Journal of Empirical Finance 9 (2002) 225–255
Journal of EMPIRICAL FINANCE
www.elsevier.com/locate/econbase

Estimation and empirical performance of Heston's stochastic volatility model: the case of a thinly traded market

Gabriele Fiorentini [a], Angel León [b], Gonzalo Rubio [c,*]

[a] Dpto. Fundamentos del Análisis Econo., Facultad de Ciencias Económicas, Universidad de Alicante, Spain
[b] Dpto. Economía Financiera, Facultad de Ciencias Económicas, Universidad de Alicante, Spain
[c] Dpto. Fundamentos del Análisis Econo., Facultad de Ciencias Económicas, Universidad del País Vasco, Avda. L. Aguirre 83, 48015, Spain

Accepted 17 October 2001

Abstract

This paper examines the stochastic volatility model suggested by Heston [Rev. Financ. Stud. 6 (1993) 327] in a thinly traded market context. We employ a time-series approach based on indirect inference to estimate the model parameters. We also discuss the potential effects of time-varying skewness and kurtosis on the performance of the model. It is found that the model tends to overprice out-of-the-money calls and underprice in-the-money calls. Moreover, the daily volatility risk premium shows a volatile behavior over time; however, our evidence suggests that the volatility risk premium has a negligible impact on the pricing performance of Heston's model. © 2002 Elsevier Science B.V. All rights reserved.

JEL classification: G12; G13
Keywords: Stochastic; Volatility; Skewness; Kurtosis; Pricing

1. Introduction

After the crash in 1987, the implied volatility of options has often been found to be decreasing and convex in the exercise price. There have been various attempts to deal with this apparent failure of the Black and Scholes (1973) (BS henceforth) formula. The deterministic volatility framework proposed by Rubin-

* Corresponding author. Tel.: +34-601-7470; fax: +34-601-7474.
E-mail address: jepruirg@bs.ehu.es (G. Rubio).

PII: S0927-5398(01)00052-4

stein (1994), Dupire (1994), Derman and Kani (1994), and the stochastic volatility models of Hull and White (1987), Scott (1987), Wiggins (1987) and Heston (1993) are two key extensions of the BS model.[1] Along these two main strands of option-pricing literature, using daily data on the S&P 500 index options, Buraschi and Jackwerth (2001) show that the underlying asset and at-the-money options do not span the whole set of prices available in the economy. They find that the returns of in- and out-of-the-money options are necessary for spanning. This implies that risk factors such as stochastic volatility, jumps or interest rates, which are not introduced by the deterministic volatility framework, are needed to correctly price options.[2] Stochastic models are therefore the most relevant approach in which option-pricing models should be tested.[3] It should be also noted that Heston (1993) provides a closed-form solution for a European call option without imposing the restrictions of zero correlation and zero price of volatility risk by using Fourier inversion methods.

That discussion motivated this research. The key objective of the paper is to evaluate the empirical performance of the stochastic volatility model proposed by Heston (1993) relative to the BS framework for options on the Spanish IBEX-35 stock exchange index. Time-series returns on the underlying index are employed to estimate the parameters of the stochastic variance process imposed by Heston's model. This is done using the indirect inference procedures of Gouriéroux et al. (1993). It is an appropriate method for adjusting for the discretization biases when estimating the parameters of the continuous variance process under the true probability measure. Moreover, this approach allows us to analyze the empirical behavior of the parameters characterizing the diffusion process assumed for the instantaneous variance. Their behavior is discussed relative to the appropriate skewness and kurtosis underlying in Heston's model. This provides us with insights towards understanding the characteristics of the stochastic volatility option-pricing model. It must be pointed out that the usual cross-sectional approach employed by Bakshi et al. (1997) to test stochastic option models may easily ignore relevant information in the original series that may not be embedded in the option prices. Our approach enables us to combine the cross-sectional information in options with the information contained in the time series returns of the underlying security. In this sense, the first contribution of this paper is the use of indirect inference to estimate Heston's model.

[1] Bates (1996) and Bakshi et al. (1997) further develop the stochastic volatility model suggested by Heston (1993) in order to account for both stochastic jumps and interest rates.

[2] See also the evidence reported by Dumas et al. (1998).

[3] The empirical evidence regarding option pricing within the stochastic models is mixed. See Bates (1996), Bakshi et al. (1997), Pan (1999), Das and Sundaram (1999), and Chernov and Ghysels (2000). In related research, Heston and Nandi (2000) propose a closed-form GARCH model in which it is possible to value options using volatilities calculated directly from the historical sample of security returns, simplifying the estimation procedure. Their empirical results are certainly promising.

The second contribution is the empirical examination of Spanish options data. This is important since it provides evidence from a thinly traded market. Empirical evidence on option pricing related to markets of this type is basically unavailable. It should be taken into account that markets frictions are substantially higher than in more actively traded markets, and that the behavior of agents participating in these markets is also likely to be different from those in large and well established option markets.

The Spanish option market shows a very limited number of exercise prices traded simultaneously relative to the US market. The lack of a sufficiently large number of contemporaneous options has serious consequences for the empirical implementation of the cross-sectional estimation of implied parameters of stochastic models. In other words, there are just not enough prices available to estimate jointly all the parameters embedded in Heston’s model. Thus, our estimation procedure has two separate steps. First, we estimate the parameters of the stochastic volatility, which are required inputs of Heston’s model, by employing the indirect inference estimation technique on a time-series of the underlying hourly return. Second, the price of the volatility risk and instantaneous variance are backed out by minimizing the sum of squared pricing errors between the option model and market prices as in Bakshi et al. (1997). Then theoretical prices are obtained for alternative specifications of Heston’s model and compared with BS option prices. Despite the fact that there is a slight overall improvement over the performance of BS, it is found that Heston’s model tends to overprice out-of-the-money calls and underprice in-the-money calls. It is also found that the daily volatility risk premium does not remain constant over time; however, our evidence suggests that the volatility risk premium has a negligible impact on the pricing performance of Heston’s model. The overall empirical evidence regarding the behavior of the volatility risk premium is the third contribution of our paper.

The remainder of the paper is organized as follows: Section 2 contains a brief summary of the Spanish options market and options data. The theoretical model employed in this paper, i.e. Heston’s model, is outlined in Section 3. The empirical results regarding the time series estimation of the stochastic volatility parameters are discussed in Section 4. The implied volatility graphs and the estimation of the volatility risk premium through option prices are shown in Section 5. In Section 6, the out-of-sample pricing performance is carried out. Finally, in Section 7, we conclude with a summary.

2. The Spanish IBEX-35 index options

2.1. Market description

The Spanish IBEX-35 index is a value-weighted index comprising the 35 most liquid Spanish stocks traded in the continuous auction market system. The official derivative market for risky assets, which is known as MEFF, trades a future contract on the IBEX-35, the equivalent option contract for calls and puts, and individual option contracts for blue-chip stocks. Trading in the derivative market started in 1992. The market has experienced tremendous growth from the very beginning. Relative to the volume traded in the Spanish continuous market, trading in MEFF represented 40% of the regular continuous market in 1992, 138% in 1996, and 249% in 1999. The number of all traded contracts in MEFF relative to the contracts traded in the Chicago Board Options Exchange (CBOE) reached 20% in 1995, and 23% in 1999.

The IBEX-35 option contract is a cash-settled European option with trading during the three nearest consecutive months and the other 3 months of the March–June–September–December cycle. The expiration day is the third Friday of the contract month. Trading occurs from 10:00 to 17:15. During the sample period covered by this research, the contract size is 100 Spanish pesetas times the IBEX-35 index, and prices are quoted in full points, with a minimum price change of one index point or 100 pesetas.[4] The exercise prices are given in 50 index point intervals.

It is important to point out that liquidity is concentrated in the nearest expiration contract. Thus, during 1995 and 1996 almost 90% of crossing transactions occurred in contracts of this type. Finally, it should be noticed that options and future contracts are directly associated. The future contract has exactly the same contract specifications as the IBEX-35 options. This will allow us to employ the future price rather than the spot price in our empirical exercise. In fact, this is what practitioners usually do.

2.2. The data

For this paper, our database is comprised of all call options on the IBEX-35 index traded daily on MEFF during the period from January 3, 1996 through April 30, 1996. Given the concentration in liquidity, our daily set of observations includes only calls with the nearest expiration day. Moreover, we eliminate all transactions taking place during the last week before expiration (to avoid the expiration-related price effects).

As usual in this type of research, our primary concern is the use of simultaneous prices for the options and the underlying security. The data, which are based on all reported transactions during each day throughout the sample period, do not allow us to observe simultaneously enough options with the same time-to-expiration on exactly the same underlying security price but with different exercise prices. In order to avoid large variations in the underlying security price, we restrict our attention to the 45-min window from 16:00 to 16:45. It turns out that on average during our sample period almost 25% of crossing transactions occur during this interval. Care was also taken to eliminate the potential problems with

[4] This has recently been changed to 10 euros, with trading taking place between 9:00 and 17:35.

Table 1
Sample characteristics of IBEX-35 future options
Average prices, average relative bid-ask spread and the number of available calls are reported for each moneyness category. All call options transacted over the 45-min interval from 16:00 to 16:45 from January 2, 1996 to April 30, 1996 are employed. K is the exercise price and F denotes the future price of the IBEX-35 index. Moneyness is defined as the ratio of the exercise price to the future price. OTM, ATM, and ITM are out-of-the-money, at-the-money, and in-the-money options, respectively.

	Moneyness (K/F)	Average price	Average bid-ask spread	Number of observations
Deep OTM calls	1.03–1.08	12.88	0.3903	116
OTM calls	1.01–1.03	28.68	0.2205	273
ATM calls	0.99–1.01	57.20	0.1491	245
ITM calls	0.97–0.99	98.54	0.1273	108
Deep ITM calls	0.90–0.97	185.42	0.0987	26
All calls	–	50.52	0.2015	768

artificial trading that are most likely to occur at the end of the day. Thus, all trades after 16:45 were eliminated so that we avoid data which may reflect trades to influence market maker margin requirements. At the same time, using data from the same period each day avoids the possibility of intraday effects.

