

Testing of an Arbitrage-free Volatility Surface

Grant Tarr

A dissertation submitted to the Faculty of Commerce, University of Cape Town, in partial fulfilment of the requirements for the degree of Master of Philosophy.

July 9, 2022

*MPhil in Mathematical Finance,
University of Cape Town.*



The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

Declaration

I declare that this dissertation is my own, unaided work. It is being submitted for the Degree of Master of Philosophy in the University of the Cape Town. It has not been submitted before for any degree or examination in any other University.

July 9, 2022

Abstract

The Ensemble Carr-Pelts surface, which is a weighted mixture of standard Carr-Pelts surfaces, is an arbitrage-free parameterization of an implied volatility surface proposed by [Antonov, Konikov and Spector \(2019\)](#). This dissertation aims to investigate the additional benefits provided by using the Ensemble Carr-Pelts surface as opposed to the standard Carr-Pelts surface. We also show its validity in comparison to stochastic volatility inspired [Gatheral \(2004\)](#) surface, which is widely used by practitioners. The approach adopted was done in three stages, with each stage calibrating to an increasingly complicated surface. Surfaces considered were a flat volatility surface, a surface changing with strike only, and a surface changing with both strike and maturity. Testing revealed that as complexity increased for the implied volatility surface, the Ensemble Carr-Pelts calibrated better than Carr-Pelts. When compared to the widely accepted stochastic volatility inspired surface; considering no-arbitrage was not enforced, the Ensemble Carr-Pelts performed adequately. However, the Ensemble Carr-Pelts takes significantly longer to calibrate.

Acknowledgements

I thank Ralph Rudd for being the best lecturer I have had the privilege of not only being taught by but also supervised by. It has been a pleasure to personally receive guidance from you, someone I acknowledge with such esteem.

I also thank Professor David Taylor for the degree he has offered, giving me such an excellent platform for growth and future endeavours.

I would like to thank my parents, Cyril and Benita Tarr, as well as my brother Dillon Tarr who were and always will be my greatest support and encouragement.

Lastly, I would like to thank Standard Bank's generous bursary, administered by StudyTrust, with which I would have never gotten through the year without.

Contents

1. Introduction	1
2. Implied Volatility Surface	3
3. Surface Parameterization	6
3.1 Stochastic Volatility Inspired	7
3.2 Carr-Pelts	8
3.2.1 No-Arbitrage Parameterization	8
3.2.2 Black-Scholes Carr-Pelts	9
3.3 Ensemble Carr-Pelts	12
4. Calibration	13
4.1 Stochastic Volatility Inspired	13
4.2 Carr-Pelts	15
4.3 Ensemble Carr-Pelts	15
4.4 Results	16
5. Local Volatility	20
5.1 Stochastic Volatility Inspired	20
5.2 Carr-Pelts	21
5.3 Ensemble Carr-Pelts	22
5.4 Results	23
6. Conclusion	26
Bibliography	28
A. Appendix	30
A.1 Flat Volatility	30
A.2 Option Pricing Under the CEV Model	31
A.3 Option Pricing Under the Heston Model	31
A.4 Little Trap Integrand	32

List of Figures

4.1	Percentage relative pricing differences under the CEV model	17
4.2	Percentage relative pricing differences under the Heston model	18
5.1	Local volatility surface under CEV compared to those of SVI, CP and ECP brought out by our calibration	23
5.2	Local volatility surface under CEV for CP and ECP with the smaller z -grid at the top and the larger z -grid below it	24
A.1	Relative pricing differences under flat volatility for $\sigma = 0.2$	30
A.2	Little Trap Convergence of Integrand	32

List of Tables

4.1	Numerical results showing error of CEV calibration	17
4.2	Numerical results showing error of Heston calibration	19
5.1	Numerical results showing error differences in local volatility	23
5.2	Numerical results showing error differences in local volatility for a larger z -grid	25

Chapter 1

Introduction

Typically, pricing an option of a specific strike and maturity requires knowing the prevailing interest rate as well as the volatility of the underlying. Assuming a constant interest rate; the interest rate, strike and maturity of the option are all known explicitly at the time of writing up the option, but volatility is not.

The [Black and Scholes \(1973\)](#) model makes the assumption of constant volatility for the underlying. If we can observe the market prices for vanilla options and assume the Black-Scholes model for the underlying, we can obtain these volatilities, known as implied volatilities. Black-Scholes proposes that the underlying that the option is written on follows geometric Brownian motion,

$$dS_t = \mu S_t dt + \sigma S_t dW_t,$$

where the volatility σ is constant. This assumes that during the life of the option the volatility of the underlying stays the same. A natural extension of this formulation, shown by [Dupire \(1994\)](#), is considering a time-dependent deterministic volatility as follows:

$$dS_t = \mu S_t dt + \sigma(S_t, t) S_t dW_t. \quad (1.1)$$

In (1.1) no additional source of randomness is introduced and volatility is assumed to move during the life of the option with its value being deterministic.

Relaxing the assumptions placed on volatility further we can arrive at a formulation in which volatility is itself stochastic in nature as represented in the [Heston \(1993\)](#) model by the following stock price and variance processes:

$$\begin{aligned} \frac{dS_t}{S_t} &= (r_t - q_t)dt + \sigma(S_t, t)\lambda(t)\sqrt{z_t}dW_t^1, \\ dz_t &= \kappa(t)(1 - z_t)dt + \gamma(t)\sqrt{z_t}dW_t^2, \\ dW_t^1 dW_t^2 &= \rho(t)dt, \end{aligned}$$

where

- λ is the overall volatility level,
- z_t is the normalized stochastic variance following a Cox-Ingersell-Ross (CIR) process,
- $\kappa(t)$ is the reversion speed,
- $\gamma(t)$ is the volatility of the variance,
- $\rho(t)$ is the correlation between the underlying and variance processes,
- r_t and q_t are the domestic and dividend rates, respectively.

The next chapter discusses the implied volatility surface and what arbitrage opportunities exist for the surface. Chapter 3 provides details on the three parameterizations of the implied volatility surface this dissertation will be investigating. We then move to the calibration of these parameterizations in Chapter 4 in an attempt to find out how each of them performs when calibrating to different implied volatility surfaces. How these calibrated surfaces can be used to obtain local volatility is shown in Chapter 5 along with a comparison between each parameterization.

Chapter 2

Implied Volatility Surface

An implied volatility surface is a mapping of strikes and maturities from a set of options to their implied volatilities. Only a finite set of these options are seen in the market. Thus, to get implied volatility data for options with corresponding strikes or maturities not seen in the market, an interpolation scheme can be used on observed data. The caveat to this method is that there is no guarantee of arbitrage-free pricing as well as positive implied volatilities.

Arbitrage from implied volatility is not easily understood. However, as mentioned in [Antonov *et al.* \(2019\)](#), it is reasonable to look towards option prices instead. With this in mind, we find ourselves facing two types of arbitrage.

Firstly, arbitrage in the maturity dimension called a calendar spread arbitrage. Calendar spread arbitrage, shown in [Gatheral and Jacquier \(2014\)](#), exploits two different maturity dates. We say calendar spread arbitrage is present if given two different maturity date time-slices in the volatility surface, total variance (volatility squared multiplied by maturity) curves for each time-slice intersect one another, i.e. $\sigma_1^2 T_1 = \sigma_2^2 T_2$ where σ and T are volatilities and maturities respectively. By extension, it is easy to see time-slice total variance curves should be strictly increasing with increasing maturity. This is captured by (2.1) below:

$$\frac{\partial w}{\partial T} > 0, \tag{2.1}$$

where w is the total variance and T is the time to maturity.

Exploiting Calendar Spread Arbitrage

To exploit this arbitrage a calendar spread is used; whereby positions in two calls with different maturities but the same strike are taken. If the total variance of the long-dated option is lower than the short-dated, the call option on the short-dated is more expensive. We sell the short-dated and buy the long-dated call and lock in an initial profit.

- If the short-dated call expires out-the-money we make a profit regardless of whether the long-dated option expires in or out the money.
- If the short-dated call expires in-the-money we short sell the stock and receive the strike.
 - If the long-dated call expires in-the-money we buy back the stock we sold with the strike we received earlier by exercising the call.
 - If the long-dated call expires out-the-money the strike is higher than the underlying, and we buy the stock with an adequate proportion of the strike resulting in a profit.

Secondly, arbitrage in the strike dimension called butterfly arbitrage. Butterfly arbitrage, as stated in [Gatheral and Jacquier \(2014\)](#), is absent if the corresponding density of a particular time-slice is strictly positive. This result was proven by [Breen and Litzenberger \(1978\)](#) whereby the second partial derivative with respect to the strike of the call options price was shown to be a valid density function captured by (2.2) below:

$$\frac{\partial^2 C}{\partial K^2} > 0, \quad (2.2)$$

where C is the call option price, and K is the strike of the call option.

Exploiting Butterfly Spread Arbitrage

One can exploit this arbitrage by making use of a butterfly spread by taking positions in three call options involving three different equally spaced strikes; one high, one low and one with a strike being the midpoint between the two strikes. The positions taken involves buying the calls with high and low strikes while selling two calls with the strike in the middle. Since the call options price is not convex, decreasing gradient, an initial positive payoff is seen at inception.

