Comput Manage Sci
DOI 10.1007/s10287-008-0083-2

ORIGINAL PAPER

American option pricing under stochastic volatility: an empirical evaluation

Farid AitSahlia · Manisha Goswami · Suchandan Guha

Abstract Over the past few years, model complexity in quantitative finance has increased substantially in response to earlier approaches that did not capture critical features for risk management. However, given the preponderance of the classical Black–Scholes model, it is still not clear that this increased complexity is matched by additional accuracy in the ultimate result. In particular, the last decade has witnessed a flurry of activity in modeling asset volatility, and studies evaluating different alternatives for option pricing have focused on European-style exercise. In this paper, we extend these empirical evaluations to American options, as their additional opportunity for early exercise may incorporate stochastic volatility in the pricing differently. Specifically, the present work compares the empirical pricing and hedging performance of the commonly adopted stochastic volatility model of Heston (Rev Financial Stud 6:327–343, 1993) against the traditional constant volatility benchmark of Black and Scholes (J Polit Econ 81:637–659, 1973). Using S&P 100 index options data, our study indicates that this particular stochastic volatility model offers enhancements in line with their European-style counterparts for in-the-money options. However, the most striking improvements are for out-of-the-money options, which because of early exercise are more valuable than their European-style counterparts, especially when volatility is stochastic.

Keywords Stochastic volatility · Indirect inference · Model calibration · American option pricing · S&P 100 index · Approximate dynamic programming

F. AitSahlia (✉) · M. Goswami · S. Guha
Department of Industrial and Systems Engineering, Weil Hall 303,
University of Florida, Gainesville, FL 32611, USA
e-mail: farid@ise.ufl.edu

1 Introduction

The Black and Scholes (1973) formula for European options is without doubt the most successful pricing formula. However, it has known biases for which remedies have been attempted in a number of different directions. They can be traced back to Black (1976) who observed that time-series of equity returns displayed fatter tails than implied by the normal distribution under the geometric Brownian model of Black and Scholes (1973). Additionally, the non-constant implied volatility across time and strike prices, together with the phenomenon of volatility clustering (sustained periods of high-variability alternating with sustained periods of low-variability) and leverage effect (high variability associated with asset price declines) as observed in many studies has provided enough evidence to seek alternative option pricing models (cf., Rossi 1996; Fouque et al. 2000; Cont and Tankov 2000). Indeed, several such models have been proposed particularly in the last two decades. They include features such as stochastic volatility, stochastic interest rates, and stochastic jumps that are considered in various specifications, separately or in combinations. A sampling of the literature includes the stochastic interest rate model of Amin and Jarrow (1992); the stochastic volatility models of Hull (1987), Scott (1987), Melino and Turnbull (1990, 1995), Stein and Stein (1991), and Heston (1993); the stochastic volatility and stochastic interest rate models of Bailey and Stulz (1989), and Amin and Ng (1993); the stochastic-volatility jump-diffusion models of Bates (1996) and Scott (1997). This list is far from being exhaustive but already points to the different orientations of the efforts to improve upon the original model of Black and Scholes (1973).

As these models become increasingly complex, requiring additional parameters to be estimated and more sophisticated numerical techniques for their pricing and hedging, it is unclear whether this complexity cost is matched by additional accuracy in pricing and hedging. In a comprehensive study based on S&P 500 data, (Bakshi et al. 1997) compare a number of popular alternative models and conclude that overall, stochastic volatility models appear to provide the main source of option pricing improvement when considering various combinations of stochastic volatility with stochastic interest rates and stochastic jumps.

Given that the volatility of an asset return is not directly observable, its modeling as a stochastic process leads to challenging issues for the estimation of the associated parameters. For example, Bakshi et al. (1997) use cross-sectional information contained in option prices with different maturities and strike prices resulting in implied volatilities in order to infer estimates for the structural parameters of the stochastic volatility model. On the other hand, it appears that these implied structural parameters deviate substantially from their time-series counterparts (e.g., (Bakshi et al. 1997) use implied volatilities to estimate the correlation coefficient of the asset return innovation with that of its stochastic volatility as 0.76 whereas their estimate based on the underlying asset time-series is 0.23.) As an alternative, (Zhang and Shu 2003) proposed to use a two-step procedure to estimate first, via the indirect inference method of Gourieroux et al. (1993), structural parameters for the underlying asset, followed by a second set of additional parameters needed for option pricing via a market price calibration based on least-squares. An example for the latter is the empirically determined volatility risk premium, which is required since the pricing model is incomplete (i.e., non-uniqueness of risk-neutral equivalent measure). Zhang and Shu (2003) apply this two-step approach in their study comparing the pricing accuracy of the stochastic volatility model of Heston (1993) against the Black–Scholes constant volatility model. They use S&P 500 data to show that the Heston model significantly outperforms the Black–Scholes model in almost all moneyness-maturity classes. As their results are limited to European options, it remains to be seen how they generalize to American-style options, which have the additional early-exercise feature that may amplify the stochastic volatility characteristic of the underlying asset. These options are notoriously difficult to price in comparison. For example, Heston (1993) derives Fourier-based expressions for European option prices that can be evaluated through standard numerical techniques. In contrast, the application of standard dynamic programming to price American options is practically infeasible due to the curse of dimensionality problem especially acute in any model that extends the standard Black–Scholes paradigm. As a result, we adopt an approximate dynamic programming method developed particularly to price American options, namely that of least-squares Monte-Carlo (LSM) attributed to Carrière (1996) and Longstaff and Schwartz (2001), which approximates the continuation value in the related optimal stopping problem by a regression for which data are generated via simulation. It should be noted that alternative methods based on numerical solutions of partial differential equations have been developed to price American options under stochastic volatility consideration (cf., Ikonen and Toivanen 2007; Zvan et al. 1998). However, we chose the LSM approach as it presents more flexibility for our approach, particularly in its anticipated extensions to higher dimensions and path-dependent option payoffs, which explain partly its wide adoption. It also avoids the stability issues associated with numerical solutions of partial differential equations.

The remainder of the paper is organized as follows: Sect. 2 reviews the stochastic volatility model of Heston (1993) that we adopt for option pricing; Sect. 3 describes the data and their sources for this study; Sect. 4 concerns the two-step statistical technique that we follow for the parameter estimation of the resulting bivariate diffusion process; in Sect. 5 we adapt the Monte Carlo algorithm of Longstaff–Schwartz to compute prices and hedging parameters for American options under stochastic volatility; Sect. 6 compares numerical results using two different estimates of volatility for pricing, namely spot volatility and long-term mean. Section 7 addresses hedging errors and compares results between constant and stochastic volatility models: Sect. 8 concludes.

