

Belledin, Michael; Schlag, Christian

Working Paper

An Empirical Comparison of Alternative Stochastic Volatility Models

Working Paper Series: Finance & Accounting, No. 38

Provided in Cooperation with:

Faculty of Economics and Business Administration, Goethe University Frankfurt

Suggested Citation: Belledin, Michael; Schlag, Christian (1999) : An Empirical Comparison of Alternative Stochastic Volatility Models, Working Paper Series: Finance & Accounting, No. 38, Johann Wolfgang Goethe-Universität Frankfurt am Main, Fachbereich Wirtschaftswissenschaften, Frankfurt a. M.,
<https://nbn-resolving.de/urn:nbn:de:hebis:30-18223>

This Version is available at:

<https://hdl.handle.net/10419/76939>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

**JOHANN WOLFGANG GOETHE-UNIVERSITÄT
FRANKFURT AM MAIN**

FACHBEREICH WIRTSCHAFTSWISSENSCHAFTEN

**Michael Belledin
Christian Schlag**

**An Empirical Comparison of Alternative
Stochastic Volatility Models**

May 1999

WORKING PAPER SERIES: FINANCE & ACCOUNTING

An Empirical Comparison of Alternative Stochastic Volatility Models

Michael Belledin

Christian Schlag

Current Version: May 28, 1999

Department of Economics and Business Administration
Goethe-University
D-60054 Frankfurt am Main
Germany
Phone: ++49 (69) 798-22674
Fax: ++49 (69) 798-22788
E-mail: belledin@wiwi.uni-frankfurt.de, schlag@wiwi.uni-frankfurt.de.

The authors thank Hermann Göppl at the University of Karlsruhe for his support in providing the data for this study.

An Empirical Comparison of Alternative Stochastic Volatility Models

Abstract

We perform an empirical comparison of stochastic volatility models using data from the German index options market. The models compared are those developed by Stein and Stein [15], Heston [8], and Schöbel and Zhu [12], while the standard Black and Scholes [4] approach is used as a benchmark.

The parameters of the four models are estimated implicitly from daily cross-sections of option prices using a simulated annealing algorithm to overcome the numerical deficiencies of standard optimization routines. The main result in terms of pricing performance is that there is a clear advantage for stochastic volatility models compared to Black and Scholes both in and out of sample.

Within the group of stochastic volatility models the more flexible approaches by Schöbel and Zhu [12] and Heston [8], allowing for a non-zero correlation between stock returns and volatility changes, are preferable to the restricted Stein and Stein model. In terms of hedging performance a slightly modified version of the Black and Scholes model is practically indistinguishable from the stochastic volatility models, and there is evidence that these more complex models are still misspecified.

Keywords: Option pricing, stochastic volatility, numerical optimization

1 Introduction

In their seminal paper Black and Scholes (B&S) [4] derived a simple closed-form solution for European options on non-dividend paying stocks. A key input to their formula is the instantaneous volatility of the log return for the underlying stock, which is assumed to be constant or at most a deterministic function of time over the remaining life of the option. Many empirical studies (e.g. Rubinstein [11]), however, document the phenomenon of a volatility smile or ‘sneer’, i.e. the existence of a systematic relationship between the implied volatility of the B&S model and the strike price or moneyness of an option.

Several attempts have been made in the literature to explain this result. Some authors investigate if transaction costs could be responsible for this phenomenon (e.g., Constantinides [5]), others try to explain the smile as the consequence of heterogeneity among investors (see Benninga and Mayshar [3]). Another line of research has focused on the use of alternative stock price distributions (like mixtures of lognormals, see Melick and Thomas [10]), but these models generally suffer from the deficiency that the distributions under consideration cannot be the result of some stochastic process for the stock price. Finally, some generalizations of the classical geometric Brownian motion of the B&S model have been suggested, e.g. the introduction of jump components, stochastic volatility, or stochastic interest rates.

In contrast to the early work on stochastic volatility by Scott [13], Wiggins [16], Hull and White [9], and Stein and Stein (S&S) [15], more recent papers like Heston [8], Scott [14], and Schöbel and Zhu (S&Z) [12] have used Fourier inversion techniques to generate closed-form pricing formulas which only involve the numerical integration of the real part of a

complex function. Bakshi, Cao and Chen [1] present the most general model in this class, incorporating stochastic volatility, jumps, and stochastic interest rates simultaneously. However, they point out that ‘... taking stochastic volatility into account is of first-order importance in improving upon the BS formula’(pp. 2042-2043.) Therefore, we restrict the empirical analysis in this paper to models whose only ‘extra feature’ is stochastic volatility, differing only in the specification of its stochastic evolution.

