

ON THE INTERCHANGEABILITY OF BARRIER OPTION PRICING MODELS

Jaco Visagie

University of Pretoria, Department of Statistics, Pretoria, South Africa
e-mail: *jaco.visagie@up.ac.za*

An important question when modelling option prices is which of the multitude of option pricing models to use. In this paper, the calculation of barrier option prices is considered. These exotic options are found in many financial markets the world over. It is demonstrated numerically that it is possible to replicate (with a high degree of accuracy) the barrier option prices obtained from one model by making use of a different model; these models are referred to as ‘interchangeable’. Tests for the interchangeability of barrier option pricing models are developed and applied. However, the tests developed are not specific to barrier option pricing models and can also be applied to the prices of other exotic options.

Key words: Barrier options, Interchangeable models, Option pricing.

1. Introduction and motivation

Schoutens, Simons and Tistaert (2004) demonstrated that it is possible to nearly perfectly replicate the observed market prices of European call options using various models. The authors demonstrated this phenomenon by calibrating the models to the observed prices which entails choosing the parameters of each model so that the model prices resemble the observed prices as closely as possible. The near identical prices achieved by the various models indicate that the models are *interchangeable*, meaning that the prices obtained using one model can be nearly perfectly replicated under another model. However, when the calibrated models are used to calculate the prices of exotic options (including barrier options) the resulting prices vary substantially. In Schoutens et al. (2004), the authors do not calibrate the models to exotic option prices. One possible reason for not considering calibrations to exotic option prices is that the market prices of these options are not as readily available as those of their European counterparts. It remains unclear whether or not the prices of exotic options calculated using a given model can be replicated using a different model.

In this paper, various option pricing models are calibrated to barrier option prices and the prices obtained under these models are compared. It is demonstrated numerically that it is possible to choose the parameters of a given model such that the option prices calculated under this model conform closely to the option prices calculated under another model. It is not surprising that the parameters of two models can be chosen so as to ensure that these models provide similar prices for a given set of options, however, this paper shows that, in certain instances, models can be calibrated such that the difference between the resulting prices are negligible (in a way made precise below). The models considered are the Heston stochastic volatility model, the Black–Scholes model, the geometric normal inverse Gaussian model and the time changed geometric lognormal-normal model.

Observed barrier option prices are not always readily available. As a result, the models considered are not calibrated to prices observed in a financial market. Instead, a hypothetical market containing three different types of barrier call options is constructed. The mentioned barrier options are digital barriers, down-and-out barriers and up-and-out barriers. The payoff of a barrier option is a function of whether or not the underlying stock price reaches some predetermined barrier level during the lifetime of the option.

In practice, barrier options are usually monitored at discrete times. This means that the stock price used to determine whether or not the barrier level has been reached is monitored at discrete intervals (often daily). Therefore, the price of a barrier option is a function of the minimum or the maximum of the stock price when observed in discrete time. The joint distributions of the stock price and its minimum or maximum observed at discrete times are generally unknown for the majority of option pricing models (including the Black–Scholes model). As a result, Monte Carlo techniques are employed to estimate barrier option prices.

The use of Monte Carlo simulation introduces Monte Carlo error into the option price estimates calculated under the models used and, because of this error, it is not feasible to require that the option prices calculated using two different models coincide exactly in order for the models to be called interchangeable. In order to determine whether or not two models are interchangeable, the concept of a *perfect* calibration is introduced. When a model is calibrated to a set of observed option prices, the difference between the observed and calculated option prices contains two components; the first is due to the misspecification of the model and the second is ascribed to the Monte Carlo error introduced by the calculation method used. In the remainder of the text, a calibration is considered to be perfect if the contribution to the discrepancy between the observed and calculated prices due to model misspecification is negligible when compared to the Monte Carlo error. This definition is made exact in Section 5. If a model can be perfectly calibrated to option prices obtained from another model, then these models are called interchangeable.

In Section 2 the various option pricing models considered are discussed. Section 3 introduces the types of barrier options used while the construction of the hypothetical financial market is discussed in Section 4. In Section 5, the calibration of the models to barrier option prices is discussed together with an analysis of the results obtained. Some conclusions are presented in Section 6.

The numerical results presented in the paper are limited to barrier options, however, the methodology developed can be applied more generally; it can also be used to compare the prices calculated under different models for other exotic options.

2. Option pricing models

The option pricing models that will be employed throughout the remainder of this paper are briefly discussed below. It is well known that, in order to ensure that option prices calculated as discounted expected payoffs under a given model are arbitrage free, the discounted value of the stock price is required to form a martingale, see Chapter 2.5 of Schoutens (2003). The calibration procedures used in Schoutens et al. (2004) enforce this restriction using the mean correcting martingale measure. This means that each of the parameters in the models are allowed to vary freely with the exception of the location parameter. The value of the location parameter is chosen so as to ensure the martingaleness of the discounted stock price using a predetermined formula. The same technique is used in this paper

and the formulae for the location parameters are included in the discussions below. The algorithms used to simulate price paths under each of the models are provided in Appendix A.