These criteria yield a final daily sample of 768 observations. Table 1 describes the sample properties of the call option prices employed in this work. Average prices, average relative bid-ask spread and the number of available calls are reported for each moneyness category. Moneyness is defined as the ratio of the exercise price to the future price. A call option is said to be deep out-of-the-money if the ratio K/F belongs to the interval $(1.03, 1.08)$; out-of-the-money if $1.03 > K/F \geq 1.01$; at-the-money when $1.01 > K/F \geq 0.99$; in-the-money when $0.99 > K/F \geq 0.97$; and deep-in-the-money if $0.97 > K/F \geq 0.90$. As indicated, there are 768 call option observations, with OTM, ATM and ITM options, respectively, accounting for 51%, 32% and 17%. The average call price ranges from 12.88 pesetas for deep OTM options to 185.42 pesetas for deep ITM options. The average relative bid-ask spread goes in exactly the opposite direction to the average price: it ranges from 0.39 for deep OTM options to 0.10 for deep ITM calls.

The implied volatility for each of our 768 options is estimated next. Note that we take as the underlying asset the average of the bid and ask price quotation given for each future contract associated with each option during the 45-min interval.[5] Recall that we are allowed to use future prices given that the expiration

[5] Lack of liquidity in the futures market might be responsible for a lack of variation in the price of the underlying asset during the 45-min window. However, this is not the case. In fact, the futures market is at least as liquid as the spot market in terms of comparable measures of trading volume.

day of the future and option contracts systematically coincides during the expiration date cycle. Moreover, note that dividends are already taken into account by the future price. These characteristics of the Spanish market facilitate our research. To proxy for riskless interest rates, we use the daily series of annualized repo T-bill rates with 1, 2 or 3 weeks to maturity. One of these three interest rates will be employed depending upon how close the option is to the expiration day. Finally, as discussed by French (1984), volatility appears to be a phenomenon that is basically related to trading days. However, interest rates are paid by the calendar day. Thus, in order to estimate the implied volatility of each option in our sample, we employ Black's (1976) option-pricing formula adjusted by two time measures to reflect both trading days and calendar days until expiration. These implied volatilities will be used later as the basis for comparison with Heston's implied instantaneous volatilities.

3. Heston's stochastic volatility option-pricing model

Heston (1993) obtains a closed-form solution for the price of a European call option on an asset with stochastic volatility. Heston works with Fourier transforms of conditional probabilities that the option expires in-the-money. The characterization of these probabilities is achieved through their characteristic function.

The stochastic volatility model proposed by Heston generalizes Geometric Brownian Motion by allowing the volatility of the return process itself to evolve stochastically over time in a square root mean-reverting fashion:

$$\begin{aligned} \mathrm{d}S_t &= \mu S_t \mathrm{d}t + \sqrt{V_t}\, S_t \mathrm{d}W_{1t} \\ \mathrm{d}V_t &= \kappa(\theta - V_t)\mathrm{d}t + \sigma\sqrt{V_t}\,\mathrm{d}W_{2t} \\ \mathrm{d}W_{1t}\mathrm{d}W_{2t} &= \rho \mathrm{d}t \end{aligned} \tag{1}$$

where μ is the instantaneous expected rate of return of the underlying asset, V_t is the instantaneous stochastic variance, θ is the long-term mean of the variance, κ governs the rate at which the variance converges to this mean, and σ represents the volatility of the variance process. The parameters of the variance process, θ, κ, and σ are all strictly positive constants. W_{1t} and W_{2t} are standard Brownian motions allowed to be instantaneously correlated. Thus, increases (decreases) in volatility could be related to the level of the underlying asset.

The risk-neutral probability measure incorporates the market price of volatility, denoted as λ, to distinguish the objective probability measure from the risk-neutral one. The volatility risk premium is assumed to be proportional to the instantaneous variance, λV_t, and its sign arises from the (sign of) correlation between the

Brownian processes assumed for the instantaneous variance and the (aggregate) consumption. The model is given by:

$$\begin{aligned}
\mathrm{d}S_t &= rS_t\mathrm{d}t + \sqrt{V_t}\,S_t\mathrm{d}W_{1t}^{*}\\
\mathrm{d}V_t &= \kappa^{*}(\theta^{*} - V_t)\mathrm{d}t + \sigma\sqrt{V_t}\,\mathrm{d}W_{2t}^{*}\\
\mathrm{d}W_{1t}^{*}\mathrm{d}W_{2t}^{*} &= \rho\mathrm{d}t
\end{aligned} \tag{2}$$

where $\kappa^{*} = \kappa + \lambda$, $\theta^{*} = \kappa\theta/(\kappa + \lambda)$.

Let $c(S,\upsilon,t)$ be the value of a European call option where $S \equiv S_t$ and $\upsilon \equiv V_t$ to abbreviate. Heston's formula is given by:

$$c(S,\upsilon,t) = S_t P_1 - K\mathrm{e}^{-r(T-t)}P_2 \tag{3}$$

where P_1 and P_2 are two risk-neutralized probabilities having the same interpretation as in the standard BS expression.

In the application below, we use future options so that the actual version of the formula we employ is given by:

$$c(F,\upsilon,t) = \mathrm{e}^{-r(T-t)}(F_t P_1 - KP_2) \tag{4}$$

where F is the future price on the underlying spot price, and

$$P_j(x,\upsilon,T-t;\ln[K]) = \mathrm{Prob}(x_T \geq \ln[K] \mid x_t = x, \upsilon_t = \upsilon) \tag{5}$$

where $x \equiv \ln[F_t]$, $j = 1$ or 2 (the probability of the event $\{x \geq \ln[K]\}$ depends on whether we chose the future for $j = 1$, or the riskless rate for $j = 2$), and P_j is, therefore, the conditional probability that the option expires in-the-money. These probabilities depend on the vector of parameters $(\kappa,\theta,\lambda,\sigma,\rho)$ given by the process assumed by Heston under the original probability. It is important to note that the expressions for these probabilities are slightly different from the original values given by Heston (1993) since we are using futures. The actual formulae employed in this paper are provided in Appendix A.

4. Estimation of the diffusion parameters

4.1. Indirect inference procedure

Consider the process under the original probability measure given by Eq. (1) and let $x_t \equiv \ln S_t$, then Eq. (1) becomes

$$\mathrm{d}x_t = (\mu - V_t/2)\mathrm{d}t + V_t^{1/2}\mathrm{d}W_{1t} \tag{6}$$

$$\mathrm{d}V_t = \kappa(\theta - V_t)\mathrm{d}t + \sigma V_t^{1/2}\mathrm{d}W_{2t} \tag{7}$$

$$\mathrm{d}W_{1t}\mathrm{d}W_{2t} = \rho\mathrm{d}t$$

where $\boldsymbol{\Omega} \equiv (\mu,\kappa,\theta,\sigma,\rho)$ denotes the vector of parameters of interest.

Since the process in Eqs. (6) and (7) is in continuous time and no exact discretization scheme exists, estimators of $\boldsymbol{\Omega}$ obtained by applying standard econometric methods such as pseudo maximum likelihood (ML), or generalized method of moments to a discrete-time approximation of the process above would be inconsistent. This effect has long been recognized and is known as the discretization bias. In this context, the indirect inference procedures of Gouriéroux et al. (1993), Smith (1993) and Gallant and Tauchen (1996) provide convenient estimation methods, which have made a substantial impact on the practice of econometrics over recent years. Specifically, Gouriéroux et al. (1993) show that, under suitable regularity conditions, consistent estimation of the parameters of interest $\boldsymbol{\Omega}$ can be based on an incorrect criterion (the auxiliary model) which, itself, depends on a vector of parameters $\boldsymbol{\Psi}$, where dim $(\boldsymbol{\Psi}) \geq$ dim $(\boldsymbol{\Omega})$. The implementation requirement of indirect inference methods is that the parameters of interests are linked to the auxiliary model parameters through the so-called binding function $\boldsymbol{\Psi} = b(\boldsymbol{\Omega})$. In practice, it is also required that the process in Eqs. (6) and (7) can be easily simulated for any given value of $\boldsymbol{\Omega}$ in the parameter space.

Gouriéroux et al. (1993) define the indirect inference estimator of $\boldsymbol{\Omega}$ as that value of $\boldsymbol{\Omega}$ which minimizes the distance between the binding function vector $b(\boldsymbol{\Omega})$ and its sample counterpart $\hat{\boldsymbol{\Psi}}_T$. Here, $\hat{\boldsymbol{\Psi}}_T$ is the pseudo ML estimate of the auxiliary model parameters based on an observed sample of length T. Notice that, analogously to the classical minimum distance estimation case, the distance is minimized under the metric defined by a weighting matrix $\boldsymbol{\Xi}$. When dim $(\boldsymbol{\Psi}) =$ dim $(\boldsymbol{\Omega})$, the model of interest is just-identified and the resulting indirect inference estimator is unaffected by the choice of the weighting matrix. In our case, as in most cases of practical relevance, no closed form expression for the binding function is available. It must therefore be computed by resorting to numerical simulation methods. For an exhaustive review of indirect inference, see Gouriéroux and Monfort (1996). Applications to the estimation of stochastic differential equations can be found, for instance, in Broze et al. (1995) and Calzolari et al. (1998).

At this stage, the most important decisions that we have to make are our choice of an auxiliary model and our definition of a simulation procedure in order to draw samples from processes (6) and (7). As mentioned, for any given value of $\boldsymbol{\Omega}$, we must be able to simulate samples from processes (6) and (7). In order to do so, we consider the Euler discretization[6] of both Eqs. (6) and (7), which is given by:

$$x_t = x_{t-\tau} + \mu\tau - \frac{V_{t-\tau}}{2}\tau + V_{t-\tau}^{1/2}\tau^{1/2}\eta_{1t} \tag{8}$$

$$V_t = \kappa\theta\tau + (1-\kappa\tau)V_{t-\tau} + \sigma V_{t-\tau}^{1/2}\tau^{1/2}\zeta_t \tag{9}$$

[6] Bakshi et al. (2000) also use Euler discretization for the method of simulated moments which is employed to estimate the structural parameters of the continuous stochastic volatility (SV) process of the underlying asset. They also use a square root process for the SV model.

where τ is the simulation frequency (typically, $\tau = 1/10$), and

$$\zeta_t = \left[\rho\eta_{1t} + \left(1-\rho^2\right)^{1/2}\eta_{2t} \right]$$

$$(\eta_{1t},\eta_{2t})' \overset{\text{iid}}{\approx} N(0,I)$$

The x_t is simulated as follows. First, we generate a sample of T/τ pseudo random errors $(\eta_{1t},\ \eta_{2t}) \sim$ i.i.d $N(0,I)$. Then, for a given value of $\boldsymbol{\Omega}$, we compute recursively x_t according to Eqs. (8) and (9). Finally, the simulated sample is obtained by retaining one every $1/\tau$ values of x_t and thus, in order to achieve a simulated sample of length T we must actually simulate T/τ values of x_t. The approximation error due to the discretization of Eqs. (6) and (7) can be made as small as desired by choosing a small enough sampling interval τ, (see for example Broze et al. (1998)).