2 Heston's stochastic volatility model

Heston's model (1993) assumes both the underlying asset and its volatility to be stochastic processes defined by the following stochastic differential equations:

$$dS_t = (\mu - \delta)S_t dt + \sqrt{v_t}S_t dW_t^{(S)} \tag{1}$$

$$dv_t = \kappa(\theta - v_t)dt + \eta\sqrt{v_t}dW_t^{(v)} \tag{2}$$

where, under the original measure P, μ and δ are the underlying mean return and dividend rates, respectively, of the asset; $\sqrt{v_t}$ is its instantaneous volatility at time t, κ is the mean reversion rate of the variance process $\{v_t\}$, θ its long-term mean, and η is the instantaneous volatility of volatility. Here $W^{(S)}$ and $W^{(v)}$ are Brownian motions modeling innovations (randomness) affecting the asset price and its volatility. They are assumed to have a correlation coefficient ρ (i.e., $dW_t^{(S)} \times dW_t^{(v)} = \rho dt$). The specification of v_t as a so-called square root process guarantees that it remains positive as long as $2\kappa\theta \geq \eta^2$ (cf. Feller 1951).

Whereas the Black–Scholes model considers only one observable process, $\{S_t\}$, which can be used to estimate its sole unknown parameter σ required for pricing, the Heston model introduces an additional process $\{v_t\}$ that is not observable and for which three parameters, κ, η, θ, must be estimated together with ρ and spot volatility v_0. This variation therefore increases immediately the complexity in the estimation procedure. Though an estimate of σ can be obtained on the basis of past observations of S, the community of financial practitioners has now grown accustomed to estimating the constant σ through the so-called implied volatility approach, which uses option data as well. It sets the Black–Scholes function evaluated for an at-the-money option equal to the corresponding observed price, $C(K, T) = C_{obs}$, from which σ is inverted, thus now becoming a function of T and K. This approach may seem unfounded since the Black–Scholes formula results from the assumption of constant σ, but it is appealing when considering the result as being a *forward-looking* estimate of σ.

An additional complication emanating from any stochastic volatility model is related to option pricing. On the one hand, the Black–Scholes model leads to a unique equivalent risk-neutral pricing measure Q, under which the expected return of the underlying asset is the prevailing riskless rate r. On the other, this measure is not unique when volatility is stochastic. An approach to select the "right" pricing measure is through the so-called market price of volatility risk, which is independent of the underlying asset, and which identifies the "right" expected instantaneous asset return. As argued in Heston (1993), this market price of volatility risk is assumed to be proportional to the instantaneous volatility; i.e.,

$$\lambda(S_t, v_t, t) = \lambda v_t,$$

where λ is a constant to be estimated. In the Heston model, the price of a European call option with strike K, time to maturity T, and spot volatility v_t is given by:

$$C(S_t, v_t, t, T) = S_t P_1(S_t, v_t, T, K) - Ke^{-r(T-t)} P_2(S_t, v_t, T, K) \tag{3}$$

where $P_j(S_t, v_t, T, K)$, $j = 1, 2$, may be interpreted as adjusted or risk-neutralized probability distributions and are obtained by inverting their corresponding characteristic functions (cf. Heston 1993):

$$P_j(x, v, T, K) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty Re\left(\frac{e^{-i\phi \ln(K)} f_j(x, v_t, T, \phi)}{i\phi}\right) d\phi \tag{4}$$

for $j = 1, 2$, with the characteristic function f_j defined by:

$$f_j(x, v, T, \phi) = exp[C_j(T - t, \phi) + D_j(T - t, \phi)v + i\phi x]$$
$$C_j(\tau, \phi) = r\phi i\tau + \frac{a}{\eta^2}(b_j - \rho\eta\phi i + d_j)\tau - 2\ln\left[\frac{1 - g_j e^{d_j\tau}}{1 - g_j}\right]$$
$$D_j(\tau, \phi) = \frac{b_j - \rho\eta\phi i + d_j}{\eta^2}\left[\frac{1 - e^{d_j\tau}}{1 - g e^{d_j\tau}}\right],$$

with $\tau = T - t$ and

$$g_j = \frac{b_j - \rho\eta\phi i + d_j}{b_j - \rho\eta\phi i - d_j}, \quad d_j = \sqrt{(\rho\eta\phi i - b_j)^2 - \eta^2(u_j\phi i - \phi^2)},$$
$$u_1 = \frac{1}{2}, \quad u_2 = \frac{-1}{2}, \quad a = \kappa\theta, \quad b_1 = \kappa + \lambda - \rho\eta, \quad b_2 = \kappa + \lambda$$

3 Data description

3.1 Data for parameter estimation

Our empirical study uses the daily closing prices of S&P 100 index options, which are of both European and American exercise styles. We use European-style call options to estimate and calibrate the parameters in the Heston model. American put options are then priced through the LSM technique, which is evaluated via an out-of-sample analysis. We acquired our data from a database maintained by the Wharton School at the University of Pennsylvania through the WRDS interface. Our entire sample period for estimation and performance evaluation is January 1, 2002 to April 28, 2006. The reported price of any given option is the average of its bid and ask prices. We employ the following exclusionary criteria to filter data: (i) to eliminate general arbitrage opportunities we consider only options that satisfy $C \geq max(0, S - K)$; (ii) options with less than 6 days and more than 120 days to expiration are excluded because they are sensitive to liquidity biases; (iii) we excluded very deep out-of-money and very deep in-the-money options for they are not traded actively (with absolute moneyness (in %) of an option defined as $|S/K - 1|$, these are options with absolute moneyness greater than 7%); (iv) quotes less than $\frac{3}{8}$th are omitted; and (v) for estimation purposes, we use European call options only from the period January 1, 2002 to December 31, 2005. Once these filter rules are applied, we are left with 33, 860 observations for 1,007 days (an average of 34 options per day).

3.2 Data for out-of-sample analysis

Once the parameters are estimated, the calibrated Heston model is used for the out-of-sample performance evaluation of American put option prices obtained via the LSM method. The option data used for this purpose are from January 1, 2006 to April 28, 2006. We apply the same exclusionary criteria (i)–(v) as above on out-of-sample data. After filtration, for 81 days in the sample period we have 8, 184 put options (an average of 101 options per day).