The models compared are S&S, S&Z, Heston, and B&S. In S&S, instantaneous volatility follows an Ornstein-Uhlenbeck process, which is forced to be uncorrelated with innovations in the stock price. S&Z present a generalization of this approach allowing for an arbitrary correlation between the two processes. Both models, S&S and S&Z, suffer from the theoretical deficiency that the instantaneous volatility, since it follows a Gaussian process, becomes negative with strictly positive probability. The basis of the Heston model is also an Ornstein-Uhlenbeck process for instantaneous volatility. However, his model is formulated in terms of the instantaneous *variance* which can be shown to follow the square-root process used by Cox, Ingersoll, and Ross [6] to describe interest rates. As a consequence, there is a certain chance that implied instantaneous volatilities are negative for both S&S and S&Z, whereas we will never *observe* this phenomenon in the Heston model, since we compute the implied volatility as the (positive) square-root of the implied variance which is non-negative by definition. The standard B&S model with constant volatility of log returns serves as the usual benchmark. The models will be compared with respect to their in sample fit and their out of sample pricing and hedging performance. Furthermore, we conduct a simple misspecification test as suggested by Bakshi, Cao and

Chen [1].

The main contributions of the paper are twofold: We conduct the first empirical test of the generalized version of the S&S model developed by S&Z against the restricted alternative and against the similar Heston specification. In addition we propose a simulated annealing algorithm to estimate implied parameters, since conventional numerical optimization routines get trapped in local minima quite frequently. In consequence, results delivered by these methods may not be reliable.

The analyses in this paper show that the parameter estimates for stochastic volatility models exhibit considerable time series variation, i.e. it is likely that the hypothesis of constant structural parameters would have to be rejected for our sample. Even without running sophisticated tests for the time series consistency of the implied stochastic processes of the stock and of the instantaneous volatility it is obvious from a simple specification test that in most cases the more complex models cannot eliminate the empirical deficiencies of the B&S model.

However, these models are superior in terms of pricing performance to the standard B&S model both in and out of sample. Moreover, it seems safe to conclude that a non-zero correlation between equity returns and volatility changes is important for option pricing, since the implied estimates for this parameter are strictly negative for each individual observation day in our sample.

The results concerning the hedging performance of the respective models are slightly different. Stochastic volatility models are still superior to B&S in standard form in most situations, but a minor modification of the B&S hedge portfolio to take into account ‘vega

risk' yields hedging results that are almost indistinguishable from those of more complex models.

One purpose of this study was to investigate whether the S&Z approach or the Heston model is the dominant alternative among stochastic volatility models with arbitrary correlation between index returns and volatility changes. For all practical purposes the two models are found to be equivalent both in terms of pricing and hedging performance. Thus, although the two approaches are not completely equivalent from a theoretical point of view, their option prices and hedge ratios seem to be close enough to yield very similar empirical results.

The rest of the paper is organized as follows: In section 2 we give a brief description of the four option pricing models used in the empirical analysis. Section 3 contains details on the data and the simulated annealing algorithm used to estimate the model parameters. The results are presented in section 4, and some closing comments and conclusions are given in section 5.

2 The Models

This section briefly introduces the models which will later be compared empirically. B&S [4] assume that the stock price follows a geometric Brownian motion with constant instantaneous volatility v , i.e.

$$dS_t = rS_t dt + vS_t dW_t,$$

where dW_t is the increment of a standard Wiener process. Note that the dynamics of the stock are given under the risk-neutralized probability measure, i.e. under probabilities that make the discounted stock price a martingale.

In the Heston [8] model the dynamics of the underlying stock and its instantaneous volatility under the risk-neutralized measure are given by

$$\begin{aligned} dS_t &= r S_t dt + v_t S_t dW_{1t} \\ dv_t &= -\beta v_t dt + \sigma dW_{2t} \end{aligned} \tag{1}$$

where the increments of the Wiener processes W_{1t} and W_{2t} are correlated with constant correlation coefficient ρ_H , i.e. $dW_{1t} dW_{2t} = \rho_H dt$. This process for v_t implies through Itô's lemma that the instantaneous variance $y_t = v_t^2$ follows the process

$$dy_t = \kappa_H(\theta_H - y_t) dt + \sigma_H \sqrt{y_t} dW_{2t}, \tag{2}$$

where κ_H , θ_H , and σ_H are the speed of adjustment, the long-run mean, and the volatility of the instantaneous variance y_t , respectively. This is analogous to the square-root process with mean reversion for the short rate in the famous model by Cox, Ingersoll, and Ross [6].