In our application, for the auxiliary model we adopt a NAGARCH (1,1) specification:[7]

$$R_t = \mu + \xi_t;\ \xi_t = h_t^{1/2}\varepsilon_t;\ \varepsilon_t \underset{7}{\overset{\text{iid}}{\approx}} N(0,1)$$

$$h_t = \omega + \beta h_{t-1} + \alpha\left(\xi_{t-1} + \gamma h_{t-1}^{1/2}\right)^2 \tag{10}$$

where $R_t \equiv x_t - x_{t-1}$ and γ is the parameter that controls the correlation between the innovations to R_t and the conditional variance. The set of parameters to be estimated by pseudo maximum likelihood is denoted by $\boldsymbol{\Psi} \equiv (\mu,\omega,\alpha,\beta,\gamma)$. Even though under processes (6) and (7), the NAGARCH (1,1) model is obviously misspecified, the choice is justified by the fact that this model is expected to capture the main characteristics of the data such as fat-tailness, time-varying conditional variance, and asymmetries in the response of volatility to the arrival of information. It must be noticed that there is no general theory for auxiliary model selection and the adequacy of a model must be judged on a case by case basis.

The indirect inference estimation of $\boldsymbol{\Omega}$ is then performed as follows. Let $\hat{\boldsymbol{\Psi}}_T \equiv (\hat{\mu},\hat{\omega},\hat{\alpha},\hat{\beta},\hat{\gamma})$ denote the pseudo maximum likelihood estimates of the NAGARCH (1,1) model computed with the observed data. Next, by using the procedure based on the Euler discretization described above, we can simulate, for a given value of $\boldsymbol{\Omega}$, N samples of size T from the continuous time processes (6) and (7). Let $\tilde{\boldsymbol{\Psi}}_i$ be the pseudo-ML estimate of the NAGARCH (1,1) model with the ith simulated sample. The resulting $\tilde{\boldsymbol{\Psi}}_i$ depends on $\boldsymbol{\Omega}$ through the simulated values, thus we can write $\tilde{\boldsymbol{\Psi}}_i(\boldsymbol{\Omega})$. The value of $\boldsymbol{\Omega}$ is calibrated in order to have

[7] See León and Mora (1999) for the behavior of alternative specifications within the GARCH family in the Spanish stock market.

$\left(1/N\sum_{i=1}^{N}\tilde{\Psi}_i(\boldsymbol{\Omega})\right)$ as close as possible to $\hat{\Psi}_T$, the indirect inference estimate of $\boldsymbol{\Omega}$ is then:

$$\hat{\Omega}=\arg\min_{\boldsymbol{\Omega}}\left(\hat{\Psi}_T-1/N\sum_{i=1}^{N}\tilde{\Psi}_i(\boldsymbol{\Omega})\right)'\left(\hat{\Psi}_T-1/N\sum_{i=1}^{N}\Psi\tilde{A}_i(\boldsymbol{\Omega})\right) \tag{11}$$

Expresion (11) can be solved through a numerical minimization method. We implement a Newton–Raphson algorithm based on numerically computed first and second derivatives of the objective function with respect to the parameters.

As shown in Gourieroux et al. (1993), since we are assuming that $\mathrm{d}x_t$ is strictly stationary and ergodic, the simulation of N samples of size T can be replaced by a single realization of the process of length NT. In this case, the indirect inference estimator is given by:

$$\hat{\Omega}=\arg\min_{\boldsymbol{\Omega}}\left(\hat{\Psi}_T-\tilde{\Psi}(\boldsymbol{\Omega})\right)'\left(\hat{\Psi}_T-\tilde{\Psi}(\boldsymbol{\Omega})\right) \tag{12}$$

where $\tilde{\Psi}_i$ is the pseudo-ML estimates of the NAGARCH(1,1) model with the simulated sample of size NT. Since in our case the model parameters are just-identified, the estimation is performed using the identity as the weighting matrix.

4.2. Estimation results

With respect to the empirical results, we first estimate the in-sample period from January 2, 1994 to January 2, 1996 with continuously compounded hourly returns from the Spanish IBEX-35 index. The calculation of returns is based on the last recorded logarithmic index levels over consecutive hourly intervals. It is well known that the first hourly return incorporates adjustments to the information which has arrived overnight, and therefore presents a higher average return variability than any other hourly return. This basically implies that this first return is not an hourly return, and we consequently delete it from the estimation. Our final sample length consists of 2450 hourly observations.

The results for the in-sample period are reported in Table 2. We consider four alternative combinations concerning the number and length of simulated samples employed in the evaluation of the binding functions and the frequency τ used in the Euler discretization scheme.[8] The results contained in the first column are obtained with N = 1 and $\tau = 1/10$. In the second column of Table 2, we simulate N = 10 samples of length T and the frequency is again $\tau = 1/10$. In the third column, we maintain N = 10 and $\tau = 1/10$ but this time we compute the binding function with a simulated series of length NT. This is the methodology employed

[8] Nowman (1997) discretization was also carried out but the results were similar to those of Euler discretization.

Table 2
Indirect inference in-sample estimation (January 1994–December 1995)
The parameters of the following processes are estimated using the indirect inference technique:

$$\mathrm{d}S_t = \mu S_t \mathrm{d}t + \sqrt{V_t}\, S_t \mathrm{d}W_{1t}$$

$$\mathrm{d}V_t = \kappa(\theta - V_t)\mathrm{d}t + \sigma\sqrt{V_t}\,\mathrm{d}W_{2t}$$

$$\mathrm{d}W_{1t}\mathrm{d}W_{2t} = \rho \mathrm{d}t$$

where μ is the instantaneous expected rate of return of the underlying asset, V_t is the instantaneous stochastic variance, θ is the long-term mean of the variance, κ governs the rate at which the variance converges to this mean, σ represents the volatility of the variance process, and ρ is the instantaneous correlation. The auxiliary model employed in the estimation is the following Nagarch(1,1) model[a]:

$$R_t = \mu + \xi_t;\ \xi_t = h_t^{1/2}\varepsilon_t;\ \varepsilon_t \overset{\text{iid}}{\approx} N(0,1)$$

$$h_t = \omega + \beta h_{t-1} + \alpha\left(\xi_{t-1} + \gamma h_{t-1}^{1/2}\right)^2$$

where R_t is the hourly rate of return of the IBEX-35 index, and γ is the asymmetric parameter of the Nagarch. The Euler discretization technique is employed in the estimation. Moreover, alternative frequencies and simulations are also used.

	N = 1; $\tau = 1/10$[b]; length = 2450	N = 10; $\tau = 1/10$[c]; length = 2450	N = 10; $\tau = 1/10$[d]; length = 2450 × N	N = 1; $\tau = 1/50$[e]; length = 2450
$\hat{\mu}$[f]	0.619	0.696	0.702	0.868
$\sqrt{\hat{\theta}}$[g]	12.09	12.77	11.92	11.73
$\hat{\kappa}$	0.029	0.034	0.034	0.025
$\hat{\sigma}$[f]	1.631	1.864	1.857	1.542
$\hat{\rho}$	0.044	0.038	0.020	0.087

[a] The estimates of the Nagarch(1,1) parameters are: $\hat{\mu} = 0.00412$; $\hat{\omega} = 0.00301$; $\hat{\alpha} = 0.07680$; $\hat{\beta} = 0.89025$; $\hat{\gamma} = -0.1225$.

[b] The process is simulated once, so that we employ 2450 data points, where 2450 is the number of hourly returns available from January 1994 to December 1995.

[c] The process is simulated N times, so that we generate N series of size 2450 data points.

[d] The process is simulated once, but now we employ 2450 × N data points. Thus, the Nagarch is calibrated with more data: 2450 × N.

[e] The process is simulated once, so that we employ 2450 data points.

[f] This is annualized and given in percentage terms.

[g] This represents the standard deviation of the long-term variance. It is annualized and given in percentage terms.

in the rolling procedure for the out-of-sample estimation which will be used in testing Heston's option-pricing model. Finally, the results reported in the fourth column refer to the case N = 1 and $\tau = 1/50$.

The estimations contained in Table 2 suggest that, regardless of which procedure is employed, the results tend to be quite similar. There seems to be some minor difference in the estimator of the volatility of the variance process, and the estimator of the correlation coefficient between the shocks when we use the frequency at $\tau = 1/50$. In any case, regardless of which procedure is used, the estimate of the correlation coefficient is surprisingly positive and close to zero.

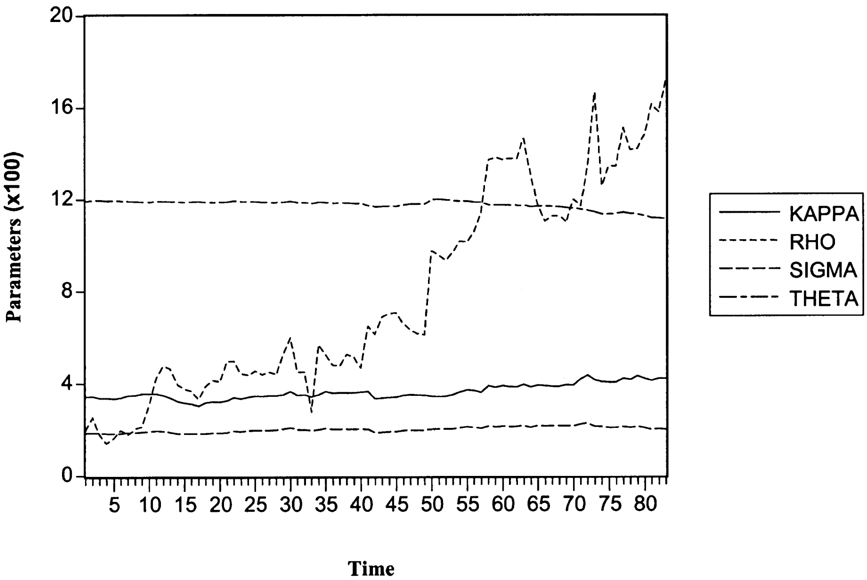

Fig. 1. Indirect inference (January 96–April 96).