4 Calibration of the Heston model

The structural parameters $(\mu, \kappa, \theta, \eta)$ of the Heston model (1)–(2), the correlation coefficient ρ and the risk premium λ are unobservable and need to be estimated. We follow a two-step procedure using first time-series data on asset returns to estimate the structural primary parameters$(\mu, \kappa, \theta, \eta)$, followed by a model calibration based on cross-sectional options data. For the latter we employ a non-linear least-squares technique using estimates from the first to evaluate the remaining parameters (λ, v, ρ), where v is the spot volatility to use at the moment of pricing (time 0, typically). Zhang and Shu (2003) use the same treatment to estimate the stochastic volatility process for S&P 500 index options, which have European exercise-style. Bakshi et al. (1997) use the least-squares method on cross-sectional option data to estimate the parameters of the Heston model. As noted in our introduction, we opt to follow instead the indirect inference method of Gourieroux et al. (1993) to estimate the structural parameters $(\mu, \kappa, \theta, \eta)$. Before proceeding any further, we remark that our study of 5-minute returns on the S&P 100 index does not show that their average is significantly different from zero, we therefore set μ to be zero.

4.1 Structural parameter estimation

The indirect inference method of Gourieroux et al. (1993) involves defining an auxiliary parameter Γ that is first estimated through an econometric approach on the basis of actual data and then used to estimate the original (structural) parameter of interest Θ through moment matching. This approach is particularly useful when likelihood functions, or any other estimation criterion, for the original model are difficult to evaluate, in contrast to the auxiliary model. The latter need not be correctly specified. If it is, then indirect inference is equivalent to maximum likelihood.

In our present case of stochastic volatility, the simulation is generated via an Euler discretization of (1)–(2):

$$r_t = \sqrt{v_t \tau}\varepsilon_{1t} \tag{5}$$

$$v_t = \kappa\theta\tau + (1 - \kappa\tau)v_{t-\tau} + \eta\sqrt{\tau v_{t-\tau}}\varepsilon_{2t} \tag{6}$$

where r_t is the return over the period $(t - \tau, t]$, with τ being the time discretization increment, and ε_{1t} and ε_{2t} are uncorrelated standard normal random variables. We should also note that ε_{1t} and ε_{2t} should be correlated by virtue of the model specification (1)–(2). However, because (5)–(6) will be matched with the GARCH model below, we can only estimate κ, θ, and η, and therefore set no correlation between ε_{1t} and ε_{2t}. This is the same approach taken by Engle and Lee (1996) and Zhang and Shu (2003) in their consideration of the Heston model for the S&P 500 index. Thus this data generating process (DGP) is fully specified when the parameter $\Theta = (\kappa, \theta, \eta)$ is fixed. On the other hand, we use a GARCH(1,1) model for the auxiliary model:

$$r_t = \varepsilon_t \tag{7}$$

$$h_t = \omega + \alpha\varepsilon_{t-1}^2 + \beta h_{t-1} \tag{8}$$

where $\varepsilon_t \sim N(0, h_t)$, with the conditional variance $h_t \equiv Var(r_t|F_{t-1})$ based on the information set F_{t-1} up to time $t-1$. This choice is predicated on the results of Nelson (1990) who argues that when the time increment τ of (5)–(6) is arbitrarily small, the first and second moments of volatility in GARCH(1,1) match those of the diffusion (2).

The parameter $\Gamma = (\omega, \alpha, \beta)$ of the GARCH model is first estimated on the basis of observed returns, yielding the estimate $\Gamma^{(0)} = \left(\omega^{(0)}, \alpha^{(0)}, \beta^{(0)}\right)$. Another set of data, now generated via the DGP (5)–(6), will be used to get another estimate $\Gamma^{(1)}$ of Γ. As the DGP (5)–(6) requires the knowledge of $\Theta = (\kappa, \theta, \eta)$, the idea of indirect inference is to get an initial estimate $\Theta^{(0)}$ of Θ to get $\Gamma^{(1)}$, which in turn will lead to the "right" estimate $\Theta^{(1)}$ of Θ via an optimality criterion. For this purpose, we follow (Engle and Lee 1996) who suggest an intermediate estimate $\Theta^{(0)}$ through the matching of the first two moments of (5)–(6) with (7)–(8), resulting in:

$$\kappa^{(0)} = \frac{\omega^{(0)}}{\tau}, \quad \theta^{(0)} = \frac{(1-\alpha^{(0)}-\beta^{(0)})}{\tau}, \quad \omega^{(0)} = \alpha^{(0)}\sqrt{(\xi-1)\tau}, \tag{9}$$

where ξ is the conditional kurtosis of the volatility shocks estimated from the actual data.

With $\Theta = \Theta^{(0)}$ now fixed, a sample of N observations is generated via the DGP (5)–(6) in order to fit a GARCH(1,1) model of the form (7)–(8), yielding an estimate $\Gamma^{(1)}$, which clearly depends on $\Theta^{(0)}$. With $R_t \equiv r_t(\Theta^{(0)})$ denoting the simulated returns and $l_t \equiv -\ln h_t - \frac{R_t^2}{(2h_t)}$ the log-likelihood function to estimate Γ, the maximum log-likelihood moment estimator

$$m(\Theta, \Gamma^{(1)})_{3\times 1} = \frac{1}{N}\sum_{t=1}^{N} \frac{\partial l_t(r_t(\Theta), \Gamma^{(1)})}{\partial \Gamma}\bigg|_{\Gamma=\Gamma^{(1)}} \tag{10}$$

is not necessarily the zero vector. As a result, the indirect inference method estimate $\Theta^{(1)}$ is that which minimizes the distance separating the above moment from zero. Engle and Lee (1996) determine that the appropriate metric to use for this purpose is defined by the matrix

$$\Omega = \left(\frac{1}{N}\sum_{t=1}^{N} \frac{\partial l_t(\tilde{R}_t, \Gamma)}{\partial \Gamma}\frac{\partial l_t(\tilde{R}_t, \Gamma)}{\partial \Gamma^T}\bigg|_{\Gamma=\Gamma^{(1)}}\right)^{-1},$$

where $\tilde{R}_t$ is the observed (market) return for period t and the superscript T stands as the transpose sign. The indirect inference estimate of Θ is then

Table 1 GARCH estimates of auxiliary parameters

	ω	α	β
Market data $\left(\Gamma^{(0)}\right)$	3.03×10^{-7}	4.26×10^{-2}	9.53×10^{-1}
Simulated data $\left(\Gamma^{(1)}\right)$	4.03×10^{-9}	0.4535	0.5457

$$\Theta^{(1)} = \arg\min_{\Theta} \left\{ m^T \left(\Theta, \Gamma^{(1)}\right) \Omega m \left(\Theta, \Gamma^{(1)}\right) \right\}, \tag{11}$$

where the function m is as in (10), with Θ no longer fixed at $\Theta^{(0)}$. As a result, the simulated observations $r_t(\Theta)$, which are generated through normal distributions in the DGP (5)–(6), will appear explicitly in terms of $\Theta = (\kappa, \theta, \eta)$. We should also note that the matrix Ω is positive definite by way of Ω^{-1} (cf. Newey and West 1987).