- If all the calls fall out-the-money the payoff at maturity is zero and we take the initial profit.

- If only the call with the lowest strike falls in-the-money we receive the higher stock value making a profit. The maximum profit reached here is the distance between strikes.
- If both the lowest and middle strikes fall in-the-money we also only make a profit since the strikes are equidistant from one another resulting in the payoff gradually moving from the maximum possible profit for the profile to zero.
- If all the calls fall in-the-money the payoff at maturity is zero since stocks cancel and since the strikes are equidistant the low and high strikes cancel with the strike in the middle.

Chapter 3

Surface Parameterization

In Chapter 2 we discussed arbitrage in an implied volatility surface. To help build an implied volatility surface that is free of arbitrage, a parameterization of a surface is created that adheres to the constraints (2.1) and (2.2) along with being strictly positive. This parameterized surface is then calibrated to market data. These functions of a finite set of parameters that parameterize the implied volatility surface have many uses in finance, as stated in [Antonov *et al.* \(2019\)](#), including

- the pricing of vanilla options and hedging,
- data cleansing and information reduction,
- detecting arbitrage,
- and the development of local volatility models.

In addition to avoiding arbitrage, there are several other desirable properties of an implied volatility surface. The following, as outlined in [Antonov *et al.* \(2019\)](#), are a few:

- The volatility surfaces in both the papers being considered are parameterizations. It should thus be parsimonious by using as few parameters as possible to parameterize the surface while still capturing the characteristics of that surface.
- Flexibility is also required for the surface to be able to fit a wide variety of option data sets.
- Dupire's formula for calculating local volatility requires partial derivatives of the call option. It is thus natural to want sufficient smoothness from the surface so that these derivatives can be obtained.

This chapter will define each of the three different implied volatility surface parameterizations we will be comparing throughout this dissertation. We will also discuss their formulation and their benefits.

3.1 Stochastic Volatility Inspired

The stochastic volatility inspired (SVI) [Gatheral \(2004\)](#) surface is a five parameter model used widely by practitioners ([Gatheral and Jacquier \(2014\)](#)). The model has two attractive features.

Firstly, the parameterization is fitted on each time-slice (a slice in the surface fixing maturity only). To eliminate calendar spread arbitrage is thus an easy task, as care must only be taken so that each time-slice does not overlap while ensuring they lie below the preceding time-slice.

Secondly, the volatility smile has a desirable shape. For a single time-slice of the surface, increasing values of absolute moneyness results in the slope of implied volatility becoming linear. This is expected, as stated in [Gatheral \(2004\)](#), since convexity increases the further out-the-money options become and convexity decreases the further in-the-money they become.

SVI is defined on total variance (implied volatility squared multiplied by maturity) by the following parameterization:

$$w(k; a, b, m, \rho, \sigma) = a + b\{\rho(k - m) + \sqrt{\sigma^2 + (k - m)^2}\},$$

where $k = \log\left(\frac{K}{F(T)}\right)$ with K and $F(T)$ representing the strike and the forward price respectively. To ensure that this surface remains positive, the following constraints on the parameters apply:

$$\begin{aligned} b &\geq 0, \\ \sigma &\geq 0, \\ |\rho| &\leq 1, \\ a + b\sigma\sqrt{1 - \rho^2} &\geq 0. \end{aligned} \tag{3.1}$$

To ensure the smile is free of butterfly arbitrage, rigorous procedures will have to be adopted. In [Gatheral and Jacquier \(2014\)](#) it is stated that this is a difficult task as general conditions to ensure no arbitrage in the strike dimension is seemingly impossible.

It is important to note that SVI is a time-slice parameterization of the volatility surface and in turn, the surface will have to be built up by many of these time-slice curves. Interpolation as well as extrapolation schemes given in [Gatheral and Jacquier \(2014\)](#), defined on these curves, allow for a surface to be built which is free of arbitrage.

It can be shown that the large maturity limit of the [Heston \(1993\)](#) model is exactly SVI, as stated in [Gatheral and Jacquier \(2011\)](#). This validates the SVI parameterization as not being an arbitrary formulation of the implied volatility surface.

This parameterization will be the benchmark we wish to compare our other parameterisations to. It is easy to implement, calibrates quickly to observed market data and provides a good fit.

3.2 Carr-Pelts

Call option prices which conform to no-arbitrage conditions are convex in the strike. The Carr and Pelts (2015a) surface (CP), is a parameterization which takes advantage of this convexity and applies convex duality. The parameterization is restricted and only holds in a static arbitrage setting as outlined in Carr and Madan (2005).

3.2.1 No-Arbitrage Parameterization

The intrinsic value of an option contract, independent of the options maturity T , is interpreted as an exchange $(N_+ - N_-)^+$. Receiving N_+ and paying N_- if $N_+ > N_-$.

The premium of the option contract, dependent on the options maturity, can be expressed as $P(N_+, N_-, T) : \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^+ \rightarrow \mathbb{R}^+$.

A static setting is described in Carr and Pelts (2015a) whereby the maturity of the option is fixed $T > 0$. Because this is a static setting, a consequence of linear homogeneity follows, for some $\lambda > 0$:

$$P(\lambda N_+, \lambda N_-, T) = \lambda P(N_+, N_-, T). \quad (3.2)$$

If (3.2) is differentiated with respect to λ and setting $\lambda = 1$ the following is obtained

$$P(N_+, N_-, T) = N_+ P_1(N_+, N_-, T) + N_- P_2(N_+, N_-, T), \quad (3.3)$$

where P_1 and P_2 are first partial derivatives of P w.r.t. N_+ and N_- respectively.

Each of these partial derivatives represent the probability that the options will end in-the-money:

$$\begin{aligned} P_1(N_+, N_-, T) &= \mathbb{Q}_+ \{N_{+T} > N_{-T}\} \equiv \Delta_+, \\ -P_2(N_+, N_-, T) &= \mathbb{Q}_- \{N_{+T} > N_{-T}\} \equiv \Delta_- \end{aligned}$$

and can be interpreted as two digital options with Δ_+ and Δ_- with payoff at maturity T if $N_{+T} > N_{-T}$ with the appropriate martingale measure represented by \mathbb{Q} .

Rearranging (3.3), $P_2(N_+, N_-, T)$ can be represented as:

$$-P_2(N_+, N_-, T) = N_+/N_- P_1(N_+, N_-, T) - P(N_+, N_-, T)/N_-. \quad (3.4)$$

Letting $\hat{P}(R, T) \equiv (P(N_+, N_-, T))/N_-$ where $R \equiv N_+/N_-$ and rearranging to get $P(N_+, N_-, T) = N_- \hat{P}(R, T)$ then differentiating w.r.t. $N_+ \implies P_1(N_+, N_-, T) = P_1(R, T)$. Substituting back in to (3.4):

$$-P_2(N_+, N_-, T) = RP_1(R, T) - \hat{P}(R, T).$$

Letting $P_1(R, T) = \Delta_+$ we have $\Delta_- = -P_2(N_+, N_-, T) = RP_1(R, T) - \hat{P}(R, T) = R\Delta_+ - \hat{P}(R, T)$. Resulting in Δ_- being an increasing function of Δ_+ . Which makes Δ_- a distortion function.

To prove that there is no calendar spread arbitrage a proof by contradiction approach is used in Carr and Pelts (2015b) whereby essentially for a fixed Δ_+ through time Δ_- only decreases. Also true is that for a fixed Δ_- through T we have that Δ_+ only increases. Hence, with increasing T price of the options increases as well leading to no calendar spread arbitrage.

From this, if a convex distortion function can be specified at each $T > 0$, then a convex function can also be specified for $\hat{P}(R, T)$ linked to R . Hence the convexity of Δ_- in Δ_+ implies convexity of \hat{P} in R which implied convexity of P in N_+ and N_- . The surface is free of butterfly spread arbitrage as shown by (2.2) since N_- represents K .

For Black-Scholes where $\Delta_- = N(d_1)$ and $\Delta_+ = N(d_2)$, it is can be shown that for $\Delta_- = N(N'(\Delta_+) - \sigma\sqrt{T})$ is indeed a distortion function being convex in Δ_+ and decreasing in $T \geq 0$.

3.2.2 Black-Scholes Carr-Pelts

The Carr-Pelts parameterization is as follows:

$$C(K, T) = D(T)(F(T)\Delta_+ - K\Delta_-), \quad (3.5)$$

with strike K and maturity T , where $C(K, T)$ is the call price, $D(T)$ is the discount factor and $F(T)$ is the forward price.

The parameterization makes use of two functions $\tau(T)$ and $h(z)$ where

$$\begin{aligned} \Delta_+ &= \Omega(z_+), \\ \Delta_- &= \Omega(z_-), \\ z_+ &= z_- + \tau(T), \\ \Omega(z) &= \int_{-\infty}^z e^{-h(z)} dz, \\ \log \frac{F(T)}{K} &= h(z_+) - h(z_-) = h(z_- + \tau(T)) - h(z_-), \end{aligned}$$

where $\tau(T)$ is non-decreasing and $h(z)$ is convex with property $\lim_{z \rightarrow \pm\infty} h'(z) = \pm\infty$. These conditions help keep Δ_- as a distortion function maintaining no-arbitrage. Along with this, it is required that $\int_{-\infty}^{\infty} e^{-h(z)} dz = 1$ since $\Omega(c)$ is a probability distribution function.

