Approximating Option Prices

Houman Younesian (337513)

Academic supervisor: Dr. M. van der Wel Co-reader: Dr. M. Jaskowski

June 20, 2013

Abstract

In this thesis we consider the method of Kristensen and Mele (2011, J. of Financial Economics) to approximate European call option prices from the constant elasticity of variance (CEV) model, for which no closed-form solution is available. The method provides a closed-form approximation that allows for direct calibration of the CEV model on option price data. Through a simulation study we find a good performance of the approximation method, both in terms of accuracy and parameter estimation. We proceed by testing the model on European call options on the S&P 500 index. The CEV model is calibrated through its closed-form approximation on option price data, using a non-linear least squares method that minimizes the sum of squared errors between the cross-sectional option price and the corresponding option price from the approximation, on a daily basis. The calibrated closed-form approximation is then used to get in-sample and out-of-sample fits of the daily option prices. When these fits are compared with those from two benchmark models, namely the Heston model and the Practitioners Black-Scholes model, the CEV model is outperformed by both. We conclude that the poor option price performance of the CEV model is due to the inaccuracy of the applied approximation.

Keywords: Option pricing; CEV model; European call options; Closed-form solution.

Contents

1 Introduction

Options are used by financial actors for the purpose of hedging and speculation. Option prices are determined using continuous-time models. The Black and Scholes (1973) (Black-Scholes, henceforth) model, specifically, has gained wide popularity among practitioners for reasons that include an elegant representation for the option price, easy applicability, and robustness. However, the model's crude assumptions (on the underlying asset) of constant volatility and log normal returns, and some of their implications, are contrary to what is observed empirically. First of all, we know that stock volatility varies over time, while the model assumes constant volatility. Second, there is the so called *leverage effect* of stock returns, which implies a negative correlation between stock returns and their volatility. This empirical feature has been reported by, for example, Black (1976). Third, there is the Black-Scholes model's implied volatility, from empirical option prices, that differs for options with different strikes. This is the so called volatility smile and is inconsistent with the model's underlying assumption, which implies that the volatility should be the same for options with different strikes. For aforementioned reasons many improvements on the model have been proposed in the literature and the original Black-Scholes model is used mostly as a benchmark model today. In this thesis, we regard a stochastic volatility model that extends the Black-Scholes model, by replacing the assumption of constant variance with a stochastic process for the variance.

The stochastic volatility model that is of our particular interest in the context of this thesis is the *constant elasticity of variance* model, or *CEV* model. The CEV model is able to cope with a lot of the stylized facts of underlying assets (like the stock index) of options, as demonstrated by, for example, Chacko and Viceira (2003). However, except in a few specific cases, there is no closed form pricing function (as there is for the Black-Scholes model) available for options from the CEV model. One specific case of the CEV model for which a closed form pricing function *is* available, is the Heston (1993) model. Under the assumptions of the Heston model, one can price European-style options using an accurate closed form pricing function in a stochastic volatility setting. Despite this, the Heston model is often rejected as an adequate model for option prices, due to shortcomings like the presence of the volatility smile and the fact that it is incapable of accounting for empirical features like the leverage effect (although the Heston model does account for the negative correlation between stock price and stock *variance*). Both the CEV and the Heston model are tested by Jones (2003), who uses a Bayesian framework from bivariate time series of S&P 100 market returns and implied volatility (VIX) data. He finds that the CEV model is able to account for the leverage effect and to a lesser extent for the volatility smile. Furthermore, as a generalization of the Heston model, the CEV model is able to account for more stylized facts of *stock returns*. However, very little is known on the *option pricing* aspect of the model.

In this thesis we regard the CEV model's application for pricing European call options. In the absence of a closed form pricing function, numerical methods like Monte-Carlo simulation are applied to determine the option price from the CEV model. Monte-Carlo simulations can be very time consuming, as one has to run many simulations to obtain an option price with decent accuracy. Also, the absence of a closed form solution makes calibration of the CEV model very cumbersome, especially when the model is calibrated directly on option price data. Another way to cope with the absence of a closed form solution of the CEV model, is to look at a closed form approximation of the solution, as proposed by Kristensen and Mele (2011).

Kristensen and Mele (2011) develop a new conceptual method to compute contingent claim prices in non-linear, multi-factor diffusion settings. One specific case is the price of a European call option with CEV model specifications for the underlying asset. The method of Kristensen and Mele (2011) (K&M method henceforth), provides a closed form approximation to the original option pricing function. In particular the K&M method provides an approximation to the (unknown) solution of a partial differential equation, of which the solution is the pricing function for a European call option from the CEV model. Once the closed form approximation is determined, a small amount of time is required to compute the option price. In their paper, Kristensen and Mele (2011) apply their method to various contingent claim prices from continuous-time models and show that it is more accurate than conventional methods, such as Monte-Carlo simulation and the finite difference method. Another method, which is mentioned by Kristensen and Mele (2011), is that of Yang (2006), who proposes a similar closed-form approximation of contingent claim prices. Kristensen and Mele (2011) show that their approximations are more accurate than Monte-Carlo simulation and the Yang method, when applied to short term maturity options from the Heston model and to some extent from the CEV model. In this thesis we elaborate on this, by also assessing the accuracy of the methods for mid and long term maturity options. Similar to Kristensen and Mele (2011), we find that the K&M method is very accurate for short term maturity options. However, the accuracy deteriorates dramatically as the maturity increases beyond the length of 1 year. The Yang (2006) method is less accurate for short term maturity options, but proves to be much more robust to maturity increases. A problem with the Yang approximation is that beyond a few corrective terms, we cannot compute the integrals that are involved and are thus left with a poor approximation that doesn't contain all the model's parameters. An important potential of the K&M approximation of the CEV model, is that it can be used for model calibration purposes; due to the accurate closed-form representation of the option price, the parameters can be estimated directly on option price data. This is not the case for the Yang approximation, as it does not include all parameters. So far the evidence of the K&M method is only based on simulated data and not on actual option price data.

The contribution of this thesis is a careful analysis of the empirical option pricing performance of the CEV model, when applied in conjunction with the K&M method. The K&M approximation of the option price function from the CEV model is used to estimate the model's parameters from empirical option price data. By doing so, the estimated model will contain information coming directly from the option market. We use prices of European call options written on the S&P500 index. To estimate the parameters we use a non-linear least squares method to imply the spot volatility, as well as the other model parameters, which, for the CEV model, are the mean reversion rate for the variance, the long run variance, the volatility of the variance, the correlation coefficient between the underlying asset return and its volatility, and the model's characteristic elasticity of variance. The non-linear least squares method is also applied by Bakshi, Cao, and Chen (1997) for the Heston model. To apply this method to a pricing function, a loss function needs to be specified. Christoffersen and Jacobs (2004) suggest "to align the estimation and evaluation loss functions". Bakshi et al. (1997) use a dollar-based loss function, namely the mean-squared absolute option pricing errors, for the estimation and evaluation. In addition, they also use a percentage-based loss function, namely the mean-squared relative option pricing errors, for the evaluation. They consider both in-sample and out-of-sample performance. Christoffersen and Jacobs (2004) show, for the Heston model (and the *Practitioners Black-Scholes model* - a slightly modified version of the Black-Scholes model), that when the evaluation is done using a percentage-based loss function, this should also be used to estimate the parameters, as it will improve the model's performance in terms of the percentage-based loss function. In our case, we are mainly interested in the out-of-sample dollar-based performance (i.e. the absolute pricing errors) of the model, therefore we leave out the percentage-based loss function; both for the estimation and for the evaluation of the model.

Although the K&M approximation does provide a closed form pricing function, it still contains an approximation error that might influence the parameter estimates. To make the estimates more stable we use the regularisation method, as suggested by Chiarella et al. (2000). By adding an extra penalty term to the loss function, for deviation from determined initial values of the parameters, Chiarella et al. (2000), suggest that the optimization problem of the *Heston model* becomes more stable. We expect the same to hold for the CEV model, as the model is also non-linear and of a similar form. The suitability of the approximation for model calibration is tested by estimating the parameters on simulated option data from the original model. The assessment is first conducted for the Heston model, for we have an accurate closed form solution, which is used to generate the data and to test how well its K&M approximation performs compared to the closed form solution. The spot prices and variances are generated using Monte-Carlo simulation. We find that the parameter estimates resulting from the calibration of the K&M approximation for the Heston model are accurate, but not as accurate as when we use the original closed form solution. This means that the approximation error does lead to deterioration in the parameter estimates, however it is nothing to serious. The errors are practically absent for the estimated spot variances. Finally, we also calibrate the approximation for another case of the CEV model, which does not have a closed form solution. In this case we generate the prices through Monte-Carlo simulation. We generate data from the model, as estimated by Jones (2003) for S&P 100 data. The Monte-Carlo simulations lead to some errors in the generated option prices, which affect the parameter estimates. Due to the errors in the generated option prices, we increase the convergence speed of the optimization scheme, by relaxing tolerance of the optimized function. Doing so, we find satisfactory estimates for the parameters of this particular case of the CEV model, when estimated with its corresponding K&M approximation.

After the K&M approximation of the CEV model is computed, we compare the empirical performance of the model to that of two benchmark models, namely the Heston model and the Practitioners Black-Scholes (PBS henceforth) model. The PBS model is a benchmark implied volatility model proposed by Dumas, Fleming, and Whaley (1998), which they refer to as the Ad-Hoc model and Christoffersen and Jacobs (2004) refer to as the Practitioners Black-Scholes model. Christoffersen and Jacobs (2004) suggest the PBS model as a benchmark for the Heston model, "both because of its simplicity and because of its use as a benchmark in the existing literature". The authors note that the PBS model is not preferred over structural models, although they show that it does outperform the Heston model for options on the S&P 500 index. We include the Heston model as a structural model benchmark and because it is a generalization of the CEV model. Doing so, we test if the generalization is beneficial for the option valuation process. Since the generalization adds another parameter, and thus more estimation uncertainty, it is natural to ask how much is to gain from this generalization. The same goes for the CEV model with respect to the PBS model, since the PBS model is a very robust and simple model. We assess the in-sample fit of each model by estimating the parameters on option prices for each day. We then use these parameters to test the 1-day and 5-day out-of-sample fits of the models. In terms of the root-means-squared error, the CEV model under-performs both the PBS and the Heston model, in-sample and both out-of-sample cases.

The rest of this thesis is organized as follows. Section 2 introduces the Black-Scholes option pricing model, the Practitioners Black-Scholes model, the Heston stochastic volatility option pricing model and the CEV stochastic volatility option pricing model. Section 3 provides al methodology used, namely the K&M method and the Yang method when applied to the CEV model, with a comparison between the two, the Monte-Carlo simulation and all estimation procedures we apply to estimate the parameters of the CEV model. Finally, Section 5 gives the application to empirical data, starting with a description of the data and a comparison of the performance of the CEV model, Heston model and the PBS model, in and out-of-sample. Section 6 concludes the thesis.

2 Stochastic Volatility Models for European Call Options

In this section we look at the various option valuation models that we use in thesis. We start with the Black-Scholes model and the Black-Scholes formula, followed by a variation of this model, namely the Practitioners Black-Scholes model. We discuss both the PBS model and its implementation. Then we discuss the Heston model and also give a pricing formula for this model. And finally, we discuss the CEV model, for which we do not have a closed-form pricing formula.

2.1 Black-Scholes Model

We start with a brief introduction of the Black-Scholes model. In the Black-Scholes model the price of an asset S_t is assumed to be the solution to the following stochastic differential equation (SDE):

$$\frac{\mathrm{d}S_t}{S_t} = r\mathrm{d}t + \sigma \mathrm{d}W_t, \tag{1}$$

where W_t is a standard Brownian motion (1) under the risk-neutral probability, r is the short term rate, taken to be constant, and σ is the asset return volatility, also taken to be constant. Consider a European call option written on this asset. The option pay-off function, at maturity time $T > 0$, is given by

$$b(S_T) \equiv \max\{S_T - K, 0\}, \tag{2}$$

where $K > 0$ is the strike price. Let $w_t = w(S_t, t; \sigma)$, denote the option price at time $t \in [0, T]$, when the the stock price is S_t with constant volatility σ. Then, at time of maturity T, the option price should satisfy the following boundary condition

$$w(S_T, T; \sigma_0) = b(S_T) \tag{3}$$

As the option price w_t is a function of the asset S_t and time t, by applying Itô's lemma for two variables, we find that it must satisfy

$$dw_t = \sigma S_t \frac{\partial w}{\partial s} \mathrm{d}W_t + (rS_t \frac{\partial w}{\partial s} + \frac{1}{2}\sigma^2 S_t^2 \frac{\partial^2 w}{\partial s^2} + \frac{\partial w}{\partial t})\mathrm{d}t, \tag{4}$$

subject to the boundary condition given by (3). As the latter expression is somewhat cumbersome, we follow Baxter and Rennie (2008) in deriving an elegant (and more common) partial differential equation (PDE) that has the pricing function of a European call option, w_t, is the solution. According to Baxter and Rennie (2008, page 89), we can find a so called replicating strategy with the same pay-off as the above European call option. That is, by investing an amount Δ_t in the stock S_t and an amount Π_t in a risk-less cash bond B_t, we obtain a portfolio of which the value at time t is equal to

$$w_t = \Delta_t S_t + \Pi_t B_t. \tag{5}$$

Baxter and Rennie (2008) state that this replicating strategy is self-financing, which means that for all t we can find a Δ_t and Π_t such that relation (5) holds. Or in other words, at *all* points in time we can redistribute the investments of our stock-bond portfolio (thus increasing the short position in stock and long position in the risk-less bond, or vice versa), such that our portfolio value is equal to the price of a European call option written on a stock S_t. Now, due to the self-financing condition, it holds that

$$\mathrm{d}w_t = \Delta_t \mathrm{d}S_t + \Pi_t \mathrm{d}B_t, \tag{6}$$

thus any changes in the value of the portfolio are due to changes in the values of the underlying assets. Furthermore, since B_t is a risk-less asset, it would see a growth of $rB_t\mathrm{d}t$ during the interval of length $\mathrm{d}t$, hence

$$\mathrm{d}B_t = rB_t\mathrm{d}t. \tag{7}$$

Substituting the latter expression for $\mathrm{d}B_t$ and expression (1) for $\mathrm{d}S_t$ in equation (6), we get

$$\mathrm{d}w_t = (\sigma\Delta_t S_t)\mathrm{d}W_t + (r\Delta_t S_t + r\Pi_t B_t)\mathrm{d}t. \tag{8}$$

Since SDE representations are unique, the volatility terms in equation (4) and (8) must match, which gives:

$$(\sigma\Delta_t S_t) = \sigma S_t \frac{\partial w}{\partial s} \Rightarrow \Delta_t = \frac{\partial w}{\partial s}, \tag{9}$$

thus the amount of stock in the replicating portfolio at any time t is the derivative of the option price with respect to the stock price. Now, if we match the drift terms (with $\mathrm{d}t$) of equations (4) and (8) and use $\Delta_t = \frac{\partial w}{\partial s}$ and $w_t = \Delta_t S_t + \Pi_t B_t$, we get the PDE of w as

$$\frac{1}{2}\sigma^2 s^2 \frac{\partial^2 w}{\partial s} + rs\frac{\partial w}{\partial s} + \frac{\partial w}{\partial t} - rw = 0. \tag{10}$$

Loosely speaking, the above expression states that growth in the option price w, due to its time value (the term with derivative to t) and the underlying stock (the terms with derivative to s), is expected to match the growth in a risk-less investment of amount w (the last term). The solution of this PDE, coupled with the boundary condition (3), gives the pricing function $w(S_t, t; \sigma)$ of a European call option, written on the asset satisfying equation (1). This pricing function $w(S_t, t; \sigma)$ is knowns as the Black-Scholes pricing function for a European call option.

We do not derive the solution of the latter equation (10), but solely state it: a derivation can be found in Baxter and Rennie (2008). The pricing equation for European Call options, denoted by $C_{BS}(S, t, T, \sigma)$ (we use this notation for the price of a European call option - instead of $w(S_t, t; \sigma)$ - as is common in the existing literature) is given by

$$C_{BS}(S, t, T, \sigma) = N(d_1) - N(d_2)K\exp(-r(T-t)), \tag{11}$$

where d_1 and d_2 are given by

$$d_1 = \frac{\ln(\frac{S_t}{K}) + (r + \frac{\sigma^2}{2})(T-t)}{\sigma\sqrt{T-t}},$$
$$d_2 = \frac{\ln(\frac{S_t}{K}) + (r - \frac{\sigma^2}{2})(T-t)}{\sigma\sqrt{T-t}}.$$

So $C_{BS}(S, t, T; \sigma)$ is the Black-Scholes price function for a European call option and a solution for the PDE (given by (10)) under boundary condition (3).

If we have the market price of an option, with spot price S_t, strike K, and time-to-maturity $\tau = T - t$, say $C_{MP}(S_t, K, \tau)$, then (in theory) we can determine the volatility σ by using the inverse of the Black-Scholes formula. This is the implied volatility and is given by

$$\hat{\sigma} = C_{BS}^{-1}(C_{MP}(S_t, K, \tau), S, K, \tau). \tag{12}$$

Because the concerning formula is not (easily) invertible, numerical methods are used for this purposes.

The Black-Scholes model assumes that the stock volatility, σ, is constant over time. Empirically it is found that the implied volatility $\hat{\sigma}$ differs greatly over time, and thus is not contstant as the model assumes. For a fixed time t (and fixed underlying stock price S_t), the implied volatility should also not vary for different maturities and strikes. Again, in practice, from empirical research, the latter is also known to be the case. This is the so called volatility smile (as mentioned in the introduction): far out-of-the-money (OTM) options (options of which the strike is far lower than the spot price) are known to have higher implied volatilities than in- the-money (ITM) options (options of which the strike price is far higher than the spot price) or at-the-money (ATM) options (options of which the strike price is equal to the spot price). One simple way to cope with the smile problem is to model the implied volatility such that it depends upon the maturity and strike price. This is the case for the model proposed by Dumas, Fleming, and Whaley (1998), which they refer to as the Ad-Hoc model and which Christoffersen and Jacobs (2004) refer to as the Practitioners Black-Scholes (PBS) model. We discuss this variation of the Black-Scholes model below.

The PBS model

We use the PBS model, as presented here, as a benchmark model for the CEV model, when it is applied to real option data. Its relevance as a benchmark model is due to the fact that it is a simple model and it is used as a benchmark model in the existing literature. Another reason is that very similar models are widely used in practice. In section 3.4 we discuss various estimation methods, for which we mainly consider their application to the (K&M approximation of the) CEV model and, to a lesser extent, to the Heston model. For the PBS model we discuss the model's specification, implementation, and estimation in this section, because these three elements are more interwoven. Also, we discuss the importance of the choice of the loss-function, as shown by Christoffersen and Jacobs (2004), in terms of its implication for the model's underlying structure.

We apply the PBS model on daily market option data (which is extensively discussed in section 5.1). For each day $t = 1, \ldots, N$, for some $N > 0$, we have a cross-sectional set of options $i = 1, \ldots, n_t$, for $n_t > 0$, with varying strikes and maturities. We focus on one day and one set of n options (without the subscript t), written on the same asset with spot price S. The implied volatilities for each option option are obtained using (12), such that we get

$$\sigma_i = C_{BS}^{-1}(C_{MP}(S, K_i, \tau_i), \tau_i, K_i, S, r), \tag{13}$$

for $i = 1, \ldots, n$.

The PBS model basically specifies the volatility in terms of an option's time-to-maturity and strike price. Dumas et al. (1998) consider various specifications for the volatility. We will limit our attention to the most general model they investigate, as done by Christoffersen and Jacobs (2004), which is of the form

$$\sigma_i = \theta_0 + \theta_1 K_i + \theta_2 K_i^2 + \theta_3 \tau_i + \theta_4 \tau_i^2 + \theta_5 K_i \tau_i + \epsilon_i. \tag{14}$$

One way to estimate the parameters $\theta = (\theta_1, \theta_2, \ldots, \theta_5)$, is by applying ordinary least squares (OLS). To apply OLS, first the implied volatilities $\hat{\sigma}_i$ have to be determined for each option $C_{MP}(S, K_i, \tau_i)$, for $i = 1 \ldots, n_t$. Second, the implied volatilities are regressed on the polynomial expression given by equation (14), such that we get the fitted values of the implied volatility as

$$\sigma_i(\theta) = \hat{\theta}_0 + \hat{\theta}_1 K_i + \hat{\theta}_2 K_i^2 + \hat{\theta}_3 \tau_i + \hat{\theta}_4 \tau_i^2 + \hat{\theta}_5 K_i \tau_i, \tag{15}$$

for $i = 1 \ldots, n_t$.

Estimating the parameters of (14) through OLS amounts to letting the estimation loss function, to minimise, be the Implied Volatility Mean Squared Error (IVMSE), such that the estimates are given by

$$\theta_{IVMSE} = \min_{\theta} IVSME(\theta) = \min_{\theta} \frac{1}{n} \sum_{i=1}^{n} (\sigma_i - \sigma_i(\theta))^2 = (Z'Z)^{-1} Z' \sigma, \tag{16}$$

where σ_i is the implied volatility obtained from (13), Z is the matrix of regressors from the implied volatility model (equation (14)), σ is the vector of length n of σ_i's, and $\sigma_i(\theta)$ is the fitted volatility from (15). Now, to obtain the option price form the PBS model, the set of estimated parameters θ_{IVMSE} is simply plugged into the fitted volatility (15), which is then plugged into the Black-Scholes formula, given by (11), yielding

$$C_{PBS} = C_{BS}(S, t, T, \sigma(\theta_{IVMSE})) \tag{17}$$

for the option price from the PBS model.

As we mentioned before, Christoffersen and Jacobs (2004) note the importance of the choice of the loss-function, when estimating the parameters of a certain option pricing model. We discuss their comments for this particular model and show how the choice of the loss-function might influence the bias in the option price. The above IVMSE loss-function, which minimizes the discrepancy between the volatilities rather than the option prices, yields an error specification of the form

$$C_i = C_{BS}(\theta_{IVMSE} + \epsilon_i), \tag{18}$$

where $C_{BS}(\theta_{IVMSE} + \epsilon_i)$ is a simplified notation for the Black-Scholes formula. Because the Black-Scholes formula is non-linear in the volatility, it is also non-linear in the error ϵ_i. This non-linearity causes a bias in the option price, given by (17), which is not present in the volatility estimate (since applying OLS for the volatility will ensure that $E(\epsilon_i) = 0$). This bias causes the inequality, given by

$$E(C_i) \neq C_{BS}(\sigma_i(\theta_{IVMSE})), \tag{19}$$

where C_i is given by (18).

Our interest is mainly in the pricing ability of the option pricing model. Therefore, we evaluate the PBS model using a dollar-based mean squared error, given by

$$MSE(\theta) = \frac{1}{n}\sum_{i=1}^{n}((C_i - C_{BS}^i(\sigma_i(\theta)))^2 \tag{20}$$

where C_i is the market price of the $i-th$ option and $C_{BS}^i(\sigma_i(\theta))$ is its corresponding PBS model price. Now, if we use θ_{IVMSE} as the estimate of θ, we use another loss function for the estimation than we use for the evaluation. If the evaluation loss function is the dollar-based MSE of (20), then Christoffersen and Jacobs (2004) state the importance of also estimating θ using the same loss function. Thus we need to use non-linear least squares (NLS) to directly estimate θ as follows:

$$\hat{\theta} = \min_{\theta} MSE(\theta) = \min \frac{1}{n}\sum_{i=1}^{n}((C_i - C_{BS}^i(\sigma_i(\theta)))^2 \tag{21}$$

The consequence of this choice of evaluation loss function is that implicitly the model under consideration is now

$$C_i = C_{BS}(\sigma_i(\hat{\theta})) + \epsilon_i \tag{22}$$

This specification for the option price differs from that given by (18). Thus, from this we can see that the change in loss function implies a change in the underlying structure of the model's dynamics. As reported by Christoffersen and Jacobs (2004), the chosen of loss function, for evaluation and estimation, is crucial for the model's performance. In this thesis we use the NLS optimization with the loss function, as presented by equation (21), because of the non-linearity of the option price given by (18) and because we are interested in prices and not in the fitting of implied volatilities.

2.2 Heston Model

The Heston model is a very popular stochastic volatility model. The model can be seen as an extension of the Black-Scholes model. The Heston model assumes a square root process for the dynamics of the instantaneous variance of the stock price , whereas the BS model takes it to be constant. The processes, assumed by Heston (1993), for the stock price S_t and its variance V_t, are given by

$$\frac{\mathrm{d}S_t}{S_t} = r\mathrm{d}t + \sqrt{V_t}\mathrm{d}W_t, \tag{23}$$

$$\mathrm{d}V_t = \kappa(\alpha - V_t)\mathrm{d}t + w\sqrt{|V_t|}\mathrm{d}W_t^v, \tag{24}$$

where W_t^v and W_t are Brownian motions correlated with instantaneous correlation ρ, α is the long run mean of the variance, w is the volatility of the variance, κ is the mean reversion rate and r is the risk free rate of return. The process for the variance, given by equation (24), is also known as the square root process.

The model's parameters determine how the distribution of the stock process S_t differs from a log-normal distribution. Kurtosis depends on the magnitude of ω relative to that of κ. If ω is relatively large, a more volatile variance will lead to fat tails. These fat tails will raise the prices of far OTM options and far ITM options. Mikhailov and Nögel (2003) state that the skewness is affected, in addition to the other parameters, by the parameter ρ. Intuitively, if $\rho > 0$, the variance will increase as the stock price increases. This will create a fat right-tail distribution. Conversely, if $\rho < 0$, then the variance will increase if the stock price decreases, leading to a fat left-tail distribution.

Consider a European call option written on an asset with specifications given by (23) and (24). The option's pay-off at maturity time $T > 0$ is given by $b(S_T)$ in equation (2), where again $K > 0$ is the strike price. We use a slightly different notation than in section 2.1. We denote the price of the option from the Heston model, at time $t \in [0, T]$, by $w(s, v, t)$, when the stock price is s and the *variance* is v. Then, at time of maturity T, the option price should satisfy

$$w(s, v, T) = b(s), \tag{25}$$

for all v. Subject to this boundary condition, the pricing function satisfies,

$$0 = L^H w(s, v, t) - rw(s, v, t) \tag{26}$$

where L^H is the differential operator associated with equations (23) and (24), given by

$$L^H w = \frac{\partial w}{\partial t} + rx\frac{\partial w}{\partial x} + \frac{1}{2}vx^2\frac{\partial^2 w}{\partial x^2} + \kappa(\alpha - v)\frac{\partial w}{\partial v} + \frac{1}{2}\omega^2 v\frac{\partial^2 w}{\partial v^2} + \rho\omega vx\frac{\partial^2 w}{\partial x \partial v}. \tag{27}$$

The form of the PDE for the Heston model, given by (26), is similar to that of the Black-Scholes model, which is given by (10). For both PDEs the term with the differential operator Lw has to match the steady growth term rw. The difference between the two is due to the stochastic process for the variance from the Heston model; the growth in the option price for the Heston model (given by the operator Lw) has additional terms that can be interpreted as the influence of changes in the variance on the option price.

The difference between the PDEs for the pricing functions of the Black-Scholes and Heston model implies that we cannot just use the Black-Scholes formula to price options from the Heston model. Contrary to the Black-Scholes model, in the Heston model the changes in volatility are expected to influence the option price. A closed form solution for the PDE of the Heston model is provided by Heston (1993). Through Fourier inversion techniques, Heston (1993) provides an analytical option pricing formula, as a solution to the PDE given by equation (26). This closed-form solution of equation (26) is the price of a European call option on an asset with stock price S_t and variance V_t, paying no dividends. Heston (1993) states this solution in the same form as that of the Black-Scholes formula, namely as[1],

$$C(S_t, V_t, t, T) = S_t P_1 - K P_2 \tag{28}$$

where P_1 and P_2 are given by

$$\begin{aligned}
P_j(x, V_t, T, K) &= \frac{1}{2} + \frac{1}{\pi}\int_0^\infty \Re\left(\frac{e^{-i\phi\ln(K)} f_j(x, V_t, T, \varphi)}{i\varphi}\right) d\varphi \qquad (29)\\
x &= \ln(S_t)\\
f_j(x, V_t, T, \varphi) &= \exp\{C(T-t, \phi) + D(T-t, \phi)V_t + i\varphi x\}\\
C(T-t, \varphi) &= r\varphi i r + \frac{a}{\omega^2}\left[(b_j - \rho\varphi i + d)\tau - 2\ln\left(\frac{1-ge^{dr}}{1-g}\right)\right]\\
D(T-t, \varphi) &= \frac{b_j - \rho\omega\varphi + d}{\omega^2}\left(\frac{1-e^{dr}}{1-ge^{dr}}\right)\\
g &= \frac{b_j - \rho\omega\varphi i + d}{b_j - \rho\omega\varphi i - d}\\
d &= \sqrt{(\rho\omega\varphi i - b_j)^2 - \omega^2(2u_j\varphi i - \varphi^2}
\end{aligned}$$

for $j = 1, 2$, where

$$u_1 = \frac{1}{2},\ u_2 = -\frac{1}{2},\ a = \kappa\alpha,\ b_1 = \kappa(1 + \omega V_t) - \rho\omega,\ b_2 = \kappa(1 + \omega V_t).$$

A large part of the above formula is easily computed with computer software like `MATLAB` (that we use in this thesis). The only part that pose a slight problem is the limit of the integral in $P_j(x, V_t, T, K)$, for j=1,2. One way to cope with this problem is to approximate the integrals, using numerical integration techniques like Gauss Lagendre or Gauss Lobatto. Numerical integration is in this case much faster than running a Monte-Carlo simulation. Carr and Madan (1999) propose a modification of the Fast Fourrier transform (FFT) that can be used to evaluate expressions of the form of (28). They show that their method is significantly faster and more accurate than standard numerical integration techniques. One downside of their approach is that it requires a certain dampening coefficient, which has to be set such that the accuracy is optimal. This coefficient depends on the model's parameters. Identifying the relationship between the model's parameters and the dampening coefficient, as Carr and Madan (1999) do for the Variance-Gamma model, is far from trivial and makes the method unusable for purpose of this thesis. Kristensen and Mele (2011) state that they use the FFT of Carr and Madan (1999), but make no mention of how they picked the dampening coefficient.