Table 1 reports the estimated parameters from GARCH fitting. The estimates on the first row are based on market data. Those on the second are based on simulated index returns with the starting structural parameters (as in set of equations in (9)) set at $\Theta^{(0)} = \left(0.0044, 6.93 \times 10^{-5}, 0.0653\right)$. These returns are at 5-minute intervals, with the starting index level and volatility set to those reported for the index level and implied volatility on January 1, 2002, that is 588.98 and 0.187377, respectively. Every day, there are approximately 77 5-min time intervals, resulting in a simulation length of 1,007 days. In other words, we generate sample paths for 77,539 time steps, with $\tau = 0.012987013$. Then the minimization problem in (11) becomes

$$\begin{aligned} \min_{\Theta} \Big\{ & 124.89 - 3245.11\kappa\theta + 0.08\kappa + 0.78\eta - 1.05\theta\kappa^2 - 10.1\kappa\theta\eta \\ & + 0.000252\kappa\eta + 21081.64\kappa^2\theta^2 + 0.000015\kappa^2 + 0.00122\eta^2 \Big\} \end{aligned}$$

resulting in the indirect structural parameter estimate $\Theta^{(1)} = (2.65, 0.029, 0.154)$.

4.2 Option pricing calibration

With the above estimates of $\Theta^{(1)}$, we proceed to evaluate (ρ, v, λ) by nonlinear least-squares, using data on S&P 100 European call option prices. For each day, there are several options available with different maturities and strike prices. For option j on day t, define the error between observed market price (p_{tj}) and model price $(p_{tj}(\rho_t, \lambda_t, v_t))$ as:

$$\Delta p_{tj} = p_{tj} - p_{tj}(\rho_t, \lambda_t, v_t), \tag{12}$$

where $p_{tj}(\rho_t, \lambda_t, v_t)$ are obtained via Heston's pricing formula for a European call (3).

Let J be the number of options on day t. The estimate $(\hat{\rho}_t, \hat{v}_t, \hat{\lambda}_t)$ for day t minimizes the sum of squares of errors $\sum_{j=1}^{J} \Delta p_{tj}^2$, which we performed through the non-linear least-squares routine of MATLAB.

Table 2 Nonlinear least-squares estimates of option pricing parameters

	Mean	Median	Std.Dev.	Minimum	Maximum
$\hat{\rho}$	−0.487	−0.644	0.4562	−0.999	0.999
$\hat{v}$	0.0349	1.5×10^{-8}	0.13	2.2×10^{-14}	0.8525
$\hat{\lambda}$	2.14	-2.2×10^{-14}	6.8338	−9.88	10

These calculations are done for each day, and for a total number T of days in the sample set, with the final estimates are reported as their averages:

$$\hat{\rho} = \frac{1}{T}\sum_{t=1}^{T}\hat{\rho}_t, \quad \hat{\lambda} = \frac{1}{T}\sum_{t=1}^{T}\hat{\lambda}_t, \quad \hat{v} = \frac{1}{T}\sum_{t=1}^{T}\hat{v}_t$$

For our data set $T = 1,007$, thus as many values of $(\hat{\rho}_t, \hat{v}_t, \hat{\lambda}_t)$ are calculated, with summary statistics reported in Table 2.

Finally, we remark that the price-based criterion (12) may lead to biased estimates in favor of out-of-the-money options as in-the-money options are likely to be priced correctly with any reasonable set of structural parameters. But this is not a concern precisely because of this latter feature. Alternatively, the calibration procedure could be based on errors based on implied volatilities instead of prices, thus accounting in some way for moneyness.

4.3 Discussion of estimated parameters

The structural parameters that control the volatility process of the Heston model are the long-term mean θ, the mean-reversion speed κ, and the volatility of volatility η. The average variance level θ is estimated to be 0.029, which corresponds to an annualized volatility of 17.03%. The average historical volatility (obtained directly from past asset prices) is 15.97%, which is very close to our estimate. We note that our estimate for θ is slightly less than the value 0.35 reported by Zhang and Shu (2003) as our observations are from a more stable subset of theirs (the S&P 100 index and S&P 500 index stocks, respectively.)

The estimated mean reversion rate κ is 2.65, i.e., every 64 days volatility reverts halfway to the long-run mean, and is very close to theirs (2.75). Our estimate of the variation of volatility η is 0.154, which is smaller than 0.425 reported by Zhang and Shu (2003) for the same reason as just described above.

The stock prices and their volatility are negatively correlated with $\rho = -0.487$, close to the value of −0.464 reported by Zhang and Shu (2003), reflecting the fact that the stocks in the S&P 100 index dominate this effect over the entire set in the S&P 500 index. This negative correlation is an expression of the so-called leverage effect, which indicates in particular that equity market reactions to negative news are more pronounced than those to positive news. Zhang and Shu (2003) also estimate this coefficient on the basis of simply the historical index returns and their associated volatility. The value they obtain, −0.23, is closer to both our value and theirs, compared to the value of −0.64 obtained by Bakshi et al. (1997), thus supporting the use of our two-step estimation procedure as discussed earlier.

Our average estimate of the volatility risk premium λ is 2.14 but varies considerably over the sample. This phenomenon is similarly observed by Zhang and Shu (2003) who obtained an average value of -0.8716. The challenge in estimating this parameter is not surprising as it is a component of the drift, which is notoriously difficult to evaluate directly due to the so-called mean blur (cf. Luenberger 1998, for example.)