The recently presented S&Z [12] model can be written as

$$\begin{aligned} dS_t &= r S_t dt + v_t S_t dW_{1t} \\ dv_t &= \kappa(\theta - v_t) dt + \sigma dW_{2t} \end{aligned} \tag{3}$$

where, again, the increments of the Wiener processes W_{1t} and W_{2t} exhibit a constant correlation ρ . As in the Heston model, v_t follows an Ornstein-Uhlenbeck process. Since the final option pricing formula is not written in terms of the variance, but in terms of volatility, this implies a strictly positive probability of observing negative implied instantaneous volatilities. The specification in (3) is also used by S&S [15] with the difference

that they impose the restriction $\rho = 0$, so that increments in W_1 and W_2 are uncorrelated. From an empirical point of view, a comparison between S&Z and S&S is thus equivalent to a test of the restriction $\rho = 0$. On the other hand, the Heston model and S&Z are not fully nested, i.e. for some parameter constellations prices from Heston models cannot be reproduced by S&Z, and vice versa.¹ Hence it remains an entirely empirical question which (if any) of the models is preferable.

3 Data and Estimation Methodology

3.1 Data

The data for this study consist of a time series of best bid and ask quotes for DAX options taken from tapes of Deutsche Terminbörse (DTB, now EUREX DEUTSCHLAND) recorded over the period from July 1 to December 31, 1996. We use the midpoint between these quotes as an estimate for the market price of the option. Put prices were converted to call prices via the standard put-call parity. This relationship between put and call prices is easily applicable here, since the DAX is a performance index with dividend reinvestment and all the options are European. The contemporaneous current level of the DAX index is taken from the tapes of the computerized trading system IBIS (now replaced by XETRA) which also provide the best bid and offer prices. The index price used in the analyses below is the midpoint between the best bid and ask quotes. To avoid systematic biases

¹In their paper S&Z describe how the parameters of their model must be set to obtain option prices from the Heston model. See their equation (22).

caused by auctions at the open or the close of the stock markets and around the noon auction on the floor of the Frankfurt Stock Exchange, only prices (for options and the DAX) recorded exactly at 11.30 a.m. were used for this study.

Observations for which the midpoint option price violated the lower boundary for European options given by the difference between current index level and present value of the strike price were eliminated from the sample. Furthermore, options with a time to maturity below ten days and above nine months were deleted from the sample due to a severe lack of liquidity. The final sample then consisted of 7,955 observations for the 126 trading days for the second half of 1996. The median of time to maturity of the options in the sample is 0.205 years, their median moneyness is 1.014, and their overall median relative spread is 5%. We compute the moneyness of an option as the ratio of the current index level to the present value of the strike price, the relative spread is given by the ratio of the absolute spread to the midpoint between bid and ask. We use the 3-month FIBOR rate as a proxy for the risk-free rate. Over the sample period the median for this rate was 3.18% p.a.

We group the options in the sample into 15 classes according to their moneyness and their time to maturity. An option was classified as short-term, if its time to maturity was less than two months, as medium-term, if time to maturity was no more than six months, and long-term otherwise. The intervals for the five moneyness categories are $[0,0.94)$, $[0.94,0.98)$, $[0.98,1.02)$, $[1.02,1.06)$, and $[1.06, \infty)$, respectively.

The characteristics of the sample with respect to price level, percentage spread and number of observations are shown in table 1. The observations are spread roughly evenly across

the moneyness groups, except for the lowest moneyness category which only contains 618 observations. The observed option prices reveal a similar pattern with respect to their time to maturity. The category containing long-term options has only 955 observations, compared to more than 3,000 for both short-term and medium-term options.