[1]Here we solely state the solution as applied in this thesis. The reader can refer to Heston (1993) for derivations and further details.

We are not able to find one certain value for the dampening coefficient, such that we get the same option prices as Kristensen and Mele (2011), for a set of options (from the Heston model) with varying strikes, maturities and spot prices, but a fixed set of parameters. There is no one value for the dampening coefficient that gets exactly the same option prices as Kristensen and Mele (2011) for all option of the concerning set. However, when we use numerical integration we get the *exact* same results as Kristensen and Mele (2011). Hence, we apply numerical integration in this thesis to determine the above closed-form pricing expression for the Heston model. In appendix C we discuss the numerical integration method we use in more detail.

Thus, the Heston model offers the great benefit of a closed form option pricing formula, while assuming a stochastic process for the variance of the option's underlying stock. Notwithstanding, the model is often rejected as a model of stock index returns. For example, Andersen, Benzoni, and Lund (2002) find that the kurtosis generated by the model is insufficient, while Pan (2002) rejects the square root specification on its complication for the term structure volatility. Jones (2003) focuses on the stochastic variance as the source of non-Gaussian return dynamics, and reports that this is not captured by the Heston model. Other rejections are reported by Chernov and Ghysels (2000) and Benzoni (2002).

Through the VIX index, Jones (2003) finds that periods of high volatility of the stock price coincide with periods where the VIX index is also more volatile, i.e. where the volatility of volatility is high. This is contradictory to the Heston model's square root formulation of the variance, which implies that the volatility of volatility is constant over time. To see this, we need to move from the instantaneous variance, as stated in equation (24), to the instantaneous volatility, given by $\tilde{v}_t = \sqrt{V_t}$.

The process for $\tilde{v}_t$ can be obtained by applying Itô's Lemma as described in the appendix A, by taking $f(X_t) = \sqrt{X_t}$ and $X_t = V_t$, such that we have

$$\begin{aligned}
\frac{\partial f(V_t,t)}{\partial V_t} &= \frac{1}{2}\frac{1}{\sqrt{V_t}} \\
\frac{\partial^2 f(V_t,t)}{\partial V_t^2} &= \frac{1}{2}\frac{1}{\sqrt{V_t}} \\
\mu_t &= \kappa(\alpha - V_t) \\
\sigma_t &= \omega\sqrt{V_t}
\end{aligned} \tag{30}$$

Plugging this into expression 113, we get the process for the stock's instantaneous volatility $\tilde{v}_t = V_t$ given by

$$\begin{aligned}
\mathrm{d}\tilde{v}_t &= (\kappa(\alpha - \tilde{v}_t^2)\frac{1}{2}\frac{1}{\tilde{v}_t} - \frac{\omega^2\tilde{v}_t^2}{2}\frac{1}{4}\frac{1}{\tilde{v}_t^2})\mathrm{d}t + \omega v_t\frac{1}{2}\frac{1}{\tilde{v}_t}\mathrm{d}W_v(t) \\
&= \frac{\kappa}{2\tilde{v}_t}(\alpha + \frac{\omega^2}{4} - \tilde{v}_t^2)\mathrm{d}t + \frac{1}{2}\omega\mathrm{d}W_v(t)
\end{aligned} \tag{31}$$

This shows that the volatility of instantaneous volatility from the Heston model is $\frac{1}{2}\omega$ and thus no longer depends on the level of the instantaneous volatility. This is in contrast to the instantaneous variance process of equation (24), which does depend on the variance level.

Another shortcoming of the Heston model, stated by Jones (2003), is that it still produces a volatility smile. Jones (2003) finds, for European call options, that the Heston model generates a volatility smile and thus is not adequate for explaining most ITM and OTM option prices.

2.3 CEV model

In this section we consider the constant elasticity of variance or CEV model, which generalizes the Heston model by replacing the square root process in the variance diffusion, as stated in equation (24), by a parameter of undetermined magnitude. The processes for the stock price, S_t, and variance, V_t, are now given by

$$\frac{\mathrm{d}S_t}{S_t} = r\mathrm{d}t + \sqrt{V_t}\mathrm{d}W_t \tag{32}$$

$$\mathrm{d}V(t) = \kappa(\alpha - V(t))\mathrm{d}t + w|V(t)|^{\xi}\mathrm{d}W_v(t) \tag{33}$$

where, as with the Heston model, we have Corr($\mathrm{d}W_v(t)$,$\mathrm{d}W_t$)=ρ. By setting $\xi = 1/2$ we obtain a square root process for the variance diffusion and thus the Heston model.

The CEV-model allows for level dependence for the stock volatility; a feature present in VIX data, as stated by Jones (2003), but not captured by the Heston model, as shown in section 2.2. This can be seen by again applying Itô's lemma to $\tilde{v}_t = \sqrt{V_t}$, for V_t from equation (33). This shows that the volatility of volatility, $\tilde{v}_t$, is now equal to $\omega\xi\tilde{v}_t^{(4\xi-2)}$, such that the level dependence of the volatility only drops (i.e. $\omega\xi\tilde{v}_t^{(4\xi-2)}$ becomes zero) when $\xi = \frac{1}{2}$, as is the case for the Heston model. When $\xi \neq \frac{1}{2}$, the volatility is level dependent, as shown by Jones (2003) using VIX data.

For a European call option written on a stock for which the price follows the process given by (32) and (33), the pricing function $w(s, v, t)$ must satisfy similar conditions as for the Heston specifications. Subject to the boundary condition $w(S_T, v, T) = b(S_T)$, at time T for all variance v, $w(s, v, t)$ should now satisfy,

$$0 = Lw(s, v, t) - rw(s, v, t) \tag{34}$$

where the differential operator for the CEV-process, associated with equations (32) and (33) is given by

$$Lw = \frac{\partial w}{\partial t} + rs\frac{\partial w}{\partial s} + \frac{1}{2}vs^2\frac{\partial^2 w}{\partial s^2} + \kappa(\alpha - v)\frac{\partial w}{\partial v} + \frac{1}{2}\omega^2 v^{2\xi}\frac{\partial^2 w}{\partial v^2} + \rho\omega v^{\xi+\frac{1}{2}}s\frac{\partial^2 w}{\partial s\partial v}. \tag{35}$$

Jones (2003) finds that the CEV model fits the empirical data (S&P 100 stock returns) much better than the Heston model; the unconditional moments from the CEV model fit the data very well, whereas the unconditional moments from the Heston model differ a lot from the empirical data. For the conditional moments, Jones (2003) finds that the CEV model outperforms the Heston model, but the fit to empirical data is not satisfactory and thus leaving room for improvement. The same goes for the volatility smile; the smile is less present for the CEV model, but it is not completely absent, as is implied by the theory. These shortcomings leave some of the conditional moments- and option pricing puzzles.

In contrast to the Heston model, the CEV model does not have an accurate closed form solution. Therefore, Jones (2003) computes the option prices from the CEV model using Monte-Carlo simulation. The shortcomings of the CEV model, as reported by Jones (2003), could be due to inaccuracies of the Monte Carlo Simulation. Kirstensen and Mele (2011) develop a new approach to approximation asset prices in the context of multifactor continuous-time models. For models, lik the CEV-model, that lack a closed-form solution, Kristensen and Mele (2011) provide a solution which relies on approximation of the intractable model through a known, auxiliary one. We suggest to apply the CEV model in conjunction with the method proposed by Kristensen and Mele (2011) for the pricing of European call options. In this project we will assess the performance of the CEV model when applied in conjunction with the K&M method.

3 Methodology

In the previous section we noted the absence of a closed form solution for the CEV model when $\xi \neq 1/2$. The empirical importance of these cases are demonstrated by Jones (2003), who estimates the parameter ξ equal to $\xi = 1.33$ and $\xi = 1.17$, using a Bayesian framework (Monte-Carlo Markov Chain), on a bivariate time series of S&P 100 returns and option implied volatility data of 1986-2000 and 1988-2000, respectively. Kristensen and Mele (2011) propose a method to approximate the price of a European call option, as a solution to the pricing equation of the CEV model, given by (34). The K&M method provides a closed form approximation to the original option price. This closed form approximation can be easily implemented and is time efficient, since the closed form expression needs to be determined only once and after implementation it requires virtually no computation time to determine the price.

In this section we describe how the K&M method is applied to the CEV model, as done by Kristensen and Mele (2011). Yang (2006) also provides a closed form approximation, which is similar to that of Kristensen and Mele (2011), for options from the CEV model. We also describe the Yang method when applied to the CEV model, as done by Yang (2006). Due to some shortcomings of the Yang method, we only use it as a benchmark for the accuracy of the K&M method in section 4. Furthermore, we describe the Euler discretization for the Monte-Carlo simulation, which we apply in section 4 to generate option prices from the CEV model. The option prices of the Monte-Carlo simulation are used as an accuracy benchmark for the prices from the K&M method and to generate prices from the CEV model when there is no closed form solution available.

Another advantage of the closed form approximation of Kristensen and Mele (2011) is that it can be used to estimate the CEV model parameters directly from option price data. The fast pricing computation and the fact that the approximations are closed form, allow for easy calibration of the concerning model. In section 5 we evaluate the empirical performance pricing of the CEV model, when applied in conjunction with the K&M method. The calibration methods used for this purpose, are described at the end of this section.

3.1 Approximation of Kristensen and Mele

In this subsection we show how the K&M method is applied to approximate the price of a European call option from the CEV model. The computation of the closed form approximation requires a so called *auxiliary model*, for which a closed form solution is available. For the CEV model we use the Black-Scholes model (for European call options) as the auxiliary model. The PDE of the Black-Scholes model is given by equation (10). We write this PDE in a similar form as that of the CEV model, given by 34, namely as,

$$0 = L_0 w^{bs}(s, t; \sigma_0) - r w^{bs}(s, t; \sigma_0), \tag{36}$$

where L_0 is the differential operator associated with the Black-Scholes model, given by

$$L_0 w = \frac{\partial w}{\partial t} + rs\frac{\partial w}{\partial s} + \frac{1}{2}\sigma^2 s^2 \frac{\partial^2 w}{\partial s^2}. \tag{37}$$

Thus, the Black-Scholes price of a European call option is the solution to (36) subject to the boundary condition $w^{bs}(s, T; \sigma_0) = \max(s - K, 0)$, where s is the price of the underlying asset with (constant) volatility σ_0, K is the strike price and T is the time of maturity.

Now, we are interested in the solution of the PDE (34), thus an expression of the price from the CEV model, but we only have an expression for the price from the Black-Scholes model. The difference between these two prices is denoted by

$$\Delta w(s, v, t; \sigma_0) \equiv w(s, v, t) - w^{bs}(s, t; \sigma_0). \tag{38}$$

where $w(s, v, t)$ and $w^{bs}(s, t; \sigma_0)$ are the CEV and Black-Scholes option prices, respectively.

The key idea of the K&M method is to subtract (37) from (35). The result is that the above price difference, $\Delta w(s, v, t; \sigma_0)$, satisfies,

$$0 = L\Delta w(s, v, t; \sigma_0) - r\Delta w(s, v, t; \sigma_0) + \delta(s, v, t; \sigma_0) \tag{39}$$

with boundary condition $\Delta w(s, v, T; \sigma_0) = 0$ for all s and v. The 'miss-pricing function', δ, is given by:

$$\delta(s, v, t; \sigma) \equiv \frac{1}{2}(v - \sigma_0^2)s^2 \frac{\partial^2}{\partial s^2} w^{bs}(s, t; \sigma_0). \tag{40}$$

Since $w^{bs}(s, t; \sigma_0)$ is known, we can compute $\delta(s, t; \sigma_0)$. By relying on the Feynman-Kac representation of the solution to (39) and recalling the definition of the price difference $\Delta w = w - w^{bs}$, the unknown pricing function w is given by

$$w(s, v, t; \sigma_0) = w^{bs}(s, t; \sigma_0) + \mathbb{E}_{s,v,t}\left(\int_t^T e^{-r(u-t)}\delta(S_u, V_u, u; \sigma_0)\,\mathrm{d}u\right). \tag{41}$$

In appendix B we describe the Feynman-Kac theorem and show how it is applied to get this expression. If we are able to compute the expectation in equation (41), we find the exact expression for the pricing function. Unfortunately this is not the case and the Feynmann-Kac representation of the pricing function does not change this, as we cannot compute the expectation. Here we encounter the most important contribution of Kristensen and Mele (2011) in solving this problem, namely the approximation for the expectation in equation (41).

Kristensen and Mele (2011) start with a power series representation of the expectation term in equation (41), such that $w(s, v, t)$ becomes

$$w(s, v, t) = w^{bs}(s, t; \sigma_0) + \sum_{n=0}^{\infty} \frac{(T-t)^{n+1}}{(n+1)!}\delta_n(s, v, t; \sigma_0), \tag{42}$$

where δ_n satisfies the recursive function:

$$\delta_{n+1}(s, v, t; \sigma_0) = L\delta_n(s, v, t; \sigma_0) - r\delta_n(s, v, t; \sigma_0), \tag{43}$$

with $\delta_0 = \delta$ and L the differential operator given by (35). This formula is truncated to finite terms for practical implementation, yielding:

$$w(s, v, t) = w^{bs}(s, t; \sigma_0) + \sum_{n=0}^{N} \frac{(T-t)^{n+1}}{(n+1)!}\delta_n(s, v, t; \sigma_0), \tag{44}$$

for some $N \geq 0$. Equation (44) gives the K&M approximation of the price of a European call option on a stock following a CEV-process. Thus, the K&M approximation of the price of a European call option from the CEV model, consists of the Black-Scholes option price together with N corrective terms.

3.2 Approximation of Yang

Similar to Kristensen and Mele (2011), Yang (2006) provides a method that yields closed form approximations for contingent claim models, using a power series expansion. Yang (2006) uses a so called 'base model' that serves a similar function as the auxiliary model used by Kristensen and Mele (2011). In this subsection we show how Yang (2006) applies his method for the CEV model. The main difference between the methods is that Yang (2006) determines the conditional expectations of the discrepancy between the CEV and the base model under the risk-neutral measure of the base model, while Kristensen and Mele (2011) use the risk-neutral measure of the CEV model for the same purpose. We give a more detailed comparison between the two methods at the end of this subsection.

The dynamics of the stock price and its variance under the CEV model are given by equation (32) and (33), respectively. For notational convenience, Yang (2006) restates the process for the stock price in logarithmic form. By taking the logarithm of the stock price $X_t = \log(S_t)$ and applying Itô's formula, we get

$$\mathrm{d}X_t = (r - \frac{1}{2}V_t)\mathrm{d}t + \sqrt{V}_t \mathrm{d}W_t, \tag{45}$$

$$\mathrm{d}V_t = \kappa(\alpha - V_t)\mathrm{d}t + \omega|V_t|^{\xi}\mathrm{d}W_{v,t}. \tag{46}$$

for the log-stock price and variance under the CEV model, respectively. In terms of the logarithmic stock price, we write $u(x, v, t)$ for the pricing function of a European call option and its pay-off function is given by $f(X) = \max\{e^X - K, 0\}$. Consequently, at time of maturity T, the pricing function should satisfy the boundary condition $u(x_T, v, T) = f(X_T)$. Subject to this boundary condition, the pricing function satisfies the following PDE:

$$\frac{\partial \mu}{\partial \tau} = \frac{1}{2}v\frac{\partial^2 u}{\partial x^2} + \rho v^{\xi+\frac{1}{2}}\frac{\partial^2 u}{\partial x \partial v} + \frac{1}{2}\omega^2 v^{2\xi}\frac{\partial^2 u}{\partial v^2} + (r - \frac{v}{2})\frac{\partial u}{\partial x} + \kappa(\alpha - v)\frac{\partial u}{\partial v} - ru. \tag{47}$$

This PDE is stated by Yang (2006) in terms of the derivative to the time-to-maturity $\tau = T - t$, instead of the time t. This expression is equivalent to expression (34), when written in terms of the log sotck price x and time-to-maturity τ.

Next, Yang (2006) divides the linear operator on the right hand side of equation (47) into two parts: L_0 and L_1, given by

$$L_0 u = \frac{1}{2}v\frac{\partial^2 u}{\partial x^2} + (r - \frac{v}{2})\frac{\partial u}{\partial x} - ru \text{ and} \tag{48}$$

$$L_1 u = \rho v^{\xi+\frac{1}{2}}\frac{\partial^2 u}{\partial x \partial v} + \frac{1}{2}\omega^2 v^{2\xi}\frac{\partial^2 u}{\partial v^2} + \kappa(\alpha - v)\frac{\partial u}{\partial v} - ru, \tag{49}$$

respectively, such that we can write equation (47) as

$$\frac{\partial u}{\partial \tau} = L_0 u + L_1 u. \tag{50}$$

By writing equation (47) as (50), we have a linear operator L_0 containing *only* derivatives with respect to x and an operator L_1 containing *all* derivatives with respect to v. If the variance is constant, we have $L_1 u(x,v,\tau) = 0$ since all derivatives with respect to v are zero. This corresponds with the assumption of constant variance of the Black-Scholes model and we are left with

$$\frac{\partial u}{\partial \tau} = L_0 u, \tag{51}$$

which is the pricing equation of the Black-Scholes model, written in terms of the log stock price and maturity τ.

Yang (2006) suggests to write the solution $u(x,v,t)$ to the CEV pricing equation (50), similar to Kristensen and Mele (2011), in power series form, as

$$u(x,v,\tau) = \sum_{m=0}^{\infty} u^{(m)}(x,v,\tau). \tag{52}$$

To determine the terms of the power series expression, we plug it into expression (50), yielding

$$\sum_{m=0}^{\infty} \frac{\partial u^{(m)}}{\partial \tau} = \sum_{m=0}^{\infty} L_0 u^{(m)}(x,v,\tau) + \sum_{m=0}^{\infty} L_1 u^{(m)}(x,v,\tau). \tag{53}$$

The first term $u^{(0)}$, is the solution to (51), thus the Black-Scholes formula.

If we take into account that $u^{(0)}$ must satisfy (51) or

$$\frac{\partial u^{(0)}}{\partial \tau} = L_0 u^{(0)} \Rightarrow \frac{\partial u^{(0)}}{\tau} - L_0 u^{(0)} = 0, \tag{54}$$

under boundary condition $u^{(0)}(x,v,0) = \max\{e^x - K, 0\}$. We can derive the following from (53)

$$\begin{aligned}
&\sum_{m=0}^{\infty} \frac{\partial u^{(m)}}{\partial \tau} = \sum_{m=0}^{\infty} L_0 u^{(m)}(x,v,\tau) + \sum_{m=0}^{\infty} L_1 u^{(m)}(x,v,\tau) \\
\Rightarrow &\frac{\partial u^{(0)}}{\partial \tau} + \sum_{m=1}^{\infty} \frac{\partial u^{(m)}}{\partial \tau} = L_0 u^{(0)}(x,v,\tau) + \sum_{m=1}^{\infty} L_0 u^{(m)}(x,v,\tau) + \sum_{m=0}^{\infty} L_1 u^{(m)}(x,v,\tau) \\
\Rightarrow &\frac{\partial u^{(0)}}{\partial \tau} - L_0 u^{(0)}(x,v,\tau) + \sum_{m=1}^{\infty} \frac{\partial u^{(m)}}{\partial \tau} = \sum_{m=1}^{\infty} L_0 u^{(m)}(x,v,\tau) + \sum_{m=0}^{\infty} L_1 u^{(m)}(x,v,\tau) \\
\Rightarrow &\sum_{m=1}^{\infty} \frac{\partial u^{(m)}}{\partial \tau} = \sum_{m=1}^{\infty} L_0 u^{(m)}(x,v,\tau) + \sum_{m=0}^{\infty} L_1 u^{(m)}(x,v,\tau),
\end{aligned} \tag{55}$$

where in the last step we use the relation given by (54). From the last expression (55), it follows that for each $m > 0$, it must hold that

$$\frac{\partial u^{(m)}}{\partial \tau} = L_0 u^{(m)}(x,v,\tau) + L_1 u^{(m-1)}(x,v,\tau). \tag{56}$$

And thus the terms $u^{(m)}$, for $m > 0$, can be found recursively as a solution to this PDE, with boundary condition $u^{(m)} = \max\{e^x - K, 0\}$. Note that at least the term $u^{(0)}$, must be known to proceed. This is why Yang (2006) refers to the model of this term as the base model. In this case it is the Black-Scholes model, similar to the auxiliary model of the K&M method.

As a general solution to (56), for $m > 0$, Yang (2006) suggests the following. First, without loss of generality, assume that $r = 0$. Otherwise, let $W(y, v, \tau) = e^{r\tau}u(x, v, \tau)$, $y = x + r\tau$, the r in L_0 will cancel out. The Green function for equation (56) is given by

$$G(x, v, \tau) = \frac{1}{\sqrt{2\pi v\tau}} e^{-\frac{(x-\frac{v\tau}{2})^2}{2v\tau}}, \tag{57}$$

using this, we get the recursive relation

$$u^{(0)}(x, v, \tau) = \int_{-\infty}^{\infty} f(x - \xi)G(\xi, v, \tau)d\xi, \tag{58}$$

$$u^{(m)}(x, v, \tau) = \int_0^{\tau} \int_{-\infty}^{\infty} L_1 u^{(m-1)}(x - \xi, v, \tau - s)G(\xi, v, s)d\xi ds, \tag{59}$$

for $m > 0$ the approximation of Yang (2006) is based on cutting of the power series expansion given by (52), with terms given by (58) and (59), for some $N > 0$. Yang's (2006) approximation of a European call option from the CEV model is given by

$$u(x, v, \tau) = \sum_{m=0}^{N} u^{(0)}(x, v, \tau) \tag{60}$$

3.2.1 Comparison Yang and K&M method

The methods of Yang (2006) and Kristensen and Mele (2011) are very much alike; both add and subtract corrective terms to a base/auxiliary price to approximate an unknown contingent claim price. In this section we will make a short comparison and show the difference between the two methods for approximating the solution of the CEV model.

We have used different notations in section 3.1 and 3.2 for the K&M and Yang method, respectively. To make the comparison more comprehensible, we use one notation to present both methods here. Yang (2006) and Kristensen and Mele (2011) start similarly with a base-model and auxiliary model, respectively, satisfying a PDE of the form

$$L_0 w_0(x, t) = r w_0(x, t). \tag{61}$$

The PDE of the unknown price of the CEV model is given by equation (34). We can write this PDE, in terms of L_0, as

$$Lw(x, t) - rw(x, t) = 0 \tag{62}$$

$$\Rightarrow L_0 w(x, t) + (L - L_0)w(x, t) - rw(x, t) = 0, \tag{63}$$

where L is the linear operator, corresponding with the CEV model, given by (35). Both methods use the Black-Scholes model as the base/auxiliary-model, such that, for L_0 we have

$$L_0 w = \frac{\partial w}{\partial t} + rs\frac{\partial w}{\partial s} + \frac{1}{2}v^2 s^2 \frac{\partial^2 w}{\partial s^2} \text{ and thus} \tag{64}$$

$$(L - L_0)w = \kappa(\alpha - v)\frac{\partial w}{\partial v} + \frac{1}{2}v^{2\xi}\frac{\partial w}{\partial v} + \rho\omega v^{\xi+1/2} s \frac{\partial^2 w}{\partial s \partial v} \tag{65}$$

This definition of L_0 closely related to that of Yang (2006), given by equation (48), with an additional term containing the derivative to t and the derivative is taken to stock price s instead of the log stock price x. The differential operator $L - L_0$ relates to L_1 of the Yang method, given by equation (49), in the same way (but without the derivative to t term). This L_0 can be obtained from that of Kristensen and Mele (2011), given by equation (37), by replacing the volatility of the Black-Scholes (auxiliary) model σ_0 with $\sqrt{v}$, and as such the corresponding pricing function becomes $w_0(x,t) = w^{bs}(s,t;\sqrt{v})$.

For the following we assume, without loss of generality, that $r = 0$. The difference between the two methods becomes clear if we look at how they go about the price difference, $\Delta w(x,t) = w(x,t) - w_0(x,t)$, between the *unknown* CEV price $w(x,t)$ and the *known* Black-Scholes price $w_0(x,t)$. Yang (2006) uses the expression in equation (63) together with the fact that $L_0 w_0(x,t) = 0$ to get $L_0 \Delta w(x,t) + (L - L_0)w(x,t) = 0$. By relying on the Feyman-Kac representation of the latter expression, we get

$$w(x,t) = w_0(x,t) + \int_t^T \mathbb{E}^0_{x,t}\left((L - L_0)w(x(\mu),\mu)\right) \mathrm{d}\mu \tag{66}$$

where $\mathbb{E}^0_{x,t}$ denotes the conditional expectation taken under the risk-neutral probability of the *base model*, which is defined by L_0. The second term on the right side of equation (66), is where the power series expansion of Yang applies to (for $m > 0$). The terms of the expansion of the expectation term are given by equation (59), which defines a recursive relation for $m > 0$, starting with the pricing function from the base model. The integral that needs to be determined for (59) is quite complicated, but can be done using computer symbolic mathematics software, e.q. Maple, like we use in this thesis.

Expression (66) looks similar to expression (41) of the K&M method. The difference lies in the measure under which the expectation is taken. This is what we look at next. The equivalent of equation (41), under the assumptions of this section (like $r = 0$) can be obtained by defining the mispricing function δ, in this case, by $\delta = (L - L_0)w_0(x,t)$. Under the assumptions and notation of this section, we have $L\Delta w(x,t) + \delta = 0$. Using the Feynman-Kac representation, we get:

$$w(x,t) = w_0(x,t) + \int_t^T \mathbb{E}_{x,t}\left((L - L_0)w_0(x(\mu),\mu\right) \mathrm{d}\mu \tag{67}$$

which is the same as expression (41). The latter expectation is taken under the risk-neutral measure of the model defined by L, i.e. the CEV model, whereas the expectation of (66) was taken under the risk-neutral probability.

So the technical difference lies in the measure under which the expectation of the error between the unknown price and the known base/auxiliary model price is taken. This difference eventually results in that for the Yang method we need to compute quite complicated integrals, as given by equation (59), repeatedly to get the terms of the approximation, whereas for the K&M method the terms in the expansion are obtained by computing derivatives of the linear operator repeatedly, as given by (43). The terms of the Yang method are much harder to compute, as these involve an integral, than it is to compute the terms of the K&M method, which only contain derivatives. In Maple, we are not able to compute more than two terms for the Yang method, while in the K&M method we are able to compute four corrective terms.

3.3 Euler Discretization

Monte Carlo simulation is a common approach to determine contingent claim prices from models that we do not have a closed form expression for, like the CEV model. To apply a Monte Carlo simulation for determining the price of a European call option from the CEV model, we need to generate a large number of trajectories from the underlying stock price process. This is done by using an Euler discretization. In this section we will explain the Euler discretization and how it is applied to generate option prices from the CEV model.

For the discretization we will use the logarithmic form of the CEV model stock price process, given by (45) and (46). The logarithmic form makes the discretization more comprehensive, because it makes the dependence of the level dependence of the stock process additive. For completeness we restate the process as

$$\mathrm{d}\ln(S_t) = (r - \frac{1}{2}V_t)\mathrm{d}t + \sqrt{V_t}\mathrm{d}W_t \tag{68}$$

$$\mathrm{d}V_t = \kappa(\alpha - V_t)\mathrm{d}t + \omega|V_t|^{\xi}\mathrm{d}W_{V,t}. \tag{69}$$

We denote the discrete counterparts of the stock price S_t and stock variance V_t, as $\dot{S}_t$ and $\dot{v}_t$, respectively, such that the Euler discretization of the above stochastic process is given by

$$\ln\dot{S}_{t+\Delta} = \ln\dot{S}_t + (r - \frac{1}{2}\dot{v}_t)\Delta + \sqrt{\dot{v}_t}\varepsilon_x\sqrt{\Delta} \tag{70}$$

$$\dot{v}_{t+\Delta} = \kappa(\alpha - \dot{v}_t)\Delta + \omega|\dot{v}_t|^{\xi}\varepsilon_v\sqrt{\Delta}. \tag{71}$$

For $t = 1, \ldots, T$, the time intervals $[t-1, t]$, of size one year, are divided into $\frac{1}{\Delta}$ increments of size $0 < \Delta \leq 1$. We compute $(\dot{S}_{t-1+\Delta}, \dot{v}_{t-1+\Delta})$, given $(\dot{S}_{t-1}, \dot{v}_{t-1})$; this process is repeated until we obtain $(\dot{S}_t, \dot{v}_t)$, for $t = 1 \ldots T$. The errors ε_x and ε_v are standard normally distributed variables, with correlation ρ. The errors are obtained in each step as follows; we draw independent uniformly distributed variables U_1 and U_2 and compute

$$\varepsilon_x = \Phi^{-1}(U_1) \tag{72}$$

$$\varepsilon_v = \rho\varepsilon_x + \sqrt{1-\rho^2}\Phi^{-1}(U_2) \tag{73}$$

where Φ denotes the cumulative distribution of the standard normal distribution.

For a European call option, with maturity time T and strike K, let $C(t)$ be the exact option price at time t. The Euler discretization is used to approximate S_T with $\dot{S}_T$, such that we can approximate the option price $C(t)$ with $\dot{C}(t)$, by computing the expectation

$$\dot{C}_t = \mathbb{E}\left(max(\dot{X}_T - K, 0).\right)$$

For the Euler discretizations, we draw N independent samples of $S_T^{(1)}, S_T^{(2)}, \ldots, S_T^{(N)}$ using an equidistant time-grid with fixed increment size Δ, such that $\dot{C}(t)$ is then estimated in a standard Monte-Carlo fashion as

$$\dot{C}(t) \approx \frac{1}{N}\sum_{i=1}^{N}\max(\dot{S}_T^{(i)} - K, 0). \tag{74}$$

Thus we use (74) as the price of a European call option, with maturity time T and strike K, from the CEV model. If N gets very large the approximation is more precise, as a lot of noise is taken out the simulation process, but the process takes longer. The contrary holds for Δ; if Δ gets very small, the approximation is more precise, but the process takes longer. Thus, a trade-off needs to be made between bias and time-efficiency. Andersen (2007) shows that a decent accuracy is found when $N = 10^6$ and $\Delta = \frac{1}{32}$, and therefore it takes very long to generate the options, especially the long term maturity options. The author notes that at this precision it takes much too long to generate an option price for the model to be of any practical use. In our case we use (almost) the same settings, since we want to generate options for the CEV model, where no closed form solution is available.