5 Pricing American options

5.1 The LSM framework

Carrière (1996) and Longstaff and Schwartz (2001) proposed simulation-based techniques with least-squares regression fitting (LSM) to price American options with constant volatility. In our paper we use the algorithm in Longstaff and Schwartz (2001) to price American options for models with both constant and stochastic volatilities.

As is standard, we assume a complete probability space $(\Omega, \mathcal{F}, \mathcal{P})$ and a finite time horizon $[0, T]$, where Ω is the set of all possible realizations of the stochastic processes, (S_t, v_t) over $[0, T]$. A generic element of Ω will be denoted ω_i. $\mathcal{F}$ and P are the associated probability filtration and original/physical measure, respectively.

One way to evaluate the price of an American option is to first determine its optimal exercise date. Since the latter is random (depending on the path followed by the underlying asset), the corresponding payoff will be random. We are therefore interested in obtaining the optimal exercise strategy, which maximizes the expected value of this payoff, in the interval $[0, T]$. Assume that the option can only be exercised at discrete time points $0 \leq t_1 \leq t_2 \leq \cdots \leq t_M \leq T$. If the option is exercised at maturity, the value of the option will simply be that of the payoff at maturity. For a given sample path ω_i at time t_m, the payoff from immediate exercise is known. If the option is not exercised at time t_m, then its expected payoff from continuation is the expectation of the remaining discounted cash flows $C(\omega_i, t : t_m, T)$, $t \in \{t_{m+1}, t_{m+2}, \ldots t_M\}$, with respect to the risk-neutral measure Q. In other words, at time t_m, the value of continuation, $F(w_i, t_m)$, is given by

$$F(w_i, t_m) = E_Q\left[\sum_{j=m+1}^{M} \exp\left(-\int_{t_m}^{t_j} r(w_i, t)dt\right) C(\omega_i, t_j : t_m, T) \mid \mathcal{F}_{t_m}\right] \quad (13)$$

where $r(w_i, t)$ is the riskless interest rate and the expectation is conditional on the information available up to time t_m. With this setup, the problem reduces to evaluating the conditional expected payoff $F(w_i, t_m)$ at every time step t_m, for every path ω_i, and comparing it with the immediate payoff. The option will then be exercised immediately if the latter is higher.

5.2 Longstaff–Schwartz algorithm

For our study we adapt the Longstaff–Schwartz algorithm to handle stochastic volatility in order to price S&P 100 put options, which are of American-style exercise. In the LSM approach the goal is to approximate the above conditional expectation at

$t_m, m = 1, 2, \ldots, M-1$. Moving backwards in time, LSM assumes that at time t_{M-1}, the unknown functional form of $F(w_i, t_{M-1})$ in Eq. (13) can be represented as a linear combination of a countable set of $\mathcal{F}_{t_{M-1}}$-measurable functions. For the purpose of this paper we choose the basis functions as the set of weighted Laguerre polynomials:

$$L_n(X) = \exp\left(\frac{-X}{2}\right) \frac{e^X}{n!} \frac{d^n}{dX^n}(X^n e^{-X})$$

For our implementation, as in Longstaff and Schwartz (2001), we use the first three Laguerre polynomials, $(L_0(X), L_1(X), L_2(X))$. Then $F(w_i, t_{M-1})$ can be approximated as:

$$F(w_i, t_{M-1}) \approx \sum_{j=0}^{3} A_j L_j(X), \tag{14}$$

where A_j are the coefficients of the regression equation. As in Longstaff and Schwartz (2001), we approximate the value of $F(w_i, t_{M-1})$ by regressing the discounted payoffs, $C(\omega_i, t : t_{M-1}, T)$ onto the basis functions for the paths where the option is in the money.

6 Empirical evaluation: pricing S&P 100 put options

To implement the LSM we generate $N = 100{,}000$ sample paths (50,000 plus 50,000 antithetic), $\{\omega_1, \omega_2, \ldots, \omega_N\}$ for stochastic processes, S_t and v_t using a second-order discretization of the Heston model (cf. Glasserman 2004). For a fixed time-step size h (we use $h = 1/365$) and letting $\sigma_1 = \rho\eta$ and $\sigma_2 = \sqrt{1-\rho^2}\eta$, we generate over time asset and variance values, s_i and v_i, respectively, through the recursions

$$\begin{aligned}
s_{i+1} &= s_i(1 + rh + \sqrt{v_i}w_1) + \left(\left(r + \frac{\sigma_1 - \kappa}{4}\right) s_i\sqrt{v_i} + \left(\frac{\kappa\theta}{4} - \frac{\eta^2}{16}\right)\frac{s_i}{\sqrt{v_i}}\right) w_1 h \\
&\quad + \frac{1}{2}r^2 s_i h^2 + \frac{s_i}{2}\left(v_i + \frac{\sigma_1}{2}\right)(w_1^2 - h) + \frac{1}{4}\sigma_2 s_i(w_1 w_2 + \varepsilon) \\
v_{i+1} &= \kappa\theta h + (1 - \kappa h)v_i + \sqrt{v_i}(\sigma_1 w_1 + \sigma_2 w_2) - \frac{1}{2}\kappa^2(\theta - v_i)h^2 \\
&\quad + \left[\left(\frac{\kappa\theta}{4} - \frac{\sigma^2}{16}\right)\frac{1}{\sqrt{v_i}} - \frac{3\kappa\sqrt{v_i}}{2}\right](\sigma_1 w_1 + \sigma_2 w_2)h \\
&\quad + \frac{1}{4}\sigma_1^2(w_1^2 - h) + \frac{1}{4}\sigma_2^2(w_2^2 - h) + \frac{1}{2}\sigma_1\sigma_2 w_1 w_2,
\end{aligned}$$

where w_1 and w_2 are independent normal random variables with mean 0 and standard deviation $\sqrt{h}$ and ε is a random variable independent of w_1 and w_2 such that $P\{\varepsilon = h\} = P\{\varepsilon = -h\} = 1/2$.

For the purpose of our study, we initialize the asset price to be the index value on the current day and volatility to be either the spot volatility or the long-term mean estimate that we obtained from the calibration above. Then successive values of (s_i, v_i) are generated with each sampling of the random variables w_1, w_2 and ε. Once the sample paths have thus been generated, we then compute the cash flows at every time step in $\{0, h, 2h, \ldots, T\}$ going backwards in time, in the same manner as in Longstaff and Schwartz (2001), followed by the sample average of discounted cash flows to compute the option price.