In the remainder of the paper stock prices are indexed daily. Let S_t denote the stock price at time t such that, if t is a whole number, then S_t denotes the stock price at the end of day t . A constant, continuously compounded risk free daily interest rate, denoted by r , is assumed throughout.

2.1 The Heston model

This model was introduced in Heston (1993). The dynamics of the stock price under the Heston model are as follows:

$$\begin{aligned}\frac{dS_t}{S_t} &= \mu dt + \sigma_t dW_t, \\ d\sigma_t^2 &= \kappa(\eta - \sigma_t^2)dt + \theta\sigma_t d\tilde{W}_t,\end{aligned}$$

where W_t and \tilde{W}_t are two correlated standard Brownian motions with $\text{Corr}(dW_t d\tilde{W}_t) = \rho dt$. In the calibration procedure used, μ is set to r in order to ensure the martingaleness of $e^{-rt} S_t$.

Under the Heston model, the volatility process σ_t^2 follows a Cox-Ingersoll-Ross process. As a result, if $2\kappa\eta \geq \theta^2$, then the volatility of the process is strictly positive and the stationary distribution of σ_t^2 is gamma($2\kappa\eta/\theta^2, \theta^2/2\kappa$), where the density function of a gamma(a, b) random variable is

$$f(x) = \frac{x^{a-1}}{b^a \Gamma(a)} \exp\left(-\frac{x}{b}\right),$$

for $a > 0$, $b > 0$ and $x > 0$. In the numerical analysis to follow, the restriction $2\kappa\eta \geq \theta^2$ is enforced, meaning that the calibration procedure chooses parameters under the constraint that $2\kappa\eta \geq \theta^2$.

2.2 The Black–Scholes model

Under this model, introduced in Black and Scholes (1973), the dynamics of the stock price are given by

$$S_t = S_0 \exp(\mu t + \sigma W_t),$$

with $\mu \in \mathbb{R}$, $\sigma > 0$ and W_t a standard Brownian motion. Enforcing the martingaleness requirement placed on $e^{-rt} S_t$ entails setting $\mu = r - \sigma^2/2$.

2.3 The geometric normal inverse Gaussian (N ◦ IG) model

A random variable $X : \Omega \rightarrow \mathbb{R}$ is said to follow a N ◦ IG distribution with parameter set $\theta = (\alpha, \beta, \mu, \delta)$ if it has density

$$f(x; \theta) = \frac{\alpha\delta}{\pi} \exp\left(\delta\sqrt{\alpha^2 - \beta^2} + \beta(x - \mu)\right) \frac{K_1\left(\alpha\sqrt{\delta^2 + (x - \mu)^2}\right)}{\sqrt{\delta^2 + (x - \mu)^2}},$$

where $\alpha > 0$, $|\beta| < \alpha$, $\mu \in \mathbb{R}$, $\delta > 0$ and K_1 denotes the modified Bessel function of the third kind with index 1.

A stochastic process $L = (L_t; t \geq 0)$ is a N ◦ IG process with parameter set $(\alpha, \beta, \mu, \delta)$ if $L_0 = 0$, L has stationary and independent increments, and

$$L_{t+s} - L_t \sim \text{N} \circ \text{IG}(\alpha, \beta, \mu s, \delta s),$$

for all $s > 0$.

Under the geometric $N \circ IG$ model the stock price is represented as

$$S_t = S_0 \exp(L_t),$$

where L_t is a $N \circ IG(\alpha, \beta, \mu, \delta)$ process. The martingaleness condition is enforced by setting

$$\mu = r + \delta \left(\sqrt{\alpha^2 - (\beta + 1)^2} - \sqrt{\alpha^2 - \beta^2} \right).$$

As a result, in order for $e^{-rt} S_t$ to form a martingale it is required that $-\alpha < \beta < \alpha - 1$. This requirement is enforced in the calibration procedure used. For more details regarding this model, see Schoutens et al. (2004).

2.4 The geometric lognormal-normal model

Clark (1973) proposed the use of time changed models in option pricing. The mentioned paper argues that changes in a stock price follow a normal distribution, but that the number of price changes per time unit is not constant. Clark proposed the use of a normal distribution to model individual price changes and a lognormal distribution to model the number of these changes per unit time. The resulting stochastic process is referred to as lognormal-normal. In Clark's 1973 paper this process is proposed as a model for the first differences in observed stock prices. However, this process is used in order to model log-returns in the remainder of this paper. Formally the model is defined as

$$S_t = S_0 \exp(L_{T_t}),$$

where T_t is a lognormal process ($T_{t+1} - T_t \sim \log N(\alpha, \beta)$) and L_t is a Brownian motion, with parameters μ and σ^2 , which is independent of T_t .

The discounted stock price forms a martingale if, and only if,

$$E[S_t] = S_0 e^{rt} \iff E[\exp(L_{T_1})] = e^{rt} \iff \int_0^\infty e^x f(x) dx = e^r, \quad (1)$$

where f is the density of L_{T_1} . Standard calculations yield

$$f(x) = \frac{1}{2\pi\sigma\beta} \int_0^\infty y^{-\frac{3}{2}} \exp\left(-\frac{(x - \mu y)^2}{2\sigma^2 y} - \frac{(\log(y) - \alpha)^2}{2\beta^2}\right) dy. \quad (2)$$

Using (1) and (2) we can solve for the value of μ that ensures the martingaleness of $e^{-rt} S_t$ numerically.