3.4 Calibration Methods

To implement the CEV model the parameters need to be estimated. We estimate the model's parameters using the (closed-form) K&M approximation, such that we can do this directly on empirical option price data. In section 2.1 we showed that OLS was not suitable for the parameter estimation, due to non-linearity in the option price function. The same goes for the price function of the CEV model, or at least the K&M approximation of the CEV model. Therefore, we use non-linear least squares to estimate the parameters for the CEV model. In this section we discuss the estimation methods we use. The adequacy of the estimation process will be assessed through simulation studies discussed in section 4.1, where we study the estimation of the approximation of the Heston model (by keeping ξ in the K&M approximation fixed) and in section 4.2 we do the same for the CEV model (by adding ξ to the set of parameters we need to estimate).

3.4.1 The Calibration Problem

For the Heston model four parameters have to be estimated, namely κ, α ,ω, ρ. The time dependent variance, v_t, also needs to be estimated, as it is not observable. The CEV model includes an additional fifth parameter ξ. For the Heston model, Bakshi et al. (1997) show that the implied parameters and their time-series counterparts (estimated on the option's underlying) are different. Such that the parameters estimated on the time-series of the underlying cannot be used to price options adequately. A solution to this problem is to find the variance and parameters which produce the correct market prices of options. This makes our problem an *inverse* problem. The inverse problem at hand can be solved through various methods. We are interested in the pricing of options in absolute terms and not in, say, relative (percentage) terms. The performance of option valuation models in absolute terms is measured by the mean squared error given by

$$\$MSE(\theta) = \frac{1}{n}\sum_{i=1}^{n}(C_i - C_i(\theta))^2, \tag{75}$$

where C_i and $C_i(\theta)$ are the data and model option prices, respectively, and n is the number of options used. In this evaluation we use a dollar based loss-function, which squares the distances between the model and market prices. Christoffersen and Jacobs (2004) argue that we should use the same loss-function for the estimation process, and not, say a percentage based loss-function, which for the purpose of evaluation can be used as

$$\%MSE(\theta) = \frac{1}{n}\sum_{i=1}^{n}((C_i - C_i(\theta)/C_i)^2. \tag{76}$$

Since we do not use a percentage-based loss function for the parameter estimation, we also do not use it to evaluate the performance of any of the models. Christoffersen and Jacobs (2004) argue and show that the conclusions drawn, on the valuation performance, by using distinct loss-functions for the estimation and evaluation make little sense. Thus, the inverse problem we want to solve is to minimise the error between market prices and model prices, as this is how we evaluate the model as well. We start with a cross-sectional set of market option prices at time t. More specifically, let $\tau_1, \tau_2, \ldots, \tau_M$ be a set of time-to-maturities and let $K_1, K_2, \ldots, K_N$ be a set of strike prices. We estimate the parameters that minimise the discrepancy between model and option price, by minimising the squared error, denoted by $SqErr(\Omega, v_t)$ between the two prices. This comes down to solving the following optimisation problem,

$$\min_{\Omega, v_t} SqErr(\Omega, v_t) = \min_{\Omega, v_t} \sum_{i=1}^{N}\sum_{j=1}^{M} w_{ij}[C_{MP}(S_t, K_i, \tau_j) - C_{\Omega_t, v_t}(S_t, K_i, \tau_i)]^2, \tag{77}$$

where $C_{MP}(S_t, K_i, \tau_j)$ denotes the market prices of an option at time t with spot price S_t, maturity τ_j and strike price K_i, $C_\Omega(S_t, v_t, K_i, \tau_j)$ is the corresponding model price with parameter set $\Omega = \{\kappa, \alpha, \omega, \rho, \xi\}$ and volatility v_t. This notation implies that we estimate the model parameters and variance using $N \times M$ option prices. If ξ is fixed ,for the Heston case when $\xi = 1/2$, it is left out of the parameter set and we have $\Omega = \{\kappa, \alpha, \omega, \rho\}$. The factors w_{ij} are the weight factors, that determine the weight of the discrepancies of various option prices in the optimisation process.

By including the weights (and not setting them equal to one), we deviate from what is suggested by Christoffersen and Jacobs (2004) , since we do not have exactly the same loss function as we use in the evaluation. This is suggested by Mikhailov and Nögel (2003) who demonstrate the importance of including the weights for the Heston model, for stabilizing the optimization problem. We follow Mikhailov and Nögel (2003), by setting the weights equal to $\frac{1}{|bid_{ij} - ask_{ij}|}$. The intuition behind this choice is that options with a greater spread carry smaller weight in the estimation process, as a large spread implies a wider range of possible mid-prices. That is, the model is allowed to imply a wider range of prices around the mid-price.

To test the estimation process for the CEV model, we will use generated option price data. This is consistent with the estimation process as presented in equation (77). However, eventually we want to calibrate the model on empirical option price data. This raises the question as to what market prices to use in the estimation process, as for any given option from empirical data there exists a bid price and an ask price. This 'problem' actually allows flexibility in the estimation process. We use the mid-price of the bid- and ask price, as the market price $C_{MP}(S_t, K_i, \tau_j)$. However, we only accept sets of estimated parameters $\{\hat{\Omega}, \hat{v}_t\}$ that satisfy the following condition

$$\begin{aligned} \min_{\Omega, v_t} SqErr(\Omega, v_t) &= \min_{\Omega, v_t} \sum_{j=1}^{M} \sum_{i=1}^{N} w_{ij} [C_{MP}(S_t, K_i, \tau_j) - C_{\Omega, v_t}(S_t, K_i, \tau_j)]^2 \\ &\leq \sum_{i=0}^{N} \sum_{j=0}^{M} w_{ij} [bid_{ij} - ask_{ij}]^2, \end{aligned} \tag{78}$$

where bid_{ij}/ask_{ij} are the bid/ask prices of an option with strike K_i and maturity τ_i. This constraint narrows the possible parameter values, making the estimation process more comprehensible, as we have a smaller parameter space (i.e. possible values that the parameters can attain). It implies that we do not require the model to match the mid-prices exactly, but to fall on average within the bid-offer spread. This is not an unreasonable relaxation, as the estimation process can only estimate parameters with a certain tolerance level. To prevent the volatility from reaching zero, we add the so called stationary condition $2\kappa\alpha > \omega^2$, as suggested by Mikhailov and Nögel (2003). This condition assures the process of the variance to be stationary by requiring that the variance of the variance, ω^2, should be half the value of the long term variance, α, times its mean reversion, κ.

The minimisation of $SqErr(\Omega, v_t)$ poses some other problems. Mikhailov and Nögel (2003) state that when the closed form solution of the Heston model is taken for $C_{\Omega, v_t}(S_t, K_i, \tau_j)$, the particular structure of $SqErr(\Omega, v_t)$ is not known. We can conclude the same for the K&M approximation of the price for the CEV model. The minimum of $SqErr(\Omega, v_t)$ is not as simple as finding those parameters that make the gradient zero, as the gradient may have multiple zero points, meaning that gradient based optimisation methods will prove to be ineffective. Therefore finding a global minimum is very difficult and strongly depends on the optimisation method used. The ideal optimization method, should deal with these and other problems faced when calibrating the CEV model through its K&M approximation.

3.4.2 Regularisation

Chiarella, Craddock, and El-Hassan (2000) discuss the method of regularisation for the closed form solution of the Heston model in detail. Instead of minimizing $SqErr(\Omega, v_t)$ as in (77), we now add a penalty function $\zeta p(\Omega)$ to (77) and we get,

$$\min_{\Omega, v_t} SqErr(\Omega, v_t) = \min_{\Omega, v_t} \sum_{j=1}^{M} \sum_{i=1}^{N} w_{ij} [C_{MP}(S_t, K_i, \tau_i) - C_{\Omega, v_t}(S_t, K_i, \tau_j)]^2 + \zeta p(\Omega). \tag{79}$$

where the parameter ζ is the regularisation parameter. Chiarella et al. (2000) argue that "the philosophy behind this strategy is one of pragmatism". As the original problem, given by (78), is too difficult to find an exact solution to, the authors argue to replace the original problem "with one which is close to the original but does not possess the ill conditioning which makes the original intractable." This means that we do not solve the original problem, but one whose solution is very close to it. In this case its equation (79) that we try to solve. This is the essence of the regularisation method. When it is applied, we also require the bid/ask spread constraint as given for (78).

For the closed-form solution of the Heston (1993) model, as expressed in section 2.2, expression (79) is convex, as shown by Chiarella et al. (2000). Intuitively, this should also hold for the closed-form approximations, as they are to represent the original function and also for the approximations to the absent closed-form solution to the CEV model. Mikhailov and Nögel (2003) suggest using $\zeta p(\Omega, v_t) = ||\{\Omega, v_t\} - \{\Omega_0, v_t^0\}||^2$ with $\{\Omega_0, v_t^0\}$ the initial estimates of the parameters and variance. This makes the optimiser very dependent on the initial values and the optimisation process more 'local' as deviation from the initial parameters is penalized.

3.4.3 MATLAB's `lsqnonlin`

We use MATLAB's least-squares non-linear optimiser `lsqnonlin(fun, x0, lb, ub, options)` for the optimization problems given by (78) and (79). This function minimises the vector-valued function, `fun`, given the vector of initial parameter values, `x0`, and vectors of lower and upper bounds of the parameters given by `lb` and `ub`, respectively. `lqnonlin` uses an interior-reflective Newton method for large scale problems. These large scale problems are characterized by MATLAB as containing bounded/unbounded parameters, where the system is not under-determined, thus where we have less parameters than the number of equations to solve. For the simulation study we have a standard of 15 options a day, so this poses no problem and for the empirical application there are always enough options present on each day. For further reference on these methods, MATLAB suggests (Coleman and Li, 1994, 1996). Due to the dependence on the initial parameter values `x0`, the optimizer is a local optimizer and not a global one. There is no way to know whether the solution is global or local, however the restriction given by (78), checks if the solution is acceptable. If this is not the case, the process has to be run with a different `x0`. We apply this optimizer in conjunction with regularisation, which in that case makes the optimization process even more local.

4 Simulation Study

In this section we give an assessment of the accuracy of the K&M approximation of the CEV model, together with an assessment of its parameter estimation performance through a simulation study. In subsection 4.1 we look at the approximation of the Heston model (which is that of the CEV model with ξ kept fixed at 1/2) for which we have an accurate closed-form pricing function available, and in subsection 4.2 we look at the approximation of the CEV model. The accuracy assessment is only done for the approximation of the Heston model, since we do not have an accurate closed form solution for the CEV model when $\xi \neq 1/2$. For both approximations we conduct a simulation on time-series data, using data from the Heston model for both cases, as well as a cross-sectional data study, where for the CEV model we use data from the Heston model and from the CEV model with $\xi = 1.33$. We do not perform a accuracy or fit test for during the simulation studies, due to the fact that the fits are, in our case, virtually always zero, even when the K&M approximation is used.

4.1 Simulation Study of the Heston Model

In the first part of this subsection we test the accuracy of the K&M approximation when applied to the Heston model. The presence of an accurate closed form solution for the Heston model enables a good assessment of the accuracy of the K&M method. As benchmarks for the accuracy, we also compute the option prices from the Yang approximation and from Monte-Carlo simulation with an euler discretization. Because the K&M approximation for the Heston (or CEV) model adds corrective terms to the Black-Scholes model price, we also include the prices from the Black-Scholes model to get an indication of the added value from the approximation. In the second part, we perform an assessment of the parameter estimation process using the K&M approximation of the Heston model on simulated data from the Heston model.

4.1.1 Numerical Accuracy

To assess the accuracy of the K&M approximation of the CEV model, we look at the Heston model case, as we have a very accurate closed form pricing function for this case, given by equation (28). This is similar to what is done by Kristensen and Mele (2011). Our contribution is that we also look at longer term maturity options, where they only assess the performance for 1-month maturity options. The K&M approximation for the CEV model is given by equation (44). We obtain the approximation for the Heston model by setting the elasticity parameter equal to $\xi = 1/2$. For the approximation we take $N = 4$, thus *five* corrective terms to compute the prices. However, for the nuisance parameter σ_0 we take $\sigma_0 = \sqrt{v}$, as done by Kristensen and Mele (2011), such that the first term becomes zero and we keep four corrective terms (thus, we keep the terms $i = 1, 2, 3, 4$).

For the remaining parameters, we follow Kristensen and Mele (2011) and use the values as estimated by Bollerslev and Zhou (2002) on intra-day five minute returns for the Deutsche Mark U.S. Dollar spot exchange rates (with a sample from December 1, 1986 through December 1, 1996, containing a total of 2,445 observations). For the long term variance, Bollerslev and Zhou (2002) find $\alpha = 0.5172$, which is close to the sample mean of the underlying five-minute quadratic variation they find for the corresponding returns. They find a mean reversion parameter of $\kappa = 0.1465$, for which they say is on the high side when compared to what is reported other literature. The volatility of variance is equal to $\omega = 0.5786$, which means that, together with the other parameter values, the stationarity condition $2\kappa\alpha > \omega^2$ is violated. Despite of this, we will stick to these values, such that we can compare our results to those of Kristensen and Mele (2011). Lastly, for the correlation parameter, Bollerslev and Zhou (2002) find $\rho = -0.0243$, which is on the low side (almost zero), but confirms the negative correlation between asset prices and variance as found in other literature.

We also compute the prices through the Yang method, as described in section 3.2. The Yang approximation is given by equation (60). We only use two corrective terms in this case, thus $N = 2$ in terms of section 3.2, since for higher order terms (e.g. $N = 3$) we were not able to compute the integral of equation (59), using computer software package Maple. Kristensen and Mele (2011) also ran into this same problem and therefore they also stick to two terms. These two corrective terms do not include the mean reversion parameter κ of the CEV model (and thus the Heston model). Therefore, we only consider the Yang approximation as a benchmark for the numerical accuracy of the K&M approximation and do not use it for further calibration or pricing purposes.

As another benchmark for the pricing accuracy, we take prices computed through Monte-Carlo simulation, as described in section 3.3. This approach is often used in practice in absence of a(n) (accurate) closed-form solution for option prices from continuous time models. The Heston model allows us to fine tune the settings (e.g. setting the amount of runs needed to compute an option price) of the Monte-Carlo, such that the obtained prices are of decent accuracy. These settings are then used (in subsection 4.2) to compute prices from the CEV model, when $\xi \neq 1/2$, for which we do not have a closed form solution and have to rely on the Monte-Carlo simulation. For the Euler discretization, as described in section 3.3, we use an increment of $\Delta = 1/36$ for all considered maturities. Note that this is a smaller increment (and thus will lead to a higher accuracy) than $\Delta = 1/32$ as suggested by Andersen (2007). To keep the standard deviation low and to obtain a high accuracy for the option prices, we run a total of $N = 10^6$ simulations to compute each price as given by equation (74).

The generated option prices, for the various methods, are reported in table 1 below. The table reports the option prices for initial stock prices of $S_0 = 950$, $S_0 = 1000$ and $S_0 = 1050$ and maturities ranging from 1 month to 15 years. The strike price is kept fixed at $K = 1000$ for all considered options. The accuracy for each method is given in terms of the method's discrepancy (% Diff) with the Heston closed form solution, given by equation (26).

For short term maturities, of one and three months, the K&M-method outperforms all others, with a maximum error of 0.0381%. The Yang-method does much worse and has errors comparable to the Black-Scholes price. However, for the at-the-money and out-of-the money options ($S_0 = 1000$ and $S_0 = 950$), the errors from the Yang method are about half the size of the Black-Scholes pricing errors. For in-the-money options ($S_0 = 1050$), the Yang-method does slightly worse than the Black-scholes price. The Yang-method shows a performance worse than the Monte-Carlo simulation, for all but the 1 month in-the-money option price. The Monte-Carlo simulation is the overall second best in this case.

If we move to mid term maturities (3 and 5 years), the Yang method is the most accurate of all three analytical methods. With errors between 1,42% and 2.86%, it is quite accurate for mid term maturity options. The K&M method shows poor results, with errors around 5% for the 3 year maturity and around 15% for the 5 years maturity options. The Monte-Carlo simulation provides the best result in this case and outperforms all other methods. For long term maturities (10 and 15 years) both the K&M-method and the Yang-method show very poor performance and both methods are even outperformed by the Black-Scholes model. In this case, only the Monte-Carlo simulation outperforms the Black-Scholes model. Thus, the approximations are not suited for long term maturities.

The K&M-method seems to provide very accurate approximations for short term maturities. As the maturity increases, its approximations deteriorate. We probably need many more terms for the approximation to be accurate at longer term maturities. The Yang expansion does seem to be more robust to maturity increase, providing a more stable performance under maturity increases. The reason for this is probably that the convergence range of the time-to-maturity (i.e. the values of the maturity for which the expansion becomes better as more terms are added) is larger for the Yang expansion than for the K&M expansion. Thus the Yang expansion shows similar performance for short and mid term maturities. However, the amount of terms are not enough to actually show a satisfactory performance. For long-term options, both approximation methods have very large errors and only the Monte-Carlo simulation provides errors that are reasonably smaller than the Black-Scholes model.

Under the current specifications it took very long to generate the prices from the Monte-Carlo simulation. For the options with 15 years maturity it took about three whole days. The long running time makes the Monte-Carlo simulation inadequate for practical purposes, at least with the current specifications. Andersen (2007) notes that if the amount of simulations is practical ($<< 10^6$), the noise in the simulations will be much higher and lead to a much greater pricing error. Also, the author shows that as the increment size is decreased (raising the implementation time), the pricing error becomes smaller. For short term maturities, the K&M method outperforms the Monte Carlo simulation, especially if we consider that the K&M method has virtually no computation time. For the mid-term maturities the Monte Carlo simulations are the most accurate, with errors of 2% at the most. Despite its bad performance for the long-term-maturity options, the K&M approximation does show satisfactory results for most empirical application, because the options that are traded most actively are short-term-maturity options.

Table 1: **Comparison of Heston (1993) model option prices under alternative computation methods.** This table compares option prices for the Heston (1993) model computed using the closed-form solution of Heston (1993), the method of Kristensen and Mele (2011) with 4 leading terms, the asymptotic expansion of Yang(2006), the Black and Scholes (1973) model, and finally a Monte Carlo simulation. All options have strike price $K = 1000$. The parameter values of equation (32) and (33) are set equal to $\kappa = 0.1465$, $\alpha = 0.5172$, $\omega = 0.5786$, $\xi = 1/2$, $\rho = -0.0243$, and $r = 0$ as in Kristensen and Mele (2011). Panel A, B and C provide option prices with initial stock value equal to 950, 1000 and 1050 respectively. Every panel contains prices with maturities ranging from 1 month to 15 years.

Panel A: $S_0 = 950$

	Closed Form	Kristensen and Mele		Yang	
Maturity	Price	Price	% Diff	Price	% Diff
1 month	57.8425	57.8449	0.0042	57.7001	-0.2461
3 months	114.5037	114.5472	0.0381	113.7612	-0.6484
1 year	241.9157	243.3274	0.5835	237.4501	-1.846
3 years	398.2931	419.7043	5.3757	391.9478	-1.5931
5 years	484.9334	563.4212	16.1853	498.4784	2.7932
10 years	614.3378	1071.9434	74.4876	778.7189	26.7575
15 years	696.3898	1891.6078	171.6306	1108.573	59.1886

	Black Scholes		Monte Carlo	
Maturity	Price	% Diff	Price	% Diff
1 month	58.0456	0.3512	57.8964	0.0932
3 months	115.5141	0.8825	114.565	0.0536
1 year	249.4865	3.1295	242.5399	0.2582
3 years	430.268	8.0281	398.0423	-0.0631
5 years	539.3976	11.2313	480.1394	-0.9886
10 years	701.0094	14.1081	558.493	-9.0902
15 years	790.4399	13.5054	613.8702	-11.850

Panel B: $S_0 = 1000$

	Closed Form	Kristensen and Mele		Yang	
Maturity	Price	Price	% Diff	Price	% Diff
1 month	82.4766	82.4797	0.0038	82.2957	-0.2193
3 months	141.6661	141.7144	0.0341	140.7799	-0.6256
1 year	273.109	274.5896	0.5421	267.7368	-1.9671
3 years	433.8391	456.0133	5.1112	423.8804	-2.2955
5 years	522.8089	603.8489	15.5009	530.2433	1.422
10 years	655.6362	1126.958	71.8877	811.1394	23.7179
15 years	739.8337	1969.607	166.2229	1145.273	54.8014

		Black Scholes		Monte Carlo	
Maturity		Price	% Diff	Price	% Diff
1 month		82.6741	0.3414	82.6468	0.2064
3 months		142.6838	0.8887	141.7907	0.0879
1 year		280.8411	3.1962	272.7199	-0.1425
3 years		466.5947	8.224	429.0991	-1.0926
5 years		578.6332	11.5118	521.7423	-0.204
10 years		744.5042	14.4657	643.53	-1.8465
15 years		836.2772	13.8491	653.0236	-11.734

Panel C: $S_0 = 1050$

	Closed Form	Kristensen and Mele		Yang	
Maturity	Price	Price	% Diff	Price	% Diff
1 month	111.9021	111.9048	0.0023	111.7139	-0.1682
3 months	171.6132	171.6598	0.0271	170.6362	-0.5693
1 year	305.7206	307.2211	0.4908	299.5354	-2.0231
3 years	470.2239	492.9466	4.8323	456.7816	-2.8587
5 years	561.3212	644.4902	14.8167	562.7763	0.2592
10 years	697.3279	1181.247	69.3962	843.9638	21.0282
15 years	783.5466	2046.153	161.1399	1182.036	50.8571

		Black Scholes		Monte Carlo	
Maturity		Price	% Diff	Price	% Diff
1 month		112.0615	0.1424	111.5867	-0.2819
3 months		172.5803	0.5636	171.4811	-0.077
1 year		313.4961	2.5433	305.1866	-0.1747
3 years		503.5813	7.0939	467.3893	-0.6028
5 years		618.318	10.1541	550.0547	-2.0071
10 years		788.2289	13.0356	654.5337	-6.1369
15 years		882.2503	12.5971	762.8913	-2.6361

4.1.2 Parameter Estimation Heston Model

Using simulated data from the Heston model, we test the accuracy of the parameter estimates obtained by calibrating the K&M approximation (of the Heston model) in conjunction with the methods discussed in section 3.4. We perform the test on two different data sets, namely a simulated time-series data set and a simulated cross-sectional data set. For the accuracy test in 4.1.1, we use the K&M approximation, given by (44), with four ($N = 4$, no nuisance parameter) corrective terms. This was possible, since we computed the prices in Maple. To calibrate the model, we use MATLAB, such that the expression needs to be transferred from Maple to MATLAB. We are not able to get more than two($N = 2$) terms in MATLAB. However, due to the fact that the approximation is not convergent (at least for the lower terms), the accuracy for $N = 2$ is pretty much the same as for $N = 4$, as we show and explain in appendix D. For these reasons, we continue our analysis using only two corrective terms for the approximation.

Simulated Time-Series Data (Heston)

We start with generating option data in time series form. More precisely, we use the euler discretization of section 3.3 to generate a time series of daily spot prices and volatilities, that we use to generate daily option prices with Heston's (1993) closed-form solution, given by expression (26). Doing so, we yield a set of option prices in panel data form; for each time $t = 1, \ldots, n$ we generate a spot price and variance that are used to obtain the corresponding 15 option prices, with varying strikes and maturities, such that we end up with a set of $n \times 15$ option prices.

The time series of $n = 1000$ spot prices with corresponding variances is denoted by $(S_{\tilde{t}}, v_{\tilde{t}})_{\tilde{t}=1}^{1000}$. We use the subscript $\tilde{t}$ for the the $\tilde{t}$-th day to circumvent notational ambiguity with respect to the notation of section 3.3, where $\Delta t = 1$ stands for a time interval of 1-*year*, whereas $\Delta\tilde{t}$ stands for a time interval of 1-*day*. We assume that there are 260 trading days in one year, such that in terms of the euler discretization of section 3.3, where $t = 1$ is first year, the first day takes place at $t = \frac{1}{260}$.

The set of daily spot prices and variances is generated as follows. First set the parameter $\Delta = 260 \cdot 10^{-8}$, such that we have $n = \frac{1/260}{260 \cdot 10^{-8}} = 10^5$ iterations per daily spot price value. This large amount of iterations is to minimize the noise in the process and make it as smooth as possible (as we generate only one time series, computationally this is very comprehensible). More specifically, we generate the data by the following iteration process:

- **Step 1** Set the value of the subscript $\tilde{t} = 1$. The initial stock price and variance are set equal to $\dot{S}_{i=0} = 1500$ and $\dot{v}_{i=0} = 0.5172$ respectively.
- **Step 2** Set $i = 1$, and generate $\varepsilon^{(i)} \sim N(0, 1)$ and determine

$$\ln \dot{S}_i = \ln \dot{S_{i-1}} + (r - \frac{1}{2})\dot{v}_{i-1}\Delta + \sqrt{\dot{v}_{i-1}}\varepsilon_x^{(i)}\sqrt{\Delta} \tag{80}$$

$$\dot{v}_i = \kappa(\alpha - \dot{v}_{i-1})\Delta + \omega|\dot{v}_{i-1}|^{\xi}\sqrt{\Delta}(\rho\varepsilon_x^{(i)} + \sqrt{1-\rho^2}) \tag{81}$$

  We repeat this for $i = 2, \ldots, n$.
- **Step 3** Let $S_{\tilde{t}} = \dot{S}_n$ and $v_{\tilde{t}} = \dot{v}_n$.
- **Step 4** If $\tilde{t} = 1000$, then stop. Otherwise, set $\dot{S}_0 = S_{\tilde{t}}$, $\dot{v}_0 = v_{\tilde{t}}$ and update $\tilde{t} = \tilde{t} + 1$ and go back to step 2.

Completing the latter process yields a set of spot prices and corresponding variances, denoted by $(S_{\tilde{t}}, v_{\tilde{t}})_{\tilde{t}=1}^{1000}$. For the parameter values we use the same values as we did for the accuracy test. Thus, $\kappa = 0.1465$, $\alpha = 0.5172$, $\omega = 0.5786$, $\xi = 1/2$, $\rho = -0.0243$, and $r = 0$.

In the following, the set $(S_{\tilde{t}}, v_{\tilde{t}})_{\tilde{t}=1}^{1000}$ is used to generate a set of daily market option prices using the Heston's (1993) closed-form solution. For an option with time-to-maturity τ, strike K, and underlying with spot price $S_{\tilde{t}}$ and variance $v_{\tilde{t}}$, the generated market price is given by $C_{MP}(S_{\tilde{t}}, K, \tau) = C(S_{\tilde{t}}, v_{\tilde{t}}, \tilde{t}, \tilde{t} + \tau)$, where we have replaced T in $C(S_{\tilde{t}}, v_{\tilde{t}}, \tilde{t}, T)$, given by (28), by $T = \tau + \tilde{t}$. For each time $\tilde{t} = 1 \dots 1000$, we generate 15 option prices, with various strikes and maturities, through the following iteration process:

- **Step 1** Let let $j = 1$ and set the following values:

  $\tau_1 = 1/12$

  $\tau_2 = 1/4$

  $\tau_3 = 1,$

  which represent maturity times of 1 month, 3 months and 1 year respectively.

- **Step 2** Set initial value $K_{j=0} = 1200$ and set $i = 1$

- **Step 3** Update $K_i = K_{i-1} + 100$, set $\tau = \tau_j$. For $\tilde{t} = 1, \dots, 1000$ set,

$$C_{MP}(S_{\tilde{t}}, K_i, \tau_j) = C(S_{\tilde{t}}, v_{\tilde{t}}, K_i, \tilde{t}, \tilde{t} + \tau_j) \tag{82}$$

- **Step 4** If $i = 5$, continue to step 4, otherwise update $i = i + 1$ and go back to step 3.

- **Step 5** If $j = 3$, then stop, otherwise update $j = j + 1$ and set $i = 1$ and go back to step 2.

Going through the latter process, we obtain a set of option prices, $(((C_{MP}(S_{\tilde{t}}, K_i, \tau_j))_{\tilde{t}=1}^{1000})_{i=1}^{5})_{j=1}^{3}$, in panel data form, consisting of 15 option prices for each point in time $\tilde{t} = 1, \dots, 1000$. Cross-sectionally (e.g. on a certain time $\tilde{t}$)), the dataset consists of options with maturities of 1 month, 3 months and 1 year and strike prices of 1300, 1400, 1500, 1600, and 1700. We choose these maturities, since we know, from the previous section, that (for the Heston model) the K&M approximation is accurate for options with these maturities and inaccurate for longer maturity options. For other contingent claim prices, Kristensen and Mele (2011) show that their approximation is accurate for short term maturities, but deteriorates as time-to-maturity increases. Therefore, we also expect this to hold for general cases of the CEV model. The choice of the strike prices is based on the choice of the initial stock price $S_0 = 1000$, for which we then can consider the options as (far) OTM, ATM, and (far) ITM options. As $S_{\tilde{t}}$ varies over time, this does not need to hold for all $\tilde{t}$, e.g. we could have $S_{\tilde{t}} = 1200$ in which all call options at that particular time are ITM. This could be adjusted by letting the strikes vary per option such that the moneyness $S_{\tilde{t}}/K$, is the same for all $\tilde{t}$. We choose not to do the latter, since in practice empirical data might be such that at a particular time all options are ITM or all options are OTM. Thus, keeping the simulated data as it is makes the tests more relevant for empirical applications.