— Insert table 1 about here —

3.2 Estimation of the Model Parameters

The parameters of the four models considered in this paper were estimated implicitly by minimizing the sum of squared differences between market prices and model prices (*SSE*). For each observation day t ($t = 1, \dots, 126$) we estimate the parameter vector Ψ_t by $\hat{\Psi}_t$ with the property

$$\begin{aligned}\hat{\Psi}_t &= \arg \min_{\Psi_t} SSE_t \\ &= \arg \min_{\Psi_t} \sum_{i=1}^{N_t} \left[C_{it} - \hat{C}_{it}(S_t, r_t, K_{it}, \tau_{it}; \Psi_t) \right]^2,\end{aligned}$$

where S_t and r_t denote the level of the DAX index and the FIBOR rate at date t , respectively, and K_{it} and τ_{it} represent the strike price and the remaining time to maturity of option i on day t . C_{it} is the observed market price for option i on day t , and $\hat{C}_{it}(S_t, r_t, K_{it}, \tau_{it}; \Psi_t)$ is the theoretical price² given the parameter vector Ψ_t . For example, in the case of the classical B&S model, we have $\Psi_t = (v_t)$ where v_t denotes the implied volatility on day t , and $\hat{C}_{it}(S_t, r_t, K_{it}, \tau_{it}; \Psi_t)$ is given by their well-known formula. For the

²The formulas for the theoretical option prices for the three stochastic volatility models and for B&S are given in the appendix.

S&Z model we would have $\Psi_t = (v_t, \kappa_t, \sigma_t, \theta_t, \rho_t)$.

It turns out that the estimation of Ψ_t poses a serious numerical problem. Conventional numerical optimization routines came to a standstill in local minima very frequently. Some of the algorithms also showed the tendency to stay too close to the starting values used for the minimization process. To overcome these problems we used a simulated annealing algorithm to solve the above minimization problem for each day t . The main property of this method is that, in contrast to standard routines, it will always find the global minimum of the objective function, since it asymptotically searches over the whole parameter space. This implies that the result of the optimization process will be independent of the starting values for the parameters. The main innovative feature of simulated annealing is that the direction in which the parameter space is searched can be changed randomly which prevents the algorithm from getting stuck in local minima.³ The only potential drawback of this method is that it is more demanding in terms of computing time than standard procedures, a disadvantage that may become less and less severe given the advances in computer development.

³An introduction to simulated annealing plus some examples of the advantages of this method over conventional optimization techniques can be found in Goffe, Ferrier, and Rogers [7]. We compared simulated annealing to various numerical optimization routines contained in the NAG Fortran library. For a sample of simulated prices from the S&Z model only the simulated annealing method found the true parameter values, independent of the starting values.

4 Empirical Results

4.1 Parameter Estimates

As expected, models with a larger number of free parameters generally show a better fit to the data than the simple B&S model as indicated by the sum of squared errors (*SSE*) shown in table 2. There is a reduction of about one third in *SSE* from B&S to S&S, the simplest stochastic volatility model. However, the restriction of uncorrelated innovations in index level and volatility is binding, since the median *SSE* for Heston and S&Z are just one fourth of that observed for S&S. Indeed, we observe strongly negative implied values for ρ and ρ_H on all of the 126 days in the sample. Our estimates for ρ_H are comparable in magnitude to those found by Bakshi, Cao and Chen [1].

— Insert table 2 about here —

In contrast to their study the median B&S implied spot volatility v in our sample is substantially larger than the estimated instantaneous volatility (or variance, respectively) for the other models. This higher implied volatility can at least partially compensate the ability of stochastic volatility models to generate the empirically observed fat tails in the distribution of the price changes of the underlying. The fact that B&S implied volatilities are larger is thus merely a technical consequence of the severe restrictions that are imposed on this model.

It is also interesting to note that the option prices in our sample imply a considerable time series variation in volatility (variance) itself, since σ (σ_H) are rather large for all models.

This supports the hypothesis that uncertainty in volatility is directly incorporated into the option pricing process by market participants. The median estimates for σ are 0.2968 for S&S and 0.2575 for S&Z. In the Heston model, we obtain a median of 0.4355 for σ_H . The fact that the latter estimate is so much larger than those for the first two models can be explained by looking at the parametrization needed for Heston to obtain the prices generated by an equivalent S&Z model. As S&Z show in their paper (equation (22), p. 7,) the volatility of variance would have to be twice as high as the volatility of volatility to obtain the same prices from the two models.

4.2 Simple Checks for Misspecification

A first simple test for misspecification is to look at the variation of the parameter estimates. Since none of the coefficients in the models are time dependent, we should observe hardly any variation at all, if the models are correctly specified. In table 2 we show the 25 percent and 75 percent quantile of the estimates for the respective parameters over the 126 days in the sample.