6.1 Options classification

A usual practice is to classify options according to maturity(T, in days) and moneyness (x), which for a put option is defined as $x = K/S - 1$. In our analysis, short-term options are those for which $T < 45$; mid-term options have $45 \leq T < 90$ and those with $T \geq 90$ are labeled long-term options. With respect to moneyness, we classify options with $x > 0.05$ as deep-in-the-money (DITM), those with $x \in (0.02, 0.05)$ as in-the-money (ITM), those with $x \in (-0.02, 0.02)$ as at-the-money (ATM), those with $x \in (-0.05, -0.02)$ as out-of-the-money (OTM), and those with $x > -0.05$ as deep-out-of-the-money (DOTM).

6.2 Pricing accuracy criteria

We compare the stochastic volatility and constant volatility models according to out-of-sample mean relative-pricing errors and absolute-pricing errors, following the option classification described below. For the ith option in a given group of n, we let C_i^M and C_i denote the observed market and model prices, respectively. Then the mean relative pricing error for the group is defined as

$$\mathrm{MRE} = \frac{1}{n} \sum_{i=1}^{n} \frac{(C_i^M - C_i)}{C_i}$$

and its mean absolute relative pricing error as:

$$\mathrm{MAE} = \frac{1}{n} \sum_{i=1}^{n} \frac{|C_i^M - C_i|}{C_i}.$$

MRE is an indicator of the pricing bias whereas MAE evaluates the magnitude of mispricing. The results in Tables 3, 4, 5 and 6 are obtained by computing these averages for each moneyness-maturity group. Since we have a choice in setting the initial volatility, namely the spot volatility or the long-term mean θ, we chose to consider both as displayed in these tables.

6.3 Analysis

Tables 3 and 4 report out-of-sample mean relative pricing error for constant volatility (CV) and stochastic volatility (SV) models, with standard deviations in parentheses.

Table 3 Out-of-sample mean relative pricing errors (initial volatility = spot volatility)

Maturity					
Moneyness	Model	Short-term	Mid-term	Long-term	Overall
DITM	CV	0.008 (0.002)	0.054 (0.004)	0.064 (0.005)	0.042 (0.004)
	SV	0.002 (0.001)	0.019 (0.002)	0.019 (0.003)	0.013 (0.002)
ITM	CV	0.183 (0.024)	0.356 (0.009)	0.395 (0.014)	0.311 (0.016)
	SV	0.133 (0.006)	0.241 (0.006)	0.228 (0.011)	0.201 (0.007)
ATM	CV	1.124 (0.099)	0.935 (0.021)	0.799 (0.042)	0.953 (0.054)
	SV	0.991 (0.032)	0.699 (0.014)	0.479 (0.016)	0.723 (0.021)
OTM	CV	2.278 (0.092)	1.649 (0.031)	1.121 (0.042)	1.683 (0.055)
	SV	2.049 (0.049)	1.217 (0.026)	0.660 (0.025)	1.309 (0.034)
DOTM	CV	1.648 (0.065)	1.477 (0.043)	1.05 (0.046)	1.392 (0.051)
	SV	1.411 (0.054)	0.969 (0.025)	0.411 (0.029)	0.931 (0.036)
Overall	CV	1.048 (0.056)	0.894 (0.022)	0.686 (0.029)	0.876 (0.036)
	SV	0.917 (0.029)	0.629 (0.014)	0.359 (0.017)	0.635 (0.019)

For each entry of constant volatility (CV) and stochastic volatility (SV), we report the average pricing error relative to the observed market price with corresponding standard errors in parentheses

Table 4 Out-of-sample mean relative pricing errors (initial volatility = long-term average θ)

Maturity					
Moneyness	Model	Short-term	Mid-term	Long-term	Overall
DITM	CV	−0.007 (0.001)	0.017 (0.002)	0.023 (0.004)	0.011 (0.002)
	SV	−0.009 (0.001)	0.000 (0.001)	0.003 (0.002)	−0.002 (0.001)
ITM	CV	0.095 (0.005)	0.208 (0.005)	0.242 (0.008)	0.182 (0.006)
	SV	0.072 (0.004)	0.146 (0.004)	0.139 (0.008)	0.119 (0.005)
ATM	CV	0.784 (0.028)	0.576 (0.009)	0.455 (0.012)	0.605 (0.016)
	SV	0.697 (0.021)	0.449 (0.008)	0.288 (0.009)	0.478 (0.013)
OTM	CV	1.233 (0.028)	0.867 (0.014)	0.581 (0.016)	0.894 (0.019)
	SV	1.149 (0.027)	0.688 (0.014)	0.347 (0.014)	0.729 (0.018)
DOTM	CV	0.487 (0.036)	0.328 (0.018)	0.103 (0.025)	0.306 (0.026)
	SV	0.484 (0.034)	0.226 (0.016)	−0.071 (0.019)	0.213 (0.023)
Overall	CV	0.518 (0.019)	0.399 (0.010)	0.281 (0.013)	0.399 (0.014)
	SV	0.479 (0.018)	0.302 (0.009)	0.141 (0.011)	0.307 (0.012)

For each entry of constant volatility (CV) and stochastic volatility (SV), we report the average pricing error relative to the observed market price with corresponding standard errors in parentheses

Correspondingly, Tables 5 and 6 are for the same evaluations in absolute value this time. Given that in-sample performance of stochastic volatility model was better than the less complex constant volatility model, one may argue that by virtue of overfitting, SV would fare less well than CV out-of-sample. However, it is evident from the results in Tables 3, 4, 5 and 6 that stochastic volatility (SV) performs better than the CV model

Table 5 Out-of-sample mean absolute relative pricing errors (initial volatility = spot volatility)