For the out-of-sample estimation, the same process is estimated 80 times systematically using 2450 past observations. This rolling procedure of the indirect inference is necessary to yield estimates of the parameters involved in Heston's expression for each day between January 3, 1996 and April 30, 1996. Thus, the daily changing estimates of these parameters are used as inputs in Heston's option-pricing formula. Fig. 1 presents the evolution of the parameters associated with the stochastic variance process throughout the out-of-sample period. As we can observe from the figure, the long-term volatility, $\sqrt{\theta}$, the rate of convergence of the instantaneous variance to the long-term average, κ, and the volatility of the variance process, σ, remain quite stable over the out-of-sample period. However, the coefficient of correlation between the shocks for the stock and the variance, ρ, increases continuously. As before, it is interesting to observe that the correlation remains positive throughout the out-of-sample period.

5. Estimating the implied variance and the volatility risk premium

Volatilities from the BS model and both the instantaneous variance and the volatility risk premium from the Heston model are estimated every day from option prices; specifically from all available call options transacted over the

45-min interval from 16:00 to 16:45 during the period January 3, 1996 to April 30, 1996.

5.1. Estimation procedure

Note that once we have estimated for each day the parameters of the variance process, κ, θ, σ and ρ, using the rolling indirect inference procedure, we still need to estimate the volatility risk premium, λ, and the instantaneous variance, V_t, before we can actually price a given call option under Heston's model. Therefore, it seems reasonable to use cross-sectional data to implicitly infer the parameters that minimize the sum of squared errors (SSE) in a given day of the sample.

Given the set of parameters of the indirect estimation procedure obtained for a particular day t in the sample, $\hat{\Omega}t=(\hat{\kappa}_t,\hat{\theta}_t,\hat{\sigma}_t,\hat{\rho}_t)$, and for each option, i ($i=1,\dots,n$) and each day t, we define the pricing error as:

$$e_{it}\left(V_t,\lambda;\hat{\Omega}_t\right)=\hat{c}_{it}(K_i)-c_{it}(K_i) \tag{13}$$

where $\hat{c}_{it}(K_i)$ is the theoretical price of call i in day t, and $c_{it}(K_i)$ is the corresponding observed market price. We then want to find the instantaneous variance, V_t, and the risk premium parameter, λ, to solve:

$$\mathrm{SSE}_t\equiv\min_{\{V_t,\lambda\}}\sum_{i=1}^{n}\left[e_{it}\left(V_t,\lambda;\hat{\Omega}_t\right)\right]^2 \tag{14}$$

Direct inspection of the quadratic form in Eq. (14) shows that it is highly valley shaped and, therefore, it is extremely flat along a direction corresponding to a nontrivial combination of λ and V_t. As a consequence, derivative-based minimization methods are expected to perform poorly since the numerically computed gradient is very unstable in a neighbourhood of a minimum. This is confirmed by some experiments that we performed with the Newton–Raphson method. In order to avoid the above problem, a derivative-free minimization algorithm is called for. Since the function in Eq. (14) only depends on two parameters, the downhill simplex method of Nelder and Mead (1965) seems to be a natural candidate. We have checked the method by choosing randomly several starting triplets: in cases of convergence to different local minima their minimum was selected as the global solution. So we obtain daily estimates of both λ and V_t from January 3, 1996 to April 30, 1996; i.e. a sample of 80 days. Finally, a series for daily implied volatilities estimated in the corresponding BS versions of Eqs. (13) and (14) is also obtained.

5.2. The volatility risk premium

In Fig. 2, we can see that the estimated daily values of λ further suggest that the volatility risk premium varies over time in the IBEX-35 future option market.

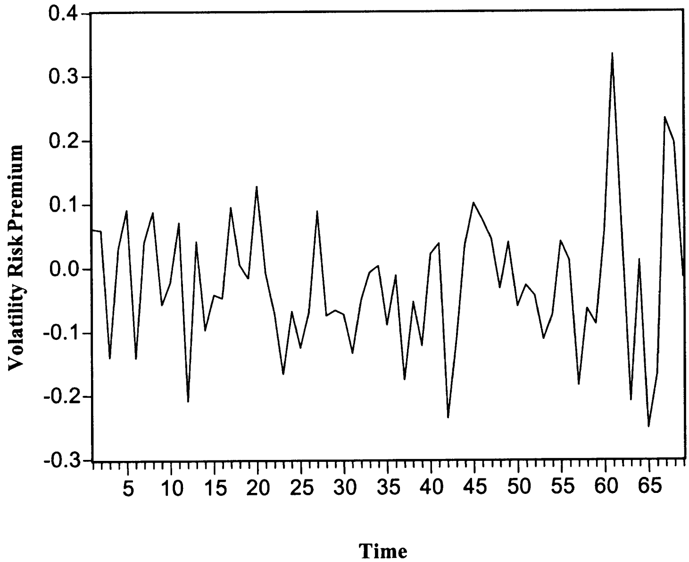

Fig. 2. Daily volatility risk premium (January 96–April 96).

Notice that this time-varying behavior of the volatility risk premium is not consistent with Heston's model since λ must be a constant value, as we can see in Eq. (2). The estimated values of λ ranges from -0.252 to 0.333. The mean (median) for the time-series of λ is -0.021 (-0.027) and the standard deviation is 0.113. The skewness and the excess kurtosis are 0.527 and 0.747, respectively. The p-value of the Jarque–Bera test (normality test) for the time-series of λ is 0.085, so the normal distribution is rejected at the 10% significance level.

The null hypothesis of $\lambda = 0$ was tested under three different methods. First, assuming that λ values were drawn from a normal distribution, the Student's t ratio was -1.58 (p-value $= 0.1183$). Second, a nonparametric sign test[9] for the population median of λ yielded a p-value of 0.1544. Third, we also test $H_0: \lambda = 0$ by obtaining accurate confidence intervals for λ based on the "bootstrap-t" confidence interval (see Efron and Tibshirani, 1993). By selecting 10 000 independent bootstrap samples, each consisting of 80 data values drawn with replacement from the series of λ, then the 90%, 95% and 99% bootstrap-t confidence intervals for the mean are, respectively: $[-0.0426; 0.0021]$, $[-0.0467; 0.007]$ and $[-0.0551; 0.0175]$. Summing up, we conclude that the evidence does not support

[9] Since the empirical distribution is skewed, the sign test is more appropriate than the popular nonparametric Wilcoxon signed-rank test because the latter assumes that the underlying distribution is symmetric (see Rohatgi, 1984).

rejection of the null hypothesis, that the volatility risk premium is zero. The finding of a zero volatility risk premium for IBEX-35 index options is opposite to the results of Guo (1998), Sarwar and Krehbiel (2000) using currency options, and those of Lin et al. (1999) using FTSE 100 index options. Nevertheless, our result is consistent with Bakshi et al. (1997), who find that the maximum likelihood estimates of κ and θ are statistically indistinguishable from their respective S&P 500 option-implied counterparts, i.e. κ^* and θ.

Since the λ parameter in Heston's model is constant, and given that we can impose that $\lambda = 0$ for the volatility risk premium embedded in Heston option prices on the Spanish IBEX-35 index according to the conclusions from the above tests, we can solve Eqs. (13) and (14) for $\lambda = 0$ and obtain a daily estimate of instantaneous variance from January 3, 1996 to April 30, 1996. We also repeat the above procedure for two alternative λ values, specifically $\lambda = 0.50$ and $\lambda = -0.50$, in order to analyze the sensitivity of the estimated instantaneous variance under different magnitudes of λ which are in accordance with our sample of λ values. In Section 6, we will test the out-of-sample pricing performance for Heston's model under the three candidate values of λ and compare each of them with BS's performance.

5.3. Implied volatility graphs

Fig. 3 contains the evolution for several series of implied volatilities from January 3, 1996 to April 30, 1996 for each version of Heston's model, i.e. for $\lambda = -0.5$, 0.0, 0.5 and for a daily-varying λ value (see Fig. 2), and also for the BS formula. It turns out that Heston's volatility, regardless of what volatility risk

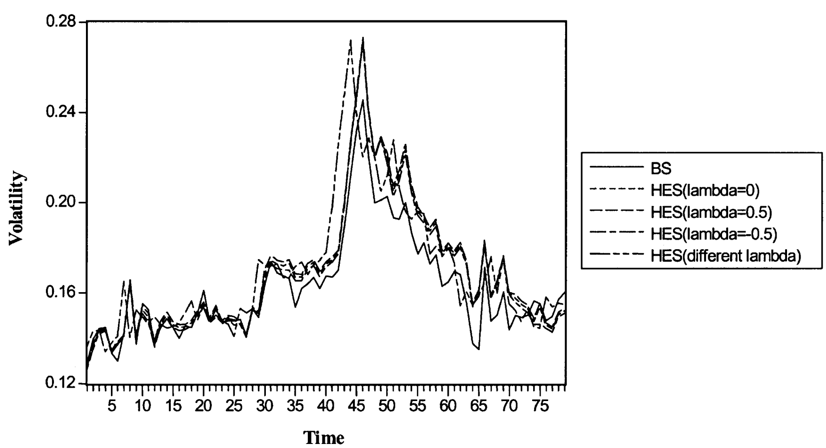

Fig. 3. Daily implied volatilities (January 96–April 96).

premium is assumed, tends to be higher than BS's volatility. A priori, this may have serious implications for pricing. In addition, it should be noted that the four implied volatility series under Heston's model are very similar. It is also the case that, in general, and for $\lambda = -0.5$, 0.0, 0.5, the larger the required volatility risk premium imposed, the (very slightly) larger the instantaneous volatility observed.

To obtain a general picture of the potential misspecifications of the option-pricing models employed in this research, we report the average pattern of implied volatilities across degrees of moneyness. The results are shown in Fig. 4.

In the BS case, we back out the implied volatility of each call option and for each day of the above period using the procedure discussed in Section 2.2. Then the equally weighted implied volatility for each moneyness category and each day in the sample period is calculated. There is a U-shaped pattern with a hump in the middle for ATM options. This suggests that the BS model tends to underprice deep OTM and deep ITM calls. This is the typical smile pattern of the Spanish option market analyzed by Peña et al. (1999). Any reasonable alternative model to BS must be able to properly price deep OTM and deep ITM call options. Of course, Heston's stochastic volatility model, given by Eqs. (4) and (5), is a potential and particularly interesting candidate.

