0.103	0.100	0.114	0.130	0.150
0.695	0.172	0.157	0.144	0.133	0.118	0.104	0.100	0.101	0.108	0.124
0.940	0.171	0.159	0.149	0.137	0.127	0.113	0.106	0.103	0.100	0.110
1.000	0.171	0.159	0.150	0.138	0.128	0.115	0.107	0.103	0.099	0.108
1.500	0.169	0.160	0.151	0.142	0.133	0.124	0.119	0.113	0.107	0.102
2.000	0.169	0.161	0.153	0.145	0.137	0.130	0.126	0.119	0.115	0.111
3.000	0.168	0.161	0.155	0.149	0.143	0.137	0.133	0.128	0.124	0.123
4.000	0.168	0.162	0.157	0.152	0.148	0.143	0.139	0.135	0.130	0.128
5.000	0.168	0.164	0.159	0.154	0.151	0.148	0.144	0.140	0.136	0.132

Table 1: Implied volatility data from S&P500 index European call options of October 1995.

5.2 Results

5.2.1 The regularized Dupire equation

Figure 7 shows the local volatility surface obtained by the calibration used in section 3.2.1. By investigating different values of the parameters λ and α, they finally were set to 10000 and 500000 respectively, since they seem to generate a plausible surface. The slope at the right side of the figure is probably not correct and may depend on numerical difficulties, since the denominator of Dupire's equation for low strike prices is very small.

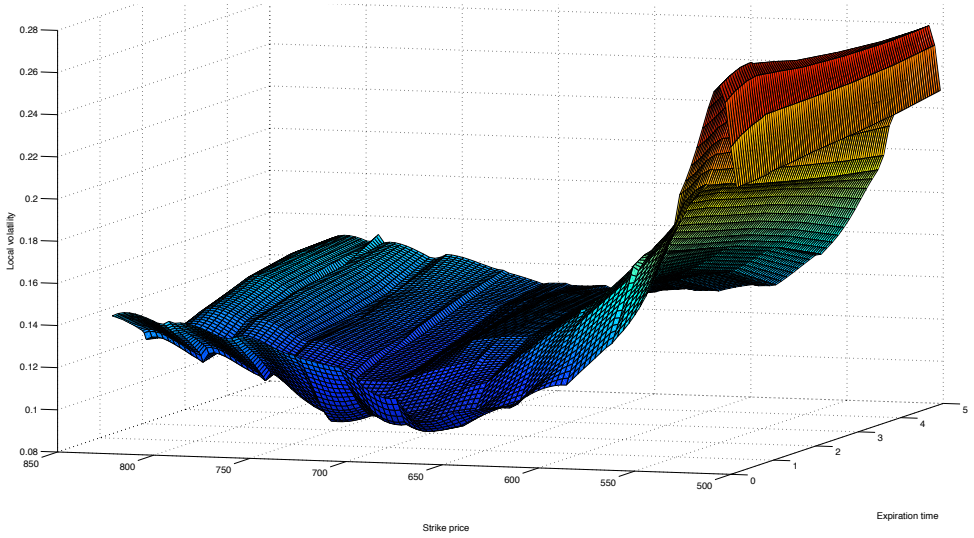

Figure 7: The linearly interpolated local volatility surface.

Figure 8 is a way of displaying which values are within the 95% confidence interval. A dot means that the corresponding value is significant. The surface seem to be very good at expiration times longer than a few months, but for short expirations it is probably not correct. Table 2 displays interpolated market prices for a few strike prices and expiration times. Table 3 displays the corresponding Monte Carlo prices from the Dupire calibrated volatility surface. Compare the values of the two tables. Values in parenthesis are out side the 95 % confidence interval.