3. Barrier options

The payoff of a barrier option depends on whether or not the stock price reaches some predetermined barrier level before the maturity of the option. In practice, barrier options are usually monitored at discrete time intervals. Throughout the paper, the barrier options considered are assumed to be monitored daily. This means that the stock price at the end of each business day is used in order to assess whether or not the barrier level has been reached. We consider three types of barrier options; digital barrier (DB) call options, down-and-out barrier (DOB) call options and up-and-out barrier

(UOB) call options. We denote by h and π respectively the payoff function and price of a barrier option.

A DB call option pays a fixed amount (set equal to 1 unit of currency throughout) at maturity if the stock price reaches the predetermined level $H > S_0$ before the maturity of the option. If this barrier is not reached the option expires worthless. The payoff and price of a DB call option with barrier level H and time to maturity T are therefore

$$h = \mathbb{I} \left(\max_{j \in 0, \dots, T} S_j \geq H \right) \quad \text{and} \quad \pi = e^{-rT} E \left[\mathbb{I} \left(\max_{j \in 0, \dots, T} S_j \geq H \right) \right].$$

A DOB barrier call option has the same payoff as a European call option if the stock price does not reach or drop below the level $H < S_0$ during the lifetime of the option. If the stock price reaches this level the option becomes worthless. The payoff and price of a DOB call option with strike price K , barrier level H and time to maturity T are therefore

$$h = \mathbb{I} \left(\min_{j \in 0, \dots, T} S_j > H \right) (S_T - K)^+ \quad \text{and} \quad \pi = e^{-rT} E \left[\mathbb{I} \left(\min_{j \in 0, \dots, T} S_j > H \right) (S_T - K)^+ \right],$$

where $(a)^+ = \max(a, 0)$.

As was the case for a DOB barrier call option, the payoff of an UOB barrier call option is the same as that of a European call option if the stock price does not achieve or exceed some level $H > S_0$ before the maturity date of the option. If the stock price reaches $H > S_0$, then the option has a payoff of 0. The payoff and price of an UOB barrier call option with strike price K , barrier level H and time to maturity T are therefore

$$h = \mathbb{I} \left(\max_{j \in 0, \dots, T} S_j < H \right) (S_T - K)^+ \quad \text{and} \quad \pi = e^{-rT} E \left[\mathbb{I} \left(\max_{j \in 0, \dots, T} S_j < H \right) (S_T - K)^+ \right].$$

The empirical results shown in Section 5 require the prices of the three types of barrier options discussed under various models. If the options are monitored continuously (meaning that the stock price is continuously observed in order to determine whether or not a barrier crossing has taken place) closed form formulae are available for the option prices under certain models, for example, see Chapter 9.1 of Schoutens (2003) for the prices of barrier options under the Black–Scholes model if the monitoring occurs in continuous time. However, in the case considered (where the stock price is monitored at discrete times), no closed form formulae are available for the option prices under the models considered. As a result, the option prices are estimated using Monte Carlo simulation as follows. A price path for the stock is simulated. The length of this price path is equal to the maximum time to maturity of the options under consideration. Next, the payoff of each of the options given this price path is calculated. The simulation of the price path is repeated 1000 times, resulting in 1000 realised payoffs for each option. The price of a given option is estimated by discounting the average of the payoffs associated with this option. Below, we refer interchangeably to the estimates calculated using this method as option prices and option price estimates. A general discussion of barrier option pricing can be found in Cont and Tankov (2004).

The calibration procedure used to obtain the numerical results is computationally expensive and therefore the number of Monte Carlo replications is fixed at 1000 in order to keep the computer time required for calculations within manageable limits while allowing for option price estimates with

small standard errors relative to their means. Various variance reduction techniques were considered in order to increase the accuracy of the estimates. However, these techniques were discarded because the associated computational cost outweighed the reduction in the variance of the option price estimates.

4. Construction of a hypothetical market

In the following discussion, rather than using observed option price data, the prices calculated from a hypothetical market are used in order to obtain the numerical results. This hypothetical market contains a risk free bond with a daily interest rate of $r = 0.1/252$ as well as a single stock. The daily price of the stock follows a Heston model with parameter set $(r, \kappa, \eta, \rho, \theta) = (0.1/252, 0.1, 0.0001, -0.7, 0.001)$. The current price of the stock is 1.

In this market, the three types of barrier options defined in Section 3 are available for various barrier levels and times to maturity. In this study, 25 options for each option type, with times to maturity ranging from 21 days to 111 days, are considered. The strike prices of each of the DOB and UOB options are fixed at 1 (the same convention is used in Schoutens et al., 2004). In order to calculate the market prices of the options defined in this hypothetical market, the option price estimates are calculated using the Monte Carlo method discussed in Section 3. These estimates are considered the known, fixed market prices of the options available in the hypothetical market. The prices, barrier levels and times to maturity of the options are reported in Appendix B. Various models are calibrated to these option prices in order to establish whether or not it is possible to accurately replicate the option prices using these models.