In the Heston case, and for $\lambda = -0.5$, 0.0, 0.5, we may analyze the pattern of implied (instantaneous) volatilities across alternative degrees of moneyness. The evidence reported in Fig. 4 suggests a rather asymmetric smile in Heston's model regardless of what volatility risk premium is imposed. It seems that Heston's approach tends to underprice (deep) ITM calls and overprice (deep) OTM calls. We now turn to formal tests of the alternative option-pricing specifications considered in this paper.

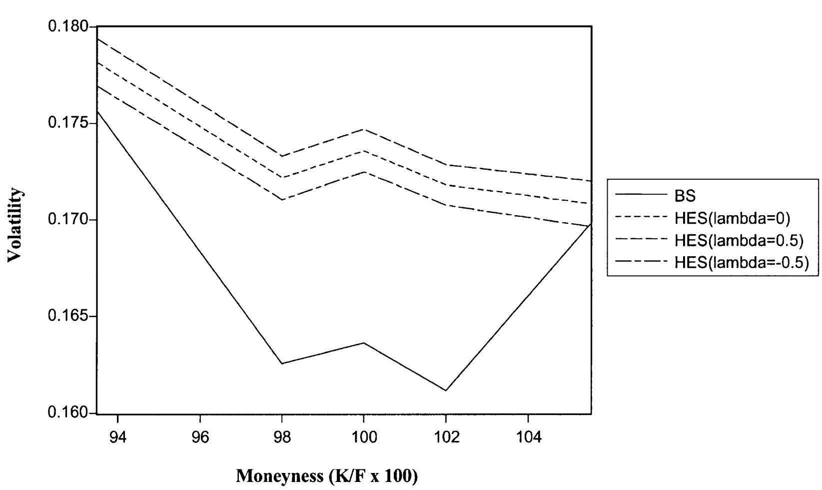

Fig. 4. Smiles (January 96–April 96).

6. Out-of-sample pricing performance

In order to test the out-of-sample pricing performance for each model analyzed in this work, we employ 2 years of rolling data to estimate, by indirect inference, the parameters of the stochastic variance process assumed in Eq. (1). Given these estimates and the volatility risk premium, λ, we use all call options available in our 45-min window to compute the instantaneous variance for each day from January 2, 1996 to April 29, 1996 that minimized the squared error between the theoretical value and the market price of the call options according to Eqs. (13) and (14). We then compute the theoretical price of each option using the previous day's instantaneous variance and the corresponding parameters of the stochastic volatility process. For the BS case, the previous day's implied volatility that minimized the squared error between the theoretical value and the market price of the options is used to obtain the theoretical BS price of each option in the sample.

In this way, we have 768 pricing errors for each of the calls available from January 3, 1996 to April 30, 1996, and for each of the models analyzed. These pricing errors are the basis for our analysis. Table 3 reports two measures of performance for the alternative model specifications. Panel A contains the absolute pricing error which is the sample average of the absolute difference between the model price and the market price for each call in the testing sample period. This statistic is reported for each moneyness category and for all calls in the sample. In Panel B, the reported percentage pricing error is the sample average of the theoretical price minus the market price, divided by the market price. Again, this statistic is calculated for each moneyness category and for all calls in the sample.

Overall, in absolute terms, Heston's option-pricing model tends to perform slightly better than BS. The absolute pricing error over all calls is approximately 2.8 pesetas for Heston's model independently of the volatility risk premium assumed, and 3.2 pesetas for the BS model. It is quite important to notice that the level of the volatility risk premium does not seem to have any influence on the performance of Heston's stochastic volatility model. The pricing errors obtained under alternative λ values are practically identical. There might be some evidence in favor of a positive risk premium, but we can safely conclude that the performance of the model does not seem to be sensitive to the volatility risk premium. This is an important empirical result. Kapadia (1998) shows that the expected value of the delta-hedged gain, under a stochastic volatility model where volatility risk is not priced, is zero. This is exactly the same result that holds under BS. Consequently, if we may assume that volatility risk is not priced, the analysis of the dynamic hedging performance of both models is clearly facilitated.[10]

[10] On the other hand, if volatility risk is priced, the average delta-hedged gain is proportional to the magnitude and sign of the volatility risk premium. Comparisons of the dynamic hedging performance between Heston's model and BS may be much more complicated than the analysis carried out by Bakshi et al. (1997).

Table 3
Out-of-sample pricing error for alternative option-pricing models: absolute pricing error and percentage pricing error
Two years of rolling daily data are employed to estimate by indirect inference the parameters of the stochastic volatility process under Heston's model. Given these estimates, and a chosen volatility risk premium, λ, we use all call options transacted over the 45-min interval from 16:00 to 16:45 to compute for each day from January 2, 1996 to April 29, 1996, the instantaneous variance that minimized the squared error between the theoretical value and the market price of the options. We then compute the theoretical price of each option using the previous day's instantaneous variance and the corresponding parameters of the stochastic volatility process. For the Black–Scholes case, the previous day's implied volatility that minimized the squared error between the theoretical value and the market price of the options is used to obtain the theoretical price of each option in the sample. The reported absolute pricing error is the sample average of the absolute difference between the model price and the market price of each call in a given moneyness category. The reported percentage pricing error is the sample average of the theoretical price minus the market price, divided by the market price. K is the exercise price and F denotes the future price of the IBEX-35 index. Moneyness is defined as the ratio of the exercise price to the futures price. OTM, ATM, and ITM are out-of-the-money, at-the-money, and in-the-money options, respectively.

Panel a: absolute pricing error

	Moneyness (K/F)	Black–Scholes (Ptas.)	Heston (Ptas.)		
			($\lambda = 0$)	($\lambda = 0.5$)	($\lambda = -0.5$)
Deep OTM calls	1.03–1.08	1.829	2.033	2.030	2.038
OTM calls	1.01–1.03	2.641	2.815	2.819	2.774
ATM calls	0.99–1.01	3.847	2.874	2.854	2.896
ITM calls	0.97–0.99	4.690	3.593	3.586	3.600
Deep ITM calls	0.90–0.97	4.357	3.809	3.795	3.823
All calls	–	3.249	2.859	2.852	2.853

Panel b: percentage pricing error

	Moneyness (K/F)	Black–Scholes (%)	Heston (%)		
			($\lambda = 0$)	($\lambda = 0.5$)	($\lambda = -0.5$)
Deep OTM calls	1.03–1.08	−7.091	5.916	6.009	5.673
OTM calls	1.01–1.03	1.189	8.429	8.583	8.310
ATM calls	0.99–1.01	−3.894	−0.053	−0.021	−0.089
ITM calls	0.97–0.99	−3.285	−1.351	−1.331	−1.371
Deep ITM calls	0.90–0.97	−2.560	−2.251	−2.243	−2.260
All calls	–	−2.439	3.607	3.689	3.513

As a way of analyzing the in-sample fitness of the models, we also calculated the pricing error for each call option on every day t in the sample using the time series estimates of parameters from the indirect inference, the instantaneous variance chosen to minimize the pricing error on the same day t, and $\lambda = 0$. The same procedure in terms of implied volatility was followed for the BS formula. We then compared the magnitude of average absolute pricing errors. As in the out-of-sample case, the average absolute error for all calls in the sample is lower for Heston's model (1.7 pesetas) than for BS (2.9 pesetas). As expected, the average absolute pricing errors are now lower than in the out-of-sample analysis. The conclusions are therefore the same in both in-sample and out-of-sample frameworks.[11]

It should be pointed out that the slightly better overall performance of Heston's model is not maintained throughout all moneyness categories. In particular, Heston's model tends to value ATM and ITM calls better than BS. However, the opposite result holds for OTM calls. This is also the case when we analyze the percentage pricing error in Panel B. Heston's model, regardless of the volatility risk premium imposed, tends to overvalue OTM calls and, at the same time, the model undervalues ITM calls. However, the percentage pricing for ATM calls is practically zero. In fact, ATM calls are clearly more consistent with a stochastic volatility model than with the well-known lognormal assumption. The BS model, on the other hand, tends to undervalue deep OTM and deep ITM calls. This is consistent with a U-shaped volatility smile. In Heston's case, the evidence points towards a sneer rather than a regular smile.

6.1. The statistical significance of the out-of-sample performance

Overall, at least over the sample period studied and regardless of λ, Heston's model seems to overvalue, while BS formula tends to undervalue call options. Unfortunately, however, the simple statistics reported above do not help in making inferences in terms of the statistical significance of improvement when we contrast one model with the other.

In this paper, the statistical significance of performance for out-of-sample pricing errors is assessed by analyzing the proportion of theoretical prices lying outside their corresponding bid-ask spread boundaries.[12] The following Z-statistic for the difference between two proportions is employed in the tests. The statistic is given by:

$$Z = \frac{p_1 - p_2}{\sqrt{p_1(1-p_1)/n_1 + p_2(1-p_2)/n_2}} \tag{15}$$

where p_1 is always the proportion of BS prices outside the bid-ask boundaries, and p_2 is the equivalent proportion for alternative Heston's specifications. n_1 and n_2 are sample sizes corresponding to these proportions. The statistic is asymptotically distributed as a standardized normal variable.

The empirical results are reported in Table 4. Given that we are also interested in knowing whether a given theoretical valuation model undervalues or overvalues

[11] The complete in-sample results are not reported for reasons of space. The results are available from the authors upon request.

[12] See Corrado and Su (1996).

Table 4

Nonparametric testing for alternative option-pricing models

The statistical significance of performance for out-of-sample pricing errors is assessed by analyzing the proportion of theoretical prices lying outside their corresponding bid-ask spread boundaries. The following Z-statistic for the difference between two proportions given by:

$$Z=\frac{p_1-p_2}{\sqrt{p_1(1-p_1)/n_1+p_2(1-p_2)/n_2}}$$

is employed in the tests, where p_1 is always the proportion of Black–Scholes prices outside the bid-ask boundaries, and p_2 is the equivalent proportion for alternative Heston's model specifications. n_1 and n_2 are sample sizes corresponding to these proportions. The statistic is asymptotically distributed as a standardized normal variable. All call options transacted over the 45-min interval from 16:00 to 16:45 from January 3, 1996 to April 30, 1996 are used in the tests below. K is the exercise price and F denotes the future price of the IBEX-35 index. Moneyness is defined as the ratio of the exercise price to the futures price. $p[c_{\text{MODEL}} \notin (\text{Bid, Ask})]$ gives the frequency in which the theoretical price does not belong to the bid-ask spread interval.