It is obvious that the implied volatility estimates for B&S are all in a rather narrow band around the median with an interquartile range of just 0.017. For the stochastic volatility models instantaneous volatility is not a structural parameter but a second underlying, so that time series variation for this parameter is to be expected. It is more disturbing to see the strong variation in κ , σ , and θ for all three models. For example, for the S&Z model the speed of adjustment parameter κ has an interquartile range of around seven which is about 1.06 times its median. Similar numbers are obtained for the other models

and parameters. This results suggest that all three stochastic volatility models are still misspecified.

Another way to perform a misspecification test is to compute model specific instantaneous implied volatilities \tilde{v}_{it} for each option i on each day t . If the stochastic volatility models are specified correctly a graph of these instantaneous volatilities against the moneyness of the options should be a flat line.

— Insert figures 1 and 2 about here —

We estimate these implied volatilities in t using the current price of the underlying and the current interest rate as well as the other structural parameters from our estimation on day $t - 1$, i.e. we find \tilde{v}_{it} as the value for the instantaneous volatility (or variance) that makes the theoretical price equal to the observed midpoint. The results showing the average implied volatilities for the different levels of moneyness and the four models are presented in figures 1 (short-term options) and 2 (long-term options).

The message in these figures is very clear. Each of the four models still exhibits a significant smile or skew for short maturities, although the graph corresponding to the S&Z model flattens out for higher moneyness values. For long-term options the picture is similar. Again the stochastic volatility models cannot eliminate the smile, and the the graph produced by B&S is surprisingly flat. However, for values of M above 0.96 the S&Z model generates a graph almost parallel to the x -axis.

4.3 Out of Sample Pricing Performance

The pricing performance for the four models is measured by the absolute DM and absolute percentage pricing error *out of sample* (table 3). In addition, table 4 shows the percentage of cases in which the respective model overprices an option and the relative frequency of observations for which the theoretical price fell outside the bid-ask band.

— Insert table 3 about here —

The DM pricing error e_{it} is defined as the difference between the observable midpoint price and the theoretical price of an option given the estimated parameter vector from day $t - 1$, i.e.

$$e_{it} = C_{it} - \hat{C}_{it}(S_t, r_t, K_{it}, \tau_{it}; \hat{\Psi}_{t-1}). \quad (4)$$

The relative pricing error is then defined as e_{it}/C_{it} .⁴

Looking at absolute DM pricing errors in table 3 first, we find that models with more degrees of freedom exhibit a better overall pricing performance than the standard B&S model. For the complete sample, B&S produces a median absolute pricing error of DM 3.16, compared to DM 2.41 for S&S and DM 1.61 for both S&Z and Heston. Again, as in the case of the goodness of fit in the estimation stage, the introduction of a non-zero correlation between index returns and volatility changes improves the pricing performance

⁴It is common practice to use the complete lagged vector Ψ_{t-1} to perform an out of sample pricing test, see for example Bakshi, Cao and Chen [1, 2]. This is not entirely consistent with the notion of volatility as a second underlying for the option, since it is implicitly assumed to remain constant from $t - 1$ to t .

of stochastic volatility models significantly so that S&Z and Heston perform better than S&S.

The ranking of the four models becomes less clear if we look at individual moneyness-maturity groups. All four models perform best for short-term options deep in the money. The main reason for this result is that theoretical option prices in this class are almost model independent, since they approach the nonparametric lower boundary given by the difference between current index level and discounted strike price.

There are six moneyness-maturity classes for which S&S performs worse than B&S, especially for options deep out of the money. For example, for options with time to maturity greater than six months and moneyness below 0.94, the median absolute error for B&S is DM 6.78 in contrast to DM 10.12 for S&S. On the other hand, we observe this phenomenon only once for the more general S&Z and Heston models. For short-term, deep out of the money options with less than two months to maturity the two stochastic volatility models produce median absolute errors of DM 1.76 each, whereas here B&S has a median error of just DM 1.54.

The numbers in table 3 also provide evidence for a time to maturity effect in pricing performance. For example, for options in the money ($1.02 \leq M < 1.06$) the difference in median absolute DM pricing errors between B&S and S&Z increases from DM 0.16 for short-term options to DM 3.42 for long-term options. A similar pattern can be found for options with $M \geq 1.06$. This maturity effect also reverses the ranking of models for options with $M < 0.94$. As we had seen for short-term options in this group B&S performed even better than S&Z or Heston, but this is no longer true for options with at least six months

to maturity. Here S&Z and Heston generate a reduction in median absolute pricing errors of more than DM 2 compared to B&S.