5. Model calibration

Calibrating an option pricing model to a set of market prices entails varying the parameters of the model freely (with the exception of the location parameter) in an attempt to minimise some distance measure between the market prices and the prices calculated using the model. The distance measure used to obtain the numerical results shown is the root mean square error (*RMSE*). In the case where $\pi_{1,j}$ and $\pi_{2,j}$, $j = 1, 2, \dots, n$ denote two sets of option prices,

$$RMSE = \sqrt{\frac{1}{n} \sum_{j=1}^n (\pi_{1,j} - \pi_{2,j})^2}.$$

The *RMSE* is often used to evaluate the fit of calibrated option pricing models (see, for example, Schoutens et al., 2004).

Each of the models is calibrated to the market prices of the DB, DOB and UOB options separately.

5.1 Optimisation method

The calibration procedure employed in this paper minimises the *RMSE* using particle swarm optimisation (PSO) (see Kennedy and Eberhart, 1995) in Matlab. PSO falls within the class of swarm intelligence techniques used for optimisation and the algorithm searches for a global minimum by evaluating several possible points within the parameter space of the model concerned. The *RMSE* is calculated at each of the points under consideration. These points, referred to as particles, are then

moved within the parameter space of the model in search of a global minimum. The movements of each of the particles are determined by a number of factors such that each particle has a tendency to move towards the most optimal position that has been located jointly by all of the particles, as well as a tendency to move to the most optimal position that has been located by that specific particle. Once a particle is moving in a specific direction, it also has a tendency to keep moving in that direction. This tendency is referred to as the inertia of the particle.

The numerical results presented in Section 5 are obtained using PSO with 50 particles. The initial location coordinates of each of these particles is obtained by specifying a range for each parameter and then simulating a uniform random number within the range specified for each parameter. (where the random numbers associated with the various parameters are independent). The ranges for the parameters are chosen so as to include the parameter values deemed likely by the user's previous experience, which is, of necessity, a subjective procedure. The ranges for the parameters associated with each of the models used are provided in Appendix C. The value of the objective function is evaluated for each particle using the simulation method described in Section 3.

The code used is an adaptation of code that can be downloaded from <http://yarpiz.com/50/yypea102-particle-swarm-optimization>.

5.2 Perfect calibrations

The aim of this paper is to ascertain the extent to which it is possible to accurately replicate the option prices calculated in the hypothetical market using the various models considered. In order to aid this determination, a calibration will be classified as either *perfect* or *imperfect*. For each individual calibration, we are interested in testing the following hypotheses:

$$\begin{aligned} H_0 &: \text{The calibration is perfect.} \\ H_A &: \text{The calibration is not perfect.} \end{aligned} \tag{3}$$

In order to explore this line of thought further, it is necessary to define what is meant by a perfect calibration. When calibrating simple models to European option prices, the *RMSE* can be computed exactly and a calibration can be defined to be perfect if the resulting *RMSE* equals 0. However, because of the Monte Carlo error contained in the barrier option prices calculated using the calibrated model, the *RMSE* associated with a given calibration is a random variable. This means that one cannot simply insist on a *RMSE* being equal to 0 in order for the calibration to be called perfect. As was explained before, the *RMSE* contains two components; the first is due to the misspecification of the model, while the second is due to the Monte Carlo error contained in the price estimates. The relative sizes of these errors will be used in order to classify a calibration as either perfect or imperfect.

In the hypothetical market under consideration, the stock price follows a Heston model. We calibrate the Heston model to the observed prices of each option type separately. Using Monte Carlo simulation, we calculate the option prices as well as the associated *RMSE* under the calibrated Heston model. This process is repeated 500 times in order to obtain as many samples of the *RMSE* under the correctly specified model. Note that the observed values of the *RMSE* are obtained in the absence of model misspecification. As a result, these values can be used in order to approximate the distribution of the *RMSE* when only Monte Carlo error is present.

Table 1. Means and standard deviations of calculated *RMSEs* (multiplied by 100).

	Heston	Black–Scholes	Exp N ◦ IG	Time changed
DB	0.428(0.131)	0.755(0.182)	0.498(0.130)	0.749(0.174)
DOB	0.056(0.028)	0.058(0.027)	0.068(0.032)	0.056(0.027)
UOB	0.044(0.015)	0.153(0.028)	0.046(0.015)	0.154(0.028)

Consider a specific calibration where the model is misspecified, for example, where the Black–Scholes model is calibrated to the option prices in the hypothetical market. In order to estimate the distribution of the *RMSE* in this case, the procedure continues in a manner similar to the one described above. The option prices and the *RMSE* under the calibrated model are obtained using Monte Carlo simulation. The process is repeated 500 times, resulting in 500 realisations of the *RMSE*. In this case, however, the misspecification of the model will contribute to the *RMSE*.