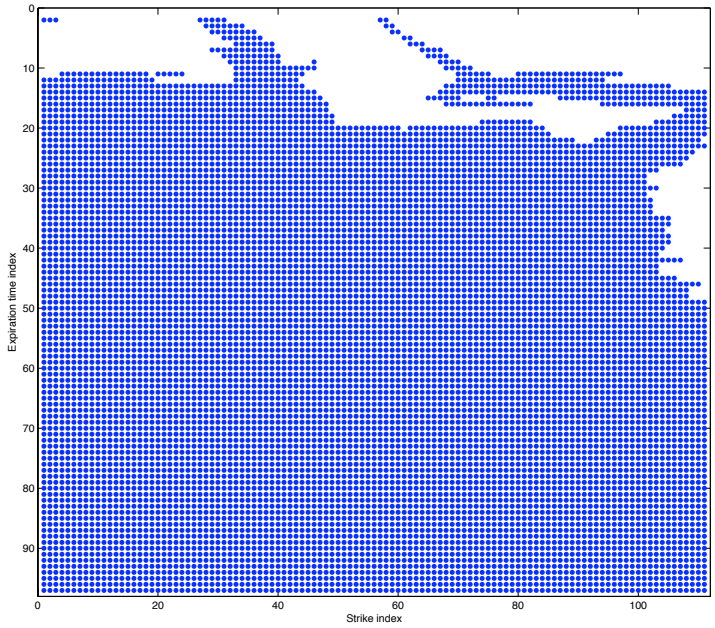

Figure 8: Dots show values within the approximate 95 % confidence interval.

$T \backslash \frac{S}{S_0}$	0.85	0.91	0.97	1.03	1.09	1.15	1.21	1.27	1.33	1.39
0.25	94.804	60.706	28.785	6.952	0.518	0.085	0.027	0.008	0.003	0.002
0.75	111.803	80.520	52.084	28.095	10.781	3.112	0.874	0.260	0.098	0.051
1.25	128.193	98.749	71.500	47.461	27.211	13.580	5.871	2.253	0.935	0.366
1.75	143.489	115.304	88.778	64.630	43.794	27.775	15.803	8.620	4.232	1.906
2.25	157.983	130.696	104.984	80.940	59.679	42.324	27.738	17.998	11.229	6.554
2.75	171.667	145.131	120.156	96.495	75.089	56.728	40.830	29.117	20.328	13.849
3.25	184.729	159.043	134.562	111.600	90.056	71.239	54.566	41.435	30.928	22.919
3.75	197.273	172.301	148.425	126.080	104.774	85.629	68.461	54.147	42.351	32.780
4.25	209.309	185.082	161.836	139.793	119.110	99.825	82.349	67.280	54.213	43.552
4.75	220.883	197.460	174.701	153.085	133.010	113.712	95.982	80.671	66.806	54.885

Table 2: Market prices in $

$T \backslash \frac{S}{S_0}$	0.85	0.91	0.97	1.03	1.09	1.15	1.21	1.27	1.33	1.39
0.25	94.312	(59.642)	(28.130)	(7.782)	(0.962)	(0.041)	(0.000)	(0.000)	(0.000)	(0.000)
0.75	110.703	79.603	(51.111)	27.827	(11.932)	(3.781)	0.951	0.220	(0.058)	(0.017)
1.25	126.963	97.856	70.764	46.788	27.402	13.832	5.976	2.426	1.039	0.474
1.75	142.351	114.492	88.047	63.960	43.316	27.105	15.571	8.405	4.441	(2.334)
2.25	157.190	130.194	104.490	80.701	59.600	41.818	27.929	18.111	11.528	(7.144)
2.75	171.317	145.161	120.103	96.577	75.221	56.639	41.301	29.565	20.756	14.327
3.25	184.390	158.908	134.373	111.222	89.909	70.854	54.455	41.225	30.864	22.839
3.75	196.648	171.844	147.837	125.050	103.884	84.732	67.924	53.770	42.189	32.850
4.25	208.282	183.970	160.429	138.002	117.088	98.010	81.097	66.450	53.974	43.343
4.75	219.986	196.317	173.241	151.194	130.505	111.463	94.309	79.172	65.992	54.564

Table 3: Monte carlo prices in $ evaluated with the Dupire calibrated surface. Values in parenthesis are not significant.

5.2.2 The most likely path calibration

Figure 9 shows the surface obtained from the most likely path calibration. It takes about 30 seconds in matlab to generate the surface. As mentioned before, it is very noisy. This is due to that the iterative process is point-wise and therefore does not take its neighbors into account. Because of the noise, a smoothing process was tested.