For at the money options we observe a rather striking result: The relative performance of the B&S model compared to its alternatives is worst for short-term options, whereas all four models exhibit a similar pricing quality for long-term options. This is surprising, since it is commonly assumed that the B&S model is most applicable for short-term options with a strike around the current index level.

The results for the absolute percentage pricing errors regarding the ranking of the four models are qualitatively very similar to those obtained for DM errors. For the cheapest options (short-term, out of the money) we observe the largest percentage errors. Here, S&Z and Heston generate average relative errors close to 90 percent of the observed option price. On the other hand, for short term options deep in the money, all four models exhibit relative pricing errors of just 0.3 to 0.5 percent. In general, the relative error is strictly decreasing with increasing moneyness for all four models. A potential explanation for this result is that the estimation method implicitly gives higher weights to options with higher prices, so that if *relative* squared errors were minimized the picture might look very different.

— Insert table 4 about here —

The overpricing statistic shown in table 4 is interesting, since it gives an idea of whether pricing errors are systematic or (more or less) just noise. In the latter case we would expect overpricing to appear just as often as underpricing, i.e. in about 50 percent of the

cases.⁵

For maturities above two months and moneyness values below 0.94 we observe a very high percentage of cases with overpricing for both B&S and S&S compared to S&Z and Heston. For example, options with a maturity between two and six months are overpriced by B&S in 80.9 percent and by S&S in 90.6 percent of the sample whereas for S&Z and Heston this share is only around 45 percent. An intuitive explanation for this result is based on the negative (implied) correlation between index returns and volatility. For options far out of the money to have a positive payoff at maturity, index returns over the remaining life of the option must be large and positive. Given a negative correlation between returns and volatility changes this implies on average a decrease in volatility, so that these options will be cheaper in the Heston and S&Z models than for B&S or S&S. The same logic applies vice versa to options far in the money with a medium or long time to maturity. For the case $M \geq 1.06$ and a time to maturity between two months and half a year the percentage of overpricing for S&Z and Heston is very close to 50 percent for both models. In contrast to this, B&S overprices these options in just 3.9 percent of the cases. S&S exhibits a very systematic pattern in mispricing, since more than 80 percent of the options in this category are underpriced.

A statistic emphasising the generally poor performance of models without a correlation component is the number of categories for which the percentage of overpricing is not even

⁵Of course, this is based on the assumption that the probability of hitting the market price exactly is equal to zero.

in the range between 35 and 65 percent⁶. For B&S this occurs in fourteen, for S&S still in ten of the fifteen classes. Both S&Z and Heston produce these systematic patterns in only four cases, which shows that the additional flexibility given by a non-zero ρ is very valuable.

Similar to the results for absolute DM and percentage pricing errors the very high percentage of overpricing produced by the B&S model for short-term at-the-money options is again somewhat surprising. The worst performance for S&Z and Heston is again observed for short-term options far out of the money. Underpricing occurs in more than 90 percent of the observations for both models, an even higher frequency than those produced by B&S and S&S.

The existence of a bid-ask spread creates a no arbitrage band around the midpoint price of an option, so that a non-zero pricing error does not automatically indicate a potential arbitrage opportunity. It is even possible to argue that a theoretical price between bid and ask constitutes no mispricing at all. Therefore, table 4 reports the relative frequency of cases in which the theoretical price is either less than the bid or greater than the ask.

Even for the two most general models the chance of obtaining prices below the bid or above the ask is still around 45 percent for the total sample, and these numbers are even worse for B&S (62.8 percent) and S&S (57.5 percent). There is a systematic tendency

⁶Even for the category with the smallest number of observations (92), the probability of obtaining less than 35 percent or more than 65 percent of cases with overpricing is much less than 1 percent if the true probability for overpricing is one half. Since this argument implicitly assumes independence, it is certainly not quite rigorous, but it still gives a good idea of whether there is a systematic tendency in the direction of mispricing.

for this percentage to decline with increasing moneyness for all four models. We obtain theoretical prices outside the bounds for almost 90 percent of the options in the category with $M < 0.94$, whereas this occurs for just twelve to 30 percent when $M \geq 1.06$. A natural explanation for this result is that DM spreads tend to get wider the more expensive the options are. The median relative spreads in table 1 do not exhibit much variation over the five moneyness categories, so that DM spreads do increase with increasing option prices.