The estimation of the Heston model parameters on panel data cannot be done directly by the methods discussed in section (3.4). The reason is that apart from the parameter set $\Omega = \{\kappa, \alpha, \omega, \rho\}$, we also need to estimate the *time-dependent* variance v_t (since we have daily option prices). The calibration methods of section 3.4 only apply to a cross-sectional set of option prices (at a certain time $\tilde{t}$). To make any of the estimation methods applicable for panel data, we add a preparatory step, leading to a two-step estimation process. The first step consists of an estimation of a time-dependent parameter set, $\Omega_{\tilde{t}} = \{\kappa_{\tilde{t}}, \alpha_{\tilde{t}}, \omega_{\tilde{t}}, \rho_{\tilde{t}}\}$ and the (time-dependent) variance parameter $v_{\tilde{t}}$. This step is not consistent with the underlying theory, since according to the model's dynamics, all parameters ($v_{\tilde{t}}$ is not a parameter) are assumed to be constant and not varying over time. In the second step we throw away all estimated parameters sets $\Omega_{\tilde{t}}$ and only keep the estimated variance time-series $\hat{v}_{\tilde{t}}$, for $\tilde{t} = 1, \ldots, 1000$. Plugging the estimated variance time-series $(\hat{v}_{\tilde{t}})_{\tilde{t}=1}^{1000}$ in the pricing formula, leaves four parameters to estimate, namely $\Omega = \{\kappa, \alpha, \omega, \rho\}$. We estimate these parameters using all available panel data.

More explicitly, to estimate the parameters we execute the following two steps,

- **Step 1**For $\tilde{t} = 1, \ldots, 1000$, we minimize

$$\min_{\Omega_{\tilde{t}}, v_{\tilde{t}}} \sum_{i=1}^{5} \sum_{j=1}^{3} [C_{MP}(S_{\tilde{t}}, K_i, \tau_j) - C_{\Omega_{\tilde{t}}, v_{\tilde{t}}}(S_{\tilde{t}}, K_i, \tau_j)]^2 \tag{83}$$

  which yields estimated sets $\{\hat{\Omega}_{\tilde{t}}, \hat{v}_{\tilde{t}}\}$. We throw away all estimated parameters and only keep the time-series $(\hat{v}_{\tilde{t}})_{t=1}^{1000}$.

- **Step 2** Using $(\hat{v}_{\tilde{t}})_{t=1}^{1000}$ we minimize

$$\min_{\Omega} \sum_{\tilde{t}=1}^{1000} \sum_{i=1}^{5} \sum_{j=1}^{3} [C_{MP}(S\tilde{t}, K_i, \tau_j) - C_{\Omega, \hat{v}\tilde{t}}(S_{\tilde{t}}, K_i, \tau_j)]^2. \tag{84}$$

This yields the estimated set of parameters $\hat{\Omega} = \{\hat{\kappa}, \hat{\alpha}, \hat{\omega}, \hat{\rho}\}$ and the set of estimated variances, from the first step. The results of this estimation process on generated option data are discussed next. We start with the estimation of the variances and close with the parameter set $\hat{\Omega}$.

Figure 1 below displays the estimated variance time-series $(\hat{v}_{\tilde{t}})_{t=1}^{1000}$, together with the generated variance time-series $(v_{\tilde{t}})_{t=1}^{1000}$. The figure shows very small errors for the estimated variances, as the plots are very much alike. This result indicates the adequacy of the estimation process.

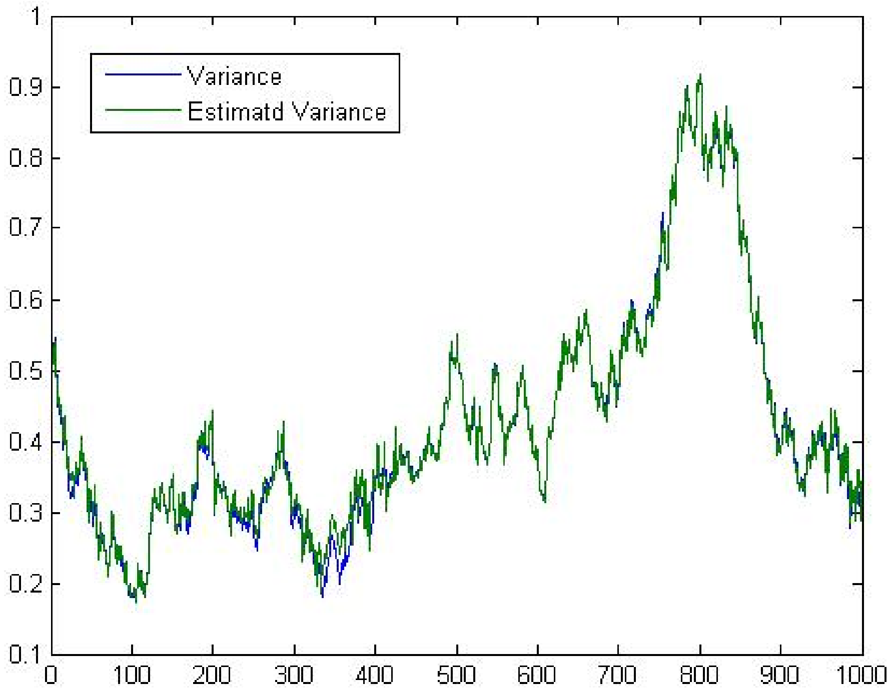

Figure 1: Time series for estimated and generated variance. The variances are estimated using the K&M approximation for the Heston model on data generated from the Heston model.

Table 6 contains the values of the estimated parameters $hat\kappa$, $hat\alpha$, $hat\omega$, and $hat\rho$ with their corresponding estimation errors. The errors are computed as

$$\%\text{ error} = 100 \times \frac{\text{Real value} - \text{Estimated value}}{\text{Real value}},$$

which is the percentage error of the estimation for each parameter. This is only done for the final value of each parameter, thus the value that we get in the second step of the estimation process. We also report the initial values used in `MATLAB`'s non-linear least squares optimizer `lsqnonlin`, to minimize the expression given by (84). We compute the initial values by taking the average of the parameters at each time $\tilde{t}$, e.g. the initial value κ_0 is computed as

$$\kappa_0 = \frac{1}{1000} \sum_{\tilde{t}=1}^{1000} \hat{\kappa}_{\tilde{t}}.$$

For the initial value of the daily parameter estimates, we use $\Omega_0 = \{0.1, 0.1, 0.1, 0.05\}$ and $v_0 = 0.1$.

The results of the estimated parameters show a slight bias; for the mean reversion parameter κ there is an estimation error around % 10, whereas for the other parameters it is around % 25. The bias could be due to the fact that we use a large set of daily data and therefore our data contains a lot of noise. Another reason that might be of importance for the distortion in our estimates, is that a lot of the generated option prices are very far OTM options, for which we know that the K&M method produces bad results, as shown in section 4.1.1 and also reported by Kristensen and Mele (2011).

Table 2: **Estimated parameters of the Heston model by calibrating its K&M approximation on a simulated set of panel data**
This table provides the estimates of the parameters of the Heston model that are obtained by calibrating its K&M approximation on simulated panel data. The parameters are estimated using the standard MSE approach. The real parameter values, used for generating the data, are given, together with the initial parameter values used in MATLAB's non-linear least squares optimizer `lsqnonlin`. The errors represent the percentage deviation of the estimated values from the real values. There are no standard deviations reported, since we only have one set of estimated parameters.

	κ	α	ω	ρ
Real value	0.1465	0.5172	0.5786	-0.0243
Initial value	0.118686	0.102557	0.136493	-0.02024
Estimated value	0.162029	0.391395	0.438066	-0.01841
% error	-10.6003	24.3242	24.28861	24.24489

Through this process we only have one estimation for each parameter. For a complete statistical inference, to check if the estimation process is appropriate, we would have to repeat this process a large amount of times, say 1000. This would mean that we need to generate 1000 sets of option prices, each consisting of a time period of 1000 days (since $\tilde{t} = 1, \ldots, 1000$). Generating one set of options for 1000 days takes a few hours and estimating the parameters through the two step process takes another few hours. Thus repeating this process as often as to perform a complete statistical inference is not feasible. And we have therefore only done the estimation once, such that we get an idea of the appropriateness of the estimation process. The errors in the parameter estimates are quite small, although still there was a bias. To perform a more complete statistical inference, in the next section we generate a large amount of cross-sectional data sets, consisting of 15 options each. This is also more relevant, since our for empirical purposes, we estimate the parameters on a daily basis, using one set of cross-sectional option prices.

Simulated Cross-Sectional Data (Heston)

The time series approach shows that we can estimate the parameters with a relatively small error, but it does not say anything about the statistics of the parameters. In this section we investigate the statistics of the estimated parameters of the Heston model, by generating many sets of independent spot prices and variances, which we use to generate corresponding sets of option price data. The simulated data consists of 1000 sets of spot prices and variances, generated by running the Euler discretization 1000 times, starting at $t = 0$ up to some fixed time $t1/12 =$(thus 1-month). Each generated spot price, at time $t = 1/12$, with corresponding variance, will be used to determine a set of 15 options with varying strikes and maturities.

We apply the Euler-discretization of section 3.3 for the Heston model (thus for $\xi = 1/2$), such that we generate a set of one month old spot prices and corresponding variances, with initial values $S_0 = 1500$ and $v_0 = 0.5172$. The increment size is set equal to $\Delta = \frac{1}{1{,}200{,}000}$, such that we end up with a total of $n = \frac{t}{\Delta} = \frac{1/12}{1/1{,}200{,}000} = 10^5$ steps to generate one 1-month old spot price and variance. The generated spot prices and variances are then, similar to the time-series case, used to generate a dataset of option prices with Heston's (1993) closed-form solution.

Setting our increment-size equal to $\Delta = \frac{1}{1{,}200{,}000}$, we need a total of $n = \frac{t}{\Delta} = \frac{1/12}{1/1{,}200{,}000} = 100,000$ steps to generate one 1-month old spot price and variance. Having these parameters set, we generate the spot price data according to the following scheme,

- **Step 1** Set the initial value $k = 1$.
- **Step 2** Set the initial values of the spot price and variance equal to $\dot{S}_0 = 1500$ and $\dot{v}_0 = 0.5172$ respectively.
- **Step 3** Set $i = 1$, and generate $\varepsilon^{(i)} \sim N(0, 1)$ and determine

$$\ln \dot{S}_i = \ln \dot{S_{i-1}} + (r - \frac{1}{2})\dot{v}_{i-1}\Delta + \sqrt{\dot{v}_{i-1}}\varepsilon_x^{(i)}\sqrt{\Delta} \tag{85}$$

$$\dot{v}_i = \kappa(\alpha - \dot{v}_{i-1})\Delta + \omega|\dot{v}_{i-1}|^{\xi}\sqrt{\Delta}(\rho\varepsilon_x^{(i)} + \sqrt{1-\rho^2}) \tag{86}$$

  This step is repeated for $i = 2, \ldots, n$.
- **Step 3** Set $S_k = \dot{S}_n$ and $v_k = \dot{v}_n$ and update $k = k + 1$.If $k < 10001$, go back to step 2 and repeat the process. Otherwise, stop.

This yields a set of independently generated spot prices and variances, denoted by $(S_k, v_k)_{k=1}^{1000}$. The option prices are generated in the same sense as with the time series approach, but the time-series variances $v_{\tilde{t}}$ and the spot prices $S_{\tilde{t}}$ are now replaced by their time *in*dependent counterparts S_i and v_i. This yields a new set of generated option prices denoted by $((C_{MP}(S_k, K_i, \tau_j)_{i=1}^{5})_{j=1}^{3})_{k=1}^{1000})$. We have the same (as the time-series case) maturities of 1-month, 3-months and 1-year and five strike prices ranging from 1300 to 1700, with a difference of 100 between two consecutive strike prices.

Now, we use the latter set of option prices to test the parameter estimation process. To estimate the parameters we start with the standard dollar-based MSE approach, given by (77) and also the regularisation approach, as described in 3.4.2, given by (79). Also, we differ from the time-series approach with the parameter values we choose to generate the data. We change the parameter value of $\kappa = 0.1465$ to $\kappa = 0.2$, such that the restriction $2\kappa\alpha > \omega^2$, as suggested by Mikhailov and Nögel (2003), holds.

We start with the normal dollar-based MSE approach (although we don't take the mean (squared error), which is just multiplying the expression below by a constant-the expression is equivalent). For each set of option prices, we minimize

$$\min_{\Omega, v_k} \sum_{i=1}^{5} \sum_{j=1}^{3} [C_{MP}(S_k, K_i, \tau_j) - C_{\Omega, v_k}(S_k, K_i, \tau_j)]^2 \tag{87}$$

yielding $\tilde{\Omega}_k = \{\tilde{\kappa}_k, \tilde{\alpha}_k.\tilde{\omega}_k, \tilde{\rho}_k\}$, for $k = 1, \ldots, 1000$. We take the mean of the estimated set of 1000 parameters, as the estimation of the corresponding parameter. For example, for the parameter κ, its estimate is given by

$$\hat{\kappa} = \frac{1}{N} \sum_{k=1}^{N} \tilde{\kappa}_k. \tag{88}$$

For the other parameters, the estimates are computed in the same way. Furthermore, we investigate the statistical properties of the estimated parameters $\hat{\Omega} = \{\hat{\kappa}, \hat{\alpha}.\hat{\omega}, \hat{\rho}\}$, by determining the percentage error from the real value and the standard deviation of the estimated parameters. The percentage error is given by

$$\text{Error } \% = 100 \times \frac{\text{Real value} - \text{Mean}}{\text{Real}}$$

For the standard deviation of the estimations, we take the sample standard deviation, for the parameter kappa, for example, this is given by

$$sd(\kappa) = \sqrt{\frac{1}{N} \sum_{i=1}^{N} (\tilde{\kappa}_i - \hat{\kappa})^2}$$

and for the other parameters it is computed in this same way. We minimize (87) using MATLAB's nonlinear optimization function `lsqnonlin`, which requires a set of initial values for the parameters and the variance.

The estimation results are shown below in table 3. Since the optimizer in MATLAB depends upon the initial parameter values, we estimate the parameters using different sets of initial values. Each of panels A, B and C show the results for a different set of initial parameter values. The initial value for the variance is kept fixed at $v^0 = 0.1$, because the real variance is different for each set of cross-sectional options. In Panel A we take the real values as the initial values, whereas in panels B and C we take the initial values deviating upwards and downwards from the real values respectively.[2] Looking at Panel A, we see that the means of the estimates are very far away from the real values. Thus, leaving no indication of unbiased estimates (in which case the % error would be around zero). Apart from the estimate for the mean reversion parameter κ, the estimates are also not of the same order as the real parameters. This holds especially for ω and ρ, of which even the sign is wrong (positive instead of negative and vice versa).

Table 3: **Summary statistics of estimated parameters of the Heston model obtained by calibrating its K&M approximation on simulated cross sectional data.**
This table provides the statistics of the parameter estimates of the Heston model obtained by calibrating the K&M approximation (of the Heston model), through the standard MSE approach, on 1000 simulated cross sectional data sets, each containing 15 option prices. The real parameter values, used for generating the data, are reported together with the initial parameters used in MATLAB's non-linear least squares optimizer `lsqnonlin`. Panels A, B and C give the estimation results for the different sets of initial parameter values. Each panel reports the mean and standard deviation that are calculated across a total of 1000 estimates for each parameter. The Error % is the percentage deviation of the mean from the real value.

	κ	α	ω	ρ
Real value	0.2	0.5172	0.5786	-0.0243
Upperbound	1	1	1	1
Lowerbound	0	0	0	-1
		Panel A		
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$
Initial value	0.2	0.5172	0.5786	-0.0243
Mean	0.116608	1.122317	-0.14439	0.016738
Error %	-41.6956	-116.999	124.9547	168.8815
Standard deviation	0.288232	1.199188	0.361775	0.050206
		Panel B		
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$
Initial value	0.1	0.1	0.1	-0.01
Mean	0.328496	0.180695	-0.16732	0.017291
Error %	64.24833	65.0628	128.9186	171.1573
Standard deviation	0.05546	0.166172	0.11846	0.074594
		Panel C		
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$
Initial value	0.9	0.9	0.9	-0.1
Mean	0.862020	0.700457	-0.00616	0.058734
Error %	-331.011	-35.4325	101.0652	341.7057
Standard deviation	1.693699	0.649912	0.394269	0.268606

[2] Because the real value of ρ is an order 10 smaller than the other parameters, we choose smaller values than for the other parameters for its initial value in panels B and C.

In the figure below we have plotted the histograms of the parameter estimates, using the initial values of Panel A of table 3. The parameters seem nicely peaked, and show the properties we would expect from the statistics reported in table, as the mean, which is at the peaks, and the standard deviations, which correspond with the width of the histograms.

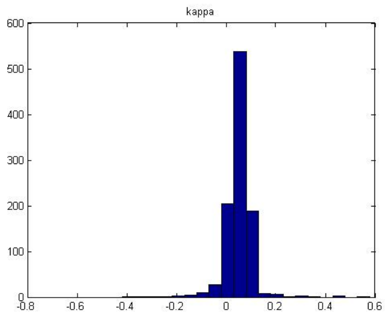

(a) parameter κ

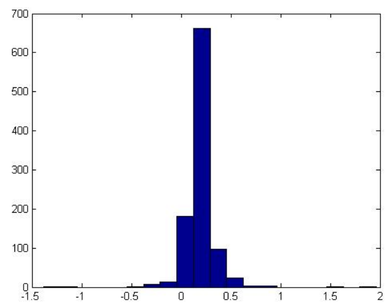

(b) parameter α

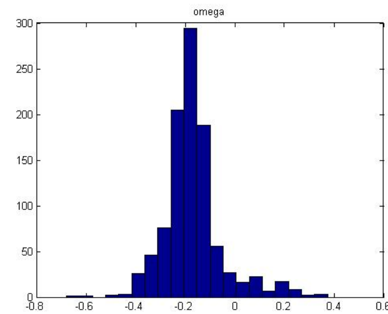

(c) parameter ω

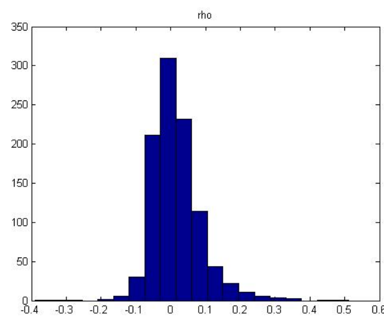

(d) parameter ρ

Figure 2: Parameter values estimated using a standard dollar-based MSE

The results in table 3 were disappointing, as the means of the estimates had very large errors and were not of the same order as the original parameters, whatsoever. Thus, the standard MSE method does not seem to be appropriate to estimate the parameters using the K&M approximation of the Heston model. For this reason we move to the regularisation method of Chiarella et al. (2000), as described in 3.4.2. As we mention in section 3.4.2, the regularisation method allows for convergence in the estimation process of the Heston model and this might also lead to better results when using the K&M approximation of the Heston model.

Applying the regularisation method, we minimise the following expression:

$$\min_{\Omega_k, v_k} \sum_{i=1}^{5} \sum_{j=1}^{3} [C_{MP}(S_k, K_i, \tau_j) - C_{\Omega_k, v_k}(S_k, K_i, \tau_j)]^2 + ||\{\Omega_k, v_k\} - \{\Omega_0, v^0\}||, \tag{89}$$

with Ω_0 and v^0 denoting the initial values of the parameters and the variance, respectively. The subscript k, denotes the estimation over the $k - th$, set of cross-sectional options. The extra term $||\{\Omega_k, v_k\} - \{\Omega_0, v^0\}||$ introduces much depends on the initial values for the parameters and variance. To make sure that the dependence upon the initial values is not too much, we estimate the model using different sets of initial parameters. We minimize the above expression (89) using MATLAB's non-linear least squares optimizer `lsqnonlin`.

The results from the parameters estimation, using the regularisation method, are shown in table 5. The table reports the real value of each parameter, together with its corresponding lower and upper bound in MATLAB's `lsqnonlin`. As in table 3, panels A, B, and C show the estimation results for the different sets of initial parameters used for the estimation. And, as in 3, we keep the initial value for the variance fixed at $v^0 = 0.1$ for all panels. Looking at panel A, we see that the mean of the estimated parameters are certainly in the same range as the real parameter values, indicating that the estimation process is adequate. The errors % show a great improvement to those obtained by using the MSE approach. Although they are far from being zero, they are very small. The error is very likely due to the error in the K&M approximation, as shown in table 1. The standard deviations are also quite small, especially for the volatility of variance ω and correlation parameter ρ, meaning that the estimates are quite consistent. The improvement could be explained by the fact that the regularisation is more stable, because of the extra penalty term. This seems helpful to deal with inaccuracies that are present in the K&M approximation.[3]

Table 4: **Summary statistics of estimated parameters of the Heston model obtained by calibrating its K&M approximation, through the regularisation approach, on generated cross-sectional data.**
This table provides the statistics of the parameter estimates of the Heston model, obtained by calibrating its K&M approximation , through the regularisation method, on 1000 simulated cross-sectional data sets, each containing 15 option prices. The real parameter values, used for generating the data, are reported, together with the lower and upper-bound values used in MATLAB's non-linear least squares optimizer `lsqnonlin`. Panels A, B and C give the estimation results for the different sets of initial parameter values used for the regularisation process. Each panel reports the mean and standard deviation calculated across a total of 1000 estimates for each parameter. Error % is the percentage deviation of the mean from the real value.

	κ	α	ω	ρ
Real value	0.2	0.5172	0.4	-0.0243
Upperbound	1	1	1	0
Lowerbound	0	0	0	-1
Panel A				
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$
Initial value	0.1	0.1	0.1	-0.01
Mean	0.2625	0.5486	0.3636	-0.0278
Error %	-31.25	-6.071152359	9.1	-14.4033
Standard deviation	0.1323	0.0416	0.0072	0.0074
Panel B				
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$
Initial value	0.2	0.5172	0.4	-0.0243
Mean	0.2661	0.5491	0.3641	-0.0279
Error %	-33.05	-6.16783	8.975	-14.8148
Standard deviation	0.133374	0.041713	0.00815	0.007398
Panel C				
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$
Initial value	0.9	0.9	0.9	-0.9
Mean	0.181397	0.47166	0.307461	-0.04654
Error %	9.301251	8.80519	23.13474	-91.5269
Standard deviation	0.098635	0.107432	0.091028	0.028252

[3] We have adjusted the upperbound of ρ from the pevious table. However, we found that changing this the upperbound for the regularisation does not lead to any different results in the estimates (not reported).

The results of panel B, of table 5, are very similar to those of panel A, implying that the shift in the initial values does not affect the results much. If we look at the results in Panel C, the change in initial values does result into a deviation that is much larger. The mean reversion parameter κ has become more accurate, whereas the error in the correlation parameter ρ has increased a lot. This last results can be expected, since the initial value used is of a completely different order for the correlation parameter. Still, the parameter value is in the same range as the real value, since we have to take into account that the Heston model is not that sensitive to changes in the parameters, although very sensitive to changes in the spot variance. Thus, although the estimates are not unbiased in any case, likely due to the inaccuracy of the K&M approximation, the results are quite good and the errors remain small when the initial values are changed.

The histogram plots of the parameters, in the figure ?? below, show the bias in the estimates: none of the modes of the estimated parameters sets, are at the true value of the parameters. For the parameter κ the histogram is even bimodal. If we take all the above into account, the regularisation method together with the K&M approximation does lead to some error, but this is less than when we combine the K&M approximation with the regular MSE approach. The question remains whether, the bias of the parameter is due to the inaccuracy of the approximation or due to the estimation process. To check this, we also estimate the parameters, using Heston's (1993) closed-form solution.

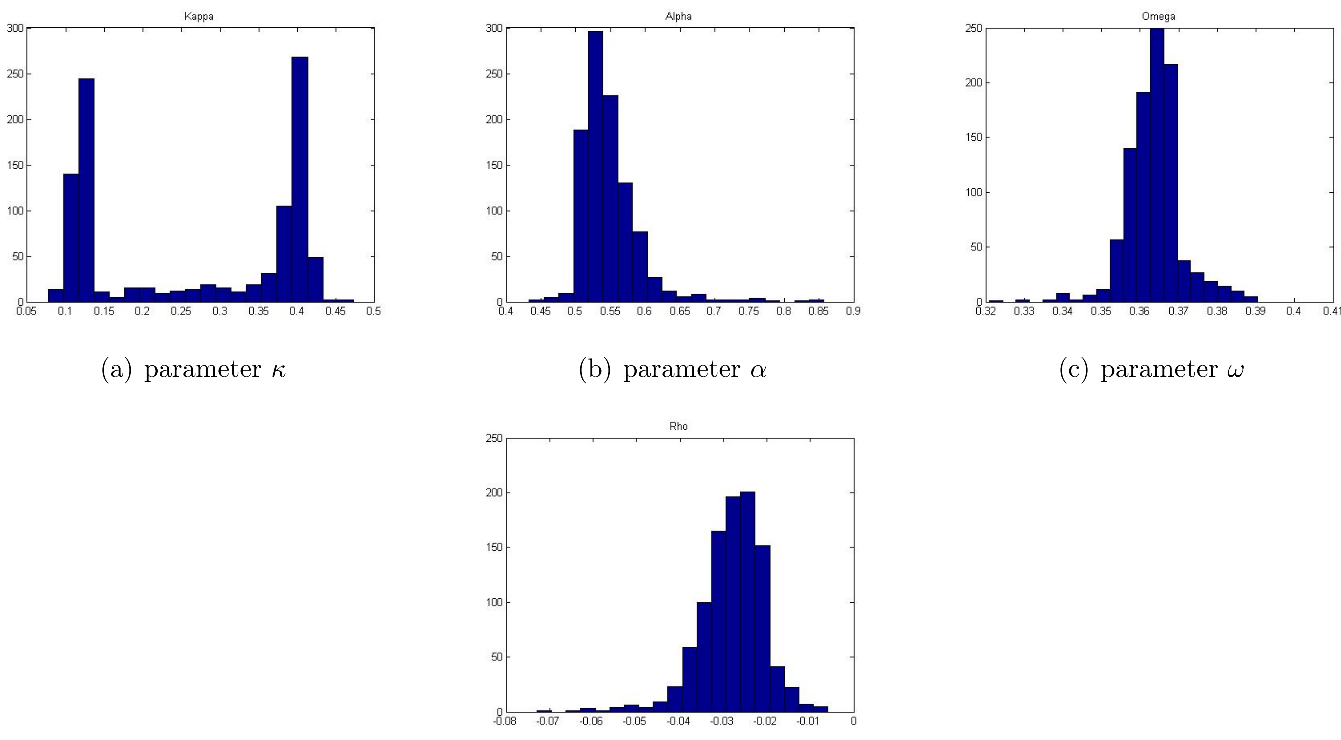

(a) parameter κ (b) parameter α (c) parameter ω

(d) parameter ρ

Figure 3: Parameter distributions using Regularisation with initial parameters of panel A in table 5.

Similarly to the above process for the estimation of the K&M approximation, we calibrate the closed-form pricing formula (26), on the cross-sectional data sets, using the regularisation approach. To keep it brief, we only use one set of initial values. The results from this estimation process are given by table 5 below. The table shows that all parameter estimates are highly accurate, which is a strong indication that the errors in the parameter estimates from the K&M approximation, are due to the inaccuracies of the approximation rather than the estimation process. One thing worth noting is the time it took to run the entire estimation process. For one set of initial parameters, for Heston's (1993) closed-form solution, it took 26720 seconds or 7 hours and (approximately) 25 minutes, while the corresponding run using the K&M approximation took 5096 seconds or 1 hour and (approximately) 25 minutes. Thus, it took more than five times longer to estimate the parameters using the closed-form solution than the K&M approximation.

Table 5: **Summary statistics of estimated parameters of the Heston model obtained by calibrating its closed form solution through the regularisation approach on generated cross-sectional data.**
This table provides the statistics of the parameter estimates of the Heston model, obtained by calibrating Heston's (1993) closed form solution, through the regularisation method, on 1000 simulated cross-sectional data sets, each containing 15 option prices. The real parameter values, used for generating the data, are reported, together with the lower and upper-bound values used in MATLAB's non-linear least squares optimizer `lsqnonlin`. Panel A gives the estimation results for one set of initial parameter values (the same as the real values) used for the regularisation estimation process. The panel reports the mean and standard deviation calculated across the total of 1000 estimates for each parameter. Error % is the percentage deviation of the mean from the real value.

	κ	α	ω	ρ
Real value	0.2	0.5172	0.4	-0.0243
Upperbound	1	1	1	0
Lowerbound	0	0	0	-1
		Panel A		
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$
Initial value	0.1	0.1	0.1	-0.01
Mean	0.201125	0.517045	0.400052	-0.02458
Error %	0.562528	-0.02992	0.012883	1.166978
Standard deviation	0.008855	0.003561	0.008593	0.00646

4.2 Simulation Study of the CEV Model

Since an accurate closed-form solution is available for the CEV model, when $\xi = 1/2$ (the Heston model), the K&M approximation becomes really useful for cases when $\xi \neq 1/2$, i.e. when we do not have a(n) (accurate) closed-form pricing function. In this section we test the performance of the K&M-approximation when $\xi \neq 1/2$. More specifically, we add ξ as a parameter to estimate in the approximation, rather than keeping it fixed at a certain value. The addition of an extra parameter will lead to more insecurity in the parameter estimates. To get an indication of this effect, we start with estimating ξ on data from the Heston model. We repeat the process for both time-series data and cross-sectional data sets. Furthermore, we use Monte-Carlo simulation to generate option prices for the case when $\xi = 1.33$ and we don not have a closed form solution available. We only generate *one* set of cross-sectional options for $\xi = 1.33$, because it takes a lot of time. For this reason also, we don not include a time-series analysis for this case.

4.2.1 Estimating ξ From The Heston Model

In section 4.1.2, we show that the calibration of the K&M approximation for the Heston model, yields small errors in the estimated parameters. Here we do the same analysis, namely calibrating the approximation on time-series and then cross-sectional data sets, with ξ as an additional estimable parameter. Thus, the set of parameters becomes $\Omega = \{\kappa, \alpha, \omega, \rho, \xi\}$. This allows us to see if the calibration method is still adequate when we need to estimate this extra parameter, as we have to for the CEV model in general.