Let $RMSE_H$ and $RMSE_M$ denote the values calculated for the *RMSEs* under the Heston and the misspecified models respectively. Table 1 contains the means and the standard deviations (in brackets) of the calculated *RMSEs* associated with each calibration. In the notation used above, the results pertaining to $RMSE_H$ are found in the first column of the table, while those associated with $RMSE_M$ are found in the remaining columns. Please note that the calculated values of *RMSEs* are close to zero. In order to aid comparison, all values in Table 1 have been multiplied by 100. The geometric N ◦ IG process model is abbreviated by “Exp N ◦ IG” in the table. The same convention is also used in the tables to follow.

Returning to the definition of a perfect calibration, a calibration is defined to be perfect if the population mean of the $RMSE_M$ does not exceed that of the $RMSE_H$. Therefore, a perfect calibration is achieved if the contribution of the misspecification of the model to the *RMSE* is negligible when compared to the contribution of the Monte Carlo error.

The hypotheses in (3) can be restated as follows;

$$\begin{aligned} H_0 : \mu_M &= \mu_H, \\ H_A : \mu_M &> \mu_H, \end{aligned} \tag{4}$$

where μ_M and μ_H denote the population means of $RMSE_M$ and $RMSE_H$ respectively. Perhaps the first idea that comes to mind when considering the hypotheses in (4) is a two sample t-test for the equality of the means. However, because of the large sample sizes used, this test will reject the null hypothesis for very small differences between the sample means. As a result, this approach is discarded in favour of the two formal tests considered next. The first utilises a mixed Bayesian-frequentist approach and the second uses a bootstrap methodology.

5.3 Mixed Bayesian-frequentist hypothesis test

The test considered below is based on an alternative to frequentist significance tests explained in Cox and Hinkley (1976). This method is outlined in general terms before describing the specifics of the test used to answer questions relating to the classification of calibrations as either perfect or imperfect.

Using Bayesian techniques to draw inference about the value of some parameter θ , it is assumed

that the parameter is a realisation of a random variable Θ with a known probability distribution. Consider the case where one has a sample \mathbf{y} from a known distribution with unknown parameter θ and the null hypothesis (H_0) to be tested specifies $\theta = \theta_0$. In order to specify the distribution of Θ completely, a probability of p_0 is assigned to the event $\{\theta = \theta_0\}$ and a prior distribution $p_A(\theta)$ is assigned to Θ under the alternative hypothesis (H_A). This means that the prior distribution of Θ consists of an atom of probability p_0 at θ_0 and a probability density function of $(1 - p_0)p_A(\theta)$ over the remaining values of θ .

In order to determine whether or not $\theta = \theta_0$, we use the posterior odds given \mathbf{y} ; this is defined as the probability of H_0 given \mathbf{y} divided by the probability of H_A given \mathbf{y} . The posterior odds is then calculated as follows:

$$\frac{P(H_0|Y = \mathbf{y})}{P(H_A|Y = \mathbf{y})} = \frac{p_0 f_{Y|\Theta}(\mathbf{y}|\theta_0)}{(1 - p_0) \int_{\theta \neq \theta_0} p_A(\theta) f_{Y|\Theta}(\mathbf{y}|\theta) d\theta}, \quad (5)$$

where $f_{Y|\Theta}$ denotes the likelihood function associated with the data for a given value of θ . If the data support H_0 , then the value of the posterior odds will be large.

In order to apply the technique outlined above, one can proceed in a way similar to Example 10.12 in Cox and Hinkley (1976). However, the approach used in the current paper deviates from the methodology described in this example in that here a mixed Bayesian-frequentist approach is chosen; this is explained below.

The Bayesian technique discussed above can be applied to test hypotheses regarding a single population. In order to apply the test to classify calibrations as perfect or imperfect, define $X = RMSE_M - RMSE_H$. The hypotheses in (4) can now be reformulated as follows:

$$H_0 : \mu_X = 0,$$

$$H_A : \mu_X > 0,$$

where μ_X denotes the expected value of X . In order to sample from X , simply subtract the realised values of $RMSE_H$ from those of $RMSE_M$, resulting in a sample of size 500. In order to apply the method described above the distribution of X needs to be specified. The normal distribution is considered as a possible candidate and the Lilliefors test is used to test the hypothesis of normality for each of the nine calibrations under consideration. Using a 1% significance level, the assumption of normality is not rejected for any of the samples considered.

Under the null hypothesis $X \sim N(0, \sigma_0^2)$, where σ_0^2 is a fixed constant. Since there is no prior knowledge regarding the value of σ_0^2 , a mixed Bayesian-frequentist approach is followed. Set σ_0^2 equal to the variance of the realisations of X . Under the alternative hypothesis $X \sim N(\mu, \sigma_0^2)$, where μ is a random variable following some known (a priori) distribution. In the interest of simplicity, we would like to assume that μ is normally distributed with mean 0 and variance $\nu > 0$. However, since the alternative hypothesis is concerned with the case where $\mu > 0$, assume that, under H_A , μ follows a half-normal distribution with density function

$$f_{H_A}(\mu) = \sqrt{\frac{2}{\pi\nu}} \exp\left(-\frac{\mu^2}{2\nu}\right),$$

for $\mu > 0$.

Table 2. Posterior odds associated with each of the calibrations. using the mixed Bayesian-frequentist hypothesis test.