A three-step procedure to find the price of the call option given strike and maturity is used:

1. We first find $\tau(T)$ and solve for z_- .
2. Using z_- , we find z_+ .
3. Use this result to compute Ω_- and Ω_+ .

With this, we have all that we require to price the call shown in (3.5).

The two functions mentioned have a relationship to each dimension of the volatility. We thus have a parameterization able to fit a wide variety of surfaces.

For the Black-Scholes model, Carr and Pelts (2015a) shows the following:

$$\tau(T) = \sigma\sqrt{T}, \quad (3.6)$$

$$\Omega(z) = \Phi(z), \quad (3.7)$$

$$h(z) = \frac{z^2}{2} + \frac{1}{2}\log(2\pi). \quad (3.8)$$

A generalization of this is formulated by Antonov *et al.* (2019) where a grid of fixed points z_j is used. Here $j = 0, \dots, 2N$ with $z_0 = -\infty$, $z_N = 0$, and $z_{2N} = \infty$. Variables of $\gamma_j > 0$, α_j and β_j is also needed in the formulation. First, our choice of h function should be defined. This choice is purely a choice but a quadratic one is proposed in Antonov *et al.* (2019), namely

$$h(z) = \alpha_j + \beta_j(z - z_{|j|}) + \frac{(z - z_{|j|})^2}{2\gamma_j} \quad (3.9)$$

where

$$|j| = (j + 1) \times 1_{j < N} + j \times 1_{j \geq N}.$$

for $z \in [z_j, z_{j+1}]$. This choice of $h(z)$ allows for $\Omega(\cdot)$ to be related to the Gaussian distribution allowing for the setup of the problem to resemble that of the Black-Scholes model.

For $h(z)$ to be continuous and differentiable the sets α_j and β_j are calculated by the following recursive procedure:

$$\alpha_{j-1} = \alpha_j + \beta_j(z_j - z_{j+1}) + \frac{(z_j - z_{j+1})^2}{2\gamma_j},$$

$$\beta_{j-1} = \beta_j + \frac{z_j - z_{j+1}}{\gamma_j},$$

for $j < N$ and

$$\alpha_{j+1} = \alpha_j + \beta_j(z_{j+1} - z_j) + \frac{(z_{j+1} - z_j)^2}{2\gamma_j},$$

$$\beta_{j+1} = \beta_j + \frac{z_{j+1} - z_j}{\gamma_j},$$

for $j \geq N$. Note that in [Antonov et al. \(2019\)](#) this is $j > N$ but this would result in a missing value of α_{N+1} and β_{N+1} .

The recursive procedures' stopping conditions are

$$\alpha_N = \alpha_{N-1} = \alpha \text{ and}$$

$$\beta_N = \beta_{N-1} = \beta.$$

Since the convexity of $h(z)$ should be preserved, γ_j values have to be strictly positive forcing the second derivative $h(z)$ to always be positive.

For the Ω function, we have

$$\Omega_0 = 0 \text{ and}$$

$$\Omega_{j+1} = \Omega_j + \sqrt{2\pi\gamma_j} e^{\frac{1}{2}\gamma_j\beta_j^2 - \alpha_j} \left[\Phi\left(\frac{z_{j+1} - z_{|j|}}{\sqrt{\gamma_j}} + \sqrt{\gamma_j}\beta_j\right) - \Phi\left(\frac{z_j - z_{|j|}}{\sqrt{\gamma_j}} + \sqrt{\gamma_j}\beta_j\right) \right]$$

We then set

$$\alpha = \log\Omega_{2N} \tag{3.10}$$

resulting in all α_j being multiplied by α and in turn all Ω_j by $e^{-\alpha}$. This is turn allow for $\Omega_{2N} = 1$ making $\Omega(z)$ a probability distribution. It is clear that

$$\Omega(z) = \Omega_j + \sqrt{2\pi\gamma_j} e^{\frac{1}{2}\gamma_j\beta_j^2 - \alpha_j} \left[\Phi\left(\frac{z - z_{|j|}}{\sqrt{\gamma_j}} + \sqrt{\gamma_j}\beta_j\right) - \Phi\left(\frac{z_j - z_{|j|}}{\sqrt{\gamma_j}} + \sqrt{\gamma_j}\beta_j\right) \right] \tag{3.11}$$

for $z \in [z_j, z_{j+1}]$.

For τ , this is defined piecewise on intervals between maturities. If we consider maturities T_i with $i = 0, 1, \dots, N_T$ and corresponding volatilities σ_i where $T_0 = 0$, T_1 is the first maturity and T_{N_T} is the last maturity, the τ function for timeslice $t \in [T_i; T_{i+1}]$ is as follows:

$$\tau(t) = \sqrt{\sum_{k=1}^{i-1} [\sigma_k^2(T_k - T_{k-1})] + \sigma_i^2(t - T_{i-1})}.$$

Carr-Pelts needs a z_- value satisfying

$$\log \frac{F(T)}{K} = h(z + \tau(T)) - h(z). \tag{3.12}$$

Since our h and Ω functions are defined piecewise, we need to not only find z but the interval to which it belongs. This is done as follows:

1. First assume $z \in [z_{j-1}, z_j]$ for each $j = 1, \dots, 2N$.
2. Find $k = \min\{z_n > z_{j-1} + \tau(T)\}, \dots, \min\{z_n > z_j + \tau(T)\}$ such that $z + \tau(T) \in [z_{k-1}, z_k]$.
3. Solve for (3.12) restricting z to $z \in [z_{j-1}, z_j]$ and $z + \tau(T)$ to $z \in [z_{k-1}, z_k]$. Repeat the process if no root is found in these intervals.

3.3 Ensemble Carr-Pelts

As mentioned, Carr-Pelts is flexible but [Antonov et al. \(2019\)](#) improves on this flexibility with the Ensemble Carr-Pelts (ECP). This is done by taking mixture of Carr-Pelts surfaces by taking a linear combination of these Carr-Pelts surfaces. This leads to the ECP parameterization with

$$\begin{aligned}
C(K, T) &= D(T)(F(T)\Delta_+ - K\Delta_-), \\
\Delta_+ &= \sum_j w_j \Omega_j(z_{+j}), \\
\Delta_- &= \sum_j w_j \Omega_j(z_{-j}), \\
z_{+j} &= z_{-j} + \tau_j(T), \\
\Omega_j(z_j) &= \int_{-\infty}^z e^{-h_j(z_j)} dz, \\
\log \frac{F(T)}{K} &= h_j(z_{+j}) - h_j(z_{-j}) = h_j(z_{-j} + \tau_j(T)) - h_j(z_{-j}),
\end{aligned}$$

where

$$\sum_j w_j = 1, \text{ and } w_j > 0.$$

Chapter 4

Calibration

Calibration in [Jäckel and Kahl \(2008\)](#) attempts to find model parameters which minimize squared error between observed data and model data with a set of strikes K_i for $i = 1, \dots, N_K$ and a set of maturities T_i for $i = 1, \dots, N_T$

$$\sum_{i,j=1} = (\text{MarketData}_{K_i,T_j} - \text{ModelOutput}_{K_i,T_j})^2.$$

This chapter uses this approach to calibrate each of the parameterizations discussed in Chapter 3 to three different implied volatility surfaces with varying complexity. We first consider flat volatility and only once we are happy with the calibration do we move forward. We then look at an implied volatility surface which is still fixed in strike but changing in maturity. Again, only once we are satisfied with this fit do we move to a more complicated surface which is one that is changing in both strike and maturity.

In the paper, [Homescu \(2011\)](#) this approach is altered into a weighting of squared errors with smaller weights being placed on shorter-dated maturities in an attempt not to overfit short-dated data points. Suggested weights include using differences in the bid and ask spreads of quoted implied volatilities in the denominator as well as other variations.

4.1 Stochastic Volatility Inspired

As stated in [Gatheral and Jacquier \(2014\)](#), market data calibration is done on prices of the call options rather than the implied volatility itself. The prices are calculated on the mid-implied volatility of bid and offer spreads by use of the Black-Scholes

formula for a call option

$$C(K, T) = S_0 N(d_1) - K e^{-rT} N(d_2), \quad (4.1)$$

$$d_1 = \frac{\ln(S_0/K) + rT + w(k, \chi)/(2T)}{w(k, \chi)}, \quad (4.2)$$

$$d_2 = d_1 - w(k, \chi), \quad (4.3)$$

where S_0 is the underlying's price, K is the strike, r is the risk-free rate, T is the maturity of the option, $N(\cdot)$ is the normal cumulative distribution function (CDF) and $w(k, \chi)$ is the SVI parameterization.

SVI does include a non-linear constraint, inequality (3.1), which will have to be dealt with carefully when considering the fitting algorithm. One cannot set bounds on non-linear constraints in an optimization problem. Bounds help for enforcing constraints for each iteration (update to the parameter set). Since each iteration does not check the inequality constraint (3.1), this could result in a negative total variance on an iteration and Black-Scholes breaks down since probabilities are always positive. To get around this we can force the parameterization into positivity when the function is called.