Simulated Time-Series Data (CEV)

We start with time-series data that we used in the previous section. Also, we will use the general estimation process in this case, not regularisation, such that we can make a fair comparison on how the extra-parameter influences the parameter estimation results. Using the same two step estimation process as in 4.1.2, we yield a set of estimated variances $(\hat{v}_{\tilde{t}})_{\tilde{t}=1}^{1}000$ and a set of estimated parameters, which is $\hat{\Omega} = \{\hat{\kappa}, \hat{\alpha}, \hat{\omega}, \hat{\rho}, \hat{\xi}\}$ in this case.

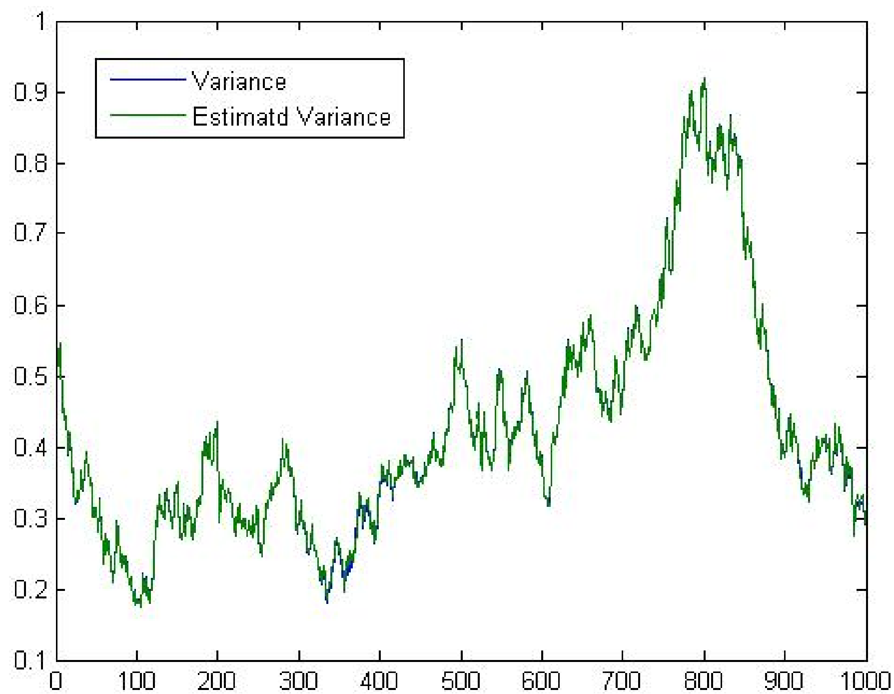

Figure 4: Time-series plot for the estimated and generated variances. The variances are estimated using the K&M approximation for the CEV model on data generated from the Heston model.

Figure 4 above displays the estimated variance time-series $(\hat{v}_{\tilde{t}})_{\tilde{t}=1}^{1}000$ together with the generated variance time series $(v_{\tilde{t}})_{\tilde{t}=1}^{1}000$, for when ξ is added as an extra parameter. If we compare this with the results in figure 1, we find that the extra freedom in the parameter ξ actually leads to better variance estimates. For comparison, the RMSE, between the estimated and generated variances, decreases from 0.0175, for when ξ was kept fixed, to 0.0034,when ξ is added to the set of parameters. The extra freedom in the parameter estimates is thus valued for the variance estimates. It is very likely that this is induced by the inaccuracy in the K&M approximation, which now went from the variance estimates to the estimates (plural, since for the time-series approach we estimate all parameters on a daily basis first, from which we also get the variances) of ξ.

In the table below we have the parameter estimates over the entire sample. As in 4.1.2, we use the averages of each parameter over time, as the initial value for the estimates over the entire panel data set. Also, the percentage error (Error %) is computed as in section 4.1.2. The estimation results are given in table 6 below. The errors are larger than what we found in table 3. This is to be expected, due to the extra parameter ξ. However, what we do note is that the initial parameter values, which are the averages over time, are much closer to the real values; much closer than the estimates for this case and also closer than the initial values of table 3. These accurate initial values, correspond to what we found for the variances, which are also estimated over time (similar to the parameter estimates of which we use the average as the intial values).

Table 6: **Estimated parameters of the CEV model by calibrating its K&M approximation on a simulated set of panel data**
This table provides the estimates of the parameters of the CEV model that are obtained when we estimate its K&M approximation on simulated panel data from the Heston model. The real parameter values, used for generating the data, are given, together with the initial parameter values used in MATLAB's non-linear least squares optimizer `lsqnonlin`. The errors represent the percentage deviation of the estimated values from the real values. No standard deviations are reported, since we only have one set of estimated parameters.

	κ	α	ω	ρ	ξ
Real value	0.1465	0.5172	0.4	-0.0243	0.5
Initial value	0.1562	0.5266	0.5557	-0.0245	0.568
Estimated value	0.116925	0.542058	0.617076	-0.02631	0.685152
Error%	20.18788	-4.80631	-54.269	-8.27442	-37.0304

Cross-Sectional Data (CEV)

In 4.1.2 we found that the parameter estimates from the K&M approximation on cross-sectional Heston model data, yielded small errors in the estimated parameters. The errors in the parameters are likely due to the inaccuracy of the approximation, since using the closed-form solution led to virtually no errors for the means of the parameter estimates (unbiased estimates), as showen by table 5. Here we use the same set of cross-sectional option prices, denoted by $((C_{MP}(S_k, K_i, \tau_j))_{i=1}^{5})_{j=1}^{3})_{k}^{1000}$, with ξ added as an extra estimable parameter, just as we did for the time-series data. We only use the regularisation method for the estimation.

The results for the parameter estimates are shown in the table 7 below. The table is similar to 5(but of course with ξ added to the parameters), with three panels with different initial parameters each. Again we keep the initial value of the variance fixed at $v^0 = 0.1$. Looking at Panel A, we see that the introduction of ξ leads to larger errors (than in (the corresponding) Panel A of table 5) in the estimates of parameters κ and α, whereas the estimates for ω and ρ have remained quite the same. The error in ξ is also quite large, indicating that we can not estimate ξ with high accuracy. If we deviate the initial values from the real values to those of Panel B, we see a large improvement for the parameter ξ, with the error almost half of what it was, whereas the other estimates have remained roughly the same. The parameter ω also deteriorates in this case. Moving to panel C, we see that all other parameters remain roughly the same as in table Panel B, although ω deteriorates again. Also, in this case the estimate of ξ is much worse than in the previous cases, which suggests that ξ is very dependent upon its initial parameter value. The results for parameters κ and α are worse than in table 5, as we add the parameter ξ. For the other two parameters it actually makes little difference. Overall, there are some difficulties in estimating the parameter ξ and the estimates are strongly dependent upon the initial value provided in this case.

Table 7: **Summary statistics of estimated parameters of the CEV model by calibrating its K&M approximation using regularisation on generated cross-sectional data from the Heston model.**
This table provides the estimates of the parameters of the CEV model, obtained by calibrating its K&M approximation on 1000 simulated cross-sectional data sets, each containing 15 option prices, from the Heston model. The real parameter values, used for generating the data, are reported, together with the lower- and upper-bound values used in MATLAB's non-linear least squares optimizer `lsqnonlin`. Panels A, B and C give the estimation results for the different sets of initial parameter values used for the regularisation process. Each panel reports the mean, standard deviation and t-statistic calculated across a total of 1000 estimates that we have for each parameter. Error % gives the percentage deviation of the mean from the real value.

	κ	α	ω	ρ	ξ
Real value	0.2	0.5172	0.4	-0.0243	0.5
Upperbound	0.3	1	1	0	1
Lowerbound	0	0	0	-1	0
Panel A					
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$	$\hat{\xi}$
Initial value	0.2	0.5172	0.4	-0.0243	0.5
Mean	0.3013	0.650906	0.394137	-0.02775	0.613583
Error %	50.6501	25.8519	-1.465807	14.185	22.7165
Standard deviation	0.248621	0.166079	0.077261	0.007207	0.297002
Panel B					
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$	$\hat{\xi}$
Initial value	0.1	0.1	0.1	-0.01	0.1
Mean	0.299645	0.64049	0.379142	-0.02768	0.551346
Error %	49.82248	23.83791	-5.21444	13.92257	10.26922
Standard deviation	0.242167	0.164426	0.085033	0.007657	0.332641
Panel C					
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$	$\hat{\xi}$
Initial value	0.9	0.9	0.9	-0.1	0.9
Mean	0.3233	0.65854	0.451463	-0.02828	0.802947
Error %	-61.6502	-27.328	-12.8657	-16.3849	-60.5895
Standard deviation	0.252857	0.175501	0.077427	0.007587	0.272114

4.2.2 The CEV Model with $\xi = 1.33$

Jones (2003) estimates the parameter ξ equal to 1.33. This means that the process is highly volatile and this could lead to problems in the parameter estimation. In this case, we do not have a closed form solution to generate the data and thus to test our model we using option prices generated with a Monte-Carlo simulation. We use a Monte-Carlo simulation in conjunction with the Euler discretization, as described in section 3.3, to generate the option prices. We generate only one set of 15 option prices, for maturities of 1 month, 3 months and 1 year and strike prices ranging from 1300 up to and including 1700, with a difference of 100 for every consecutive pair of strike prices. Contrary to section 4.1.2, we cannot do a study on time-series data or 1000 sets of cross-sectional data, because it takes a long time to generate the options, with a decent accuracy, using Monte-Carlo simulation. Therefore, we rely on the results of one set of 15 option prices.

We use the same settings for the simulation as in 4.1.2. Namely, we generate a total of $N = 10^6$ option prices and set $\Delta = \frac{1}{36}$. For the various maturities, we need the following amount of simulation steps when $\Delta = \frac{1}{36}$. For the one month maturity $n_1 = \frac{1/12}{1/36} = 3$, for the three months maturity $n_2 = \frac{3/12}{1/36} = 9$ and for the maturity time of 1 year $n_3 = \frac{1}{1/36} = 36$. This process is implemented as follows.

- **Step 1** Set the initial stock price and variance equal to $\dot{S}_{i=0} = 1500$ and $\dot{v}_{i=0} = 0.5172$ respectively. And set $j = 1$.
- **Step 2** Set $i = 1$, and generate $\varepsilon^{(i)} \sim N(0,1)$ and determine

$$\ln \dot{S}_i = \ln \dot{S_{i-1}} + (r - \frac{1}{2})\dot{v}_{i-1}\Delta + \sqrt{\dot{v}_{i-1}}\varepsilon_x^{(i)}\sqrt{\Delta} \tag{90}$$

$$\dot{v}_i = \kappa(\alpha - \dot{v}_{i-1})\Delta + \omega|\dot{v}_{i-1}|^{\xi}\sqrt{\Delta}(\rho\varepsilon_x^{(i)} + \sqrt{1-\rho^2}) \tag{91}$$

  Repeat this for $i = 2, \ldots, 36$.
- **Step 3** We save the spot prices at the following maturities

$$S_{n_1}^j = \dot{S}_3$$
$$S_{n_2}^j = \dot{S}_9$$
$$S_{n_3}^j = \dot{S}_{36}$$

  and throw away all other $\dot{S}_i$
- **Step 4** If $j = 100,000$, then continue to step 5. Otherwise, reset $\dot{S}_0 = 1500$, $\dot{v}_0 = 0.5172$, update $j = j + 1$ and go back to step 2.
- **Step 5** The generated stock prices $(S_{n_1}^j, S_{n_2}^j, S_{n_3}^j)_{j=1}^{10^5}$ are used to determine the option prices given by

$$C_{MPMC}(S, K, \tau_{n_i}) = \frac{1}{10^5}\sum_{j=1}^{10^5} \max(S_{n_k}^j - K_k, 0) \tag{92}$$

  for $K = 1300, 1400, 1500, 1600, 1700$ and $\tau_{n_1} = 1/12$, $\tau_{n_2} = 3/12$, $\tau_{n_3} = 1$. The prices are computed for an option with spot price $S = 1500$.

Going through this process, we end up with a set of 15 options through Monte-Carlo simulation. This set is denoted by $((C_{MPMC}(1500, K_i\tau_j))_{i=1}^5)_{j=1}^3)$. The set of parameters we need to estimate is given by $\Omega = \{\kappa, \alpha, \omega, \rho, \xi\}$. The estimation process is similar to that of section 4.1.2, namely we minimise the expression

$$\min_{\Omega, v_0} \sum_{i=1}^{5}\sum_{j=1}^{3}[C_{MPMC}(S_0, K_i, \tau_j) - C_{\Omega,v_0}(S_0, K_i, \tau_j)]^2 + ||\Omega, v_0 - \Omega^0, v_0^0||, \tag{93}$$

For the estimation of the parameters, we must also take the extra noise, due to the Monte-Carlo simulation, into account. At the beginning of this section we saw that Monte-Carlo simulations have some errors, when compared to the K&M approximation, for example. For this reason we relax the convergence criteria of the optimizer `lsqnonlin` used in `MATLAB`. The relaxation of the criteria will prevent over-fitting.

The results for the estimation process are given by table 8. Note that the standard deviations are missing, due to the fact that we only have *one* cross-sectional set of 15 option prices here. Because of this same reason, we included the initial variance, v_0, used to get the option prices. We also adjusted the upper and lower bound for the parameter ξ, as we have generated the data with $\xi = 1.33$, which lies outside of the previous bounds. If we look at panel A, we see relatively good results; most estimates seem to be in the same range as their real counterparts. And especially the estimate of ξ is of high accuracy. In the other panels, we see quite large errors for the parameter ρ and in Panel C for α, meaning that the process is sensitive to changes in the initial parameter values. However, the estimates of v_0 remain quite good over all panels. Another thing one must realize is the errors in the generated option prices, due to the use of Monte-Carlo simulation. In the accuracy tests, we saw, for the Heston model, that the Monte-Carlo simulation had relatively large errors for options with maturities that we consider here. These errors in the option prices, in turn, lead to errors in the parameters that are estimated on these option prices.

Table 8: **Estimated parameters of the CEV model, for $\xi = 1.33$, by calibrating its K&M approximation using regularisation on one set of cross-sectional options generated from the CEV model.**
This table provides the estimates of the CEV model that are obtained by calibrating its K&M approximation, using regularisation, on one set of 15 option prices, with various strikes and maturities. The option prices are generated from the CEV model through Monte-Carlo simulation. The real parameter values, used for generating the data, are reported, together with the lower and upper-bound values used in MATLAB's non-linear least squares optimizer, `lsqnonlin`. Panels A, B and C give the estimation results for different sets of initial parameter values, used for the regularisation process. Each panel reportes the estimates, the initial values it used and the percentage errors (Error %) in the estimations.

	κ	α	ω	ρ	v_0	ξ
Real value	0.2	0.5172	0.4	-0.0243	0.5172	1.33
Upperbound	1	1	1	0	1	2
Lowerbound	0	0	0	-1	0	0
			Panel A			
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$	$\hat{v_0}$	$\hat{\xi}$
Initial value	0.2	0.5172	0.4	-0.0243	0.5172	1.33
Estimated value	0.265886	0.647077	0.629676	-0.02864	0.523125	1.397678
Error %	-32.943	-25.1115	-57.4191	-17.8601	-1.14554	-5.08857
			Panel B			
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$	$\hat{v_0}$	$\hat{\xi}$
Initial value	0.1	0.1	0.1	-0.01	0.1	1.1
Estimated value	0.195401	0.60639	0.360655	-0.00589	0.502494	1.497039
Error %	2.299563	-17.2448	9.836317	75.77259	2.843443	-12.5593
			Panel C			
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$	$\hat{v_0}$	$\hat{\xi}$
Initial value	0.9	0.9	0.9	-0.09	0.9	1.9
Estimated value	0.162021	0.97767	0.591677	-0.06066	0.527137	1.574047
Error %	18.9895	-89.0314	-47.9192	-149.609	-1.92124	-18.3494

5 Application To Empirical Data

In the previous sections we show that the K&M approximation is able to yield satisfactory results for data generated from the CEV model and thus conclude that we have an accurate enough closed form approximation for the option price from the CEV model. This leaves the testing of the closed form approximation on empirical data, which we conduct in this section. In the following we start with a description of the data, followed by an assessment of the performance of the CEV model compared to the Heston model and the PBS model.

5.1 Data Description

We use closing prices of S&P 500 call index options, obtained from OptionMetrics, which are the focus of many existing investigations, including, among others, Bakshi, Cao, and Chen (1997) and Christoffersen and Jacobs (2004). These options are the most actively traded European-style contracts. Our sample period dates from January 1, 2002 to December 30, 2011 and is much more recent than the sample periods used in former studies, which focus on data around 1992. Our data sample contains both a relative stable period, from 2002 to 2007, and a highly volatile period between 2008 to 2011, which contains a lot of crises.

The S&P 500 index options expire on the third Friday of each contract month. Their strike price intervals are 5 points and minimum tick is 1/8. On each trading day, we have a set of options with various bid-ask quotes, strike prices, and maturity dates. The time-to-maturity is measured from the trading day to the Thursday immediately preceding the Friday of the option's expiration, because S&P 500 index options expire immediately at the opening of the trading day. As the spot price we take the closing price level of the S&P 500 index. The index levels are obtained from DataStream.

For the construction of our data sample we apply several filters, based on the following criteria that are set up by Bakshi et al. (1997). First, quotes that do not satisfy the arbitrage condition

$$C \geq \max(0, S - K, S - e^{r\tau}) \tag{94}$$

must be eliminated, since this could lead to negative implied volatilities. Second, options with very long and very short time-to-maturities are removed. Options with time-to-maturity longer than 120 days are excluded, since these trade at a very high-premium. Due to their high-premium, there will be a lot of weight put on these contracts in the non-linear least squares estimation method. Options with less than 6 days to expiration are excluded, because they are very sensitive to liquidity biases and their prices are extremely volatile. Third, very far in-the-money and very far out-of the money options are excluded. These options are highly illiquid and therefore their price quotes may not reflect the true option value. We say that options are very high in- and out-of-the-money, if their absolute moneyness is greater than 7%. The option's moneyness is defined by the percentage difference between the S&P 500 index level, S, and the strike price, K, i.e. $Moneyness(\%) = S/K - 1$. Lastly, we remove all options with bid prices below \$ 3/8, since these prices are also very volatile and contain too much noise.

Applying the filters yields a sample of daily option prices consisting of 190834 observations (option prices) for 2519 days. The amount of options on each day varies a lot, but the average is about 76 for one day. The CEV model also requires the input of the risk-free interest rate. We proxy the risk-free interest rate by Treasury bill yields collected from the resource center of the US department of treasury website. Option prices are known not to be very sensitive to changes in interest rates and interest rates do not change much on a daily basis. Therefore, we use the T-bill rate, with a maturity closest to the option's expiration, to represent the interest used in the option pricing model. Our analysis is restricted to options with time-to-maturity up to 120 days, therefore we only use the 1-month and 3-month T-bill yields to represent the risk-free interest rate for our options.

As option prices are very sensitive to their exercise prices and their times to maturity, we divide the options in our data set into different maturity and moneyness bins to study of their price behaviour. We say that a call option is at-the-money (ATM) if the moneyness is between -0.02 and 0.02, in-the-money (ITM) if the moneyness is between 0.02 and 0.05, out-of-the-money (OTM) if the moneyness is between -0.02 and -0.05, and deep in-the-money or far out-of-the-money if the moneyness is greater than 0.05 or smaller than -0.05, respectively. For the maturities of the options we define short-term, middle-term and long-term options as option with a maturity between 7 and 45 days, 45 to 90 days and 90 to 120 days, respectively.

Table 9: **Summary statistics of S&P 500 Index Call Options (2002-2011).**
This table contains the summary statistics of daily closing prices of European call options written on the S&P 500 index. The data is tabulated in maturity and moneyness bins. The sample starts on January 2, 2002 and ends on December 30, 2011, consisting of 2519 days with a total of 190834 option prices. DTM denotes the days to maturity and S/K denotes the moneyness. Panel A displays the amount of option prices per bin, panel B displays the average call option price per bin and panel C displays the standard deviation of option prices over each bin.

Panel A. Number of call option contracts				
	DTM<45	45<DTM<90	90<DTM<120	Total
S/K<0.95	8128	11404	3228	22760
0.95<S/K<0.98	20040	19112	5085	44237
0.98<S/K<1/02	30324	25334	6659	62317
1.02<S/K<1.05	19315	15804	4391	39510
1.05<S/K	10388	9063	2559	22010
Total	88195	80717	21922	190834

Panel B. Average call price				
	DTM<45	45<DTM<90	90<DTM<120	Total
S/K<0.95	4.74	10.45	20.47	9.83
0.95<S/K<0.98	7.48	18.42	31.22	14.94
0.98<S/K<1/02	22.16	38.15	52.14	31.87
1.02<S/K<1.05	51.17	65.58	78.19	59.94
1.05<S/K	75.50	87.19	97.89	82.92
Total	29.86	40.44	53.18	37.01

Panel C. Standard deviation of call prices				
	DTM<45	45<DTM<90	90<DTM<120	Total
S/K<0.95	3.23	5.39	7.53	5.22
0.95<S/K<0.98	4.66	7.57	9.29	6.83
0.98<S/K<1/02	10.58	11.62	14.63	10.51
1.02<S/K<1.05	12.16	12.76	17.48	11.37
1.05<S/K	13.57	13.96	18.87	11.90
Total	24.19	24.59	27.85	26.86

Table 9 provides the descriptive statistics of the entire data set. In Panel A we see that about 33% of all the options are ATM, indicating that these options are the most actively traded in the market. The ITM and OTM options account for a total of 23% and 21% respectively. The deep ITM and far OTM show the least presence, with a total of about 12% each. Over this sample period, we see that there are more OTM options than ITM options. This indicates that investors are optimistic about future market increases, since OTM options will be valuable only if the market increases significantly in the future. For the time-to-maturity groups, the short-term options are the most dominant in our sample, accounting for about 46% of the data set, closely followed by mid-term maturity options with a presence of 42%, and the long-term options account for only 12%.

If we look at Panel B of table 9, we see a large difference in prices across maturities and moneyness. This has an important implication for the pricing process, as in our dollar based loss function of section 3.4, expensive contracts will receive more weight than cheap contracts. The pricing pattern is as we would expect, with the deep OTM options being the cheapest and the deep ITM options being the most expensive, as these are very likely to give a pay-off at maturity. Also, we note that options with a longer maturity trade at a higher premium; as time-to-maturity increases, this also increases the risk and hence the premium increases. The overall average call option price in our sample period is \$37 with a standard deviation of about \$27.

The data used by Bakshi et al. (1997) (and by Christoffersen and Jacobs (2004)) contains prices on 755 days in the period from June 1, 1988 through May31, 1991. Our data contains prices on 2519 days and of course covers an entirely other period. It is more recent and covers a much longer time period than the data used by Bakshi et al. (1997). The amount of options traded has increased in the recent years, which are contained in our data. However, Bakshi et al. (1997) do use a wider range of options in terms of moneyness and time-to-maturity. Since our data contains a lot of days, it also contains a lot of different regimes. Two main periods that we differentiate are the pre-crisis period, containing prices on 1446 days in the period from January 2, 2002 through September 28, 2007, and the crisis period, containing prices on 1073 days in the period from October 1, 2007 through December 30, 2011.

In appendix E we consider the two different time periods distinctively. The distribution of the options in the different bins changes over time as a result of the difference in regimes from one period to another. Also, the second period (crisis) contains a lot more options than the first period (pre-crisis), which indicates the increase of amount of options traded. As the different regimes seem to affect the pricing of the options, it is likely that this will cause problems for our model, since the CEV model assumes constant parameters for the process over the entire sample period.

5.2 Empirical Results

Using the option price data described in the previous subsection, we test the empirical option valuation performance of the CEV model when applied in conjunction with the K&M method. Once the option prices for the CEV model are computed, we compare the empirical performance of the CEV model to two benchmark models, namely the PBS and Heston model. The models are tested by their in sample fits and (1-day and 5-day) out-of-sample fits. The parameters estimates are conducted by using the regularisation method described in 3.4.2. We mainly focus on the CEV model here, but the estimates and performance computations are done similarly for the other two models. For the Heston model, we use Heston's (1993) closed-form solution, given by (28) and described in 2.2, and for the PBS model we use the Black-Scholes formula, given by (11) with volatility specification given by (14), described in 2.1.

Similar to Bakshi et al. (1997), we estimate the parameters, $\Omega = \{\kappa, \alpha, \omega, \rho, \xi\}$, and variance, v_t, using all available option prices for each day. Because we estimate the parameters on a daily basis, we denote the daily parameter set by $\Omega_t = \{\kappa_t, \alpha_t, \omega_t, \rho_t, \xi_t\}$. The estimates are obtained by minimizing the sum of squared errors between the observed and model price. For the regularisation method this is given by

$$\min_{\Omega_t, v_t} SqErr_t(\Omega_t, v_t) = \min_{\Omega_t, v_t} \sum_{j=1}^{M} \sum_{i=1}^{N} w_{ij} [C_{MP}(S_t, K_i, \tau_i) - C_{\Omega_t, v_t}(S_t, K_i, \tau_j)]^2 + ||\{\Omega_t, v_t\} - \{\Omega_0, v_t^0\}||^2, \tag{95}$$

which is the expression given by (79), with an additional subscript t. The subscript t is added to $SqErr_t$, in order to clarify that it concerns the squared errors on day t. Minimizing this expression yields a set of estimated parameters, denoted by $\hat{\Omega}_t = \{\hat{\kappa_t}, \hat{\alpha_t}, \hat{\omega_t}, \hat{\rho_t}, \hat{\xi_t}\}$ and estimated variance $\hat{v}_t$. The estimation process requires a set of initial parameters Ω_0 and initial variance, v_t^0. As the default set of initial values we take

$$\Omega_0 = [3, 0.05, 0.5, -0.1, 1] \text{ and } v_t^0 = 0.1. \tag{96}$$

To prevent some locality issues with the optimization, the estimated parameters $\hat{\Omega}_t$ and $\hat{v}_t$ need to satisfy the condition given by (78), or

$$SqErr_t(\hat{\Omega}_t, \hat{v}_t) \leq \sum_{i=0}^{N} \sum_{j=0}^{M} w_{ij} [bid_{ij} - ask_{ij}]^2, \tag{97}$$

where $SqErr_t(\Omega_t, v_t)$ is given by (95). If the condition is not satisfied, a new set of initial parameters is chosen and the process is repeated. The new set of initial values are determined by increasing or decreasing the old initial values through a bounded grid. For reasons of practicality, the process is stopped if the condition is not satisfied at the $10-th$ set of initial parameter values. In this case we choose the estimates inducing the smallest squared errors of all ten trials. Furthermore, we use MATLAB's `lsqnonlin` optimizer and set the lower and upper bounds equal to

$$\Omega_{lb} = [0, 0, 0, -1, 0] \text{ and } v_t^{lb} = 0 \tag{98}$$

$$\Omega_{ub} = [10, 1, 1, 0, 2] \text{ and } v_t^{ub} = 0, \tag{99}$$

respectively.

We evaluate the pricing performance by comparing the (average) root mean-squared errors (RMSEs) for each model. The RMSEs are conducted for each day, both in and out-of-sample, as

$$RMSE_{t+h} = \sqrt{\frac{1}{M_t + N_t} \sum_{j=1}^{M_t} \sum_{i=1}^{N_t} [C_{MP}(S_{t+h}, K_i, \tau_j) - C_{\Omega_t, \hat{v}_{t+h}}(S_{t+h}, K_i, \tau_j)]^2}, \tag{100}$$

where N_t and M_t, denote the amount of distinct strike prices K and maturities τ on day t, respectively, and h denotes the forecast horizon. For the forecast of the variance, we take

$$\hat{v}_{t+h} = \hat{\alpha}_t + e^{-\hat{\kappa}_t \frac{h}{360}}(\hat{v}_t - \hat{\alpha}_t), \tag{101}$$

as suggested by Christoffersen and Jacobs (2004) , where $\hat{v}_t$ is the estimate for the variance at time t. The above expressions are computed for in-sample fits by setting $h = 0$ and for the 1-day and 5-days out-of-sample forecasts by setting $h = 1$ and $h = 5$ respectively.

For the Heston model all computations are done similarly, except the parameter ξ is not present and the prices $C_{\Omega_t, v_t}(S_t, K_i, \tau_j)$ are given by Heston's (1993)closed-form pricing function (28). The procedure is a little different for the PBS model, e.g. the variance v_t is not present. We denote the corresponding option price of the PBS, by $C_{\sigma(\theta)}(S_t, K_i, \tau_j)$, which is just the Black-Scholes option price, given by (11), with volatility $\sigma(\theta)$ given by (14), which we'll write here as

$$\sigma(\theta) = \theta_0 + \theta_1 K_i + \theta_2 K_i^2 + \theta_3 \tau_j + \theta_4 \tau_j^2 + \theta_5 K_i \tau_j + \epsilon, \tag{102}$$

with $\theta = \{\theta_0, \ldots, \theta_5\}$. We add a subscript t, such that we have θ_t as the parameter set of day t. We estimate the parameters through regularisation, by minimizing the squared errors given by

$$\min_{\theta_t} SqErr_t^{PBS}(\theta_t) = \min_{\theta_t} \sum_{j=1}^{M} \sum_{i=1}^{N} w_{ij}[C_{MP}(S_t, K_i, \tau_i) - C_{\theta_t}^{PBS}(S_t, K_i, \tau_j)]^2 + ||\theta_t - \theta^0||^2, \tag{103}$$

yielding a set of parameters $\hat{\theta}_t = \{\hat{\theta}_{0,t}, \ldots, \hat{\theta}_{5,t}\}$. The estimated volatility for time t becomes

$$\hat{\sigma_t}(\hat{\theta}_t) = \hat{\theta}_{0,t} + \hat{\theta}_{1,t} K_i + \hat{\theta}_{2,t} K_i^2 + \hat{\theta}_{3,t} \tau_j + \hat{\theta}_{4,t} \tau_j^2 + \hat{\theta}_{5,t} K_i \tau_j. \tag{104}$$

The initial parameters, together with the lower and upper bounds, for the PBS model are given by

$$\theta_0 = \{0.5, -0.1, 0.3, 0.7, -0.5, 0.1\}, \tag{105}$$

$$\theta_{lb} = \{-1, -1, -1, -1, -1, -1\}, \tag{106}$$

$$\theta_{ub} = \{0, 0, 0, 0, 0, 0\}, \tag{107}$$

respectively. Similarly as for the CEV model, for the PBS model we demand the parameter estimates to satisfy the condition given by

$$SqErr_t^{PBS}(\hat{\theta}_t) \leq \sum_{i=0}^{N} \sum_{j=0}^{M} w_{ij}[bid_{ij} - ask_{ij}]^2. \tag{108}$$

If the condition is not satisfied, a new set of initial values is chosen and the process is repeated ten times at the most. Lastly, the root mean-squared error for the PBS model, is computed as

$$RMSE_{t+h}^{PBS} = \sqrt{\frac{1}{M_t + N_t} \sum_{j=1}^{M_t} \sum_{i=1}^{N_t} [C_{MP}(S_{t+h}, K_i, \tau_j) - C_{\hat{\theta}_{t+h}}^{PBS}(S_{t+h}, K_i, \tau_j)]^2}. \tag{109}$$

5.2.1 The implied parameters

By implementing the above procedure we use all options on each given day as inputs to estimate that day's relevant structural parameters and spot variance. This estimation is done separately for all three models, on each day in the period from January 2, 2002 to December 30, 2011. Since we use the K&M approximation, and not a precise closed-form solution, for the CEV model, we look at its parameter estimates to see if they are reasonable. For this reason, we add the estimation results of the Heston model for comparison. In addition it can give an idea on how the additional parameter ξ influences the other estimates. We do not add the parameter estimates of the PBS model, since it is only a benchmark model and not of our interest. Table 10 reports the daily average and standard deviations of each estimated parameter/volatility series, for the CEV model and the Heston model, respectively. The table also reports the initial values used for the regularisation and the additional minimum and maximum values used for MATLAB's `lsqnonlin`.