Maturity					
Moneyness	Model	Short-term	Mid-term	Long-term	Overall
DITM	CV	0.024 (0.002)	0.058 (0.004)	0.066 (0.005)	0.049 (0.004)
	SV	0.019 (0.001)	0.028 (0.002)	0.029 (0.003)	0.026 (0.002)
ITM	CV	0.185 (0.023)	0.356 (0.09)	0.395 (0.014)	0.312 (0.015)
	SV	0.136 (0.006)	0.241 (0.006)	0.228 (0.011)	0.202 (0.008)
ATM	CV	1.126 (0.099)	0.932 (0.023)	0.798 (0.051)	0.952 (0.058)
	SV	0.991 (0.032)	0.699 (0.014)	0.479 (0.016)	0.723 (0.021)
OTM	CV	2.283 (0.073)	1.636 (0.031)	1.169 (0.043)	1.696 (0.049)
	SV	2.048 (0.049)	1.217 (0.026)	0.660 (0.025)	1.309 (0.034)
DOTM	CV	1.684 (0.062)	1.492 (0.042)	1.074 (0.043)	1.417 (0.049)
	SV	1.419 (0.053)	1.002 (0.023)	0.514 (0.022)	0.979 (0.033)
Overall	CV	1.060 (0.052)	0.895 (0.022)	0.885 (0.030)	0.947 (0.035)
	SV	0.923 (0.028)	0.638 (0.014)	0.382 (0.016)	0.648 (0.019)

For each entry of constant volatility (CV) and stochastic volatility (SV), we report the average absolute pricing error relative to the observed market price with corresponding standard errors in parentheses

Table 6 Out-of-sample mean absolute relative pricing errors (initial volatility = long-term average θ)

Maturity					
Moneyness	Model	Short-term	Mid-term	Long-term	Overall
DITM	CV	0.014 (0.000)	0.028 (0.002)	0.036 (0.004)	0.026 (0.002)
	SV	0.015 (0.000)	0.017 (0.001)	0.019 (0.001)	0.017 (0.001)
ITM	CV	0.099 (0.004)	0.208 (0.005)	0.242 (0.017)	0.183 (0.009)
	SV	0.079 (0.004)	0.147 (0.004)	0.140 (0.007)	0.122 (0.005)
ATM	CV	0.784 (0.028)	0.576 (0.009)	0.455 (0.012)	0.605 (0.016)
	SV	0.697 (0.022)	0.449 (0.008)	0.288 (0.009)	0.478 (0.013)
OTM	CV	1.233 (0.028)	0.867 (0.014)	0.581 (0.016)	0.894 (0.019)
	SV	1.149 (0.027)	0.689 (0.014)	0.347 (0.014)	0.729 (0.018)
DOTM	CV	0.589 (0.028)	0.485 (0.011)	0.377 (0.013)	0.484 (0.017)
	SV	0.571 (0.027)	0.394 (0.009)	0.281 (0.011)	0.416 (0.016)
Overall	CV	0.544 (0.018)	0.433 (0.008)	0.338 (0.012)	0.438 (0.013)
	SV	0.503 (0.016)	0.339 (0.007)	0.215 (0.009)	0.352 (0.011)

For each entry of constant volatility (CV) and stochastic volatility (SV), we report the average absolute pricing error relative to the observed market price, with corresponding standard errors in parentheses

while both tend to underestimate the actual market price. The latter is likely due to the LSM approach, which is applied to both models, and which is known to have a downward bias (cf. Longstaff and Schwartz 2001.) However, these results are in line with the European option results obtained by Zhang and Shu (2003) and do not necessarily point to any systematic bias due to model misspecification as supported by a comparison between relative and absolute errors.

A few additional observations emerge as well. First, we note that for all scenarios SV results are less variable than those of CV. Second, in contrast to Zhang and Shu (2003) we also assess the effect of using the long-term average (θ) as the volatility estimate when pricing. As shown in Tables 4 and 6, this choice induces a significant reduction in error, both relatively and absolutely, as compared to using the spot volatility. Third, the SV model improves upon the CV model on both criteria, particularly for deep-in-the-money (DTIM)) and deep-out-of-the-money (DOTM) options, with mid-to long-term maturities. In fact, whereas Zhang and Shu (2003) note that significant improvement is mostly for deep-in-the-money European options, our study points out that a similar pattern emerges for both deep-in-the-money and out-of-the-money or deep-out-of-the money American-style options. The former was to be expected, given the immediate early exercise opportunity, but the latter can also be explained by the fact that stochastic volatility may lead to an early exercise opportunity despite the option currently being out-of-the-money or deep-out-of-the-money. In conclusion, we stress that the improvement provided by the Heston model is relative. In an absolute sense (i.e., in direct comparison with actual prices), both the Black–Scholes and Heston models are significantly off the mark, particularly for out-of-the-money options. Fortunately, a better improvement, both relatively and absolutely, is noticeable with respect to hedging as discussed next.

7 Hedging performance

In this section, we compare single instrument hedging errors for the CV and SV models, in which only the underlying asset is used for the two models. An argument supporting this strategy is given in Bakshi et al. (1997). For the SV model, our aim is to hedge a short position in the put option with τ periods to expiration. Let $X_S(t)$ be the number of shares of stock (index) to be purchased and $X_0(t)$ be the residual cash, resulting in a time-t portfolio value $X_0(t) + X_S(t)S(t)$. Adapting formula (21) from Bakshi et al. (1997) to our setting with no jump, the standard minimum variance hedging problem under the SV model yields:

$$X_S(t) = \Delta_S(t, \tau) + \frac{1}{S}\rho\eta\Delta_V(t, \tau) \tag{15}$$

where $\Delta_S(t, \tau)$ and $\Delta_V(t, \tau)$ are the delta and vega parameters, which are estimated using backward finite differences. The resulting cash position is then

$$X_0(t) = C(t, \tau) - X_S(t)S(t), \tag{16}$$

where $C(t, \tau)$ is the (observed) time-t price of an option expiring at time $t + \tau$. Formula (15) shows that if volatility is stochastic and correlated with the underlying asset returns, then the position in the asset is governed not only by the impact of asset price changes but also by volatility changes. For the CV model we need only consider the first term on the right hand side of formula (15), corresponding to the classical delta-hedge in the Black-Scholes context.

Whereas the hedging strategy (15) assumes continuous rebalancing in response to changing market conditions, it is clear that in practice such rebalancing can only be effected at discrete intervals, say of length Δt. We can then construct a self-financing portfolio as follows. At time-t short the put and go long $X_S(t)$ shares of the asset, with the remainder $X_0(t)$ in a risk-free asset with rate $R(t)$. At time $t + \Delta t$ the resulting hedging error is

$$H(t + \Delta t) = X_S(t)S(t + \Delta t) + X_0(t)e^{R(t)\Delta t} - C(t + \Delta t, \tau - \Delta t) \tag{17}$$

Proceeding in the same manner at times $t+2\Delta t$, $t+3\Delta t$, $\ldots$, $t+\tau$, yields a collection of hedging errors $H(t + j\Delta t)$, for $j = 1, \ldots, M$, where M is an integer such that $M\Delta t = \tau$, which are averaged as

$$H(\Delta t) = \frac{1}{M}\sum_{j=1}^{M} |H(t + j\Delta t)|.$$

To obtain the hedging results in Table 7, we use day $t - 1$ data to determine the model parameters, which are then used on day t with the current day's index and interest values to construct the desired hedge as given in Eq. (15) (for both CV and SV models.) Finally, since we are rebalancing the hedge daily, the hedging error is computed on day $t + 1$ for this strategy. These steps are repeated for each option and every trading day in the month of January 2006. The average absolute hedging errors for each moneyness-maturity group are reported in Table 7.