We calibrate SVI to market data on each time-slice since it is defined as a curve and not a surface. There are two important notes to make about the setup and calibration with regards to the SVI surface built in this dissertation. Firstly, [Gatheral and Jacquier \(2014\)](#) proposes a method for the surface to be built in such a way to eliminate arbitrage. As mentioned by [Antonov et al. \(2019\)](#), this results in a reduction of parameters from five to three. We lose calibration quality as a result. For this dissertation, an arbitrage-free surface for SVI is not considered as [Antonov et al. \(2019\)](#) explicitly states refraining from creating such a surface as it leads to a lesser quality fit.

Secondly, [Antonov et al. \(2019\)](#) does not state their calibration approach; as a result, an initial assumption was made that calibration should be done on market prices. They do not, however, enforce positivity which should imply that our calibration should not be done to market prices but on total variance directly to reproduce their results. This is obvious since by looking at the Black-Scholes equation (4.1), we see that negative total variance in (4.2) and (4.3) leads to complex numbers, and the normal CDF breaks down. Clearly, [Antonov et al. \(2019\)](#) calibrates to total variance instead, which is the calibration approach used in this dissertation for SVI.

SVI suffers from a poor initial guess and a failure to perform well at short-dated maturities. We ignore the first issue, as this is tedious and not in-line with the goal of this dissertation, and take advantage of the latter for better calibration. With this

in mind and considering SVI is calibrated to each time-slice; we can break down our calibration into two parts: One involving the longest dated-maturity and the other involving all other maturities.

We first calibrate to the furthest maturity with an initial guess of zero for all five parameters. The parameters producing the best fit here are then used as an initial guess for the second furthest maturity. We work our way forward to shorter-dated maturities using the previously calibrated parameter set as our initial guess. In this fashion, we build up our surface calibrating to each time-slice resulting in $5 \cdot N_T$ parameters where N_T is the number of maturities.

4.2 Carr-Pelts

When calibrating, an intuitive initial guess is of high importance and can reduce calibration time significantly. We consider the initial guess suggested by [Antonov et al. \(2019\)](#) in which $\gamma_j = 1$ for all $j = 1, \dots, 2N$, $\alpha = 0$ and $\beta = 0$ while stripping σ_i from market at-the-money implied volatilities.

To obtain the σ_i values the following recursive procedure is used: Given maturities $\{T_i\}$ and corresponding market observed implied volatilities $\sigma_{imp}(T_i)$ where $i = 1, \dots, N_T$ we have

$$(T_i - T_{i-1}) \cdot \sigma_i = T_i \cdot \sigma_{imp}(T_i) \quad \text{for } i = 1$$

$$(T_i - T_{i-1}) \cdot \sigma_i = T_i \cdot \sigma_{imp}(T_i) - \sum_{k=1}^{i-1} (T_k - T_{k-1}) \cdot \sigma_k \quad \text{for } i = 2, \dots, N_T$$

allowing for τ to be piecewise constant matching (3.6) in the Black-Scholes setup at and between each maturity.

Looking at the h and Ω functions for this initial guess we find that it is precisely the Black-Scholes model's h and Ω given by (3.8) and (3.7). Since the Black-Scholes model would assume the same volatility across all maturities and strikes, it is clear that the initial guess proposed in [Antonov et al. \(2019\)](#) should fit flat volatility, without the need for calibration. Appendix A.1 shows this result.

4.3 Ensemble Carr-Pelts

A 3-factor ECP is considered (3 Carr-Pelts surfaces) for this dissertation matching the amount in the paper [Antonov et al. \(2019\)](#). This results in a total number of parameters calculated as $3 \cdot (18 + 1 + 11)$. The 3 represents the factor number, 18 the number of strikes, 18 the number of maturities and the 1 being the weighting on the surface.

For CP and ECP, we now have our parameter set, $\{\gamma_0, \dots, \gamma_{2N}, w, \sigma_1, \dots, \sigma_{N_T}\}$.

4.4 Results

We have shown that our initial guess in CP and ECP fit flat volatility almost perfectly, shown in Appendix A.1. No calibration was done here besides for SVI, proving the setup of an initial guess put forth by Antonov *et al.* (2019) fits Black-Scholes well for at-the-money volatilities. Considering that the initial guess does not cater to smiles, which need calibrating; we next consider fitting implied volatility surfaces which are not flat. The layout of the volatility surface is as follows:

- $K = 1238, 1486, 1733, 1981, 2229, 2476, 2724, 2972, 3219, 3467, 3715$ and
- $T = 0.05, 0.12, 0.22, 0.30, 0.37, 0.62, 0.87, 1.14, 1.39, 1.89, 2.39, 3.38, 4.38, 5.38, 6.38, 7.39, 8.39, 9.39$.

Firstly, an implied volatility surface which is unchanging in maturity but changing in the strike is considered. A model which satisfies the Schroder (1989) constant elasticity of variance (CEV) model whose stock price process is given by

$$dS_t = rS_t dt + \sigma S_t^\alpha dW_t. \quad (4.4)$$

We can obtain prices under this model easily since the solution is known and shown in Appendix A.2. Following this price generation approach for CEV with parameters ($\sigma = 0.6, \alpha = 0.85, r = 0.06$), the implied volatility surface is backed out from the pricing surface. The plot in Figure 4.1 highlights an implied volatility surface which is changing only in the strike, as desired.

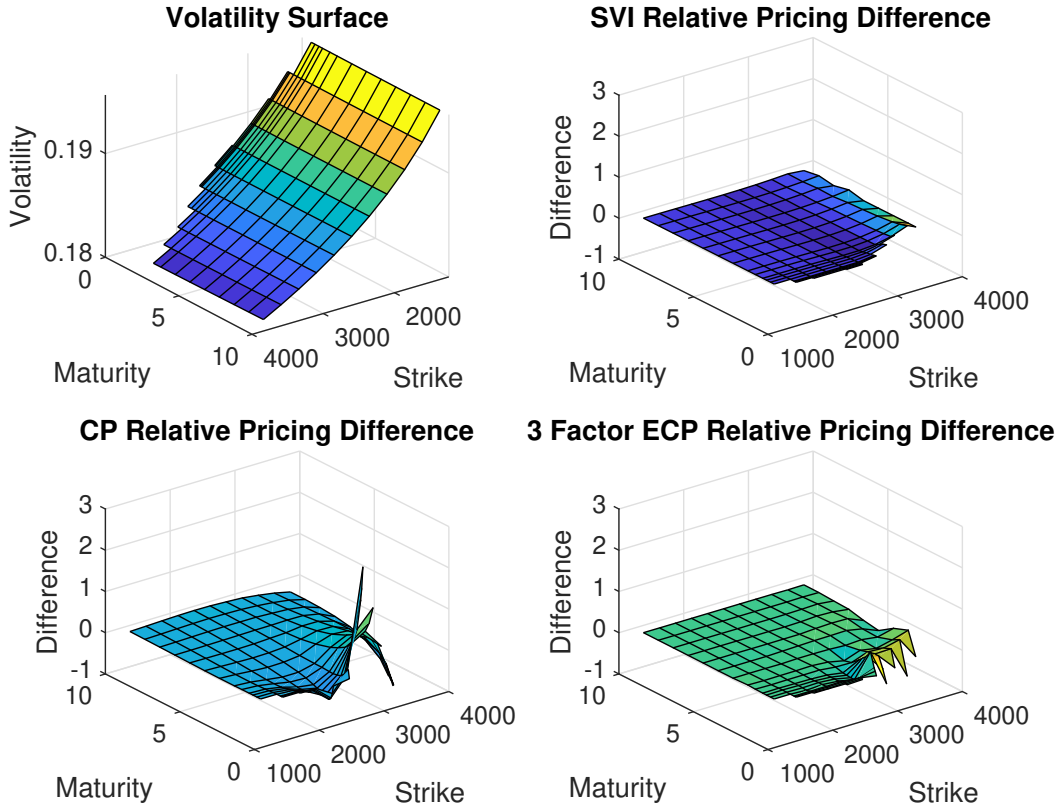


Fig. 4.1: Percentage relative pricing differences under the CEV model

	Min	Mean	Max	Runtime (seconds)
SVI	7.834e-11	0.0218	0.2345	3.2792
CP	2.0512e-7	0.2134	2.5837	794.3239
ECP	1.0136e-8	0.0447	0.6257	1495.6848

Tab. 4.1: Numerical results showing error of CEV calibration

Both CP and ECP had identical fits for flat volatility. Moving away from flat volatility, we find ECP beginning to fit more complicated options data better. We also find that ECP calibrates reasonably well compared to SVI. Considering the complexity of CP and ECP, it is of no surprise that computation time is slower than SVI.

The performance of ECP is expected to be better than CP due to how many more parameters it has. However, it should be noted that when increasing the granularity of the z-grid on CP does not result in large improvements to fit, even resulting in overfitting on low strike and high maturity.