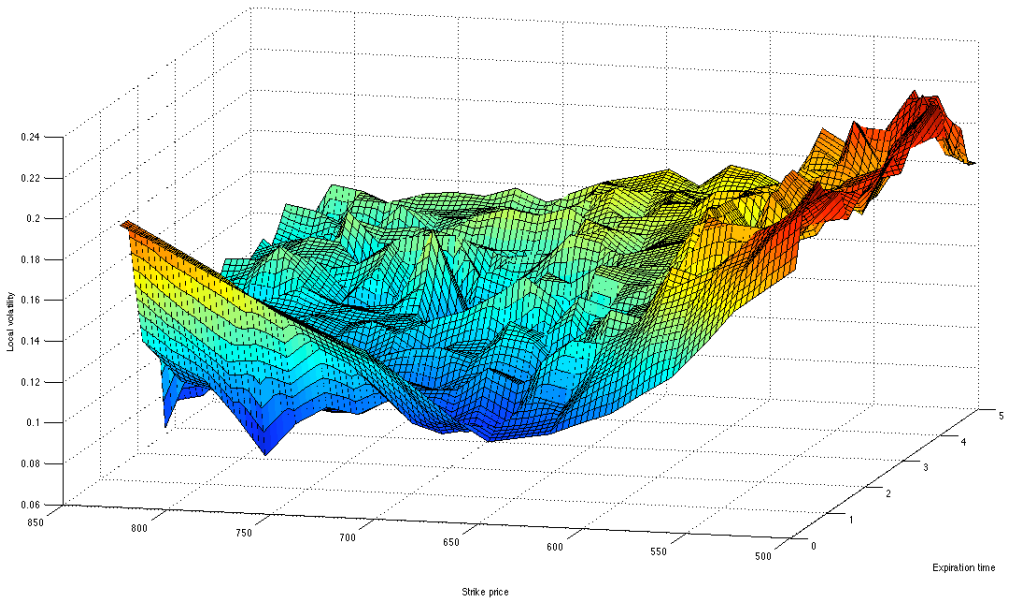

Figure 9: The linearly interpolated local volatility surface.

Figure 10 shows the significance of the most likely path calibrated surface.

This surface is not as good as the Dupire-calibrated surface. The values for low strikes are good but for high strikes a lot of values are not significant.

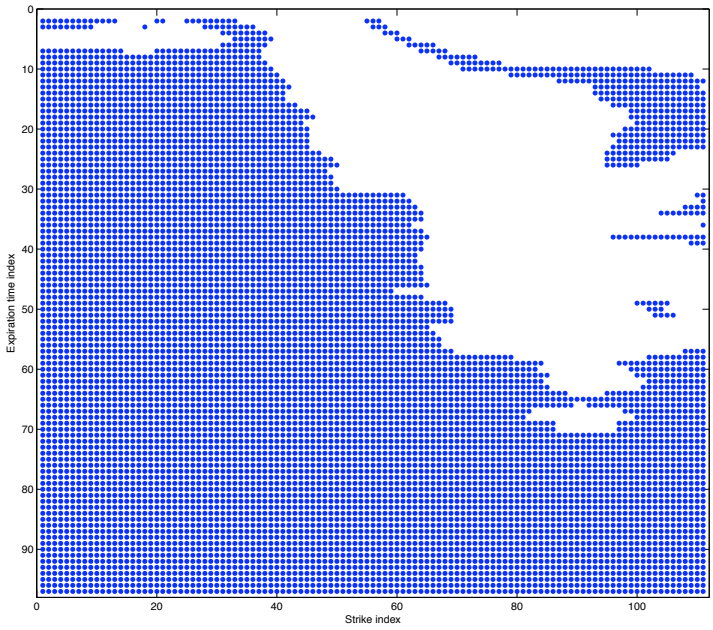

Figure 10: Dots show values within the approximated 95% confidence interval.

$T \backslash \frac{S}{S_0}$	0.85	0.91	0.97	1.03	1.09	1.15	1.21	1.27	1.33	1.39
0.25	94.900	60.161	28.314	(7.503)	(0.794)	(0.029)	(0.002)	(0.000)	(0.000)	(0.000)
0.75	111.003	79.636	51.262	28.109	(12.158)	(4.063)	(1.192)	(0.357)	0.110	0.037
1.25	127.828	98.406	71.192	47.488	(28.441)	(15.014)	(6.984)	(2.830)	1.094	0.458
1.75	142.872	114.685	88.383	64.662	44.505	28.594	(17.094)	(9.478)	(4.806)	(2.159)
2.25	157.724	130.614	105.133	81.647	60.866	43.409	(29.671)	(19.463)	(12.241)	(7.314)
2.75	170.969	144.674	119.824	96.749	75.880	57.633	(42.502)	(30.714)	(21.673)	(14.859)
3.25	183.652	158.331	134.131	111.445	90.693	72.192	56.090	(42.939)	(32.364)	23.881
3.75	196.051	171.346	147.824	125.644	105.071	86.420	69.899	55.555	43.586	33.701
4.24	208.066	184.056	161.047	139.293	119.029	100.381	83.468	68.437	55.373	44.292
4.75	219.678	196.409	174.008	152.728	132.776	114.236	97.282	81.980	68.393	56.565