Categories	BS	Heston ($\lambda = 0$)	Z-stat. (p-value)	Heston ($\lambda = 0.5$)	Z-stat. (p-value)	Heston ($\lambda = -0.5$)	Z-stat. (p-value)
All options							
$p[c_{\text{MODEL}} \notin (\text{Bid, Ask})]$	0.4761	0.4249	1.825 (0.068)	0.4249	1.825 (0.068)	0.4233	1.882 (0.060)
$p(c_{\text{MODEL}} < \text{Bid})$	0.3487	0.1837	6.730 (0.000)	0.1821	6.804 (0.000)	0.1901	6.434 (0.000)
$p(c_{\text{MODEL}} > \text{Ask})$	0.1274	0.2412	− 5.253 (0.000)	0.2428	− 5.319 (0.000)	0.2332	− 4.919 (0.000)
OTM options ($K/F > 1$)							
$p[c_{\text{MODEL}} \notin (\text{Bid, Ask})]$	0.4821	0.4347	1.421 (0.155)	0.4347	1.421 (0.155)	0.4324	1.421 (0.155)
$p(c_{\text{MODEL}} < \text{Bid})$	0.3184	0.1351	6.694 (0.000)	0.1329	6.791 (0.000)	0.1419	6.399 (0.000)
$p(c_{\text{MODEL}} > \text{Ask})$	0.1637	0.2995	− 4.864 (0.000)	0.3018	− 4.940 (0.000)	0.2905	− 4.566 (0.000)
ITM options ($K/F < 1$)							
$p[c_{\text{MODEL}} \notin (\text{Bid, Ask})]$	0.4615	0.4011	1.166 (0.244)	0.4011	1.166 (0.244)	0.4011	1.1656 (0.244)
$p(c_{\text{MODEL}} < \text{Bid})$	0.4231	0.3022	2.418 (0.015)	0.3022	2.418 (0.015)	0.3077	2.303 (0.021)
$p(c_{\text{MODEL}} > \text{Ask})$	0.0385	0.0989	− 2.294 (0.022)	0.0980	− 2.294 (0.022)	0.0934	− 2.123 (0.034)
ATM options ($1.01 > K/F \geq 0.99$)							
$p[c_{\text{MODEL}} \notin (\text{Bid, Ask})]$	0.4583	0.3750	1.762 (0.078)	0.3749	1.763 (0.078)	0.3735	1.765 (0.077)
$p(c_{\text{MODEL}} < \text{Bid})$	0.3981	0.1806	5.134 (0.000)	0.1800	5.395 (0.000)	0.1893	5.189 (0.000)
$p(c_{\text{MODEL}} > \text{Ask})$	0.0602	0.1944	− 4.272 (0.000)	0.1987	− 4.341 (0.000)	0.1915	− 4.231 (0.000)

market prices, the Z-statistic is also calculated to obtain the proportion for which the theoretical model yields a price below the bid quote, and the proportion for which the model gives a price above the ask quote. If a theoretical model tends to undervalue market prices, it should yield a higher proportion of prices below the bid quote. If, on the other hand, the model tends to overvalue market prices, it should have a higher proportion of prices above the ask quote.

When we consider all call options together, the p-value for statistical improvement of Heston's model over the BS formula is 0.068. The proportion of BS prices lying outside the bid-ask boundaries is 48%, while the proportion of Heston's model, regardless of the volatility risk premium, is about 42%. As we see, there is a slight improvement. However, given the costs of implementation, it is not clear whether Heston's approach is worthwhile in terms of practical applications.

It is also the case that 35% of BS prices are below the bid quote. Given that Heston's model yields approximately 18% of prices below the bid, we can conclude that BS, on average, significantly undervalues market prices relative to Heston's approach. At the same time, 24% of Heston's prices are above the ask quote. We can also conclude that Heston's model tends to significantly overvalue market prices relative to the BS formula. Thus, mispricing associated with BS is basically related to the tendency of the model to yield prices below market values. However, Heston's mispricing is a consequence of the model's tendency to offer prices above their corresponding market values. This is consistent with the results reported in Panel B of Table 3, where the overall percentage pricing error associated with BS is negative, while the corresponding percentage error for Heston's model is positive.

If we classify call options by moneyness, similar conclusions are obtained. A call option is now said to be OTM if $K/F > 1$, and ITM otherwise. In general, within a given category of moneyness, neither model is statistically superior to the other. However, once again, BS mispricing comes from the tendency of the model to undervalue either OTM or ITM calls relative to Heston's model (and relative to market prices). The opposite is true for Heston's formula relative to BS (and market prices). It should be pointed out that in the latter case the main problem arises when we value OTM with Heston's formula. There seems to be a strong tendency in Heston's model to overvalue this type of options. This is also reflected in Table 3.

It was mentioned above that the percentage pricing error of Heston's model for ATM calls turns out to be quite small. It may be the case that Heston's formula works particularly well for ATM options. Hence, calls are now said to be ATM when $1.01 > K/F \geq 0.99$. The exercise described above is repeated for this type of options. Overall, as expected, Heston's model tends to yield lower proportions of prices lying outside the bid-ask spread than in previously analyzed cases. However, the p-value for the difference relative to BS proportions is just 0.078. As before, BS tends to yield a statistically significant higher proportion of prices lower than the bid quote relative to Heston's prices, and Heston's model presents a statistically significant higher proportion of prices above the ask quote relative to BS prices. Now, however, Heston's formula has similar proportions above the ask or below the bid. Regarding ATM calls, Heston's misspecification cannot be explained by either overvaluation or undervaluation of call prices.

6.2. Skewness, kurtosis and the parameters of the stochastic volatility process

The time-varying behavior of the correlation coefficient found above may have a serious impact on the capacity of Heston's model to explain option-pricing data. It suggests that, in this case, we may have a problem similar to the one we have when assuming constant volatility in the BS context. If the correlation between prices and volatility changes continuously over time, skewness of the underlying asset may also exhibit time-varying behavior. This is clearly a potential and relevant problem for models with stochastic volatility.

On the other hand, according to the estimates shown in Fig. 1, it seems that the behavior of the volatility of volatility is rather stable over time. This suggests that accounting for changing kurtosis of the underlying asset may not be as crucial as taking into consideration changing skewness.

Das and Sundaram (1999) obtain closed-form expressions for conditional and unconditional skewness and kurtosis under Heston's stochastic volatility model. Their expressions, for the frequency of $\Delta t = 1$, can be seen in Appendix B.

Given our rolling estimates of ρ and σ from January 1996 to April 1996, we can daily evaluate (B1) and (B2) in Appendix B, so that we may observe how the characteristics of the distribution of the underlying asset—skewness and kurtosis — change with both the correlation coefficient and the volatility of volatility.

Fig. 5 depicts the conditional and unconditional skewness over the sample period. As we can easily observe, their behavior closely follows the pattern of the

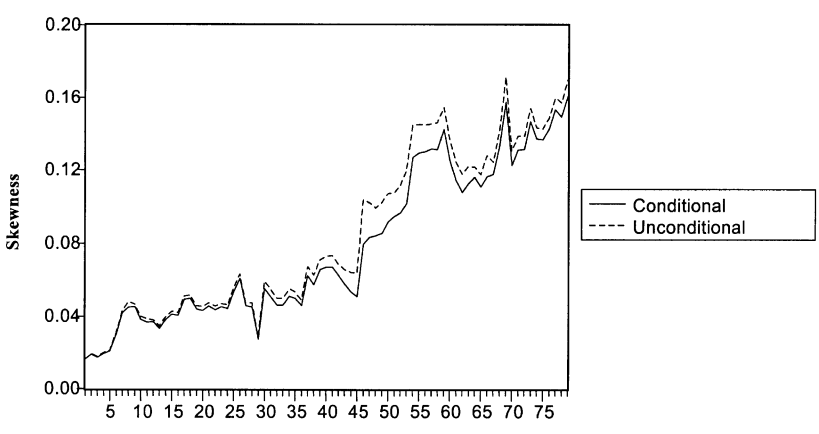

Fig. 5. Time-varying skewness (January 96–April 96).

Fig. 6. Time-varying kurtosis (January 96–April 96).

correlation coefficient of Fig. 1. As expected, skewness mainly arises from the correlation between changing prices and stochastic volatility of the underlying asset. The problem is that, of course, Heston's model assumes a constant unconditional skewness over time.[13]

In León and Rubio (2000), it is shown that, all else being constant, the relationship between skewness and the correlation coefficient is positive; i.e. $\partial \mathrm{SKEW}/\partial\rho > 0$ for both the conditional and unconditional cases. This explains the behavior of skewness in Fig. 5. Therefore, if the correlation coefficient presents a time-varying behavior (as seems to be the case in practice), an option-pricing model with a stochastic differential equation for correlation is a desirable extension of Heston's model. Given this evidence, we should not expect to find a good performance of the stochastic volatility model when we compare observed market prices with theoretical prices.

Fig. 6 contains the same evidence for the kurtosis case. In principle, the impact of time-varying kurtosis on the misspecification of Heston's model seems to be much less severe than the influence of the correlation coefficient. Its pattern over time is much more stable than ρ. This is related to the behavior of the volatility of volatility, σ, in Fig. 1. This is the parameter that allows for kurtosis in the stochastic volatility option-pricing model, and it does not seem to change much over time. Once again, in León and Rubio (2000), it is shown that the relationship

[13] Given the similarity between the behavior of unconditional and conditional skewness, the impact of the instantaneous stochastic variance on the conditional skewness must be negligible relative to the effect of the correlation coefficient.

between kurtosis and σ is positive, i.e. $\partial \mathrm{KURT}/\partial\sigma > 0$, for both the conditional and unconditional cases.

As a summary, time-varying skewness may be the key issue to analyze if we want to understand the failure of the tests we report in previous sections. Its consequences may be much more serious than the potential effects of changing kurtosis.