Table 10: **Summary statistics about the implied parameters of the Heston and CEV model (2002-2011).** This table gives the summary statistics of the implied parameters from S&P 500 Index Call Options for the Heston, using Heston's (1993) closed-form solution, and the CEV model, using its K&M approximation. Each day in the data sample, the structural parameters of a given model are estimated using the regularisation method together with MATLAB's `lsqnonlin`, which requires the minimum and maximum boundaries and initial value as reported. The daily average of the estimated parameters is reported first, followed its median and the standard deviation. The parameters are estimated using all available data on each day. The data sample used, starts on January 2, 2002 and ends on December 30, 2011 (2519 days). The parameters, κ,α,ω,ρ and ξ are respectively the mean-reversion rate, the long-run mean, the volatility of variance, the correlation between stock price and variance, and the elasticity of variance and v_t is the implied variance for the options on day t.

	Parameters					
	κ	α	ω	ρ	v_t	ξ
Initial value	3.0	0.1	0.5	-0.1	0.1	1.0
Min	0.0	0.0	0.0	-1.0	0.0	0.0
Max	10.0	1.0	1.0	0.0	1.0	2.0
	CEV model					
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$	$\hat{v}_t$	$\hat{\xi}$
Mean	3.0669	0.0779	0.4454	-0.9029	0.0417	0.5784
Median	3.0687	0.0542	0.4716	-0.9921	0.0238	0.6436
Stanard deviation	1.1236	0.0780	0.1817	0.1583	0.0581	0.2590
	Heston model					
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$	$\hat{v}_t$	
Mean	4.7825	0.0518	0.6151	-0.7096	0.0445	
Median	4.4919	0.0418	0.6273	-0.7029	0.0250	
Stanard deviation	1.6414	0.0444	0.2806	0.1724	0.0620	

The parameters are quite informative. As such several observations are in order. First of all, the correlation parameter ρ for the CEV model is pretty much on its boundary value; the median is virtually equal to -1, which is the lower boundary for this parameter. Such high correlation (between the stock price and its variance) is very unlikely and therefore indicates the possible presence of an error in the estimation of this parameter. As a matter of fact, we saw in the simulation section that the K&M approximation had a hard(est) time estimating the parameter ρ. A close inspection of the closed form approximation (not reported) shows that the parameter ρ does not appear as often in the approximation as the other parameters do. It is likely that more corrective terms are needed for the $K\&M$ approximation of the CEV model to get a good estimation for ρ. For the Heston model, the correlation parameter takes on the average and median value of -0.71, which is very close the the -0.64 Bakshi et al. (1997) report. For the period between 1988 and 1991 Bakshi et al. (1997) report that the correlation coefficient between the historical index return and historical index return volatility is about −0.23 (whereas −0.64 is from option price data). The reason for this difference is, as Bakshi et al. (1997) conclude, "the models with stochastic volatility rely on implausible levels of correlation and volatility to rationalize the observed option price". Unfortunately, the CEV model (in conjunction with the K&M approximation, as applied here) cannot improve upon this.

Furthermore, the implied spot variance v_t, is very similar for both models. The means (medians) of the corresponding implied volatilities, $\sqrt{v_t}$, are 20.4% (15.8%)and 21.1% (15.4%) for the CEV model and Heston model, respectively. The difference between the average volatilities is less than 0.70%. It should be noted that option prices are very sensitive to differences in volatility inputs, as shown by Figlewski (1989). Thus the small differences in volatilities have big impacts on the prices. In figure 5 we have plotted the implied volatilities from both models in one figure. The volatilities look very similar for both models. An outstanding difference is the peak during the 2008 crises, which is higher for the Heston model than for the CEV model. The maximum values of the volatilities are 79.0% and 83.0% for the CEV and Heston model, respectively. However, the latter are not volatilities of the same day.

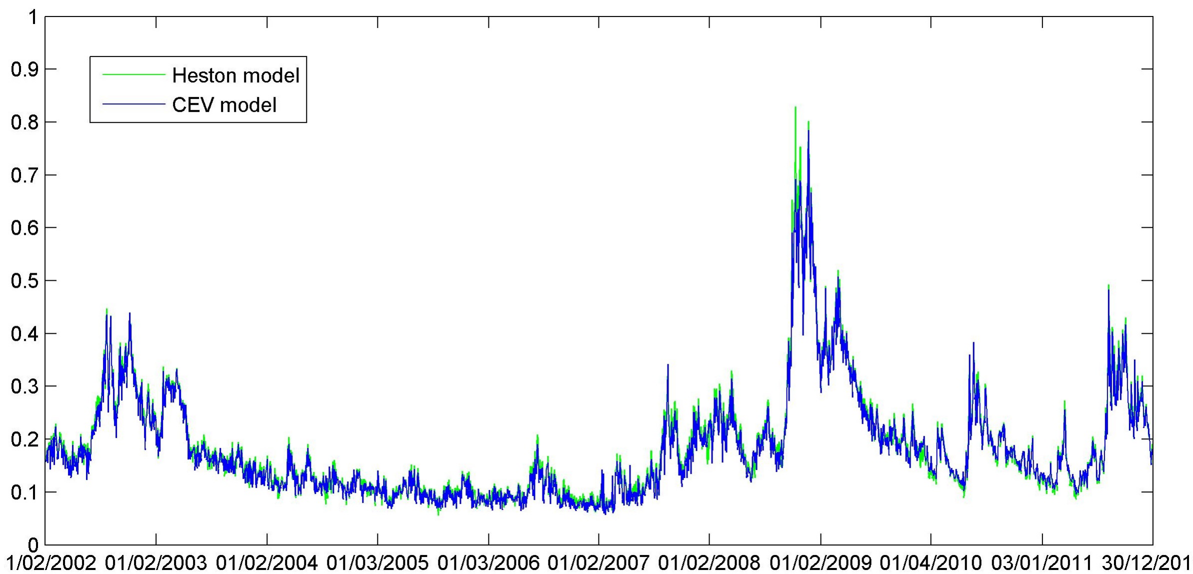

Figure 5: Implied volatilities for the Heston model and the CEV model, from S&P 500 Index Options. The sample period covers a total of 2519 days from January 2, 2002 to December 30, 2011.

Another notable result is the estimate of the variance elasticity parameter ξ, which, on average, is equal to 0.58 for the CEV model. This estimate is very close to $\xi = 1/2$, which corresponds to the Heston model, implying that we estimate the CEV model quite close to the Heston model. This is much lower than the level dependence found by Jones (2003), who estimates $\hat{\xi} = 1.33$ for the S&P 100. It is difficult to interpret the other parameters of the CEV model, since they might contain large errors for the estimates, due to the errors in ρ. For example, the average of the long run mean of the variance α, for the CEV model, is rather high with a value of 0.078, whereas for the Heston model, it is 0.052, which is more realistic when compared to the mean of the (corresponding) variance in table 10. The average of the mean reversion parameter is estimated at 3.1 and 4.8 for the CEV model and the Heston model, respectively, both of which are higher than the value of 1.15 estimated by Bakshi et al. (1997). The difference between the two models is large, especially since these estimates imply that the time it takes for the variances to revert halfway to their long run mean is 81 days and 53 days, for the variance of the CEV model and that of the Heston model, respectively. Since the implied variance series of both models are very much alike, we expect the same also to hold for the estimates of the long run mean. Lastly, the volatility of variance ω, is estimated at 0.45 and 0.62 on average for the CEV model and Heston model, respectively. Bakshi et al. (1997) estimate this at 0.39. The value of 0.62

sounds more realistic for our data sample, which is highly volatile at the second half. The ω estimate of the CEV model is closer to that of Bakshi et al. (1997) than that of the Heston model is. This is of course not of much importance, since we consider a completely different time-span.

5.2.2 The sample fit

To compare the model performances, we compute the average RMSEs as

$$\overline{RMSE}_h = \frac{1}{T-h} \sum_{t=1}^{T-h} RMSE_{t+h}, \tag{110}$$

where $h = 0$ for the in-sample, and $h = 1$ and $h = 5$ for the 1-day and 5-day out-of-sample, respectively. We compute the RMSEs over the various maturity and moneyness bins, that we used in table 9. We also compute the standard deviation of the RMSEs as

$$sd(RMSE_h) = \sqrt{\frac{1}{T-h} \sum_{t=1+h}^{T} (RMSE_{t+h} - \overline{RMSE}_h)^2}. \tag{111}$$

The results for the RMSEs for the various moneyness and maturity bins are shown in table 11, 12, and 13 for the in-sample, 1-day out-of-sample, and 5-days out-of-sample, respectively. The tables give the average RMSEs with their corresponding standard deviation.

For the in-sample fits, shown in table 11, we see that the average RMSEs of the CEV model are by far the largest in all maturity-moneyness groups, implying that the model shows a worse in-sample performance than the benchmark models. Within each maturity group, all models have much smaller RMSEs for the (deep) OTM, than for the (deep) ITM options. For example, for the short-term maturity options the average RMSEs of the CEV, Heston, and PBS model for the OTM options are 0.54, 0.28, and 0.34, respectively, whereas for the ITM options (in the same maturity group) the average RMSEs are 0.81, 0.48, and 0.46, respectively. Furthermore the Heston model has the smallest average RMSE for the deep OTM options and the short term OTM, ATM, and deep ITM options. In all other cases the PBS model has the smallest average RMSE. The CEV model shows the worst in-sample performance by far. The performance of the models in-sample might not be very representative, as the regularisation method is used for the estimation of the parameters and this involves minimizing the squared errors including the extra term for the parameters' distance from their initial values. Thus, the optimal in-sample fit is not obtained, such that we have stable parameter estimates and thus better out-of-sample fits. However, using regular optimization doesn't change anything to the hierarchy for the in-sample fit, as we have tried (not reported).

Tables 12 and 13 report the RMSEs for the 1-day and 5-day out-of-sample, respectively. We start with the 1-day out-of-sample RMSEs. The PBS model shows the smallest errors in absolutely all cases. Its difference with the Heston model has increased compared to the in-sample RMSEs. Within the maturity groups, the ITM options are the most mispriced, whereas the deep OTM options are the least mispriced. For example, the average RMSEs for the short-term ITM options are 1.65, 1.63, and 1.42 for the CEV model, Heston model and PBS model, respectively, whereas the corresponding deep OTM average RMSEs are 1.07, 0.89 and 0.70, respectively. Between the moneyness groups, the long-term maturity options are the most mispriced and the short-term maturity options the least. For the long-term ATM options, the models' average RMSEs are equal to 2.16, 2.17 and 1.56 for the CEV, Heston and PBS model, respectively, while for the short-term options we have 1.67, 1.65 and 1.35, respectively. Overall the CEV model has the worst performance and the benchmark PBS model has the best performance. The results for the CEV and Heston model are closer to each other, compared to the PBS model, which has the best performance by far. For the 5-days out-of-sample (table 13) average RMSEs, the results are even more pronounced, with very large errors for the CEV model, closely followed by the errors' of the Heston model, whereas the errors from the PBS model are the smallest in all moneyness-maturity groups.

As we mentioned, the first and second half of the data sample are quite distinct, as one is a pre-crisis period and the other a highly volatile crisis period. In appendix E, we show that the composition of the first and second half of the data sample varies. This difference in composition may change the models' relative performances. In appendix F we assess the performances of the models during the first and second part of the data sample, to see how the regime change influences the model performances. We find that although all models show a better performance during the first half of the sample (as we could expect), the performance hierarchy remains the same for both the first and second half as over the data sample. Lastly, we plot the $RMSE_t$ over time to get a closer look of the performances over time.

Figure 6 presents the in-sample, the 1-day, and 5-days out-of-sample RMSEs over the entire data set, spanning a total of 2519 days. For the 1-day and 5-day out-of-sample performance, the first 1 and 5 days, respectively, are not available for obvious reasons. From the plots it is clear that the CEV model show the worst performance, especially for the in-sample fit. From 11 we know that the overall in-sample performance of the Heston model (0.42) was slightly worse than that of the PBS model (0.40). From the figure we can see that the RMSE of the PBS model is much more volatile than that of the Heston model. This is also displayed in table 11, where the standard deviation of the RMSE is substantially higher for the PBS model (0.54) than for the Heston model (0.38). Thus, the Heston model does show a more consistent performance in-sample, with a slightly worse overall fit. For the PBS model, there are two peaks for the observations corresponding with May 4, 2004 and November 18, 2004 in-sample (the peaks are also present for their out-of-sample counterparts, both 1-day and 5-days). The RMSEs of the PBS model take values of 9 and 11, on the respective dates, without any particular (economic) event taking place. The other models are not affected by this and have quite low RMSEs on the concerning days. The reason for the peaks in the PBS model is that the options on these days, had a low variation in strike prices and time-to-maturities. This low variation in the concerning option properties, leads to bad estimates of the volatility specification parameters of the PBS model; since its volatility is written in terms of the strike price and the time-to-maturity, option prices with a variety of strikes and maturities are required to obtain good parameter estimates. For the crisis period the performance of the models deteriorates. With extreme peaks around the fall of Lehman brothers, which reach values around 30 for the stochastic volatility models. This effect is much larger for the CEV and Heston model, than for the PBS model, for which the peaks remain under 20. The performance of the models remain very bad for the crisis period, with peaks around the Greek bailout in 2010 and the European debt crisis, at the end of 2011.

The CEV model underperforms for both the Heston model and the PBS model. Overall, the PBS model shows the best performance by far. The out-performance of the PBS model compared to the Heston model, is also stated by Christoffersen and Jacobs (2004). The under-performance of the CEV model compared to the Heston model, suggests that the generalization of the variance process (adding a volatility of volatility), deteriorates the pricing accuracy significantly. However, the inaccuracy of the K&M approximation could also play a part. For the parameters we showed that the parameter ρ was hard to estimate and often reached its boundary value for the CEV model. And we showed in the simulation study, that the inaccuracies, although small in pricing terms, resulted in relatively large errors in the parameter estimates (for the Heston model, when compared with the results obtained by using Heston's (1993) the closed-form solution). The inaccuracies are more pronounced for long-term maturity options than for mid-term options, as shown in 4.1.1. Therefore, estimating the parameters on short-term options only, could result in a better performance (especially a *relatively* better performance) for the CEV model.

In table 14 we present the RMSEs for the models, when the parameters are estimated using only short-term options, with maturities up to 45 days. Also, these are the only options the RMSEs are computed for. The table shows the in-sample, 1-day out-of-sample, and 5-days-out-of-sample results. Again, the CEV model shows an under performance when compared to the other models, both in and out-of-sample. The only difference we see is that, overall, the Heston model shows the best performance. The reason for the relatively bad performance of the PBS model, is probably due to the low variation in maturities. As mentioned earlier, the PBS model, due to its specification, requires a wide range of both strike prices and maturity date options to obtain good parameter estimates. Thus when one only uses options with a certain maturity for the calibration, the PBS model is bound to underperform.

The parameter estimates of table 10 show that the parameter ρ is very poorly estimated. This is likely due to the shortcomings of the approximation of the CEV model, because we did not incorporate that many terms. By keeping the elasticity parameter fixed at $\xi = 1/2$, we have the K&M approximation for the Heston model, which is also used in the simulation study of section 4.1. In appendix G we estimate the Heston model using its K&M approximation. When the parameter estimates are compared to those of table 10, we find that the estimates, given by table 24, are close to those of the CEV model than those of the Heston model (using Heston's (1993) closed form solution). Again, the estimate of the parameter ρ is very close to -1. The fits of the Heston model in conjunction with the K&M approximation, are worse than those of the CEV model, although the RMSEs are very close, both in and out-of-sample. Lastly, we also compute the K&M approximation prices for the Heston model, using the parameter estimates we obtained through the closed form function of the Heston model. This led to even much worse results, with a large increase in RMSEs for all cases. The latter result implies that the approximation is not accurate at all. Since we are using the same input as we did in Heston closed form solution, and assuming that our approximation is close to that, we expect to get roughly the same result, while the errors were tremendously larger for when we use the same parameters in the K&M approximation.

Table 11: **In-sample average RMSEs of the CEV model, Heston model and PBS model (2002-2011).**
This table contains the averages of the daily in-sample root mean-squared errors (RMSEs) of the K&M approximation of the CEV model, the Heston model and the PBS model, for pricing S&P 500 index options, within various moneyness-maturity groups. The RMSEs are computed between the market option price and the model price. The model prices are computed by using the parameters estimates on each day of the sample. The sample period extends from January 2, 2002 to December 30, 2011 for a total of 190,834 European call options on 2519 days. DTM denotes the days to maturity and S/K denotes the moneyness (spot price divided by strike price). The smallest average RMSE within each moneyness-maturity bin is bold.

Moneyness S/K		Maturity			Overall
		Short term DTM<45	Mid term 45<DTM<90	Long term 90<DTM<120	
Deep out-of-the-money S/K<0.95	CEV	0.57	0.68	0.58	0.63
		(0.53)	(0.61)	(1.16)	(0.61)
	Heston	**0.31**	**0.32**	**0.51**	**0.34**
		(0.39)	(0.44)	(0.97)	(0.47)
	PBS	0.33	0.40	0.56	0.41
		(0.27)	(0.67)	(1.52)	(0.72)
Out-of-the-money 0.95<S/K<0.98	CEV	0.54	0.46	0.76	0.54
		(0.44)	(0.35)	(1.01)	(0.45)
	Heston	**0.28**	0.32	0.49	**0.33**
		(0.28)	(0.3)	(0.93)	(0.37)
	PBS	0.34	**0.32**	**0.40**	0.35
		(0.31)	(0.61)	(1.36)	(0.62)
At-the-money 0.98<S/K<1.02	CEV	0.51	0.38	0.80	0.50
		(0.44)	(0.28)	(0.75)	(0.37)
	Heston	**0.37**	0.36	0.46	0.38
		(0.27)	(0.26)	(0.58)	(0.28)
	PBS	0.39	**0.33**	**0.32**	**0.37**
		(0.4)	(0.57)	(1.02)	(0.5)
In-the-money 1.02<S/K<1.05	CEV	0.69	0.36	0.61	0.59
		(0.6)	(0.24)	(0.58)	(0.43)
	Heston	0.48	0.40	0.48	0.46
		(0.47)	(0.4)	(0.53)	(0.41)
	PBS	**0.46**	**0.32**	**0.37**	**0.41**
		(0.42)	(0.53)	(0.83)	(0.46)
Deep in-the-money 1.05<S/K	CEV	0.75	0.65	1.22	0.78
		(0.69)	(0.48)	(0.82)	(0.54)
	Heston	0.61	0.49	0.52	0.58
		(0.64)	(0.52)	(0.64)	(0.54)
	PBS	**0.57**	**0.40**	**0.46**	**0.52**
		(0.54)	(0.52)	(0.88)	(0.53)
Overall CEV		0.62	0.51	0.84	0.60
		(0.47)	(0.34)	(0.79)	(0.43)
Overall Heston		**0.42**	0.39	0.53	0.42
		(0.36)	(0.33)	(0.7)	(0.38)
Overall PBS		0.43	**0.35**	**0.43**	**0.40**
		(0.38)	(0.55)	(1.18)	(0.54)

Table 12: **1-day out-of-sample average RMSEs of the CEV model, Heston model and PBS model (2002-2011).**
This table contains the averages of the daily 1-day out-of-sample root mean-squared errors (RMSEs) of the K&M approximation of the CEV model, the Heston model and the PBS model, for pricing S&P 500 index options, within various moneyness-maturity groups. The RMSEs are computed between the market option price and the model price. The model prices are computed by using the parameters estimates of one day earlier for each day of the sample. The sample period extends from January 2, 2002 to December 30, 2011 for a total of 190,834 European call options on 2519 days. DTM denotes the days to maturity and S/K denotes the moneyness (spot price divided by strike price). The smallest average RMSE within each moneyness-maturity bin is bold.

Moneyness S/K		Maturity: Short term DTM<45	Mid term 45<DTM<90	Long term 90<DTM<120	Overall
Deep out-of-the-money S/K<0.95	CEV	1.07	1.42	1.84	1.31
		(1.56)	(1.85)	(2.49)	(1.69)
	Heston	0.89	1.17	1.81	1.11
		(1.59)	(1.8)	(2.36)	(1.64)
	PBS	**0.70**	**0.78**	**1.23**	**0.81**
		(1.98)	(1.08)	(3.69)	(1.8)
Out-of-the-money 0.95<S/K<0.98	CEV	1.28	1.63	2.02	1.55
		(1.62)	(1.95)	(2.54)	(1.8)
	Heston	1.12	1.54	2.06	1.43
		(1.66)	(1.94)	(2.45)	(1.81)
	PBS	**0.89**	**1.00**	**1.38**	**1.03**
		(1.6)	(1.22)	(3.65)	(1.79)
At the money 0.98<S/K<1.02	CEV	1.67	1.83	2.16	1.80
		(1.81)	(2.08)	(2.53)	(1.94)
	Heston	1.65	1.85	2.17	1.80
		(1.83)	(2.04)	(2.39)	(1.92)
	PBS	**1.35**	**1.34**	**1.56**	**1.40**
		(1.45)	(1.39)	(3.42)	(1.64)
In the money 1.02<S/K<1.05	CEV	1.65	1.94	2.23	1.88
		(1.84)	(2.05)	(2.39)	(1.91)
	Heston	1.63	1.89	2.13	1.81
		(1.81)	(2.03)	(2.34)	(1.9)
	PBS	**1.42**	**1.51**	**1.67**	**1.53**
		(1.8)	(1.52)	(3.18)	(1.75)
deep in the money 1.05<S/K	CEV	1.56	2.03	2.49	1.95
		(1.8)	(2.07)	(2.43)	(1.86)
	Heston	1.50	1.83	2.07	1.73
		(1.78)	(2.01)	(2.3)	(1.85)
	PBS	**1.36**	**1.53**	**1.70**	**1.55**
		(2.21)	(1.55)	(3.15)	(2.05)
Overall CEV		1.58	1.84	2.26	1.77
		(1.62)	(1.93)	(2.38)	(1.79)
Overall Heston		1.49	1.74	2.13	1.68
		(1.66)	(1.92)	(2.33)	(1.79)
Overall PBS		**1.27**	**1.30**	**1.58**	**1.35**
		(1.62)	(1.29)	(3.41)	(1.71)

Table 13: **5-days out-of-sample average RMSEs of the CEV model, Heston model and PBS model (2002-2011).**
This table contains the averages of the daily 5-days out-of-sample root mean-squared errors (RMSEs) of the K&M approximation of the CEV model, the Heston model and the PBS model, for pricing S&P 500 index options, within various moneyness-maturity groups. The RMSEs are computed between the market option price and the model price. The model prices are computed by using the parameters estimates of five days earlier for each day of the sample. The sample period extends from January 2, 2002 to December 30, 2011 for a total of 190,834 European call options on 2519 days. DTM denotes the days to maturity and S/K denotes the moneyness (spot price divided by strike price). The smallest average RMSE within each moneyness-maturity bin is bold.

		Maturity			
Moneyness S/K		Short term DTM<45	Mid term 45<DTM<90	Long term 90<DTM<120	Overall
Deep out-of-the-money S/K<0.95	CEV	1.60	2.34	3.31	2.13
		(2.29)	(2.89)	(3.72)	(2.57)
	Heston	1.39	2.05	3.26	1.89
		(2.19)	(2.77)	(3.66)	(2.46)
	PBS	**1.00**	**1.30**	**2.20**	**1.31**
		(2.05)	(1.73)	(3.81)	(2.03)
Out-of-the-money 0.95<S/K<0.98	CEV	1.99	2.83	3.65	2.56
		(2.43)	(3.1)	(3.81)	(2.79)
	Heston	1.79	2.68	3.70	2.41
		(2.39)	(3.05)	(3.8)	(2.76)
	PBS	**1.29**	**1.70**	**2.53**	**1.67**
		(1.84)	(1.91)	(3.81)	(2.03)
At the money 0.98<S/K<1.02	CEV	2.55	3.11	3.89	2.92
		(2.69)	(3.26)	(3.94)	(2.98)
	Heston	2.50	3.08	3.92	2.89
		(2.67)	(3.23)	(3.94)	(2.96)
	PBS	**1.87**	**2.11**	**2.83**	**2.12**
		(1.96)	(2.13)	(3.83)	(2.13)
In the money 1.02<S/K<1.05	CEV	2.32	3.08	3.89	2.83
		(2.61)	(3.18)	(3.81)	(2.87)
	Heston	2.31	3.00	3.79	2.76
		(2.55)	(3.17)	(3.81)	(2.87)
	PBS	**1.85**	**2.21**	**2.84**	**2.16**
		(2.09)	(2.21)	(3.67)	(2.22)
deep in the money 1.05<S/K	CEV	2.09	3.02	3.93	2.75
		(2.49)	(3.11)	(3.75)	(2.74)
	Heston	2.05	2.81	3.55	2.54
		(2.45)	(3.04)	(3.65)	(2.71)
	PBS	**1.67**	**2.16**	**2.73**	**2.05**
		(1.9)	(2.28)	(3.64)	(2.3)
Overall CEV		2.29	2.99	3.86	2.75
		(2.4)	(3.05)	(3.75)	(2.76)
Overall Heston		2.19	2.86	3.76	2.65
		(2.38)	(3.02)	(3.77)	(2.74)
Overall PBS		**1.69**	**2.00**	**2.74**	**1.98**
		(1.83)	(1.97)	(3.69)	(2.05)

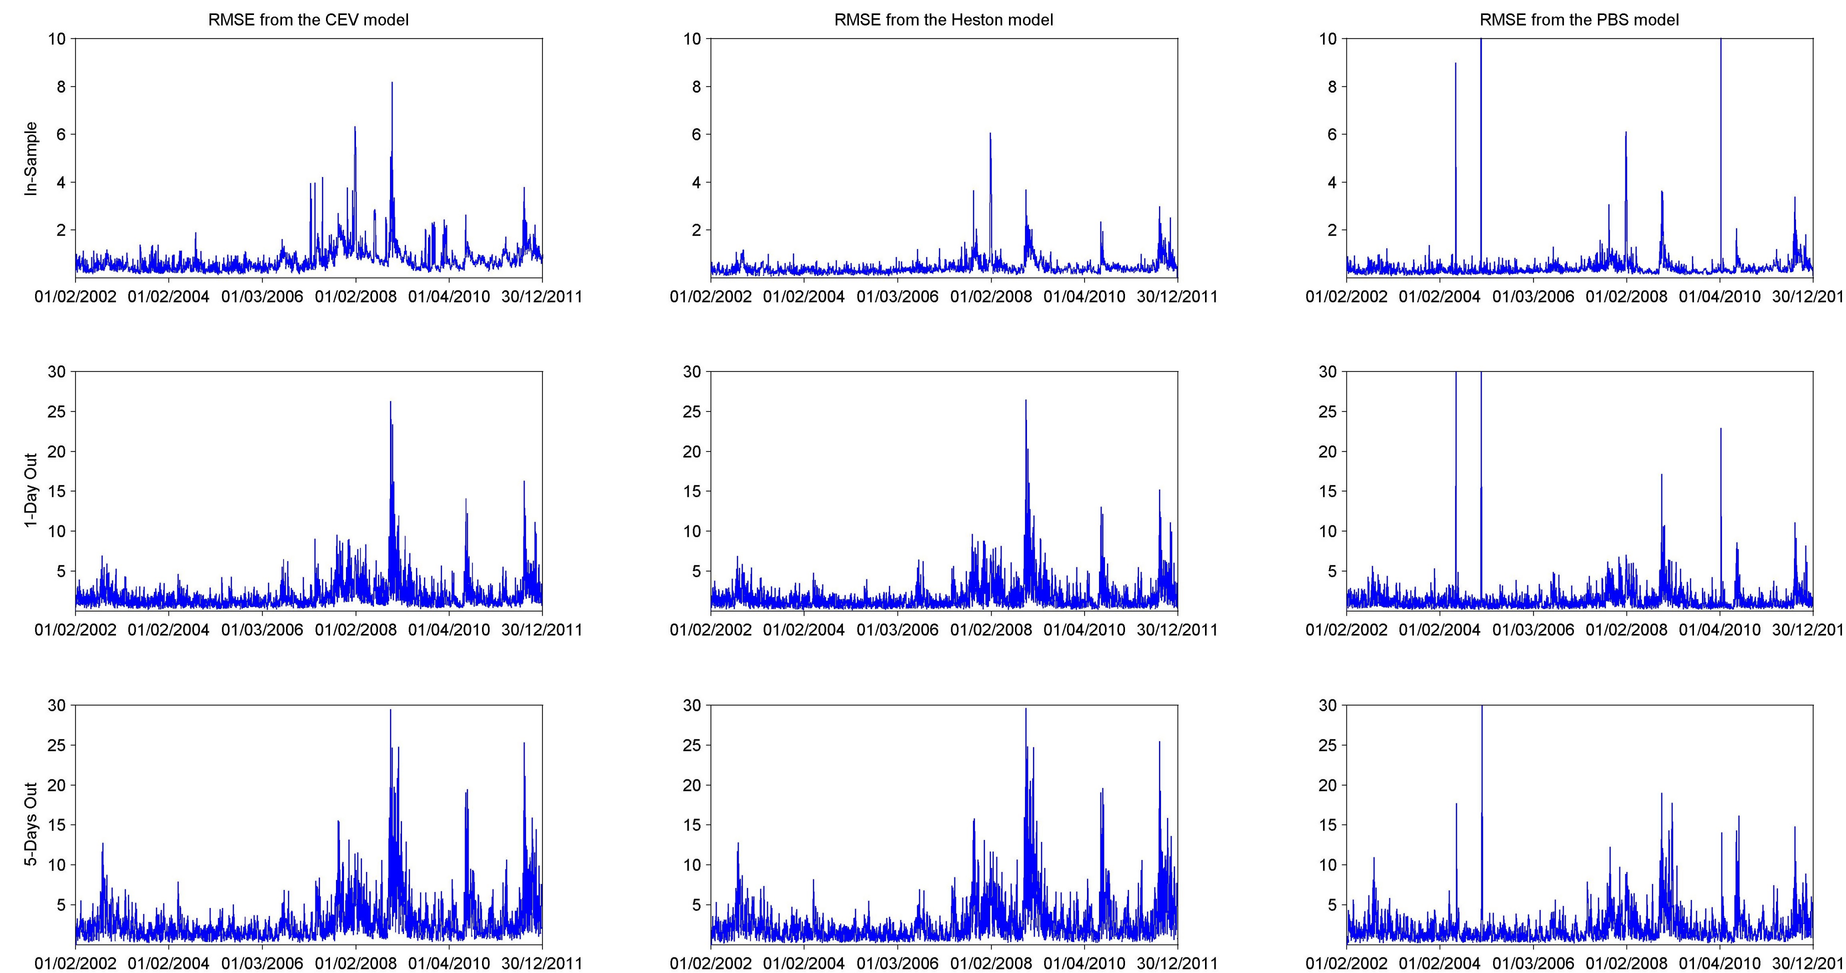

Figure 6: The root mean-squared error (RMSEs) plotted for the CEV model, the Heston model and the PBS model across 2519 days from January 2, 2002 through December 30 2011. The first column shows the RMSEs from the CEV model, the second column shows the RMSEs from the Heston model and the third column shows the RMSEs for the PBS model. The first row shows the in sample RMSEs, the second and third the 1-day and 5-day out of sample RMSEs, respectively.