As a common characteristic for the variables related to pricing performance (except DM errors) we find that options out of the money with a very short time to maturity are priced worst by all four models. This may in part be due to a tick-size problem, since the model prices may be so low in this category that the midpoint price simply has to be far off in percentage terms. The problem becomes less severe when time to maturity increases, since option prices then increase as well.

We also observe a related time-to-maturity effect for options in the money, i.e. all those with $M \geq 1.02$. For longer maturities the performance of all the models deteriorates, and a potential explanation is again that for short maturities the choice of the pricing model is almost irrelevant. Any model will yield prices that are reasonably close to the European lower boundary, but this no longer holds when time to maturity increases. The dynamics of the underlying variables become more and more important, and prices from different models will be systematically different.

4.4 Hedging Performance

It is generally accepted that the hedging performance of a model is more important than its pricing quality since option traders are more worried about the changes in the value of their positions than the absolute value itself. The success or failure of a given model depends on its ability to properly describe the price changes of options given a change in the value of the underlyings.

— Insert table 5 about here —

We test the hedging performance of the four models using the same technique as Bakshi, Cao, and Chen [1]. Using parameters estimated on day $t - 1$ we compute on day t the sensitivities of the option price with respect to the risk factors, i.e. the usual ‘delta’ for the B&S formula and the partials $\partial C / \partial S_t$ and $\partial C / \partial v_t$ ($\partial C / \partial y_t$) for the stochastic volatility models.⁷ Next, a portfolio is formed which should exhibit no sensitivity to any of the risk factors and which should thus earn the risk-free rate over the next time interval. In the case of the stochastic volatility models the portfolio consists of one option short, and the underlying and another option (with different strike and/or different time to maturity) long.⁸

⁷The hedging test was also run using parameters estimated on day t to determine the hedge ratios for the same day. The results are almost indistinguishable from those reported above.

⁸If possible, we always use an option with the same time to maturity, but with a different strike price. If there is no such option we use an option of the next available maturity date. Although theoretically irrelevant the choice of the hedging instrument will almost certainly have an impact on the empirical results. We therefore also used options with the same strike, but a different maturity to form the hedge

For B&S the hedge portfolio just consists of a fraction of the underlying. As Baskhi, Cao, and Chen [1] point out the comparison between B&S and the other models with their two-instrument hedges is implicitly unfair against B&S, since the second instrument will always hedge part of what is called convexity or Γ -risk, i.e. the risk that the hedge ratio of the option with respect to the underlying index may change over time. Therefore we adopt the methodology of these authors and introduce an extended version of B&S which incorporates both a delta and a ‘vega’ component. Of course, given the assumptions of B&S this extension is inherently inconsistent, since there is no volatility risk in the model, but in practice it seems to be a widely accepted way of hedging option positions.

If the proceeds from the short sale of the option do not cover the costs of acquiring the hedge portfolio, the necessary funds are borrowed at the risk-free rate. The position is unwound completely on day $t + 1$ (which may be more than a day after t , e.g. if the position is held over a weekend.) The net wealth remaining after liquidation (positive or negative) is recorded as the DM hedging error, since under ideal conditions this wealth should always equal zero (neglecting a potential discretization error, since the position is held longer than just an infinitesimal time span.)⁹

The most important overall result is that for the purpose of hedging the choice of a model is almost irrelevant, as long as the hedge position contains two instruments. Extended B&S, S&S, S&Z, and Heston produce median DM hedging errors of 0.94, 0.94, 0.93, and

portfolio. The results shown here were the most favorable for all the models.

⁹A natural restriction for the options used in the hedging test is that observations for these contracts must be available both for day t and day $t + 1$. The total number of hedges is thus lower than the total number of observations in the sample (around 5,600).

0.92, respectively. The simple B&S model generates hedges that are considerably more unreliable, since its median DM hedging error for the complete sample is 1.36. Compatible with the results presented for pricing performance above the hedges of all four models work best for short-term options deep in the money, for which the option price and, consequently, also the hedge ratios become independent of the model being used. In this category we find the same ranking of models as in the overall sample, with B&S performing significantly worse than the other four approaches between which there are hardly any differences.