Table 4: Monte carlo prices in $ evaluated with the most likely path calibrated surface. Values in parenthesis are not significant.

Since the first surface was pretty noisy, a Tikhonov smoother was applied with $\zeta = 0.005$, according to (3.35) in the calibration section 3.2.1. Then the surface becomes like figure 11. It is much smoother but still has the overall shape of the original one.

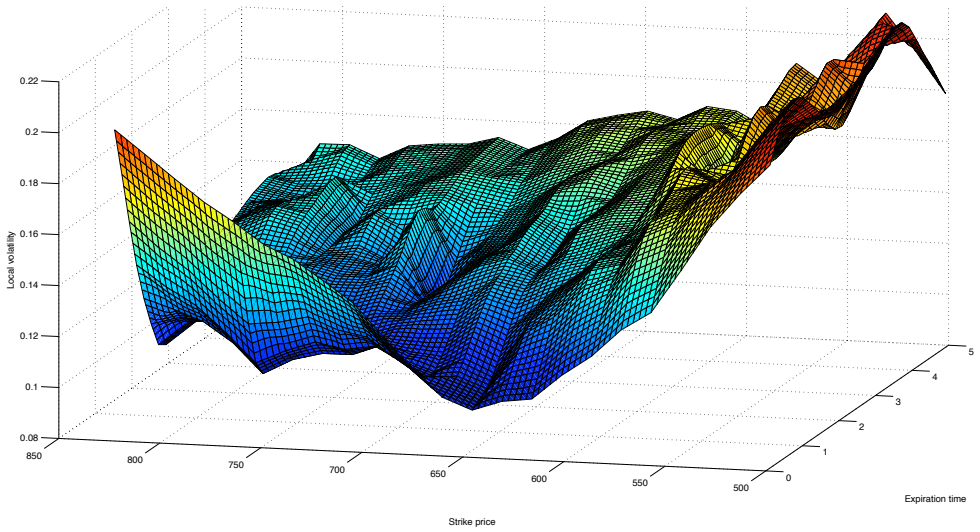

Figure 11: The linearly interpolated local volatility surface.

Figure 12 shows the significance of the smoothed surface. Even though the local volatility surface looks smoother, it unfortunately does not give the correct prices. The surface seem equally good as the unsmoothed surface.

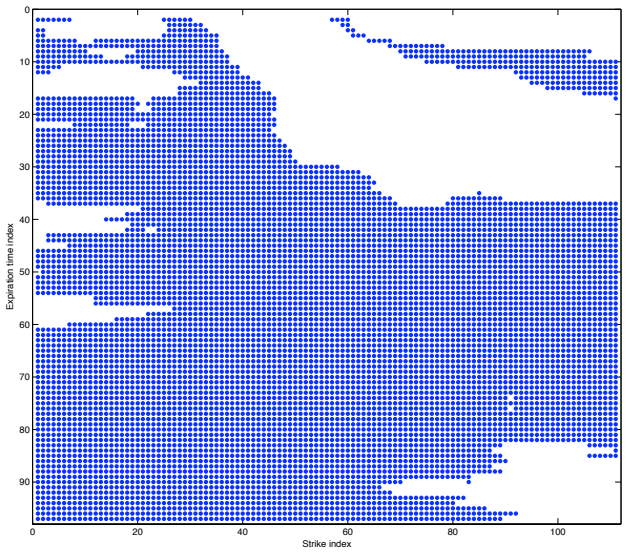

Figure 12: Difference between the confidence bounds and the market price. Values greater than 0 is within the approximated 95 % confidence interval.