Table 14: **In-sample and out-of-sample RMSEs of the CEV model, Heston model, and PBS model calibrated on option with maturities up to 45 days.**
This table contains the averages of the daily in-sample and the 1 and 5-day out-of-sample root mean-squared (pricing) errors (RMSEs) of the K&M approximation of the CEV model, the Heston model and the PBS. The parameters are estimated using only option with maturities up to 45 days. The RMSEs are only computed in different moneyness bins, using only the short-term maturity options on each day of the sample. The sample is from January 2, 2002 to December 30, 2011, consisting of 2519 days with a total of 88,195 European call options. S/K denotes the moneyness (spot price divided by strike price).The smallest average RMSE within each moneyness bin is bold.

Moneyness S/K		In Sample	1-day out	5-days out
Deep out-of-the-money S/K<0.95	CEV	0.56	1.16	1.61
		(4.92)	(4.66)	(2.38)
	Heston	**0.26**	**0.91**	**1.53**
		(0.36)	(1.6)	(2.36)
	PBS	0.36	1.00	2.04
		(0.54)	(2.07)	(4.64)
Out-of-the-money 0.95<S/K<0.98	CEV	0.46	1.38	2.15
		(3.77)	(3.84)	(2.85)
	Heston	**0.23**	**1.15**	**1.97**
		(0.24)	(1.68)	(2.57)
	PBS	0.29	1.16	2.56
		(0.47)	(2.27)	(4.41)
At-the-money 0.98<S/K<1.02	CEV	0.49	1.88	2.89
		(3.73)	(4.01)	(3.43)
	Heston	**0.28**	**1.70**	**2.76**
		(0.17)	(1.86)	(2.87)
	PBS	0.32	1.74	3.30
		(0.86)	(2.62)	(4.36)
In-the-money 1.02<S/K<1.05	CEV	0.46	1.90	2.64
		(3.4)	(3.68)	(3.07)
	Heston	0.43	**1.67**	**2.53**
		(0.38)	(1.83)	(2.73)
	PBS	**0.38**	1.87	3.09
		(2.04)	(3.3)	(4.63)
Deep in-the-money 1.05<S/K	CEV	0.59	1.83	2.37
		(3.55)	(3.67)	(2.63)
	Heston	0.57	**1.52**	**2.19**
		(0.58)	(1.79)	(2.59)
	PBS	**0.47**	1.83	2.88
		(2.49)	(4.25)	(5.38)
Overall CEV		0.50	1.76	2.56
		(3.69)	(3.84)	(2.95)
Overall Heston		**0.36**	**1.54**	**2.40**
		(0.31)	(1.68)	(2.55)
Overall PBS		0.37	1.65	3.01
		(1.37)	(2.79)	(4.43)

6 Conclusion

In this thesis we study the method proposed by Kristensen and Mele (2011) to approximate the price of derivative assets in the context of multi-factor continuous-time models. The idea behind their method is simple: given a model with no closed-form solution, an "auxiliary" model, which has a closed-form solution, is selected and the unknown price is expanded around the auxiliary one. We apply this method to the CEV model, for which no closed form solution form solution exists, to price European call options. The approach yields a closed-form approximation, which requires no simulation, and once implemented, requires a small amount of computation time. The closed-form approximation also allows for estimation and calibration of the CEV model, in this case.

Similar to Kristensen and Mele (2011), we test the accuracy of their approximation, by considering options from a specific case of the CEV model, namely the Heston (1993) model, for which Heston's (1993) closed-form solution is available. The presence of an accurate closed-form solution allows for a good assessment of the approximation's accuracy, since we can compute the discrepancy between the 'real' option price (which we wish we had) and its approximation. The accuracy of the K&M approximation is then compared to conventional Monte-Carlo simulation and the method of Yang (2006), who proposes a similar closed-form approximation. For prices of option with time-to-maturities up to one year, we find that the K&M method is much more accurate than Monte-Carlo simulation and the Yang method. However, as the option maturity increases, the accuracy of the K&M method quickly deteriorates and becomes much worse than the accuracies of the Monte-Carlo simulation and the Yang method. In a simulation study we test the K&M approximation's parameter estimation performance; both for the Heston model and the CEV model. We find small errors when the parameters are estimated using the regularisation method, which is a small extension of the standard non-linear least squares method. Although the errors in the parameter estimates are not that large, there is a significant bias present. When we compare the estimates from the K&M approximation of the Heston model with the estimates obtained when using Heston's (1993) closed-form solution, we find the errors in the estimates of the approximation are much larger, while the closed-form solution yields rather well unbiased results. This indicates that the approximation leads to errors in the parameter estimates, rather than that these errors are caused by the estimation procedure.

We test the empirical pricing performance of the CEV model on S&P 500 index options. By using the closed-form approximation we can estimate the model's parameters directly from cross-sectional option data, via a non-linear least squares method, or the regularisation method in our case. This calibration procedure uses information directly from the option market, instead of information from the model's underlying, which is the S&P 500 index in this case. The estimated structural parameters and the implied variance are then used to compute the CEV model price through its approximation. The model's structural parameters and spot variance are estimated on a daily basis. Furthermore, the same is done for the Heston model, using Heston's (1993) closed form solution, and the Practitioners Black-Scholes model, using its corresponding pricing formula (which is a practical variation of the Black-Scholes formula). The in-sample and out-of-sample fits of the latter two models are used as a benchmark for the corresponding fits of the CEV model. We compare the parameter estimates of the CEV model to those of the Heston model, since the CEV model has only one parameter more. The corresponding parameter estimates from both models differ quite a lot; the estimates for the Heston model seem to be more sensible, e.g. the correlation parameter of the CEV model is very close to its boundary value of -1, which is unlikely to be its correct value. This is presumably due to the inaccuracy of the approximation. The implied variances of the two different models, however, show much resemblance as the discrepancy between the two series is small. Furthermore, we test the performances, of all three models, both in-sample, and 1-day and 5-day out-of-sample. We find that the benchmark models outperform the CEV model in all maturity-moneyness groups. Restricting the estimation and evaluation of the models on short term maturity options does not change the under performance of the CEV model compared to the benchmark models. Lastly, when we apply the K&M approximation to the Heston model (by fixing one parameter) and test its empirical performance (similarly to how we did for the CEV model), we find results worse than for the CEV model. Furthermore, when we plug the estimated parameters from Heston's (1993) closed-form solution, into the approximation, the obtained fits are even less accurate than in any other before case. Thus, from this we conclude that the under-performance of the CEV model is due to the inaccuracy of the K&M approximation.

The under performance of the K&M approximation of the CEV model, is probably due to the low amount of corrective terms used in the approximation, which makes the approximation inaccurate. This inaccuracy leads to errors in the estimation of the model's parameters and the model's pricing performance. For the accuracy test we used two more corrective terms than for the simulation/calibration study and the empirical analysis. However, we find that the accuracy is not affected a lot by these two additional terms and thus leaving them out does not change much in the pricing performance. Still, the absence of these terms in the calibration process might play a different role. However, probably even more than two terms are needed to get an approximation suitable for empirical purposes. We were not able to compute the (extra) terms due to the computational power/ability available; the approximation is computed in Maple, for a certain amount of corrective terms, and then copy-pasted to MATLAB for the simulation and empirical application. The maximum amount of terms we could compute in Maple was five. Interestingly, Kristensen and Mele (2011) also compute the price for the Heston model using five corrective terms (at the most) for their accuracy test. Although Maple can compute five corrective terms, it can only give the approximation explicitly (such that we can transfer it to MATLAB) for three terms. The reason is that the extra terms are too long for Maple to display: the fourth term increases the amount of characters of the approximation to over 1,000,0000.

For further research we suggest the development of a method to compute the approximation with more corrective terms, such that the accuracy increases. One could look for computational software for handling derivatives of the sort encountered in the above method. Also, applying various mathematical properties of partial derivatives might lead to a shorter expression of the approximation, with more corrective terms nonetheless. Kristensen and Mele (2011) mention Kimmel (2008), who proposes a method to determine the convergence rate of similar approximations and even to increase this rate for all maturities. Kimmel's (2008) method allows to determine after how many corrective terms approximations, like that of Kristensen and Mele (2011), converge and whether the convergence is uniform or not. Unfortunately, the method of Kimmel (2008) is still restricted to one-dimensional models, whereas the CEV model is a two dimensional model. And thus, as Kimmel (2008) points out, his method needs to be generalized for multi-dimensional models. All of this cannot be applied in full if we do not have the computational power and ability necessary for generating more corrective terms. Thus our main suggestion for further research is to find a way to compute the approximation with many corrective terms in such a way that it can be implemented in programs like MATLAB. If more corrective terms can be computed, such that the approximation is really accurate, the CEV model can be assessed properly. Furthermore, the method of Kristensen and Mele (2011) could then be applied for a whole range of contingent claim prices from continuous models, since these models require a process similar to that of the CEV model and will therefore in many cases encounter the same problem of having an inaccurate closed form solution due to the low amount of corrective terms. So, although the theory of the approximation is there, the implementation should be investigated further.

A Itô's Lemma

In this appendix we shortly state Itô's Lemma, as presented by Baxter and Rennie (2008).

For a stochastic process X, with drift μ_t at time t and volatility σ_t at time t, the infinitesimal change is given by

$$dX_t = \sigma_t dW_t + \mu_t dt. \tag{112}$$

Then, Itô's lemma states that for any twice differentiable function f. $Y_t := f(X_t)$ is also a stochastic process and is given by

$$dY_t = (\sigma_t f'(X_t))dW_t + (\mu_t f'(X_t) + \frac{1}{2}\sigma_t^2 f''(X_t))d_t \tag{113}$$

B Feynmann-Kac

In this appendix we give the Feynman-Kac Theorem for a two-dimensional stochastic process. We do not discuss the theorem in detail (we omit the conditions that need to be satisfied), but only give the necessary aspects for its application. The Feynman-Kac Theorem is explained in detail in textbooks such as the one by Karatzas and Shreve (1991). We illustrate the use of the theorem by deriving the Feynman-Kac representation of the difference function $\Delta w(s,v,t;\sigma_0) = w(s,v,t) - w^{bs}(s,t;\sigma_0)$ given by equation (41).

Suppose that $\mathbf{X}_t$ follows a two-dimensional stochastic process

$$d\mathbf{X}_t = \boldsymbol{\mu}_t(\mathbf{X}_t)dt + \boldsymbol{\sigma}_t(\mathbf{X}_t)d\mathbf{W_t} \tag{114}$$

where $\mathbf{X}_t$ and $\boldsymbol{\mu}_t(\mathbf{X}_t)$, the drifts, and $\mathbf{W}_t$, risk-neutral Brownian motions, are each two-dimensional vectors and $\boldsymbol{\sigma}(\mathbf{X}_t,t)$ is a 2×2 matrix of volatilities. Thus, we can write this expression as

$$\begin{pmatrix} X_t^{(1)} \\ X_t^{(2)} \end{pmatrix} = \begin{pmatrix} \mu_t^{(1)}(\mathbf{X}_t) \\ \mu_t^{(2)}(\mathbf{X}_t) \end{pmatrix} d_t + \begin{pmatrix} \sigma_t^{(1,1)}(\mathbf{X}_t) & \sigma_t^{(1,2)}(\mathbf{X}_t) \\ \sigma_t^{(2,1)}(\mathbf{X}_t) & \sigma_t^{(2,2)}(\mathbf{X}_t) \end{pmatrix} \begin{pmatrix} dW_t^{(1)} \\ dW_t^{(2)} \end{pmatrix} \tag{115}$$

The theorem states that the PDE, for a function $w(\dot{X}_t, t)$, given by For notational convenience we write $X_t^(i) = x_i$ $\mu_i = \mu_t^{(i)}(\mathbf{X}_t)$, $\boldsymbol{\sigma} = \boldsymbol{\sigma}_t(\mathbf{X}_t)$ and $(\boldsymbol{\sigma}\boldsymbol{\sigma}^T)_{ij}$ is element (i,j) of the 2×2 matrix $\boldsymbol{\sigma}\boldsymbol{\sigma}^T$. The Feynman-Kac theorem states that a function on the stochastic process $\dot{X}_t$, say $w(\dot{X}_t, t)$, the PDE given by

$$\frac{\partial w}{\partial t}(\mathbf{x},t) + \sum_{i=1}^{2}\mu_i \frac{\partial w}{\partial x_i}(\mathbf{x},t) + \frac{1}{2}\sum_{i=1}^{2}\sum_{j=1}^{2}(\boldsymbol{\sigma}\boldsymbol{\sigma}^T)_{ij}\frac{\partial^2 w}{\partial x_i \partial x_j}(\mathbf{x},t) - k(\mathbf{x},t)w(\mathbf{x},t) + f(\mathbf{x},t) = 0 \tag{116}$$

with boundary condition $w(\mathbf{x},T) = d(\mathbf{x})$ has solution

$$w(\mathbf{x},t) = \mathbb{E}_{\mathbf{x}}\left(d(\mathbf{X}_t)e^{-\int_t^T k(\mathbf{X}_\tau,\tau)d\tau} + \int_t^T f(\mathbf{X}_u,u)e^{-\int_t^u k(\mathbf{X}_\tau,\tau)d\tau}du \right) \tag{117}$$

Next, we apply this Theorem for the difference function $\Delta w(s,v,t;\sigma_0) = w(s,v,t) - w^{bs}(s,t;\sigma_0)$, where s and v denote the stock-price and variance at time t and σ_0 is the constant volatility of the Black-Scholes model. The underlying of this function is the stock-process of the CEV model given by equation (32) and (33). This process $\mathbf{x} = (S_t, V_T)$ can be written in terms of two independent Brownian motions Z_1 and Z_2 as

$$d\begin{pmatrix} s \\ v \end{pmatrix} = \begin{pmatrix} r \\ \kappa(\alpha - v) \end{pmatrix} dt + \begin{pmatrix} s\sqrt{v} & 0 \\ \omega|v|^\xi \rho & \omega|v|^\xi\sqrt{1-\rho^2} \end{pmatrix} \begin{pmatrix} dZ^{(1)} \\ dZ^{(1)} \end{pmatrix}, \tag{118}$$

where the relation with the original risk-neutral brownian motions is $dW_t = dZ_t^{(1)}$ and $dW_t^{(v)} = \rho dZ_t^{(1)} + \sqrt{1-\rho^2}dZ_t^{(2)}$. We now have that the process for the underlying $\mathbf{x} = (s, v)$ is of the same form as (115). Using this we can find the Feynman-Kac representation of $\Delta w(s, v, t; \sigma_0)$ from equation (39). Furthermore, we have $f(\mathbf{x}, t) = \delta(s, v, t; \sigma_0)$,$\boldsymbol{\mu} = \begin{pmatrix} r \\ \kappa(\alpha - v) \end{pmatrix}$, $k(\mathbf{x}, t) = r$ and and $\boldsymbol{\sigma}\boldsymbol{\sigma}^T$ is given by

$$\boldsymbol{\sigma}\boldsymbol{\sigma}^T = \begin{pmatrix} s^2 v & s\omega\rho|v|^{\xi+\frac{1}{2}} \\ s\omega\rho|v|^{\xi+\frac{1}{2}} & \omega^2|v|^{2\xi} \end{pmatrix}. \tag{119}$$

Equation (116) is therefore equal to the PDE (39) for the difference function. The boundary condition for the difference function Δw, is the difference between the boundary function of the Black-Scholes model and the CEV model, namely $d(S_T, V_T) = b(S_T, V_T) - b_0(S_T)$. For both models we consider a European call option and thus their boundary conditions are identical, and we have $b(S_T) = b_0(S_T)$ such that $\Delta w(S_T, V_T, T; \sigma_0) = 0$ boundary condition or $d(S_T, V_T) = 0$ is the boundary function for $\Delta w(s, v, t; \sigma_0)$. Plugging all of this into equation (117) gives the Feynman-Kac representation of $\Delta w(s, v, t; \sigma_0)$, namely

$$\Delta w(s, v, t; \sigma_0) = \mathbb{E}_{s,v,t}\left(\int_t^T e^{-r(u-t)}\delta(S_u, V_u, u; \sigma_0)du\right), \tag{120}$$

which is the same as equation (41). This solution or representation is only valid if certain boundary conditions are satisfied. In this case, Kristensen and Mele (2011) show that all of the necessary boundary conditions of the Theorem are satisfied. To be precise, they state the following: "All the conditions in Karatzas and Shreve (1991, Theorem 5.7.6) are, therefore, met".

C Numerical Integration of the Heston integral

To evaluate expression 28, i.e. the European call-option price from the Heston model, we need to compute the integrals given by

$$P_j(x, V_t, T, K) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \Re \left(\frac{e^{-i\phi \ln(K)} f_j(x, V_t, T, \varphi)}{i\varphi} \right) d\varphi \tag{121}$$

for $j = 1, 2$. This is done by using the MATLAB function `quadl(@fun,a,b)` which approximates the integral of the function `@fun`, from a to b, within an error of 10^{-6} using an adaptive Gauss Lobatto quadrature rule. This technique requires the limits a and b to be finite, which is not the case for the above integral.

In figure 7 we have plotted the price of an option with a fixed set of parameters, using MATLAB's `quadl(@fun,a,b)` with different values for the upper boundary b. The figure shows that the price converges very early, namely at $b = 6$. Therefore, we conclude that choosing a finite value for b should not lead to much trouble in this case. To keep the error very low, we set the upper boundary equal to 1000, just to make sure.

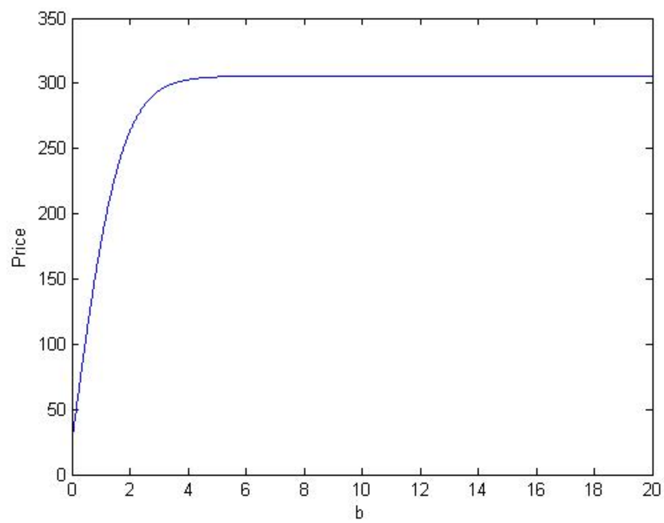

Figure 7: Plot of option price from the Heston model computed using MATLAB's `quadl[@fun,a,b]` with different finite upper boundaries b. The model parameters are kept fixed at $\kappa = 0.1465$, $\alpha = 0.5172$, $\omega = 0.5786$,$\rho = 0.0243$, $r = 0$, $t = 0$, $T = 1$ and spot variance $v_0 = 0.5172$. The strike price and spot price are set at $K = 1000$ and $S_0 = 1050$ respectively.

D Number of Corrective Terms for the K&M approximation

In section 4.1.1 we compute the K&M approximation, given by (44), for $N = 4$. We set the nuisance parameter equal to the spot variance, such that $N = 4$ corresponds to four corrective terms. Doing so, we compute the prices in Maple. However, the expression of the option price approximation for $N = 4$ contains too many characters for us to transfer it to MATLAB, such that we can use it for calibration and simulation purposes. In MATLAB, we have to repeat the entire differential procedure each time we want an option price. This takes 130 second, which is too long if we want to conduct a simulation study (or calibrate the model for that matter). However, for the case when $N = 2$ we do have a closed form solution from Maple[4] that we can transfer to a MATLAB file. Also, we are able to conduct the $N = 3$ case using MATLAB's `matlabFunction`, the problem with this is that matlab seems to determine some components numerically (replacing fractions with decimal numbers for example) and this makes the price deviate a from the actual K&M-approximation for $N = 3$. The deviation is very small and we therefore to implement this for testing purposes.

In table 15 below we have made a comparison for the cases when $N = 2$, $N = 3$ (of which we have the closed form from MATLAB), and when $N = 4$ for short term options. Option prices are conducted for options with a spot price of $S_0 = 1500$ and strike prices $K = 1300$, $K = 1400$, $K = 1500$, $K = 1600$ and $K = 1700$ and the three short term maturities of section 4.1.1. We have also added Black-Scholes model price, such that a comparison can be made when no corrective terms are added. Comparing the cases when different amounts of terms are added, we see that the improvement from adding three or four terms when compared to adding two terms deviates a lot, if there is any improvement at all. There are cases (e.g. in panel C, when $\tau = 1$ and $K = 1800$) where adding two terms outperforms the cases with three and four terms. So an increase in terms does not imply an increase in the performance of the method. Although Kristensen and Mele (2011) show that their approximation is convergent, we do not know if the convergence is uniform for the amount of corrective terms we use.

If the convergence is not uniform in a certain region, adding more terms does not necessarily lead to an improvement, and can even lead to a deterioration of the approximation as we have seen in table 15. Kimmel (2008) shows that we can modify power-series expansion, like that of the K&M-method, such that we get uniform convergence for all maturities. The modification involves a so called 'time-transformation'. However, Kimmel (2008) only discusses the case where the price of the contingent claim depends on one state-space equation. For European call options with CEV-specifications we have two state space equations (namely, (32) and (33)), such that we cannot apply Kimmel (2008) time-transformation. Kimmel (2008) does state that the generalization of the time-transformation to higher order state space equations is possible, but he does not explain how it is done. For further research it is interesting to look at how one can apply the K&M method in conjunction with Kimmel's (2008) time-transformation. We leave this as a suggestion for further research, as it is beyond the scope of this thesis.

[4]Unfortunately, the $N = 4$ (and even the $N = 3$ case) is too large for maple to display, as it contains more than $1000,000$ characters

Table 15: **Comparison of accuracy performance of K&M approximation of the Heston model for different numbers of corrective terms.**
This table compares option prices for the Heston model computed with the K&M-method using two, three and four leading terms. The errors are computed with respect to the closed form option price. All option prices are computed with a strike price $S_0 = 1500$. The parameter values of (34) and (33) are set equal to the values used in table 1. Also we have computed the percentage deviation from the correct closed form price. Panel A, B and C provide option prices with maturities equal to 1 month, 3 months and 1 year.

Panel A: $\tau = 1$ month

	Closed form	N=2		N=3	
Strike	Price	Price	% error	Price	% error
1300	242.0178	242.0155	-0.0009	242.0176	-0.0001
1400	176.2205	176.2175	-0.0017	176.2176	-0.0016
1500	123.7149	123.7121	-0.0023	123.7106	-0.0034
1600	83.9444	83.9418	-0.0031	83.9408	-0.0043
1700	55.2284	55.2259	-0.0045	55.2266	-0.0033
1800	35.3588	35.3567	-0.0059	35.3587	-0.0001

	N=4		Black Scholes	
Strike	Price	% error	Price	% error
1300	242.0171	-0.0003	242.0502	0.0134
1400	176.2235	0.0017	176.4101	0.1076
1500	123.7196	0.0038	124.0111	0.2395
1600	83.9477	0.0039	84.2628	0.3792
1700	55.2286	0.0002	55.4946	0.4819
1800	35.3561	-0.0074	35.5332	0.4935

Panel B: $\tau = 3$ months

	Closed form	N=2		N=3	
Strike	Price	Price	% error	Price	% error
1300	314.2324	314.1912	-0.0131	314.1959	-0.0116
1400	259.3434	259.3018	-0.0161	259.2884	-0.0212
1500	212.4992	212.4584	-0.0192	212.4356	-0.0299
1600	173.0745	173.0337	-0.0235	173.0127	-0.0357
1700	140.2836	140.2421	-0.0296	140.2320	-0.0368
1800	113.2769	113.2348	-0.0371	113.2403	-0.0323

	N=4		Black Scholes	
Strike	Price	% error	Price	% error
1300	314.2676	0.0112	315.1697	0.2983
1400	259.4051	0.0238	260.6455	0.5021
1500	212.5717	0.0341	214.0257	0.7183
1600	173.1407	0.0383	174.6729	0.9235
1700	140.3303	0.0333	141.8170	1.0931
1800	113.2968	0.0176	114.6400	1.2033

Panel C: tau=1 year

	Closed form	N=2		N=3	
Strike	Price	Price	% error	Price	% error
1300	490.7879	489.4835	-0.2658	489.0096	-0.3623
1400	448.1465	446.8051	-0.2993	446.1367	-0.4485
1500	409.6636	408.2889	-0.3356	407.5217	-0.5228
1600	374.9629	373.5541	-0.3757	372.7811	-0.5819
1700	343.6825	342.2382	-0.4202	341.5418	-0.6228
1800	315.4812	314.0028	-0.4686	313.4502	-0.6438

	N=4		Black Scholes	
Strike	Price	% error	Price	% error
1300	492.6156	0.3724	500.5911	1.9974
1400	450.2397	0.4671	459.0118	2.4245
1500	411.8844	0.5421	421.2616	2.8311
1600	377.1813	0.5916	386.9804	3.2050
1700	345.7848	0.6117	355.8357	3.5362
1800	317.3750	0.6003	327.5234	3.8171

E Data Composition During Pre-Crisis (2002-2007) and Crisis (2007-2011)

In section 5.1 we have option price data on 2519 days, from 2002 through the end of 2011. This is a very long period during which a lot happened, that might change the composition of the option data set, during different time periods. Especially, the occurrence of the financial crisis in 2007/2008, changed the market from bull to bear. We divide the sample into before and after the credit crunch in 2007, such that we can compare the composition of the option price data for the two time periods. More specifically, we separate the sample from January 2, 2002 through December 30, 2011 into the the pre-crisis period from January 2, 2002 through September 28 2007, and the crisis period from October 1 2007 through December 30, 2011. The descriptive statistics of the option price data for the pre-crisis and crisis period are given by table 16 and table 17, respectively.

Panel A of both tables shows a large increase in the amount of European call options from the pre-crisis period in table 16 to the crisis period in table **??**. The tendency towards out-of-the-money options remains for both sample periods, however there is a shift from the at-the-money options to the deep in and far out-of-the-money options. This is likely due to the high volatility during the crisis period, where investors favoured deep-out-of-the-money options, because the sharp decreases were seen as a market sentiment which would restore due to the remaining intrinsic value of the underlying companies on the US equity market.

If we look at the average option prices in Panel B of tables 16 and 17, we see an increase in the average option price, from \$31.80 during the pre-crisis to \$40.72 during the crisis period. The overall increase is to be expected due to the higher volatility during the crisis period. The price increase is particularly sharp for (deep) out-of-money options for which the price more than doubled from the pre-crisis to the crisis period. This again confirms that investors believed that the sharp declined would restore shortly after.

F Model Fits During Pre-Crisis (2002-2007) and Crisis (2007-2011)

In this appendix we state the sample fit results of the CEV model (using the K&M approximation), the Heston model and the PBS model for pre-crisis period and crisis period, as described in appendix E. The fits are computed through the root-means-squared-error(RMSE), of which the average over time in different moneyness-maturity bins is taken. The results for the in-sample fits are given by table 18 and table 19, respectively, for the 1-day out-of-sample by tables 20 and table 21, respectively, and for the 5-days out-of-sample by table 22 and table 23, respectively.