6.3. The structure of pricing errors

Given the poor empirical performance of both the BS formula and Heston's stochastic volatility model, a further analysis trying to understand the structure of pricing errors of these models would seem to be called for. Following the evidence reported by Bakshi et al. (1997) and Peña et al. (1999), we use a simple regression framework to study the relation between the percentage pricing errors and factors that are either option-specific or market dependent. We first take as given an option-pricing model, and let e_{it} be the ith call option's percentage pricing error on day t. Finally, we run the following regression for the whole sample period:

$$e_{it} = \alpha + \beta_1 X_{it} + \beta_2 \tau_{it} + \beta_3 \mathrm{SP}_t + \beta_4 \mathrm{VOL}_t + \beta_5 \mathrm{TERM}_t + \beta_6 \mathrm{SKEW}_t + \beta_7 \mathrm{KURT}_t + \omega_{it} \tag{16}$$

where X_{it} is the moneyness of the ith call at time t as defined by the ratio between the exercise price (K) and the future price (F); τ_{it} is the annualized time to expiration of the ith call on day t; SP is the average relative bid-ask spread of all calls and puts transacted between 16:00 and 16:45 on date t; VOL_t is the annualized daily standard deviation of the IBEX-35 index returns computed from 1-min intradaily returns; TERM_t is the yield differential between the annualized 10-year government bond and the annualized 1-month repo Treasury bill; SKEW_t and KURT_t are the daily implicit skewness and (excess) kurtosis obtained from option prices.[14] Table 5 contains the regression results based on the entire sample period and 768 call options.[15] The explanatory variables employed in the regressions tend to be statistically significant. However, there are important differences between the percentage errors associated with either BS or Heston.

In accordance with our analysis in Section 6.2 above, Heston's formula has a (slight) skewness bias. It should be noted that this is not the case for the BS formula, where both the traditional moneyness bias and the volatility bias seem to

[14] In order to estimate the daily skewness and kurtosis implicit in option prices, we use a Gram–Charlier series expansion of the normal density function truncating it in the fourth moment. See Jarrow and Rudd (1982) and Corrado and Su (1996). The authors thank Gregorio Serna (Universidad Castilla la Mancha) for the estimation of these series.

[15] The standard error for each coefficient estimate is adjusted by the White (1980) heteroskedasticity-consistent estimator. The same regression was run using 1-day lagged values for the market dependent variables. Very similar results were found.

Table 5

Percentage pricing errors and explanatory variables

For a given option pricing model, the following regression is employed to explain the percentage pricing errors of all call options transacted over the 45-min interval from 16:00 to 16:45 from January 3, 1996 to April 30, 1996:

$$e_{it} = \alpha + \beta_1 X_{it} + \beta_2 \tau_{it} + \beta_3 \mathrm{SP}_t + \beta_4 \mathrm{VOL}_t + \beta_5 \mathrm{TERM}_t + \beta_6 \mathrm{SKEW}_t + \beta_7 \mathrm{KURT}_t + \omega_{it}$$

where e_{it} is the percentage pricing error of the ith call on date t; X is the moneyness of the ith call at time t as defined by the ratio between the strike price (K) and the future price (F); τ_{it} is the annualized time to expiration of the ith call on day t; SP is the average relative bid-ask spread of all calls and puts transacted between 16:00 and 16:45 on date t; VOL_t is the annualized standard deviation of the IBEX-35 index returns computed from 1-min intradaily returns; TERM_t is the yield differential between the annualized 10-year government bond and the annualized 1-month repo Treasury bill; SKEW_t is the implicit skewness found in option prices and KURT_t is the implicit (excess) kurtosis obtained from option prices. The t-statistic reported in parentheses is based on White's heteroskedasticity consistent estimator of standard errors. A total of 768 call options are employed in the regressions.

Coefficient	Black–Scholes	Heston ($\lambda = 0$)	Heston ($\lambda = 0.5$)	Heston ($\lambda = -0.5$)
Constant	0.775 (1.98)	−1.094 (−2.68)	−1.105 (−2.69)	−1.069 (−2.64)
Moneyness (X)	−0.822 (−2.16)	0.901 (2.25)	0.915 (2.27)	0.875 (2.19)
Time to expiration (τ)	0.024 (2.64)	0.022 (2.34)	0.022 (2.30)	0.022 (2.35)
Spread (SP)	0.228 (1.38)	0.429 (2.66)	0.431 (2.67)	0.430 (2.68)
Volatility (VOL)	−1.423 (−2.00)	−1.050 (−1.52)	−1.062 (−1.54)	−1.061 (−1.47)
Term structure (TERM)	0.016 (0.53)	0.075 (2.63)	0.075 (2.61)	0.074 (2.58)
Skewness (SKEW)	−0.073 (−0.96)	−0.128 (−1.72)	−0.127 (−1.70)	−0.126 (−1.70)
Kurtosis (KURT)	−0.002 (−0.31)	−0.003 (−0.61)	−0.002 (−0.46)	−0.003 (−0.59)

be capturing the role played by skewness.[16] On the other hand, neither Heston's model nor BS formula presents a kurtosis bias.

As expected, given the assumption of constant volatility under the BS model, the annualized standard deviation of the IBEX-35 index explains its percentage pricing errors. On average, BS pricing errors tend to be lower the higher the volatility of the index. On the other hand, the coefficients associated with the alternative specifications of Heston's model are not statistically significant. This is, of course, a crucial distinction between the two models, and it suggests how important it is to employ option-pricing models with time-varying volatility.

Traditional biases are not corrected for any of the models. The bias associated with moneyness has the opposite sign in the two types of models. As expected, given previous results, Heston tends to price OTM options worse than ITM calls. However, on average, the opposite result is found for BS. Moreover, the bias related to time to expiration seems to be relevant for both types of models. On average, the percentage pricing errors are larger the longer the time to expiration.

[16] The coefficient of skewness becomes significant when we run the regressions without moneyness and volatility.

The influence on percentage pricing errors of both the interest rates yield differential and the average relative bid-ask spread for all calls and puts transacted over the 45-min window is, interestingly enough, different for Heston and BS expressions. In the stochastic volatility model, the higher the long-term yield relative to the short-term rate, the higher the percentage pricing error. However, this result becomes statistically insignificant when we consider the BS model. To understand this result, it should be noted that neither of these two models incorporates a dynamic term structure setup, so that both models should present a term structure bias. At the same time, however, we argue below that the effect of changing the yield differential should be stronger in Heston's model than in BS. Let TERM be equal to $r_L - r$, the difference between the long and short interest rates. Assuming that the short-term interest rate is more volatile than the long rate (as was actually the case during the sample period), it has to be the case that the yield differential is (mainly) a function of the short term interest rate, so that $\mathrm{dTERM} = -\mathrm{d}r$. Then, $\mathrm{d}c/\mathrm{dTERM} = (\mathrm{d}c/\mathrm{d}r)(\mathrm{d}r/\mathrm{dTERM}) = (\mathrm{d}c/\mathrm{d}r)(-1)$. By recalling either Heston's formula for options on future or the BS expression, it is the case that $\partial c/\partial r = -(T-t)c$. Therefore, the elasticity of option prices relative to the yield differential (TERM) is positive. Given that Heston's model tends to overvalue market prices, while the opposite result is obtained for the BS framework, and taking into account our definition of pricing error—theoretical price minus market price—, the effect of changing the yield differential should be stronger under Heston's framework than in the BS case.

Regarding the bid-ask spread variable, it should be mentioned that regressions of a similar type were run including the relative bid-ask spread at date t of the call i. This is, contrary to the results reported in Table 5, a contract-specific variable. Surprisingly, the estimated coefficients are never significant regardless of the model considered. By doing this, we are really incorporating a transaction cost variable directly associated with the liquidity of each individual option. Again, this variable does not seem to be significant in explaining percentage pricing errors. However, when we include the average relative bid-ask spread over all call and put options for a particular day t, the estimated coefficient becomes positive and significant in Heston's case. This aggregate variable may indicate the average consensus about the uncertainty of trading throughout the option market. It may be understood as the average adverse selection confronting market makers in trading options on the Spanish index. As we see from Table 5, the larger the average spread—larger average adverse selection among traders—the higher the percentage pricing error in Heston's stochastic volatility models. The impact of this type of uncertainty does not seem to be relevant in explaining the percentage pricing errors of BS.

In short, the pricing errors from Heston's framework have some moneyness, maturity, average (aggregate) bid-ask spread, skewness, and yield differential related biases. On the other hand, BS presents some moneyness, maturity, and volatility associated biases. Neither model seems to capture appropriately the underlying distribution characteristics of the underlying asset. Further research is clearly justified.

7. Conclusions

This paper introduces a two-step procedure, based on both time series and cross-section data, for estimating all the parameters that we need to compute Heston call price. This estimation approach should be particularly useful in thinly traded markets where a single cross-sectional estimation approach may be difficult to implement, as is the case in the Spanish options data on the IBEX-35 stock index. Moreover, to employ just a cross-sectional procedure may ignore relevant information that may be included in the original series but not in the option prices. By using the indirect inference procedure to adjust for discretization biases, our approach combines the time-series information contained in the underlying security with the cross-sectional information embedded in option prices.

On average, over all call options available in our sample, Heston's model improves the (poor) performance of BS just marginally. It is clear that this extremely limited improvement barely justifies the implementation costs involved in the estimation of Heston's approach.[17] The overall rejection of Heston's model coincides with recent findings by Bakshi et al. (1997) and Chernov and Ghysels (2000) for options written on the S&P 500 index.

It is also found that the daily volatility risk premium behaves in a quite volatile fashion over time. However, our evidence suggests that the volatility risk premium has a negligible impact on the pricing performance of Heston's model.

The findings of this paper suggest possible directions for future research. Presumably, the ultimate reasons behind the shortcomings of Heston's model are closely related to the time-varying skewness found in the data. In particular, the assumption of a constant correlation coefficient between returns and stochastic volatility should be relaxed if we want to have a richer model. It should be clear that our results are generally applicable regardless of what market is used in the estimation. The behavioral restrictions imposed on the parameters of the diffusion process should be understood as a limitation of the stochastic volatility model. Unfortunately, of course, the complexities needed to price options seem to increase without bounds. Another potentially interesting area of research might be related to endogenously incorporating liquidity costs in option-pricing models with either stochastic volatility, stochastic jumps or both. Unfortunately, once again, this approach may be extremely demanding from a theoretical point of view.

[17] Hedging performance of alternative models is not analyzed in this paper. It is possible that the hedging performance under Heston's stochastic volatility model might be clearly superior to BS.