	Black–Scholes	Exp N ◦ IG	Time changed
DB	0.000	0.000	0.000
DOB	1.045	0.000	4.059
UOB	0.000	6.761	0.000

It remains to specify the value of ν . Again, we opt for a mixed approach by estimating the value of ν from the data as follows. The average of X is calculated, this can be viewed as a realisation of μ . Let $Y = RMSE_M^{(1)} - RMSE_H$ and $Z = RMSE_M^{(2)} - RMSE_H$, where $RMSE_M^{(1)}$ and $RMSE_M^{(2)}$ denote the $RMSE$ s associated with the calibration of the remaining two misspecified models (DB, DOB and UOB options are treated separately). In the same way as before, sample from Y and Z , so that the averages of these samples can be viewed as a second and third realisation of μ . As a result, a total of three realisations from the distribution of μ are obtained. The value of ν is set equal to the variance of this (admittedly small) sample.

The last quantity required for the test that remains unspecified is p_0 (the prior probability of H_0 being true). Since one would have no prior knowledge, and little in the way of intuition, regarding whether or not the null hypothesis holds, the value is arbitrarily set to of $p_0 = 0.5$ for the purposes of this discussion.

If the realised values of X are denoted by $\mathbf{x} = (x_1, x_2, \dots, x_{500})$, then the posterior odds defined in (5) can be calculated using

$$\frac{P(H_0|X = \mathbf{x})}{P(H_A|X = \mathbf{x})} = \frac{p_0 \sqrt{\pi\nu} \exp\left(-\frac{\sum x_j^2}{2\sigma_0^2}\right)}{\sqrt{2}(1-p_0) \int_0^\infty \exp\left(-\frac{\sum (x_j - \mu)^2}{2\sigma_0^2} - \frac{\mu^2}{2\nu}\right) d\mu}.$$

The posterior odds associated with each of the nine calibrations under consideration are reported in Table 2. Reject H_0 for posterior odds smaller than 1.

Table 2 indicates the presence of three perfect calibrations; that of the Black–Scholes and the time changed models to the DOB options, as well as that of the geometric N ◦ IG process model to the UOB options.

Consider, for example, the calibration of the time changed model to the DOB options. This calibration is deemed perfect. In order to gain some understanding of the level of precision involved in this calibration, the market prices are compared to a single set of price estimates obtained from the calibrated model. Figure 1 shows the market prices as circles and the option price estimates calculated under the calibrated model as crosses.

Table 2 classifies one calibration relating to each of the three alternative models as perfect. Therefore, the mixed Bayesian-frequentist hypothesis test provides some evidence that, for the specific market used, there are instances where the models considered are interchangeable. However, several imperfect calibrations serve to indicate that this is not always the case.




Figure 1. Market prices (circles) and prices under the calibrated model (crosses).

5.4 Bootstrap based hypothesis test

Efron and Tibshirani (1993) explain how to test hypotheses of the form specified in (4) using the bootstrap. For ease of notation, let $X = RMSE_M$ and let $Y = RMSE_H$; let x_1, \dots, x_{500} and y_1, \dots, y_{500} represent the samples taken from X and Y respectively. Let \bar{x} and \bar{y} be the means of the two samples and let $s_x^2 = \frac{1}{500} \sum_{j=1}^{500} (x_j - \bar{x})^2$ and $s_y^2 = \frac{1}{500} \sum_{j=1}^{500} (y_j - \bar{y})^2$. μ_M and μ_H respectively denote the population means of X and Y . The test statistic used to test the hypotheses in (4) is

$$T = \frac{\bar{x} - \bar{y}}{\sqrt{\frac{s_x^2}{500} + \frac{s_y^2}{500}}}.$$

In order to test the hypotheses in question, calculate the realised value of T and use the bootstrap to estimate the distribution of T under H_0 as follows. Under the null hypothesis $\mu_X = \mu_Y$ and this equality is mimicked when using bootstrap hypothesis testing procedures by applying the following transformation to the sample data: $r_j = x_j - \bar{x}$ and $s_j = y_j - \bar{y}$ for $j = 1, \dots, 500$. A simple random sample of size 500 is then drawn with replacement from r_1, \dots, r_{500} ; denote the elements of this sample by x_1^*, \dots, x_{500}^* . We also draw a simple random sample of size 500 with replacement from s_1, \dots, s_{500} ; denote the elements of this sample by y_1^*, \dots, y_{500}^* . Using notation similar to those defined above, calculate

$$T^* = \frac{\bar{x}^* - \bar{y}^*}{\sqrt{\frac{s_{x^*}^2}{500} + \frac{s_{y^*}^2}{500}}}.$$

T^* is a realisation of the test statistic under the null hypothesis. This process is repeated 100 000 times to obtain as many realisations of T^* . The p-value associated with the test is equal to the proportion of the realised values of T^* that are larger than or equal to the value of T realised from the original sample. Table 3 reports the p-value associated with each of the nine calibrations.

Using a 1% significance level, the same three calibrations that were deemed perfect by the mixed Bayesian-frequentist test are deemed perfect using the Bootstrap methodology. The results

Table 3. Bootstrap p-values associated with each of the calibrations.