Next, we would like to form an implied volatility surface which is changing in

both strike and maturity. To achieve this, we model prices under the [Heston \(1993\)](#) model. Prices are thus obtained in closed-form using a characteristic function. We use the “little trap” specification of the characteristic function from [Albrecher et al. \(2007\)](#) discussed in [Appendix A.3](#) which is preferred over others for its stability. Using the parameters $\{v_0 = 0.06, \kappa = 9, \theta = 0.06, \sigma = 0.4, \rho = -0.4, r = 0.06, N_m = 10, N_q = 100, u_{\max} = 80\}$ prices are generated and the implied volatility surface is obtained. Since a discretization of an integral is used we will want convergence and have the integral be properly defined. If we do not obtain convergence it means the integral does not exist, possibly due to the curve blowing up. We show the convergence of the discretization in [Appendix A.4](#). For these prices, [Figure 4.2](#) shows their corresponding implied volatility surface, which highlights the meeting of our conditions for this particular test (varying implied volatility with both strike and maturity).

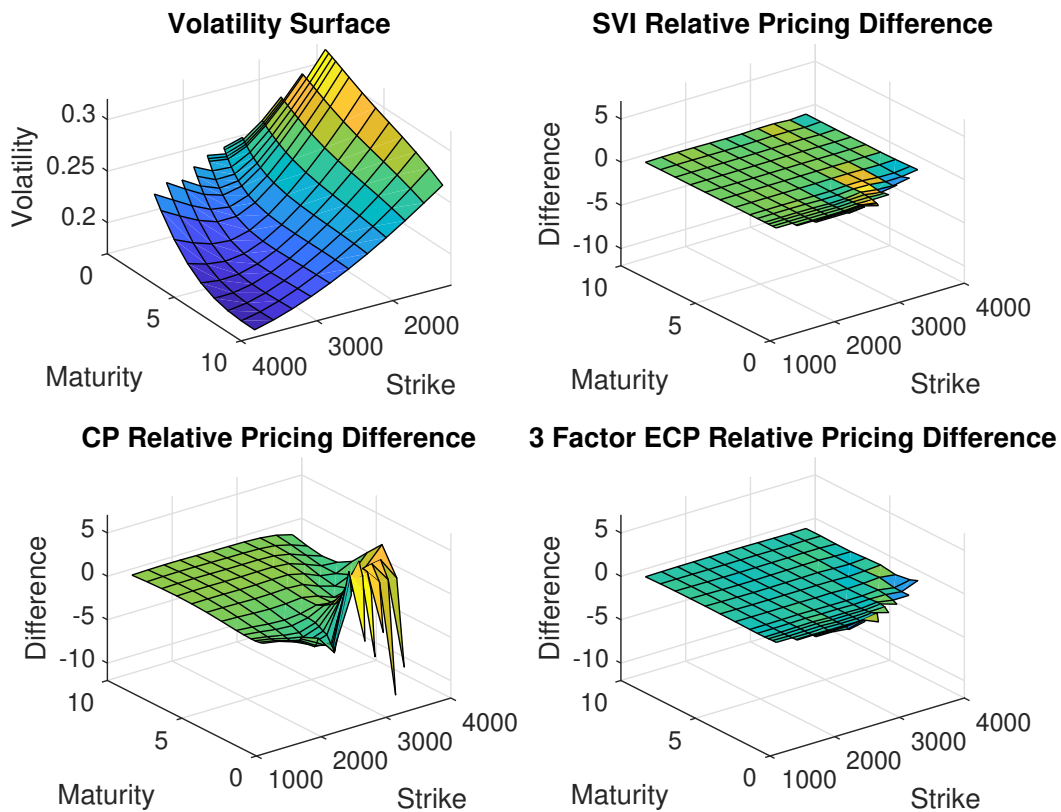


Fig. 4.2: Percentage relative pricing differences under the Heston model

	Min	Mean	Max	Runtime (seconds)
SVI	7.4767e-06	0.0363	0.3687	3.3244
CP	3.7439e-03	1.2391	11.6427	7055.0519
ECP	1.768e-04	0.0595	0.5843	37517.9062

Tab. 4.2: Numerical results showing error of Heston calibration

We see the same result for the Heston case as with the CEV case with ECP fitting better than CP showing support for ECP parameterization. This pricing structure does begin to show that CP and ECP are starting to become very computationally heavy as the data moves further from the flat implied volatility with SVI performing similar to CEV.

Runtime is far worse than CEV for this Heston case. For both methods, runtime is evenly spread between the functions with most of the runtime spent solving for (3.12) and setting up the h function.

The testing for runtime was done using the following device specifications with MATLAB 2019a:

- Processor: Intel(R) Core(TM) i3-6100 CPU @ 3.70GHz
- Installed RAM: 8.00 GB (7.87 GB usable)

Chapter 5

Local Volatility

Local volatility, as described in [Kotzé et al. \(2015\)](#), is seen as an average of instantaneous volatilities observed over the life of an option given that the underlying falls on a particular strike at maturity. This was formulated by [Dupire \(1994\)](#) as follows:

$$\sigma^2(K, T) = \frac{\frac{\partial C}{\partial T} + \mu_T K \frac{\partial C}{\partial K} + q_T C}{\frac{1}{2} K^2 \frac{\partial^2 C}{\partial K^2}}, \quad (5.1)$$

where C is the price of the call option, K and T is the strike and maturity of the option respectively, μ_T and q_T is the drift and dividend yield of the underlying respectively.

This section shows how to obtain this local volatility surface from the calibrated surfaces obtained in Chapter 4. A comparison is also done between the local volatility surfaces backed out by the parameterizations in an attempt to highlight any shortfalls and benefits of each. To highlight their effectiveness we compare them to the analytical solution for local volatility in the CEV case.

5.1 Stochastic Volatility Inspired

Our parameterization for SVI was done with respect to total variance. A formulation of the Dupire formula in this context will be ideal. In [Gatheral and Lynch \(2004\)](#) a change of variables is considered on (5.1) leading to the result we desire

$$v_L = \frac{\frac{\partial w(k; \chi)}{\partial T}}{1 - \frac{k}{w(k; \chi)} \frac{\partial w(k; \chi)}{\partial k} + \frac{1}{4} \left(-\frac{1}{4} - \frac{1}{w(k; \chi)} + \frac{y^2}{w^2(k; \chi)} \right) \left(\frac{\partial w(k; \chi)}{\partial k} \right)^2 + \frac{1}{2} \frac{\partial^2 w(k; \chi)}{\partial k^2}}. \quad (5.2)$$

Since our surface was calibrated for each maturity we can retrieve the strike derivatives of (5.2) using

$$\begin{aligned} \partial_k w(k; \chi) &= b\rho + b(k - m)(\sigma^2 + (k - m)^2)^{-1/2}, \\ \partial_k^2 w(k; \chi) &= b(\sigma^2 + (k - m)^2)^{-1/2} - b(k - m)^2(\sigma^2 + (k - m)^2)^{-3/2}. \end{aligned}$$

For the maturity derivative of (5.2) we employ the finite difference formula,

$$f'(t) \approx \frac{f(t+dt) - f(t-dt)}{2dt},$$

where as we make dt finer our estimate for the derivative of f with respect to t gets better. To find these differenced value points we have to employ an interpolation scheme on our parameterized surface.

5.2 Carr-Pelts

In the case of the Carr-Pelts parameterization calibration was done to market prices. Therefore, logically, we would use the Dupire formula from equation (5.1).

Following [Antonov et al. \(2019\)](#), we consider the first order strike derivative

$$\begin{aligned} \partial_K C &= \partial_K [D(T)(F(T)\Delta_+ - K\Delta_-)] \\ &= [D(T)[F(T)\partial_K \Delta_+ - K\partial_K \Delta_-]] \\ &= D(T)[F(T)\Omega'(z_+)\partial_K z_+ - K\Omega'(z_-)\partial_K z_-] \\ &= D(T)[F(T)e^{-h(z_+)}\partial_K z_+ - Ke^{-h(z_-)}\partial_K z_-] \\ &= D(T) \left([F(T)e^{-h(z_+)} - Ke^{-h(z_-)}]\partial_K z_- - \Delta_- \right) \\ &= D(T) \left([F(T)e^{-h(z_+)} - Ke^{-h(z_-)}]\partial_X z_- \left(-\frac{1}{K} \right) - \Delta_- \right) \\ &= D(T) ([e^{-h(z_-)} - e^X e^{-h(z_+)}]\partial_X z_- - \Delta_-) \\ &= D(T) ([e^{-h(z_-)} - e^{h(z_+) - h(z_-)} e^{-h(z_+)}]\partial_X z_- - \Delta_-) \\ &= -D(T)\Delta_-. \end{aligned}$$

Then, for the second order strike derivative, we have

$$\partial_K^2 C = \frac{D(T)}{K} \Omega'(z_-) \partial_X z_-. \quad (5.3)$$

To find $\partial_X z_-$ we note that

$$1 = [h'(z_- + \tau) - h'(z_-)] \partial_X z_-,$$

and substitute this into (5.3) to obtain

$$\partial_K^2 C = \frac{D(T)e^{-h(z_-)}}{K(h'(z_- + \tau) - h'(z_-))}.$$

Lastly, we have the first order maturity derivative

$$\begin{aligned} \partial_T C &= D(T)(r_T K \Omega(z_-) - q_T F(T) \Omega(z_+) + F(T) \Omega'(z_+) \partial_T z_+ - K \Omega'(z_-) \partial_T z_-) \\ &= D(T) \left(r_T K \Omega(z_-) - q_T F(T) \Omega(z_+) + F(T) e^{-h(z_+)} \tau'(T) \right), \end{aligned}$$