Based on the results in Table 7, the SV model is the better of the two for all groups except for short-term deep-in-the-money options. But even in this case, SV model is essentially doing better since in the CV model the option is immediately exercised and thus no hedge is created, which leads to a reported zero hedging error. In contrast, in the SV model the option is allowed additional exercise opportunities till maturity in most cases, thus resulting in higher profits.

8 Conclusions

In this paper we assessed the performance of a stochastic volatility model to price American-style options on the basis of actual market data. Using S&P 100 index options data we compared the pricing and hedging accuracy of the stochastic volatility model of Heston (1993) against the widely adopted model of Black and Scholes (1973). The specification of the unobservable volatility as a stochastic process leads to the challenge of estimating its associated parameters, which we address through the use of a two-step procedure of indirect inference and non-linear optimization for calibration. The additional complexity of American option pricing is then tackled through the adaptation of the least-squares Monte-Carlo algorithm of Longstaff and Schwartz (2001) to our context.

Table 7 Out-of-sample average single instrument hedging errors

Maturity					
Moneyness	Model	Short-term	Mid-term	Long-term	Overall
DITM	CV	0.284	0.492	0.437	0.404
	SV	0.328	0.238	0.161	0.242
ITM	CV	0.607	0.682	0.624	0.637
	SV	0.552	0.572	0.593	0.572
ATM	CV	0.487	0.423	0.385	0.432
	SV	0.462	0.397	0.353	0.404
OTM	CV	0.214	0.249	0.207	0.223
	SV	0.194	0.218	0.171	0.194
DOTM	CV	0.046	0.0663	0.059	0.057
	SV	0.029	0.045	0.038	0.037
Overall	CV	0.328	0.382	0.342	0.351
	SV	0.313	0.294	0.263	0.290

For the CV model, this strategy is the usual delta hedge. For the SV model, the strategy involves both delta and vega parameters (cf. formula (15))

Using criteria that pertain to both bias and mispricing magnitude, our study shows that the Heston model is a better alternative for both pricing and hedging purposes, particularly for out-of-the-money and deep-out-of-the-money options. The latter is clearly attributable, when volatility is stochastic, to the early exercise opportunity that may materialize for American-style options, in contrast to European-style exercise.

This study can be further extended in a number of different ways: (i) by including more parameters in the model to account for features such as jumps and stochastic interest rates; (ii) by conducting the empirical assessment using other underlying assets and options; and (iii) by additional comparisons with alternative stochastic volatility models such as those of Hull (1987), Scott (1987), and Stein and Stein (1991).

References

Amin K, Jarrow R (1992) Option pricing on a risky asset in a stochastic interest economy. Math Finance 2:881–910

Amin K, Ng V (1993) Option valuation with system stochastic volatility. J Finance 48:881–910

Bailey W, Stulz R (1989) The pricing of stock index options in a general equilibrium model. J Financ Quant Anal 24:1–12

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

Bates D (1996) Jumps and stochastic volatility: exchange rate processes implicit in deutschemark options. Rev Financ Stud 9:69–108

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. J Polit Econ 81:637–659

Carrière J (1996) Valuation of the early-exercise price for derivatives using simulations and splines. Insurance Math Econ 19:19–30

Cont R, Tankov P (2000) Financial modelling with jump processes, 1st edn. Chapman and Hall/CRC Press, Boca Raton

Engle R, Lee G (1996) Estimating diffusion models of stochastic volatility. Modeling stock market volatility. Academic Press, Inc, New York, pp 333–355

Feller W (1951) Two singular diffusion problems. Ann Math 54:173–182

Fouque J, Papanicolaou G, Sircar R (2000) Derivatives in finacial markets with stochastic volatility. Cambridge University Press, Cambridge

Glasserman P (2004) Monte carlo methods in financial engineering, 1st edn. Springer, Heidelberg

Gourieroux C, Monfort A, Renault E (1993) Indirect inference. J Appl Economet 8:s85–s118
Heston S (1993) A closed-form solution for options with stochastic volatility with applications to bond and currency options. Rev Financ Stud 6:327–343
Hull JC, White A (1987) The pricing of options on assets with stochastic volatilities. J Finance 42:281–300
Ikonen S, Toivanen J (2007) Efficient numerical methods for pricing american options under stochastic volatility. Numer Methods Partial Differ Equat 24(1):104–126
Longstaff FA, Schwartz ES (2001) Valuing american options by simulation: a simple least-squares approach. Rev Financ Stud 14:113–147
Luenberger D (1998) Investment science. Oxford University Press, Oxford
Melino A, Turnbull S (1990) Pricing foreign currency options with stochastic volatility. J Economet 45:239–265
Melino A, Turnbull S (1995) Misspecification and the pricing and hedging of long-term foreign currency options. Int Money Finance 14:373–393
Nelson D (1990) Arch models as diffusion approximations. J Economet 45:7–38
Newey W, West K (1987) A simple positive definite heteroscedasticity and autocorrelation consistent covariance matrix. Econometrica 55:703–708
Rossi P (1996) Modeling stock market volatility. Academic Press, New York
Scott L (1987) Option pricing when the variance changes randomly: theory, estimators, and applications. J Financ Quant Anal 22:419–438
Scott L (1997) Pricing stock options in a jump-diffusion model with stochastic volatility and interest rates: an application of fourier inversion methods. Math Finance 7:413–426
Stein E, Stein J (1991) Stock price distributions with stochastic volatility. Rev Financ Stud 4:727–752
Zhang J, Shu J (2003) Pricing S&P 500 index options with heston's model. In: IEEE Computational intelligence for financial engineering proceedings, pp 85–92
Zvan R, Forsyth PA, Vetzal KR (1998) Penalty methods for american options with stochastic volatility. J Comput Appl Math 91:199–218