	Black–Scholes	Exp N \circ IG	Time changed
DB	0.000	0.000	0.000
DOB	0.118	0.000	0.569
UOB	0.000	0.036	0.000

presented in Tables 2 and 3 coincide, which reinforces the conclusion that the models considered are interchangeable to a certain extent in the market under consideration. As was the case before, the majority of the calibrations are not deemed perfect, which indicates that the models are not always interchangeable.

6. Conclusions

This paper is concerned with the calibration of various option pricing models to barrier option prices. The price of a barrier option depends not only on the terminal stock price, but also on whether or not the stock price reaches some predetermined level during the lifetime of the option. In practice, the stock prices at the end of each business day are often used in order to determine whether or not the barrier has been reached; the same convention is used in the paper. In the discussion, three types of barrier options are considered: digital barriers, down-and-out barriers and up-and-out barriers.

Since the market prices of barrier options are not readily available, a hypothetical financial market is constructed which contains a single stock, the price of which follows a Heston model with a known parameter set. The market also contains all three types of barrier option types discussed, the market prices of which are calculated using Monte Carlo simulation.

In the paper, the degree to which it is possible to replicate the market prices of the options (calculated under the Heston model) using three other models is ascertained. The models include the Black–Scholes model, the geometric normal inverse Gaussian model and the geometric lognormal-normal model. No closed-form formulae exist for the barrier options considered under these models and, as a result, the prices of the options are estimated using a Monte Carlo procedure.

The four option pricing models considered are calibrated to the market prices of the options. The option types are considered separately, meaning that each of the four models is calibrated to the three different sets of market option prices separately. The calibration procedure minimises the root mean square error (*RMSE*) between the market option prices and the prices calculated using the model considered using a particle swarm optimisation algorithm. In order to determine whether or not the models are interchangeable, the *RMSE* associated with the Heston model calibration and the *RMSE* obtained using the remaining models are compared. Since the calculated option prices contain Monte Carlo errors, one cannot simply require that the option prices calculated using the calibrated model exactly match the market option prices in order for the models to be interchangeable. Instead, the concept of a *perfect* calibration is introduced to overcome this obstacle. The *RMSE* associated with a given calibration contains two components; the first is due to the misspecification of the model and the second is due to Monte Carlo error. A calibration is deemed perfect if the population average of the *RMSE* associated with the calibration does not exceed the population average of the *RMSE* associated with the calibrated Heston model (which is correctly specified). As is to be expected,

the *RMSEs* obtained using the method described above are small (indicating that the option prices obtained correspond closely). However, it is shown that, in several cases, the contribution of the model misspecification to the *RMSE* is negligible when compared to that of the Monte Carlo error.

Two formal hypothesis tests are employed in order to test the hypothesis that a given calibration is perfect; a mixed Bayesian-frequentist test and a Bootstrap based test. Both of the hypothesis testing procedures used classify several of the calibrations considered as perfect. This demonstrates numerically that, in the theoretical Heston market considered, it is possible to replicate, with one model, the barrier option prices calculated using another model, with a high degree of accuracy. It is concluded that, when the underlying market follows a Heston model, there are instances where the models considered are interchangeable. However, the findings presented also suggest that there are instances where these models are not interchangeable.

Acknowledgements. This research was done as part of the author's doctoral studies under the supervision of Prof. F. Lombard. The author would like to sincerely thank Prof. Lombard for his guidance. The financial assistance of the DST-NRF Centre of Excellence in Mathematical and Statistical Sciences (CoE-MaSS) is hereby acknowledged. Opinions expressed and conclusions arrived at are those of the author and are not necessarily to be attributed to the CoE-MaSS.

A. Algorithms for the simulation of price paths

Let n represent the number of days for which we wish to simulate the stock price and let N be the number of simulated price paths that are required. The algorithms used to simulate price paths for the various models discussed in Section 2 are provided below.

A.1 The Heston model

1. Generate σ^2 from a $\text{gamma}(2\kappa\eta/\theta^2, \theta^2/2\kappa)$ distribution.
2. Generate W_1 from a $N(0, 1)$ distribution.
3. Calculate $S(1, k) = 1 + \mu + \sigma W_1$.
4. Generate W_2 and W_3 independently from a $N(0, 1)$ distribution.
5. Calculate $W_4 = \rho W_2 + \sqrt{1 - \rho^2} W_3$.
6. Calculate $S(j, k) = (1 + r + \sigma W_2)S(j - 1, k)$.
7. Calculate $\sigma^2 = \max(\sigma^2 + \kappa(\eta - \sigma^2) + \sigma\theta W_4, 0)$.
8. Repeat steps 4 to 7 for $j = 2$ to n .
9. Repeat steps 1 to 8 for $k = 1$ to N .

A.2 The Black-Scholes model

1. Generate an n by N matrix Δ where each element of Δ follows a $N(\mu, \sigma^2)$ distribution.
2. Calculate $\Delta^*(j, k) = \sum_{l=1}^j \Delta(l, k)$.

Table 4. Digital barrier call option prices.