Substituting these derivatives into (5.1) we get

$$\begin{aligned}
\sigma^2(T, K) &= \frac{r_T K \Omega(z_-) - q_T F(T) \Omega(z_+) + F(T) e^{-h(z_+) \tau'(T)}}{\frac{K e^{-h(z_-)}}{2(h'(z_+) - h'(z_-))}} \\
&+ \frac{-\mu_T K \Omega(z_-) + q_T F(T) \Omega(z_+) - q_T K \Omega(z_-)}{\frac{K e^{-h(z_-)}}{2(h'(z_+) - h'(z_-))}} \\
&= 2\tau'(T)(h'(z_+) - h'(z_-)) \\
&= \frac{(\tau^2(T))'(h'(z_+) - h'(z_-))}{\tau(T)}.
\end{aligned}$$

In the generalization of the Carr-Pelts surface described in Chapter 3 we have

$$\begin{aligned}
(\tau^2(t))' &= \sigma_i^2 \quad \text{for } T_{i-1} < t \leq T_i, \\
h'(z) &= \beta_j z + \frac{z - z_j}{\gamma_j} \quad \text{for } z \in [z_j, z_{j+1}].
\end{aligned}$$

For the Black-Scholes case $(h'(z_+) - h'(z_-)) = z_- + \tau(T) - z_- = \tau(T)$ resulting in $\sigma^2(T, K) = \sigma^2(T)$ which is just the Black-Scholes volatility. We expect this by the definition of local volatility.

5.3 Ensemble Carr-Pelts

We follow the same method as with Carr-Pelts and arrive at the following:

$$\begin{aligned}
\sigma^2(T, K) &= \frac{F(T) \sum_j w_j e^{-h_j(z_{+j})} \tau_j'(T)}{K \sum_j w_j \frac{e^{-h_j(z_{-j})}}{2(h_j'(z_{+j}) - h_j'(z_{-j}))}} \\
&= \frac{F(T) \sum_j w_j e^{-h_j(z_{+j})} \frac{(\tau^2(T))'}{\tau(T)}}{K \sum_j w_j \frac{e^{-h_j(z_{-j})}}{(h_j'(z_{+j}) - h_j'(z_{-j}))}} \\
&= \frac{e^X \sum_j w_j e^{-h_j(z_{+j})} \frac{(\tau^2(T))'}{\tau(T)}}{\sum_j w_j \frac{e^{-h_j(z_{-j})}}{(h_j'(z_{+j}) - h_j'(z_{-j}))}} \\
&= \frac{\sum_j w_j e^{-h_j(z_{-j})} \frac{(\tau^2(T))'}{\tau(T)}}{\sum_j w_j \frac{e^{-h_j(z_{-j})}}{(h_j'(z_{+j}) - h_j'(z_{-j}))}},
\end{aligned}$$

where again

$$\begin{aligned}
(\tau^2(t))' &= \sigma_i^2 \quad \text{for } T_{i-1} < t \leq T_i, \\
h'(z) &= \beta_j z + \frac{z - z_j}{\gamma_j} \quad \text{for } z \in [z_j, z_{j+1}].
\end{aligned}$$

5.4 Results

For the CEV model we know what the local volatility is. By simple inspection of (4.4) it is obvious that

$$\sigma(t, S_t) = \sigma S_t^{(\alpha-1)}.$$

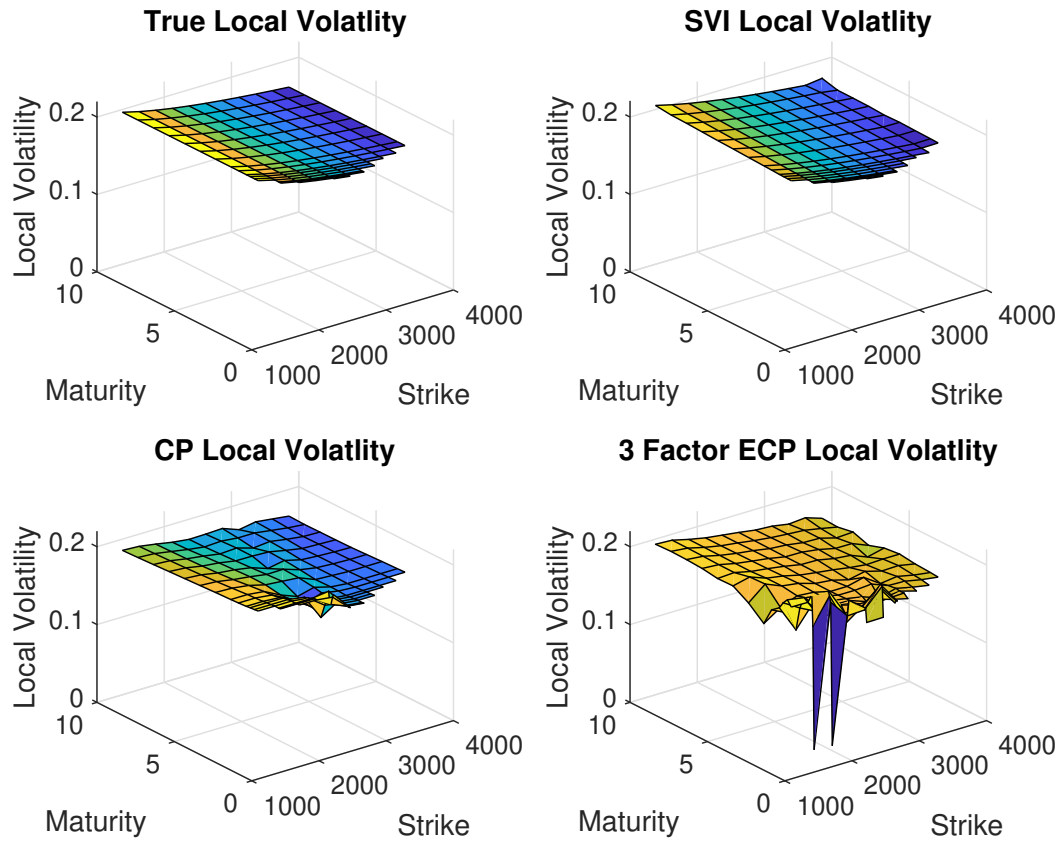


Fig. 5.1: Local volatility surface under CEV compared to those of SVI, CP and ECP brought out by our calibration

	Min	Mean	Max
SVI	2.4257e-06	0.0034	0.0113
CP	1.7346e-07	0.0035	0.0194
ECP	1.3295e-06	0.0061	0.1750

Tab. 5.1: Numerical results showing error differences in local volatility

Given our calibration, these values are to be expected. SVI has a noticeable increase in the surface with increasing maturity and is subject to interpolation issues.

SVI is defined for each time-slice, making it reasonable for the local volatility to suffer from changing maturity. ECP and CP fail at lower end strikes and maturities. We note that at these points we find our z points solving equation (3.12) falling outside the range of our z -grid.

This also highlights a benefit to parameterizing the implied volatility surface. A suitable parameterization of implied volatility leads to a local volatility surface without the need to calibrate. This is usually done by the use of a finite difference scheme, as shown in [Andreasen and Huge \(2010b\)](#), where the local volatility surface is treated as a parameter set in calibration to market prices.

To help with the lower maturity and strike issue seen in Figure 5.1 for ECP we set up a larger grid for our z -grid and consider $\{-\infty, -1000, -100, -10, -1, 0, 1, 10, 100, 1000, \infty\}$ instead of $\{-\infty, -5, -1, -0.5, -0.1, 0, 0.1, 0.5, 1, 5, \infty\}$ which was considered previously. This produces a new set of results shown in Figure 5.2 below.