Acknowledgements

We thank Stewart Mayhew, Eduardo Schwartz, Enrique Sentana, Ignacio Peña, Gregorio Serna, Rafael Salinas, José Luis Fernández, Eliseo Navarro, Alejandro Balbás, Rafael Santamaría, two anonymous referees, and one of the editors of the Journal of Empirical Finance for their helpful comments and useful discussions. Earlier versions of this paper were presented at the European Finance Association meeting in Helsinki (1999), the European Financial Management Association conference in Paris (1999), Universidad Carlos III, II Foro Finanzas-Segovia, and MEFF Institute for Derivative Assets. Gabriele Fiorentini, Angel León and Gonzalo Rubio acknowledge the financial support provided by Dirección Interministerial Científica y Técnica (DGICYT) grants PB96-0339, PB98-0979 and PB97-0621, respectively. All three authors benefited from a research grant received from the Instituto Valenciano de Investigaciones Económicas (IVIE). The contents of this paper are the sole responsibility of the authors.

Appendix A. Heston's stochastic volatility model using future options

Given that we use future options, the actual version of the formula we employ is given by:

$$c(F,\upsilon,t) = \mathrm{e}^{-r(T-t)}(F_t P_1 - KP_2)$$

where F is the future price on the underlying spot price, K is the exercise price and the probabilities are given by:

$$P_j = \frac{1}{2} + \frac{1}{\pi}\int_0^{\infty} \mathrm{Re}\left(\frac{\mathrm{e}^{-i\phi \ln[K]} f_j}{i\phi}\right)\mathrm{d}\phi; \quad j = 1,2$$

where $\mathrm{Re}(y)$ is the real part of the function y; i is the imaginary number $i = \sqrt{-1}$, and

$$f_j(x,\upsilon,T-t;\phi) = \exp\left[C(T-t;\phi) + D(T-t;\phi)\upsilon + i\phi x\right]$$

where,[18]

$$C(T-t;\phi) = \frac{a}{\sigma^2}\left\{\left(b_j - \rho\sigma\phi i + d\right)(T-t) - 2\ln\left[\frac{1 - g\mathrm{e}^{\mathrm{d}(T-t)}}{1-g}\right]\right\}$$

$$D(T-t;\phi) = \frac{b_j - \rho\sigma\phi i + d}{\sigma^2}\left[\frac{1 - \mathrm{e}^{\mathrm{d}(T-t)}}{1 - g\mathrm{e}^{\mathrm{d}(T-t)}}\right]$$

[18] As we have already pointed out, the expression for $C(T-t;\phi)$ below is slightly different from the original value given by Heston (1993), since we are using futures.

$$g = \frac{b_j - \rho\sigma\phi i + d}{b_j - \rho\sigma\phi i - d}$$

$$d = \sqrt{\left(\rho\sigma\phi i - b_j\right)^2 - \sigma^2\left(2u_j\phi i - \phi^2\right)}$$

$$a = \kappa\theta$$

$$b_1 = \kappa + \lambda - \rho\sigma$$

$$b_2 = \kappa + \lambda$$

$$u_1 = 1/2; u_2 = -1/2$$

verifying that $C(0) = D(0) = 0$, and where $C(T-t;\phi)$ and $D(T-t;\phi)$ (and therefore the probabilities P_j; $j = 1, 2$) depend on the vector of parameters, $(\kappa, \theta, \lambda, \sigma, \rho)$ given by the processes assumed by Heston under the original probability.

Appendix B. (Un)conditional skewness and kurtosis

(a) The conditional case:

$$\mathrm{SKEW} = \left(\frac{3\sigma\rho e^{\kappa/2}}{\sqrt{\kappa}}\right)\left[\frac{\theta(2 - 2e^{\kappa} + \kappa + \kappa e^{\kappa}) - \upsilon(1 + \kappa - e^{\kappa})}{\left[\theta(1 - e^{\kappa} + \kappa e^{\kappa}) + \upsilon(e^{\kappa} - 1)\right]^{3/2}}\right]$$

$$\mathrm{KURT} = 3\left[1 + \sigma^2\left(\frac{\theta A_1 - \upsilon A_2}{B}\right)\right] \tag{B1}$$

where

$$A_1 = (1 + 4e^{\kappa} - 5e^{2\kappa} + 4\kappa e^{\kappa} + 2\kappa e^{2\kappa}) + 4\rho^2(6e^{\kappa} - 6e^{2\kappa} + 4\kappa e^{\kappa} + 2\kappa e^{2\kappa} + \kappa^2 e^{\kappa})$$

$$A_2 = 2(1 - e^{2\kappa} + 2\kappa e^{\kappa}) + 8\rho^2(2e^{\kappa} - 2e^{2\kappa} + 2\kappa e^{\kappa} + \kappa^2 e^{\kappa})$$

$$B = 2\kappa\left[\theta(1 - e^{\kappa} + \kappa e^{\kappa}) + \upsilon(e^{\kappa} - 1)\right]^2$$

and where $\upsilon \equiv V_t$.

(b) The unconditional case:

$$\mathrm{SKEW} = 3\left(\frac{\rho\sigma}{\sqrt{\kappa\theta}}\right)\left[\frac{1 - e^{\kappa} + \kappa e^{\kappa}}{\kappa^{3/2}e^{\kappa}}\right]$$

$$\mathrm{KURT} = 3\left[1 + \frac{\sigma^2}{\kappa\theta\kappa^2 e^{\kappa}}\left(1 - e^{\kappa} + \kappa e^{\kappa} + 4\rho^2[2 - 2e^{\kappa} + \kappa + \kappa e^{\kappa}]\right)\right] \tag{B2}$$

References

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

Bakshi, G., Cao, C., Chen, Z., 2000. Pricing and hedging long-term options. Journal of Econometrics 94, 277–318.

Bates, D., 1996. Jumps and stochastic volatility: exchange rate processes implicit in Deutsche mark options. Review of Financial Studies 9, 69–107.

Black, F., 1976. The pricing of commodity contracts. Journal of Financial Economics 3, 167–179.

Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. Journal of Political Economy 81, 637–659.

Broze, L., Scaillet, O., Zakoian, J.M., 1995. Testing for continuous time models of the short-term interest rate. Journal of Empirical Finance 2, 199–223.

Broze, L., Scaillet, O., Zakoian, J.M., 1998. Quasi-indirect inference for diffusion processes. Econometric Theory 14, 161–186.

Buraschi, A., Jackwerth, J., 2001. The price of a smile: hedging and spanning in option markets. Review of Financial Studies 14, 495–527.

Calzolari, G., Di Iorio, F., Fiorentini, G., 1998. Control variates for variance reduction in indirect inference: interest rate models in continuous time. Econometrics Journal 1, C100–C112.

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

Corrado, C., Su, T., 1996. Skewness and kurtosis in S&P 500 index returns implied by option prices. Journal of Financial Research 19, 175–192.

Das, S., Sundaram, R., 1999. Of smiles and smirks: a term-structure perspective. Journal of Financial and Quantitative Analysis 34, 211–239.

Derman, E., Kani, I., 1994. Riding on the smile. Risk 7, 32–39.

Dumas, B., Fleming, J., Whaley, R., 1998. Implied volatility functions: empirical tests. Journal of Finance 53, 2059–2106.

Dupire, B., 1994. Pricing with a smile. Risk 7, 18–20.

Efron, B., Tibshirani, R., 1993. An Introduction to the Bootstrap. Chapman & Hall, London, United Kingdom.

French, D., 1984. The weekend effect on the distribution of stock prices: implications for option pricing. Journal of Financial Economics 13, 547–559.

Gallant, A.R., Tauchen, G., 1996. Which moments to match? Econometric Theory 12, 657–681.

Gouriéroux, C., Monfort, A., 1996. Simulation-Based Econometric Methods. Oxford Univ. Press, Oxford, United Kingdom.

Gouriéroux, C., Monfort, A., Renault, E., 1993. Indirect inference. Journal of Applied Econometrics 8, S85–S118.

Guo, D., 1998. The risk premium of volatility implicit in currency options. Journal of Business and Economic Statistics 16, 498–507.

Heston, S., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. Review of Financial Studies 6, 327–344.

Heston, S., Nandi, S., 2000. A closed-form GARCH option valuation model. Review of Financial Studies 13, 585–625.

Hull, J., White, A., 1987. The pricing of options on assets with stochastic volatilities. Journal of Finance 42, 281–300.

Jarrow, R., Rudd, A., 1982. Approximate option valuation for arbitrary stochastic processes. Journal of Financial Economics 10, 347–369.

Kapadia, N., 1998. “Do equity options price volatility risk? Empirical evidence”, Working Paper, University of Massachusetts.

León, A., Mora, J., 1999. Modelling conditional heteroskedasticity: application to the IBEX-35 stock return index. Spanish Economic Review 1, 215–238.

León A., Rubio, G., 2000. "Smiling under stochastic volatility", unpublished manuscript, University of Alicante and University of País Vasco.

Lin, Y., Strong, N., Xu, X., 1999. "Pricing FTSE 100 index options under stochastic volatility", working paper 99/018, The Management School, Lancaster University.

Nelder, J.A., Mead, R., 1965. A simplex method for function minimization. Computer Journal 7, 308–313.

Nowman, K., 1997. Gaussian estimation of single-factor continuous time models of the term structure of interest rates. Journal of Finance 52, 1695–1706.

Pan, J., 1999. "Integrated time series analysis of spot and option prices", Manuscript, Stanford University.

Peña, I., Serna, G., Rubio, G., 1999. Why do we smile? On the determinants of the implied volatility function. Journal of Banking and Finance 23, 1151–1179.

Rohatgi, V.K., 1984. Statistical Inference. Wiley, Chichester, United Kingdom.

Rubinstein, M., 1994. Implied binomial trees. Journal of Finance 49, 771–818.

Sarwar, G., Krehbiel, T., 2000. Empirical performance of alternative pricing models of currrency options. The Journal of Futures Markets 20, 265–291.

Scott, L.O., 1987. Option pricing when the variance changes randomly: theory, estimation and an application. Journal of Financial and Quantitative Analysis 22, 419–438.

Smith, A., 1993. Estimating nonlinear time-series models using simulated vector autoregressions. Journal of Applied Econometrics 8, S63–S84.

White, H., 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. Econometrica 48, 817–838.

Wiggins, J.B., 1987. Option values under stochastic volatilities: theory and empirical estimates. Journal of Financial Economics 19, 351–372.