Barrier	$T = 21$	$T = 46$	$T = 68$	$T = 111$
1.050	0.2730	0.4970	0.5940	0.6894
1.075	0.1020	0.3130	0.4312	0.5637
1.100	0.0282	0.1764	0.2933	0.4453
1.125		0.0882	0.1861	0.3394
1.150		0.0390	0.1101	0.2495
1.175			0.0603	0.1771
1.200			0.0310	0.1210
1.225				0.0799
1.250				0.0508
1.275				0.0312

Table 5. Down-and-out barrier call option prices.

Barrier	$T = 21$	$T = 46$	$T = 68$	$T = 111$
0.99	0.0149	0.0188	0.0212	0.0248
0.98	0.0193	0.0262	0.0302	0.0363
0.97	0.0213	0.0309	0.0366	0.0450
0.96	0.0222	0.0337	0.0409	0.0516
0.95		0.0353	0.0437	0.0564
0.94			0.0455	0.0599
0.93			0.0465	0.0622
0.92				0.0639
0.91				0.0649

3. Calculate $S(j, k) = \exp(\Delta^*(j, k))$.
4. Repeat steps 2 and 3 for $j = 1$ to n .
5. Repeat steps 2 to 4 for $k = 1$ to N .

A.3 The geometric $N \circ IG$ model

1. Generate an n by N matrix Δ where each element of Δ follows a $N \circ IG(\alpha, \beta, \mu, \delta)$ distribution.
2. Calculate $\Delta^*(j, k) = \sum_{l=1}^j \Delta(l, k)$.
3. Calculate $S(j, k) = \exp(\Delta^*(j, k))$.
4. Repeat steps 2 and 3 for $j = 1$ to n .
5. Repeat steps 2 to 4 for $k = 1$ to N .

A.4 The time changed geometric Brownian motion model

1. Generate an n by N matrix T where each element of T follows a $\log N(\alpha, \beta)$ distribution.
2. Generate an n by N matrix Δ where each element of Δ follows a $N(0, 1)$ distribution.
3. Calculate $\Delta_1(j, k) = \mu T(j, k) + \sigma \sqrt{T(j, k)} \Delta(j, k)$.
4. Calculate $\Delta^*(j, k) = \sum_{l=1}^j \Delta_1(l, k)$.
5. Calculate $S(j, k) = \exp(\Delta^*(j, k))$.
6. Repeat steps 3 to 5 for $j = 1$ to n .
7. Repeat steps 3 to 6 for $k = 1$ to N .

B. Options available in the hypothetical market

Tables 4, 5 and 6 respectively provide the details of the DB, DOB and UOB call option prices calculated in the hypothetical Heston market. Each table shows market prices of the options considered. The tables also provide the barrier levels and times to maturity associated with the options. All of the DOB and UOB options have a strike price of 1.

Table 6. Up-and-out barrier call option prices.

Barrier	$T = 21$	$T = 46$	$T = 68$	$T = 111$
1.100	0.0196	0.0172	0.0131	
1.125	0.0219	0.0250	0.0215	0.0149
1.150	0.0225	0.0307	0.0296	0.0230
1.175	0.0227	0.0342	0.0364	0.0316
1.200		0.0359	0.0411	0.0398
1.225		0.0366	0.0442	0.0470
1.25			0.0460	0.0529
1.275				0.0574
1.3				0.0606

Table 7. Ranges for the starting values of the parameters.

Black–Scholes	$\sigma = (0, 0.02)$
Heston	$\kappa = (0, 0.3)$; $\eta = (0, 0.0005)$; $\rho = (-0.95, 0.95)$; $\theta = (0, 0.005)$
Exp N \circ IG	$\alpha = (5, 100)$; $\beta = (-50, 50)$; $\delta = (0, 0.1)$
Time changed	$\alpha = (-0.2, 0.2)$; $\beta = (0, 0.2)$; $\sigma = (0, 0.2)$

C. Starting values for the optimisation procedure

Table 7 contains the ranges of the possible starting values for each of the parameters under the various models considered.

References

- BLACK, F. AND SCHOLES, M. (1973). The pricing of options and corporate liabilities. *The Journal of Political Economy*, **3**, 637–654.
- CLARK, P. (1973). A subordinated stochastic process with finite variance for speculative prices. *Econometrica*, **41**, 135–156.
- CONT, R. AND TANKOV, P. (2004). *Financial Modelling With Jump Processes*. Chapman & Hall, London.
- COX, D. R. AND HINKLEY, D. V. (1976). *Theoretical Statistics*. Chapman & Hall, London.
- EFRON, B. AND TIBSHIRANI, R. J. (1993). *An Introduction to the Bootstrap*. Chapman & Hall, New York.
- 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.
- KENNEDY, J. AND EBERHART, R. C. (1995). Particle swarm optimization. *Proceedings of the IEEE International Conference on Neural Networks*, 1942–1948.
- SCHOUTENS, W. (2003). *Lévy Processes in Finance*. Wiley & Sons, Chichester.
- SCHOUTENS, W., SIMONS, E., AND TISTAERT, J. (2004). A perfect calibration! Now what? *Wilmott magazine*, **8**, 66–78.