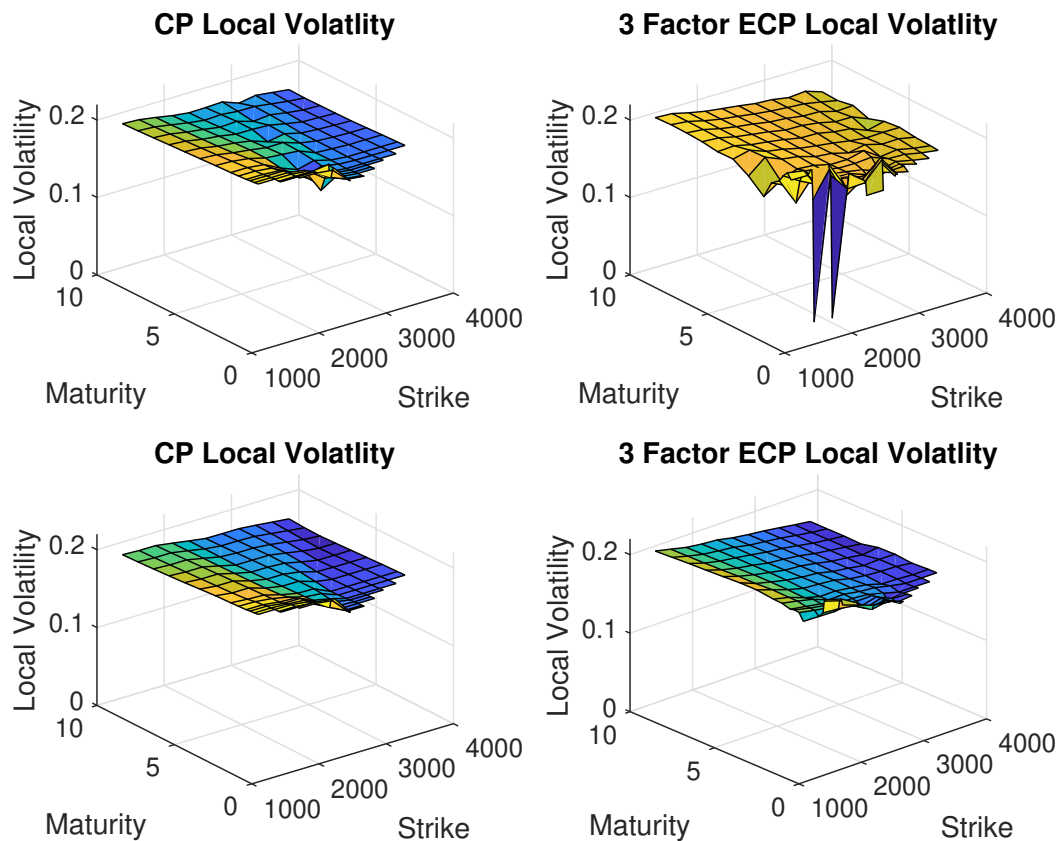


Fig. 5.2: Local volatility surface under CEV for CP and ECP with the smaller z -grid at the top and the larger z -grid below it

	Min	Mean	Max
CP	1.3718e-07	0.0031	0.0178
ECP	2.6765e-06	0.0022	0.0185

Tab. 5.2: Numerical results showing error differences in local volatility for a larger z -grid

We see the results for ECP to be far more satisfactory in figure 5.2. This does come at a cost, however. The larger z -grid leads to worse calibration quality with an increase in average percentage relative pricing difference from 0.044688 to 0.091034 for ECP.

We can also test any local volatility surface by use of finite difference as described in [Andreasen and Huge \(2010b\)](#). This can be extended to local stochastic volatility (LSV) where [Gyöngy \(1986\)](#) theorem and a finite difference scheme can be employed as shown in [Andreasen and Huge \(2010a\)](#) and [de Graaf \(2012\)](#) to find these local stochastic volatilities at set points as well as testing of these LSV values. Both these finite difference methods should return whatever prices are initially used in the parameterization of implied volatility.

Chapter 6

Conclusion

We tested three parameterizations of the implied volatility surface for three different surfaces of increasing complexity. Comparisons between the three surface parameterizations were done for both implied volatility surface calibration as well as a local volatility surface. All testing was done to best mimick the same environment as in [Antonov *et al.* \(2019\)](#), particularly their strike distribution and maturity layout resulting in the same number of parameters mentioned in the paper.

No testing was done on the actual implied volatility surface, only on the corresponding prices. Extracting these implied volatilities from the pricing surface seemed very sensitive to low strikes and maturities as well as the initial guess used in the `fsolve` and `fzero` functions from MATLAB. Experiments could be run on other methods of obtaining implied volatility such the `bslimpv` function in MATLAB. Further research can be done on a weighting of different maturities in our objective function which could alleviate overfitting at these lower end maturities.

There is a fast calibration scheme mentioned in [Antonov *et al.* \(2019\)](#), which involves the fixing of the h parameters. CP and ECP use at the money implied volatilities as an initial guess for the τ parameters. Since we only calibrating σ_i no optimization is needed since this guess fits perfectly for at-the-money strike points. We are thus left with a surface which is changing in maturity and constant in the strike. If we see extreme smiles in the market, the performance of the ECP will decrease greatly as we move further away from the at-the-money strike.

Building the Ω function had its problems for implementation. MATLAB often returned NaN values for certain input parameters which are to do with how MATLAB handles infinity multiplied by zero. This results in Ω_j values being NaN and cascades into an unsolvable problem for (3.12) and an error is thrown. To deal with this issue, we use exception handling, which allows the optimization to continue.

We first considered the simplest case of flat volatility where CP and ECP performed identically, as expected. This was also used as a check to see if our CP parameterization was working correctly. Moving to a more complicated pricing sur-

face, ECP begins to calibrate better than CP. ECP's calibration ability was seen as adequate when comparing to SVI. As the surface got more complicated, we found SVI starting to calibrate better than ECP. This is evidenced by ECP performing worse than SVI in the Heston case while performing similarly in the CEV case.

When considering local volatility, all three parameterizations performed similarly barring the ECP which performed poorly for very low strikes and short maturities. This was a result of ECP solving (3.12) for z out of the range of the specified z -grid, disregarding end infinity points. This problem can be avoided by setting a large enough z -grid; however, calibration quality is lost by doing so. A simple interpolation scheme for SVI led to negative local volatilities closer to maturity when testing on the Heston prices. This is not present in ECP and CP as a result of their piecewise construction.

Better fitting on ECP over CP should also be scrutinized as ECP had a much larger number of parameters. This not only increases computation time but also increases chances for overfitting of the volatility surface which stands in contrast to the added flexibility proposed by the method. Increasing the z -grid for CP only resulted in overfitting.

Taking note that ECP is an arbitrage-free parameterization, these results are impressive. Especially when finding that ECP calibrates almost as well as SVI, which did not have arbitrage-free conditions enforced. However, we do have to note that calibration time is far longer with ECP than SVI and will have to be run overnight.

For further investigations into the plausibility of the method itself, a wider range of surface constructions should be considered besides SVI as with the dissertation. A better examination into the setup of the z -grid would be highly beneficial in elevating any unwanted results such as with the local volatility fit in this dissertation on ECP.

Bibliography

- Albrecher, H., Mayer, P., Schoutens, W. and Tistaert, J. (2007). The little Heston trap, *Wilmott* (1): 83–92.
- Andreasen, J. and Huge, B. N. (2010a). Finite difference based calibration and simulation, *Available at SSRN 1697545* .
- Andreasen, J. and Huge, B. N. (2010b). Volatility interpolation, *Available at SSRN 1694972* .
- Antonov, A., Konikov, M. and Spector, M. (2019). A new arbitrage-free parametric volatility surface, *Available at SSRN 3403708* .
- Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities, *The Journal of Political Economy* **81**(3): 637–654.
- Breeden, D. T. and Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices, *Journal of Business* **51**(4): 621–651.
- Carr, P. and Madan, D. B. (2005). A note on sufficient conditions for no arbitrage, *Finance Research Letters* **2**(3): 125–130.
- Carr, P. and Pelts, G. (2015a). Duality, deltas, and derivatives pricing, *Presented at the Conference Dedicated to Steve Shreve's 65th Birthday* .
- Carr, P. and Pelts, G. (2015b). Game of vols, *Available at SSRN 3422540* .
- de Graaf, C. (2012). Finite difference methods in derivatives pricing under stochastic volatility models, *Master's Thesis, Leiden University* .
- Dupire, B. (1994). Pricing with a smile, *Risk* **7**(1): 18–20.
- Gatheral, J. (2004). A parsimonious arbitrage-free implied volatility parameterization with application to the valuation of volatility derivatives, *Presentation at Global Derivatives & Risk Management, Madrid* .
- Gatheral, J. and Jacquier, A. (2011). Convergence of heston to SVI, *Quantitative Finance* **11**(9): 1129–1132.
- Gatheral, J. and Jacquier, A. (2014). Arbitrage-free SVI volatility surfaces, *Quantitative Finance* **14**(1): 59–71.

- Gatheral, J. and Lynch, M. (2004). Lecture 1: Stochastic volatility and local volatility, *Case Studies in Financial Modeling Notes, Courant Institute of Mathematical Sciences* .
- Gyöngy, I. (1986). Mimicking the one-dimensional marginal distributions of processes having an Itô differential, *Probability Theory and Related Fields* **71**(4): 501–516.
- 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.
- Homescu, C. (2011). Implied volatility surface: Construction methodologies and characteristics, *Available at SSRN 1882567* .
- Jäckel, P. and Kahl, C. (2008). Hyp hyp hooray, *Wilmott Magazine* **34**: 70–81.
- Kotzé, A., Oosthuizen, R. and Pindza, E. (2015). Implied and local volatility surfaces for south african index and foreign exchange options, *Journal of Risk and Financial Management* **8**(1): 43–82.
- Schroder, M. (1989). Computing the constant elasticity of variance option pricing formula, *The Journal of Finance* **44**(1): 211–219.

Appendix A

Appendix

A.1 Flat Volatility

This is the graphic comparison between the three different parameterizations calibrating to the flat volatility mentioned in Chapter 4.