$T \backslash \frac{S}{S_0}$	0.85	0.91	0.97	1.03	1.09	1.15	1.21	1.27	1.33	1.39
0.25	94.499	(59.831)	28.314	(7.868)	(1.029)	(0.054)	(0.003)	(0.000)	(0.000)	(0.000)
0.75	110.711	(79.368)	(51.141)	28.266	(12.599)	(4.471)	(1.401)	(0.410)	0.119	0.034
1.25	(126.668)	97.440	70.425	47.003	(28.372)	(15.251)	(7.349)	(3.189)	(1.319)	(0.538)
1.75	142.102	114.124	88.024	64.385	44.267	28.446	16.857	9.330	4.886	2.451
2.25	(155.715)	(128.700)	103.269	79.982	59.416	42.148	28.424	18.278	11.187	6.511
2.75	169.434	143.367	118.625	95.644	74.802	56.591	41.339	29.456	20.562	13.937
3.25	182.499	157.221	133.145	110.526	89.756	71.244	55.277	42.029	31.426	23.126
3.75	195.647	171.057	147.530	125.348	104.763	86.112	69.691	55.490	43.505	33.655
4.25	208.737	184.844	161.863	140.117	119.781	100.994	84.022	69.038	56.018	44.981
4.75	221.380	198.208	175.840	154.503	134.451	115.829	(98.700)	(83.265)	(69.505)	(57.458)

Table 5: Monte carlo prices in $ evaluated with the smoothed most likely path calibrated surface. Values in parenthesis are not significant.

6 Discussion

Our first goal was to find a couple of models which in a fast and stable manner can give an implied volatility surface. The test has shown that both the SABR and SVI serve as good models. The SABR model was really stable and fast. The same initial point was used on all test data and a solution was always found. This is probably due to that the SABR model just has a few parameters, especially when the α is chosen according to the article, and just two parameters are calibrated.

The SVI was slightly more difficult to calibrate, since some initial points did not lead to a solution. Therefore it was necessary to use random initial points. When using 10 initial points it was enough to find a solution for every day and option from the test data set. However it is no guarantee that it will always find a solution with exactly 10 restarts. For each restart, the calibration time increases, so that number of restarts was chosen as a compromise of stability and calculation time. With 10 restarts, the calculation time was about a few seconds, but in the worst case it could be 20 seconds.

If one model was to be chosen, the SABR model might be preferable, but both models could be interesting to use and as was said earlier, they both perform well.

A minor point which not has been discussed earlier in the thesis is that the SABR model can be used to price exotic options as well by using the SABR dynamics, with the parameters obtained from the calibration, in a Monte Carlo pricer. It is actually often the case that banks in Stockholm rather use stochastic volatility to price exotic options than local volatility, especially Heston and Bates model, which were mentioned earlier. But from an optimization point of view it was more interesting to work with the local volatility.

The second goal was to investigate whether a fairly fast and stable method could extract a local volatility surface from quoted European option prices or not. The short answer is, probably not. One of them is fast and could generate a plausible surface, but requires a lot of tweaking to set the parameters right. The other one is fairly quick to calibrate and does not need any initial guesses but it is doubtful whether the obtained surface is significant or not. To further give a measure of how good the surfaces are, the original implied volatility surface was used as the local volatility in the Monte Carlo pricer. Figure 13 shows the significance of using implied volatility as local volatility. Notice that almost all points are insignificant. Points for greater expiration time index than 40 are not displayed, since they are all insignificant. This shows that the calibrated surfaces obtained before, are a lot more accurate than when using the implied volatility surface, but maybe not good enough.

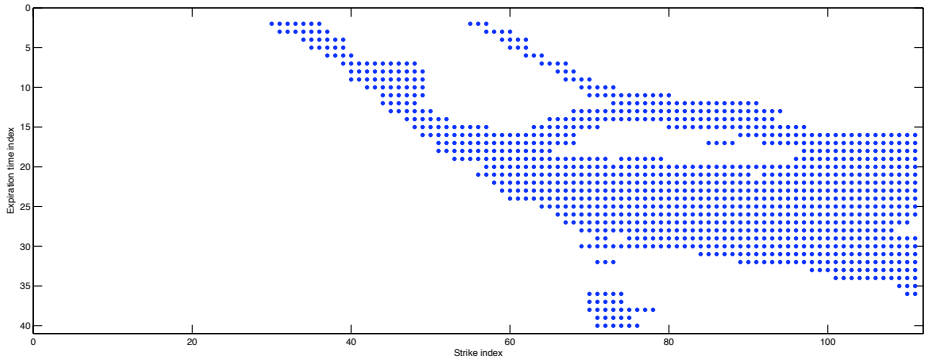

Figure 13: Dots show values within the approximated 95 % confidence interval.