The distinction between the first and second part of the sample, denoted by pre-crisis and crisis period, respectively, that we pointed out in appendix E, barely changes the performance of the models relative to each other. All model's do show a better performance during the pre-crisis period than during the crisis period, as could be expected, since the crisis period is much more volatile.

Table 16: **Summary statistics of S&P 500 Index Call Options (2002-2007)**
This table contains the summary statistics of daily closing prices of European call options written on the S&P 500 index.The data is tabulated in maturity and moneyness bins. The sample starts on January 2, 2002 and ends on September 28, 2007, consisting of 1446 days with a total of 79,238 option prices. DTM denotes the days to maturity and S/K denotes the moneyness. Panel A displays the amount of option prices per bin, panel B displays the average call option price per bin and panel C displays the standard deviation of the call prices over each bin.

Panel A. Number of call option contracts

	DTM<45	45<DTM<90	90<DTM	Total
S/K<0.95	2195	4095	720	7010
0.95<S/K<0.98	8750	8275	1278	18303
0.98<S/K<1/02	15134	11712	2095	28941
1.02<S/K<1.05	8962	6516	1347	16825
1.05<S/K	4115	3286	758	8159
Total	39156	33884	6198	79238

Panel B. Average call price

	DTM<45	45<DTM<90	90<DTM	Total
S/K<0.95	2.82	5.15	11.72	5.10
0.95<S/K<0.98	4.50	11.61	21.77	8.92
0.98<S/K<1/02	18.32	31.28	44.69	25.48
1.02<S/K<1.05	48.43	60.23	75.07	55.13
1.05<S/K	74.55	83.78	96.86	80.34
Total	27.33	34.60	50.48	31.80

Panel C. Standard deviation of call prices

	DTM<45	45<DTM<90	90<DTM	Total
S/K<0.95	3.23	5.39	7.53	5.22
0.95<S/K<0.98	4.66	7.57	9.29	6.83
0.98<S/K<1/02	10.58	11.62	14.63	10.51
1.02<S/K<1.05	12.16	12.76	17.48	11.37
1.05<S/K	13.57	13.96	18.87	11.90
Total	24.19	24.59	27.85	26.86

Table 17: **Summary statistics of S&P 500 Index Call Options (2007-2011)**
This table contains the summary statistics of daily closing prices of European call options written on the S&P 500 index.The data is tabulated in maturity and moneyness bins. The sample starts on October 1, 2007 and ends on December 30, 2011, consisting of 1,073 days with a total of 111,596 option prices. DTM denotes the days to maturity and S/K denotes the moneyness. Panel A displays the amount of option prices per bin, panel B displays the average call option price per bin and panel C displays the standard deviation of the call prices over each bin.

Panel A. Number of call option contracts

	DTM<45	45<DTM<90	90<DTM	Total
S/K<0.95	5933	7309	2508	15750
0.95<S/K<0.98	11290	10837	3807	25934
0.98<S/K<1/02	15190	13622	4564	33376
1.02<S/K<1.05	10353	9288	3044	22685
1.05<S/K	6273	5777	1801	13851
Total	49039	46833	15724	111596

Panel B. Average call price

	DTM<45	45<DTM<90	90<DTM	Total
S/K<0.95	5.45	13.42	22.98	11.94
0.95<S/K<0.98	9.79	23.63	34.39	19.18
0.98<S/K<1/02	25.99	44.06	55.56	37.41
1.02<S/K<1.05	53.54	69.34	79.57	63.50
1.05<S/K	76.13	89.13	98.32	84.44

Panel C. Standard deviation of call prices

	DTM<45	45<DTM<90	90<DTM	Total
S/K<0.95	6.80	10.40	12.11	9.69
0.95<S/K<0.98	9.05	12.35	13.05	11.44
0.98<S/K<1/02	13.17	14.26	14.27	13.19
1.02<S/K<1.05	12.70	13.72	13.97	12.32
1.05<S/K	12.25	13.61	14.10	11.71
Total	24.37	24.49	23.90	28.53

Table 18: **In-sample average RMSEs of the CEV model, Heston model and PBS model (2002-2007).**
This table contains the averages of the daily in-sample root mean-squared errors (RMSEs) of the K&M approximation of the CEV model, the Heston model and the PBS model, for pricing S&P 500 index options, within various moneyness-maturity groups. The RMSEs are computed between the market option price and the model price. The model prices are computed by using the parameters estimates on each day of the sample. The sample period extends from January 2, 2002 to September 28, 2007 for a total of 79,238 European call options on 1146 days. DTM denotes the days to maturity and S/K denotes the moneyness (spot price divided by strike price). The smallest average RMSE within each moneyness-maturity bin is bold.

Moneyness S/K		Maturity: Short term DTM<45	Maturity: Mid term 45<DTM<90	Maturity: Long term 90<DTM<120	Overall
Deep out-of-the-money S/K<0.95	CEV	0.35	0.44	0.45	0.42
		(0.41)	(0.45)	(0.48)	(0.4)
	Heston	**0.22**	**0.24**	**0.35**	**0.24**
		(0.18)	(0.2)	(0.34)	(0.18)
	PBS	0.29	0.37	0.39	0.35
		(0.19)	(0.67)	(0.37)	(0.55)
Out-of-the-money 0.95<S/K<0.98	CEV	0.41	0.37	0.41	0.41
		(0.37)	(0.28)	(0.29)	(0.3)
	Heston	**0.20**	0.26	0.39	**0.25**
		(0.16)	(0.2)	(0.36)	(0.17)
	PBS	0.28	**0.28**	**0.28**	0.29
		(0.24)	(0.66)	(0.23)	(0.42)
At the money 0.98<S/K<1.02	CEV	0.39	0.33	0.56	0.39
		(0.23)	(0.28)	(0.45)	(0.24)
	Heston	**0.29**	0.32	0.40	0.32
		(0.19)	(0.24)	(0.34)	(0.2)
	PBS	0.35	**0.30**	**0.23**	**0.33**
		(0.39)	(0.66)	(0.15)	(0.46)
In the money 1.02<S/K<1.05	CEV	0.56	0.34	0.41	0.49
		(0.4)	(0.24)	(0.3)	(0.3)
	Heston	0.40	0.34	0.36	0.38
		(0.32)	(0.29)	(0.29)	(0.28)
	PBS	**0.42**	**0.30**	**0.29**	**0.38**
		(0.3)	(0.55)	(0.24)	(0.39)
Deep in the money 1.05<S/K	CEV	0.64	0.60	1.01	0.68
		(0.51)	(0.46)	(0.79)	(0.44)
	Heston	0.53	0.46	0.43	0.52
		(0.47)	(0.38)	(0.38)	(0.38)
	PBS	**0.55**	**0.41**	**0.41**	**0.51**
		(0.43)	(0.49)	(0.39)	(0.43)
Overall CEV		0.49	0.41	0.58	0.47
		(0.32)	(0.28)	(0.38)	(0.29)
Overall Heston		**0.34**	**0.33**	0.41	**0.34**
		(0.23)	(0.22)	(0.28)	(0.22)
Overall PBS		0.39	0.33	**0.31**	0.36
		(0.32)	(0.59)	(0.22)	(0.42)

Table 19: **In-sample average RMSEs of the CEV model, Heston model and PBS model (2007-201).**
This table contains the averages of the daily in-sample root mean-squared errors (RMSEs) of the K&M approximation of the CEV model, the Heston model and the PBS model, for pricing S&P 500 index options, within various moneyness-maturity groups. The RMSEs are computed between the market option price and the model price. The model prices are computed by using the parameters estimates on each day of the sample. The sample period extends from October 1, 2007 to December 30, 2011, for a total of 111,596 European call options on 1,146 days. DTM denotes the days to maturity and S/K denotes the moneyness (spot price divided by strike price). The smallest average RMSE within each moneyness-maturity bin is bold.

		Maturity			
Moneyness S/K		Short term DTM<45	Mid term 45<DTM<90	Long term 90<DTM<120	Overall
Deep out-of-the-money S/K<0.95	CEV	0.77	1.00	0.64	0.91
		(0.56)	(0.64)	(1.38)	(0.71)
	Heston	**0.39**	**0.43**	**0.59**	**0.47**
		(0.49)	(0.61)	(1.15)	(0.66)
	PBS	0.37	0.43	0.64	0.48
		(0.32)	(0.66)	(1.84)	(0.91)
Out-of-the-money 0.95<S/K<0.98	CEV	0.72	0.58	0.95	0.73
		(0.47)	(0.41)	(1.18)	(0.54)
	Heston	**0.40**	0.40	0.55	**0.45**
		(0.35)	(0.39)	(1.11)	(0.51)
	PBS	0.42	**0.37**	**0.46**	0.44
		(0.37)	(0.53)	(1.67)	(0.82)
At the money 0.98<S/K<1.02	CEV	0.66	0.45	0.93	0.65
		(0.58)	(0.27)	(0.84)	(0.45)
	Heston	**0.48**	0.41	0.50	0.47
		(0.32)	(0.27)	(0.67)	(0.34)
	PBS	0.45	**0.36**	**0.36**	**0.42**
		(0.42)	(0.41)	(1.25)	(0.54)
In the money 1.02<S/K<1.05	CEV	0.87	0.40	0.72	0.71
		(0.77)	(0.24)	(0.65)	(0.54)
	Heston	0.59	0.48	0.54	0.56
		(0.6)	(0.49)	(0.61)	(0.52)
	PBS	**0.52**	**0.35**	**0.41**	**0.46**
		(0.54)	(0.5)	(1)	(0.55)
Deep in the money 1.05<S/K	CEV	0.86	0.71	1.33	0.92
		(0.84)	(0.51)	(0.81)	(0.61)
	Heston	0.70	0.53	0.57	0.65
		(0.77)	(0.65)	(0.73)	(0.68)
	PBS	**0.59**	**0.38**	**0.48**	**0.52**
		(0.64)	(0.56)	(1.04)	(0.63)
Overall CEV		0.80	0.65	0.98	0.77
		(0.57)	(0.36)	(0.9)	(0.52)
Overall Heston		**0.53**	0.46	0.59	0.52
		(0.46)	(0.42)	(0.83)	(0.5)
Overall PBS		0.49	**0.39**	**0.49**	**0.46**
		(0.43)	(0.49)	(1.44)	(0.67)

Table 20: **1-day out-of-sample average RMSEs of the CEV model, Heston model and PBS model (2002-2007).**
This table contains the averages of the daily 1-day out-of-sample root mean-squared errors (RMSEs) of the K&M approximation of the CEV model, the Heston model and the PBS model, for pricing S&P 500 index options, within various moneyness-maturity groups. The RMSEs are computed between the market option price and the model price. The model prices are computed by using the parameters estimates of one day earlier for each day of the sample. The sample period extends from January 2, 2002 to September 28, 2007 for a total of 79,238 European call options on 1146 days. DTM denotes the days to maturity and S/K denotes the moneyness (spot price divided by strike price). The smallest average RMSE within each moneyness-maturity bin is bold.

Moneyness S/K		Maturity: Short term DTM<45	Maturity: Mid term 45<DTM<90	Maturity: Long term 90<DTM<120	Overall
Deep out-of-the-money S/K<0.95	CEV	0.61	0.85	1.05	0.79
		(0.78)	(0.94)	(1.16)	(0.84)
	Heston	0.51	0.68	1.00	0.65
		(0.65)	(0.84)	(1.21)	(0.74)
	PBS	**0.54**	**0.55**	**1.07**	**0.61**
		(2.55)	(0.75)	(5.63)	(2.03)
Out-of-the-money 0.95<S/K<0.98	CEV	0.86	1.14	1.24	1.05
		(0.9)	(1.09)	(1.27)	(0.97)
	Heston	0.71	1.07	1.32	0.95
		(0.82)	(1.09)	(1.4)	(0.95)
	PBS	**0.67**	**0.77**	**1.25**	**0.81**
		(1.64)	(0.9)	(5.52)	(1.91)
At the money 0.98<S/K<1.02	CEV	1.34	1.41	1.47	1.40
		(1.12)	(1.3)	(1.45)	(1.18)
	Heston	1.29	1.46	1.56	1.41
		(1.11)	(1.3)	(1.48)	(1.18)
	PBS	**1.18**	**1.20**	**1.58**	**1.25**
		(1.12)	(1.09)	(5.2)	(1.55)
In the money 1.02<S/K<1.05	CEV	1.39	1.61	1.65	1.55
		(1.2)	(1.4)	(1.48)	(1.24)
	Heston	1.37	1.55	1.58	1.49
		(1.17)	(1.35)	(1.44)	(1.22)
	PBS	**1.31**	**1.42**	**1.72**	**1.43**
		(1.6)	(1.21)	(4.73)	(1.59)
Deep in the money 1.05<S/K	CEV	1.30	1.73	2.02	1.64
		(1.17)	(1.54)	(1.84)	(1.33)
	Heston	1.25	1.50	1.56	1.43
		(1.14)	(1.32)	(1.43)	(1.2)
	PBS	**1.25**	**1.45**	**1.77**	**1.46**
		(2.15)	(1.22)	(4.67)	(2)
Overall CEV		1.25	1.43	1.59	1.37
		(0.97)	(1.17)	(1.35)	(1.06)
Overall Heston		1.19	1.36	1.50	1.31
		(0.96)	(1.15)	(1.34)	(1.04)
Overall PBS		**1.14**	**1.16**	**1.57**	**1.21**
		(1.53)	(0.98)	(5.12)	(1.71)

Table 21: **1-day out-of-sample average RMSEs of the CEV model, Heston model and PBS model (2007-2011).**
This table contains the averages of the daily 1-day out-of-sample root mean-squared errors (RMSEs) of the K&M approximation of the CEV model, the Heston model and the PBS model, for pricing S&P 500 index options, within various moneyness-maturity groups. The RMSEs are computed between the market option price and the model price. The model prices are computed by using the parameters estimates of one day earlier for each day of the sample. The sample period extends from October 1, 2007 to December 30, 2011 for a total of 111,596 European call options on 1,146 days. DTM denotes the days to maturity and S/K denotes the moneyness (spot price divided by strike price). The smallest average RMSE within each moneyness-maturity bin is bold.

Moneyness S/K		Maturity: Short term DTM<45	Maturity: Mid term 45<DTM<90	Maturity: Long term 90<DTM<120	Overall
Deep out-of-the-money S/K<0.95	CEV	1.49	2.16	2.24	2.01
		(1.92)	(2.41)	(2.86)	(2.21)
	Heston	1.22	1.82	2.23	1.72
		(2.04)	(2.41)	(2.67)	(2.21)
	PBS	**0.85**	**1.07**	**1.32**	**1.09**
		(1.25)	(1.34)	(2.12)	(1.41)
Out-of-the-money 0.95<S/K<0.98	CEV	1.85	2.30	2.43	2.21
		(2.11)	(2.56)	(2.91)	(2.35)
	Heston	1.68	2.17	2.44	2.07
		(2.24)	(2.57)	(2.77)	(2.4)
	PBS	**1.19**	**1.30**	**1.45**	**1.34**
		(1.5)	(1.51)	(2.1)	(1.57)
At-the-money 0.98<S/K<1.02	CEV	2.13	2.39	2.52	2.34
		(2.38)	(2.7)	(2.87)	(2.54)
	Heston	2.12	2.38	2.49	2.33
		(2.41)	(2.65)	(2.7)	(2.51)
	PBS	**1.58**	**1.54**	**1.55**	**1.61**
		(1.77)	(1.7)	(1.91)	(1.72)
In-the-money 1.02<S/K<1.05	CEV	2.01	2.38	2.54	2.33
		(2.41)	(2.62)	(2.7)	(2.47)
	Heston	1.98	2.34	2.42	2.24
		(2.38)	(2.63)	(2.65)	(2.47)
	PBS	**1.58**	**1.63**	**1.64**	**1.68**
		(2.04)	(1.84)	(1.92)	(1.93)
Deep in-the-money 1.05<S/K	CEV	1.85	2.40	2.72	2.35
		(2.27)	(2.52)	(2.65)	(2.31)
	Heston	1.78	2.23	2.32	2.12
		(2.26)	(2.55)	(2.6)	(2.39)
	PBS	**1.48**	**1.63**	**1.67**	**1.66**
		(2.27)	(1.87)	(2)	(2.1)
Overall CEV		2.01	2.40	2.61	2.30
		(2.14)	(2.52)	(2.71)	(2.35)
Overall Heston		1.90	2.26	2.46	2.18
		(2.23)	(2.54)	(2.66)	(2.38)
Overall PBS		**1.45**	**1.48**	**1.59**	**1.53**
		(1.73)	(1.61)	(1.99)	(1.7)

Table 22: **5-days out-of-sample average RMSEs of the CEV model, Heston model and PBS model (2002-2007).**
This table contains the averages of the daily 5-days out-of-sample root mean-squared errors (RMSEs) of the K&M approximation of the CEV model, the Heston model and the PBS model, for pricing S&P 500 index options, within various moneyness-maturity groups. The RMSEs are computed between the market option price and the model price. The model prices are computed by using the parameters estimates of five days earlier for each day of the sample. The sample period extends from January 2, 2002 to September 28, 2007 for a total of 79,238 European call options on 1146 days. DTM denotes the days to maturity and S/K denotes the moneyness (spot price divided by strike price). The smallest average RMSE within each moneyness-maturity bin is bold.

		Maturity			
Moneyness S/K		Short term DTM<45	Mid term 45<DTM<90	Long term 90<DTM<120	Overall
Deep out-of-the-money	CEV	0.93	1.41	1.78	1.30
S/K<0.95		(1.31)	(1.61)	(1.84)	(1.43)
	Heston	0.83	1.23	1.81	1.15
		(1.14)	(1.48)	(1.95)	(1.29)
	PBS	**0.73**	**0.92**	**1.99**	**0.96**
		(2.31)	(1.37)	(5.29)	(2.03)
Out-of-the-money	CEV	1.33	1.94	2.21	1.72
0.95<S/K<0.98		(1.53)	(1.91)	(1.99)	(1.69)
	Heston	1.15	1.85	2.37	1.61
		(1.4)	(1.84)	(2.31)	(1.63)
	PBS	**0.92**	**1.29**	**2.44**	**1.26**
		(1.55)	(1.46)	(5.32)	(1.8)
At-the-money	CEV	1.90	2.26	2.47	2.12
0.98<S/K<1.02		(1.78)	(2.11)	(2.12)	(1.91)
	Heston	1.83	2.27	2.60	2.10
		(1.71)	(2.05)	(2.41)	(1.87)
	PBS	**1.51**	**1.72**	**2.77**	**1.75**
		(1.45)	(1.65)	(5.24)	(1.78)
In-the-money	CEV	1.74	2.29	2.52	2.10
1.02<S/K<1.05		(1.68)	(2.09)	(2.2)	(1.85)
	Heston	1.72	2.20	2.45	2.02
		(1.6)	(1.99)	(2.25)	(1.78)
	PBS	**1.53**	**1.83**	**2.65**	**1.81**
		(1.7)	(1.66)	(4.82)	(1.89)
Deep in-the-money	CEV	1.54	2.24	2.73	2.04
1.05<S/K		(1.6)	(2.12)	(2.54)	(1.84)
	Heston	1.48	1.99	2.23	1.82
		(1.49)	(1.87)	(2.12)	(1.67)
	PBS	**1.33**	**1.74**	**2.39**	**1.69**
		(1.28)	(1.75)	(4.74)	(2.09)
Overall CEV		1.70	2.16	2.49	1.99
		(1.52)	(1.91)	(2.02)	(1.71)
Overall Heston		1.62	2.08	2.45	1.91
		(1.44)	(1.84)	(2.19)	(1.65)
Overall PBS		**1.38**	**1.62**	**2.59**	**1.63**
		(1.45)	(1.5)	(4.99)	(1.79)

Table 23: **5-days out-of-sample average RMSEs of the CEV model, Heston model and PBS model (2007-2011).**
This table contains the averages of the daily 5-days out-of-sample root mean-squared errors (RMSEs) of the K&M approximation of the CEV model, the Heston model and the PBS model, for pricing S&P 500 index options, within various moneyness-maturity groups. The RMSEs are computed between the market option price and the model price. The model prices are computed by using the parameters estimates of five days earlier, for each day of the sample. The sample period extends from October 1, 2007 to December 30, 2011 for a total of 111,596 European call options on 1,146 days. DTM denotes the days to maturity and S/K denotes the moneyness (spot price divided by strike price). The smallest average RMSE within each moneyness-maturity bin is bold.

Moneyness S/K		Maturity: Short term DTM<45	Maturity: Mid term 45<DTM<90	Maturity: Long term 90<DTM<120	Overall
Deep out-of-the-money S/K<0.95	CEV	2.19	3.55	4.09	3.23
		(2.77)	(3.66)	(4.17)	(3.26)
	Heston	1.88	3.13	3.98	2.90
		(2.72)	(3.57)	(4.08)	(3.2)
	PBS	**1.24**	**1.80**	**2.31**	**1.78**
		(1.74)	(2)	(2.77)	(1.95)
Out-of-the-money 0.95<S/K<0.98	CEV	2.88	4.03	4.40	3.69
		(3.06)	(3.9)	(4.29)	(3.5)
	Heston	2.67	3.79	4.39	3.49
		(3.07)	(3.9)	(4.22)	(3.51)
	PBS	**1.78**	**2.26**	**2.59**	**2.21**
		(2.06)	(2.27)	(2.72)	(2.19)
At the money 0.98<S/K<1.02	CEV	3.42	4.27	4.64	4.00
		(3.38)	(4.09)	(4.44)	(3.73)
	Heston	3.41	4.18	4.62	3.95
		(3.38)	(4.09)	(4.39)	(3.73)
	PBS	**2.37**	**2.64**	**2.86**	**2.64**
		(2.4)	(2.55)	(2.83)	(2.44)
In the money 1.02<S/K<1.05	CEV	3.09	4.14	4.61	3.83
		(3.33)	(4)	(4.26)	(3.61)
	Heston	3.11	4.07	4.49	3.76
		(3.28)	(4.03)	(4.26)	(3.64)
	PBS	**2.28**	**2.73**	**2.94**	**2.64**
		(2.46)	(2.71)	(2.89)	(2.53)
deep in the money 1.05<S/K	CEV	2.71	3.95	4.53	3.65
		(3.09)	(3.79)	(4.11)	(3.38)
	Heston	2.69	3.79	4.22	3.46
		(3.07)	(3.8)	(4.05)	(3.44)
	PBS	**2.05**	**2.65**	**2.90**	**2.52**
		(2.36)	(2.7)	(2.93)	(2.48)
Overall CEV		3.08	4.11	4.58	3.79
		(3.05)	(3.85)	(4.22)	(3.48)
Overall Heston		2.98	3.91	4.45	3.65
		(3.08)	(3.87)	(4.22)	(3.5)
Overall PBS		**2.11**	**2.50**	**2.82**	**2.45**
		(2.17)	(2.39)	(2.76)	(2.28)

G Empirical Test of K&M Approximation of the Heston model

In this appendix we apply the Heston model on empirical data by using its K&M approximation (which is that of the CEV model, used e.g. in section 5.2, with the elasticity parameter fixed at $\xi = 1/2$), instead of its closed-form pricing formula given by (26). By comparing the fit of the K&M approximation with that of the corresponding closed-form solution, we get an impression on how the approximation (as applied in this thesis) influences the fit of the model. The K&M approximation of the Heston model is calibrated on empirical data (described in section 5.1) similarly to the calibration of the CEV model described in section 5.2(but of course with $\xi = 1/2$). Table 24 gives the summary statistics of the estimated parameters.

The table shows quite similar results as in the case for the CEV model, reported in table 10 With the mean of the mean reversion parameter κ being equal to 2.75 below and 3.0 for the CEV model. Also, in this case, similarly to the CEV model, we have that the correlation parameter is equal to -0.9 and with a median very close to -1. Indicating a bad estimation of the correlation parameter for the approximation. If we compare these parameter statistics with the statistics for the parameter from the closed-form solution of the Heston model, reported in table 10, we find that there is large discrepancy between the estimates. This was to be expected as the parameter ξ of the CEV model was estimated very close to 0.5, the value of the Heston model. Also, the variance parameter is much closer to that of the CEV model than that of the Heston model closed-form.

For the fit of the corresponding K&M approximation of the Heston model we compute the RMSEs in two ways: first, we use the estimates obtained by calibrating the K&M approximation on empirical data, second we use the estimates obtained by using the closed-form solution in section 5.2. The results for the entire data sample are given in table 24.[5] First of all, when we use the corresponding parameters, we find that the all fits underperform the CEV model, although the difference is not that large. If we use the parameters from the (more accurate) closed-form solution we see that the errors increase tremendously, in all three cases (in-sample, 1-day out-of-sample and 5-days out-of-sample). Again, this indicates the inaccuracy of the approximation we use.

[5]The results for different moneyness-maturity bins, as we have done for the other fits, are left out, because these had no added value for our conclusion in this case.

Table 24: **Summary statistics Implied parameters of the Heston model, obtained through its K&M approximation, on S&P 500 Index Call Options (2002-2011).** This table gives the summary statistics of the implied parameters for the Heston using its K&M approximation, on S&P 500 Index Call Options. For each day of the data sample, the structural parameters of the model are estimated using the regularisation method. We use MATLAB's `lsqnonlin`, which requires the minimum and maximum boundaries and initial value as reported. The daily average of the estimated parameters is reported first, followed by its median and its standard deviation. The parameters are estimated using all available data on each day. The data sample used, starts on January 2, 2002 and ends on December 30, 2011. The parameters, κ,α,ω and ρ are respectively the mean-reversion rate, the long-run mean, the volatility of variance and the correlation between stock price and the corresponding variance v_t.

	Parameters				
	κ	α	ω	ρ	v_t
Initial value	3.0	0.1	0.5	-0.1	0.1
Min	0.0	0.0	0.0	-1.0	0.0
Max	10.0	1.0	1.0	0.0	1.0
	Heston model using the K&M approximation				
	$\hat{\kappa}$	$\hat{\alpha}$	$\hat{\omega}$	$\hat{\rho}$	$\hat{v}_t$
Mean	2.7330	0.0807	0.3182	-0.9031	0.0416
Median	2.7095	0.0562	0.2521	-0.9929	0.0239
Standard deviation	1.1690	0.0808	0.2186	0.1615	0.0566

Table 25: **In-sample and out-of-sample average root mean squared errors (RMSEs) of the K&M approximation of the Heston model.(2002-2011).** This table contains the averages and standard deviations (between parentheses) of daily root mean-squared errors (RMSEs) of the K&M approximation of the Heston model, both in and out-of-sample. The data sample is from January 2, 2002 to December 30, 2011, consisting of 2519 days with a total of 190,834 European call options.

K&M approximation of the Heston model		
In-sample	1-day out	5-days out
0.63	1.85	2.78
(0.50)	(1.80)	(2.76)
Heston's (1993) closed-form solution		
In-sample	1-day out	5-days out
2.66	3.04	3.50
(3.84)	(4.01)	(4.26)

References

Andersen, L. B. G., 2007. Efficient simulation of the heston stochastic volatility model. Bank of America Merrill Lynch, Working paper .

Andersen, T. G., Benzoni, L., Lund, J., 2002. An empirical investigation of continuous-time equity return models. Journal of Finance 57, 1239–1284.

Bakshi, G., Cao, C., Chen, Z., 1997. Empirical performance of alternative option pricing models. Journal of Finance 52, 2003–2049.

Baxter, M. W., Rennie, A. J. O., 2008. Financial Calculus: An introduction to derivative pricing. Cambridge University Press, 17th ed.

Benzoni, L., 2002. Pricing options under stochastic volatility: an empirical investigation. Unpublished Manuscript, Carlson School of Management .

Black, F., 1976. Studies of stock price volatility changes. In: Proceedings of the 1976 meetings of the american statistical association pp. 171–181.

Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. Journal of Political Economy 81, 637–654.

Bollerslev, T., Zhou, H., 2002. Estimating stochastic volatility diffusion using conditional moments of integrated volatility. Journal of Econometrics 109, 33–65.

Carr, P., Madan, D. B., 1999. Option valuation using the fast fourier transform. Journal of Computational Finance 2, 61.

Chacko, G., Viceira, L. M., 2003. Spectral gmm estimation of continuous-time processes. Journal of Econometrics 116, 259–292.

Chernov, M., Ghysels, E., 2000. A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. Journal of Financial Economics 56, 407–458.

Chiarella, C., Craddock, M., El-Hassan, N., 2000. The calibration of stock option pricing models using inverse problem methodology. Research Paper Series 39, Quantitative Finance Research Centre, University of Technology, Sydney.

Christoffersen, P., Jacobs, K., 2004. Empirical performance of alternative option pricing models. Journal of Financial Economics 72, 291–318.

Coleman, T. F., Li, Y., 1994. On the convergence of reflective newton methods for large-scale non-linear minimization on bounds. Mathematical Programming 67, 189–224.

Coleman, T. F., Li, Y., 1996. An interior, trust region approach for nonlinear minimization subject to bounds. SIAM Journal on Optimization Programming 6, 418–445.

Dumas, B., Fleming, J., Whaley, R. E., 1998. Implied volatility functions: empirical tests. Journal of Finance 53, 2059.

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, 327–343.

Jones, C. S., 2003. The dynamics of stochastic volatility: evidence from underlying and options markets. Journal of Econometrics 116, 181–224.

Karatzas, I., Shreve, S. E., 1991. Brownian Motion and Stochastic Calculus. Springer, second ed.

Kimmel, R. L., 2008. Changing times: accurate solutions to pricing and conditional problems in non-affine continuous-time models. Ohio State University, Working paper .

Kristensen, D., Mele, A., 2011. Adding and subtracting black-scholes: A new approach to approximating derivative prices in continuous-time models. Journal of Financial Economics 102, 390–415.

Mikhailov, S., Nögel, U., 2003. Heston's stochastic volatility model implementation, calibration and some extensions. Wilmott p. 74.

Pan, J., 2002. The jump-risk premia implicit in options: evidence from an integrated time-series study. Journal of Financial Econometrics 63, 3–50.

Yang, J., 2006. Stochastic Volatility Models: Option Price Approximations, Asymptotics and Maximum Likelihood Estimation. Ph.D. thesis, University of Illinois at Urbana-Champaign.