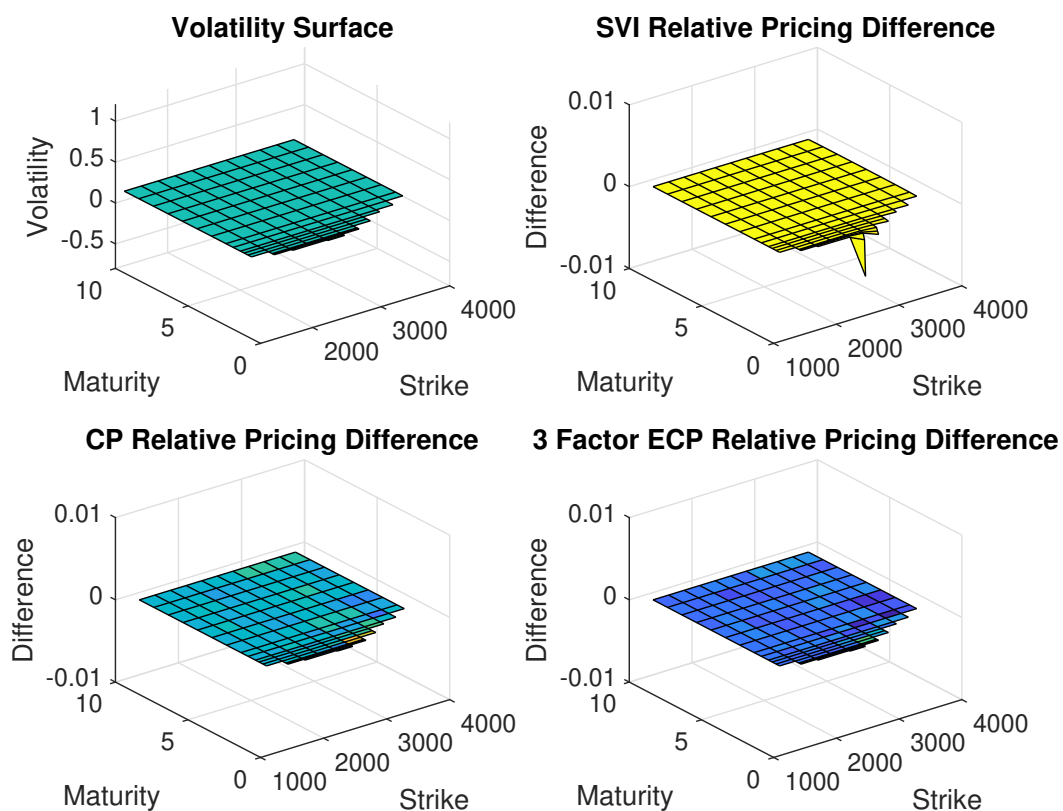


Fig. A.1: Relative pricing differences under flat volatility for $\sigma = 0.2$

A.2 Option Pricing Under the CEV Model

European call options written on a underlying modeled with the CEV model can be priced with the following formula from [Schroder \(1989\)](#):

$$C(S_0, K, \sigma, \alpha, T) = S_0[1 - \mathcal{P}(y; z, x)] - Ke^{-rT}\mathcal{P}(x; z - 2, y),$$

where S_0, K and T are the initial stock price, strike and maturity of the option respectively, $\mathcal{P}(\cdot; d, \lambda)$ is the noncentral chi-squared distribution function with d degrees of freedom and non-centrality parameter λ with

$$x = \kappa S_0^{2(1-\alpha)} e^{2r(1-\alpha)T}, \quad y = \kappa K^{2(1-\alpha)}, \quad z = 2 + \frac{1}{1-\alpha},$$

where

$$\kappa = \frac{2r}{\sigma^2(1-\alpha)(e^{2r(1-\alpha)T} - 1)}.$$

A.3 Option Pricing Under the Heston Model

We present the “little trap” formulation of the characteristic function and how to use it in option pricing as shown in [Albrecher *et al.* \(2007\)](#). We first note that option prices can take the following form by change of numeraire:

$$\begin{aligned} C(K) &= e^{-rT} \mathbb{E}^{\mathbb{Q}}[(S_T - K)^+] \\ &= S_0 P_1 - Ke^{-rT} P_2. \end{aligned}$$

Here K, T and S_T are the strike, maturity and stock price at maturity of the option. By use of Euler’s identity, characteristic functions and discretizing the resulting integral we arrive at

$$\begin{aligned} P_1 &= \frac{1}{2} + \frac{1}{\pi} \sum_{n=1}^N \operatorname{Re} \left[\frac{e^{-iu_n k} \phi_{s_T}(u_n - i)}{iu_n \phi_{s_T}(-i)} \right] \Delta u, \\ P_2 &= \frac{1}{2} + \frac{1}{\pi} \sum_{n=1}^N \operatorname{Re} \left[\frac{e^{-iu_n k} \phi_{s_T}(u_n)}{iu_n} \right] \Delta u, \end{aligned}$$

where ϕ_{s_T} is the characteristic function of the log stock price, $u_n = (n - \frac{1}{2})\Delta u$ with integration limits on the interval $[0, u_{\max}]$ and $\Delta u = u_{\max}/N$.

In this formulation no assumption was made on the dynamics of the stock price process. The “little trap” characteristic function captures these dynamic for the Heston model shown by

$$\phi_{s_T}(u) = \exp(C + Dv_0 + iu \log(S_0)),$$

where

$$\begin{aligned} C &= rTiu + \theta\kappa \left(Tx_- - \frac{1}{\alpha} \log \left(\frac{1 - ge^{-Td}}{1 - g} \right) \right), \\ D &= \frac{1 - e^{-Td}}{1 - ge^{-Td}} x_-, \end{aligned}$$

with

$$a = \frac{\sigma^2}{2}, \quad b = \kappa - \rho\sigma iu, \quad c = -\frac{u^2 + iu}{2}, \quad d = \sqrt{b^2 - 4ac},$$

$$x_{\pm} = \frac{b \pm d}{2a}, \quad \text{and} \quad g = \frac{x_-}{x_+}.$$

A.4 Little Trap Integrand

We graph the convergence of the discretized integral of the “little trap” characteristic function formulation for the Heston model with parameters discussed in Chapter 4.

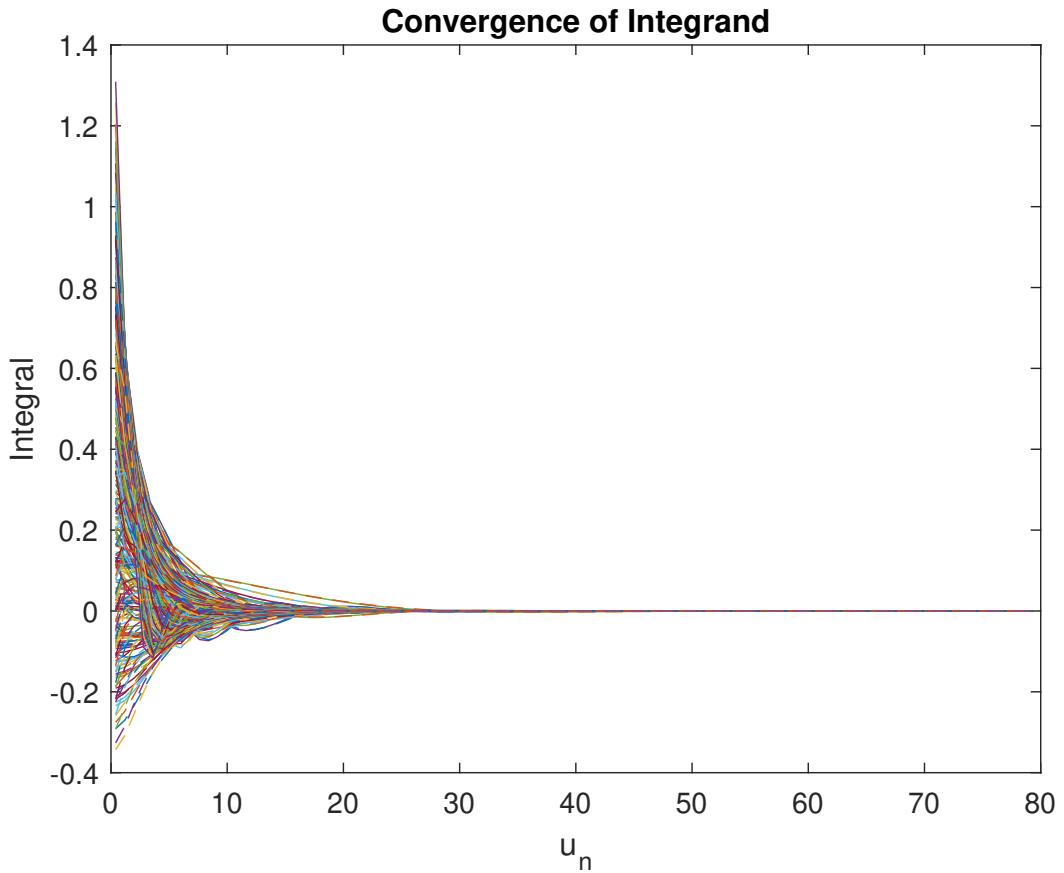


Fig. A.2: All integrands of P_1 and P_2 showing convergence.