To illustrate the problem of setting the right parameters, look what happens if λ decreases with 0.1 %. Two humps appear in the middle of the surface. This is due to that the second derivative, and therefore the denominator of (3.25), locally is relatively small. By punishing the splines sufficiently, the variations are treated, which yields a smoother surface. There is not yet a strategy to automatically set these parameters in advance.

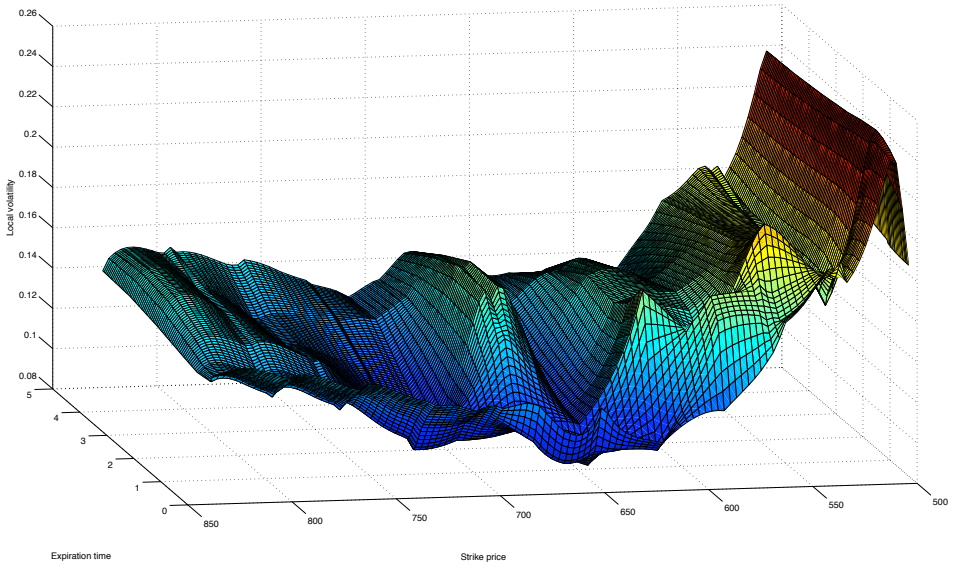

Figure 14: The local volatility when λ changes from 10 000 to 1 000.

It seems that the local volatility is difficult to extract and probably a more rigorous method is needed. Perhaps a pde-solver based on finite elements.

References

[1] Bates, D. *Jump and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options*,
The Review of Financial Studies, 9: 69-107, 1996

[2] Black, F. & Scholes, M. *"The Pricing of Options and Corporate Liabilities*,
Journal of Political Economy 81, 3: 637-654, 1973

[3] Fengler, M. *Arbitrage-Free Smoothing of the Implied Volatility Surface*,
Humboldt-Universität zu Berlin, 2005

[4] Gatheral, J. *The volatility surface; A Practitioner's Guide*,
John Wiley & Sons, Inc., 2006

[5] Gatheral, J. *A parsimonious arbitrage-free implied volatility parameterization with application to the valuation of volatility derivatives*,
Global Derivatives & Risk Management 2004

[6] Hagan, P. S. *Managing smile risk*,
Wilmott Magazine, 2002

[7] Hanke, M. & Rösler, E. *Computation of Local Volatilities from Regularized Dupire Equations*,
Fachbereich Mathematik, Johannes Gutenberg-Universität, 2004

[8] Heston, S. L. *A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bonds and Currency Options*,
The Review of Financial Studies, 1993

[9] Hirsa, A. & Pender, P. *Local Volatility Calibration using the Most Likely Path*,
Computational Methods in Finance, 2006

[10] Levenberg, K. *A Method for the Solution of Certain NonLinear Problems in Least Squares*,
The Quarterly of Applied Mathematics 2: 164-168, 1944

[11] Marquardt, D. *An Algorithm for Least-Squares Estimation of Nonlinear Parameters*,
SIAM Journal on Applied Mathematics 11:431-441, 1963

[12] Rebonato, R. *Volatility and Correlation*,
John Wiley & Sons Inc., 1999