Again, this aggregate analysis tells only part of the story, since there are cases in which even the simplest model beats all or at least some of the others. Consider, for example, options far out of the money ($M < 0.94$) with short time to maturity. Here B&S has a median DM hedging error of 1.14, compared to 1.54, 1.41, 1.96, and 2.16 for the other four models. Simple B&S also performs best for long-term options in general. When time to maturity is longer than six months B&S exhibits the best overall hedging performance with a median DM error of 1.95 compared to around DM 3 for the stochastic volatility models. Note that in this case S&S performs better than the correlation models, too. On the other hand, for options in the money with a short or medium time to maturity the models with two hedging instruments clearly beat the simple B&S, which generates hedging errors about twice as large. For example, for $1.02 \leq M < 1.06$ and up to two months to maturity the median DM error is 1.10 for B&S, compared to 0.56, 0.56, 0.51, and 0.49 for the other four models. This result seems a little counterintuitive in that there should be no significant convexity risk in this category because the hedge portfolio

(approximately) consists just of one unit of the underlying.

The results for options out of the money ($M < 0.98$) are in general rather mixed. No model dominates the others, and it is interesting to note the large differences (with varying signs) between S&Z and Heston in these categories.

5 Summary

Theoretically, stochastic volatility models have the ability to mend the deficiencies of the B&S model. However, it has to be investigated empirically if they actually perform better than simple models. This study compares three stochastic volatility models and the standard B&S model using a sample of stock index option prices from the German market. It presents the first empirical test of the S&Z model which is a generalization of the S&S approach, augmenting the latter by introducing an arbitrary correlation between the stock and instantaneous volatility.

The parameters are estimated implicitly using a cross-section of option prices for every day in the sample. Due to numerical problems with standard optimization routines a simulated annealing algorithm is used which is much more likely to find the global optimum of the objective function. It turns out that the most important feature determining the pricing quality of stochastic volatility models is the correlation between index returns and volatility changes. The S&S with a zero correlation performs much worse than S&Z and Heston in terms of the average sum of squared errors in the estimation stage. The estimated correlation is strongly negative for both Heston and S&Z for all days in the

sample, indicating that a decrease in volatility is expected in bull markets, and vice versa. Still, the S&S model is much more flexible than B&S and easily outperforms this simple approach in terms of goodness of fit in sample.

Concerning pricing performance out of sample stochastic volatility models are again clearly superior to B&S, although there are some cases in which the simpler models perform better than more sophisticated approaches. This ranking is not reproduced in our hedging test, where an extended B&S model eliminating not only stock price, but also ‘vega’ risk, yields hedging results that are almost indistinguishable from those of stochastic volatility models.

In summary, the greater flexibility provided by stochastic volatility models clearly improves the pricing of index options on the German market. However, these models are not able to consistently outperform simpler approaches, especially when the models are compared with respect to their hedging performance. This might be considered as a slightly disappointing result at first sight. Nevertheless, these models still have the potential to perform better than B&S or related models. Improvements can be made in terms of parameter estimation. In a very recent paper Baskhi, Cao, and Chen [2] suggest to use the Method of Simulated Moments to estimate the parameters of stochastic volatility models. This approach yields just one estimate of the parameter vector for a time series of cross-sections of option prices and thereby eliminates the inherent inconsistency of daily changing parameter estimates. These authors still do not use sampling information from both the time series of *stock and option prices*, since they do not include the stock price in their set of moment conditions. Furthermore, as long as we do not know the distribution

of hedging errors, we can only compare two models in terms of the average errors they produce. There is a definite need for further research (both theoretical and empirical) to obtain more precise answers to these problems.

A Appendix

- Black and Scholes:

$$\begin{aligned}
C(S, K, \tau) &= SN(d_1) - Ke^{-r\tau}N(d_2) \\
d_1 &= \frac{\ln(S/K) + (r + 0.5v^2)\tau}{v\sqrt{\tau}} \\
d_2 &= d_1 - v\sqrt{\tau} \\
N(y) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^y \exp(-x^2/2) dx
\end{aligned}$$

- Heston:

$$\begin{aligned}
C(S, K, y, \tau) &= S_t F_1 - Ke^{-r\tau} F_2 \\
F_j &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left(\frac{f_j(\phi) \exp(-i\phi \ln K)}{i\phi} \right) d\phi \quad (j = 1, 2) \\
f_j(\phi) &= \exp[C_j(a, b_j, d_j, g_j, \tau, \phi) + D_j(b_j, d_j, \tau, \phi)y_t + i\phi \ln S_t] \quad (j = 1, 2) \\
C_j(a, b_j, d_j, g_j, \tau, \phi) &= r\phi i\tau \\
&\quad + \frac{a}{\sigma_H^2} \left[(b_j - \rho_H \sigma_H \phi i + d_j)\tau - 2 \ln \left(\frac{1 - g_j \exp(d_j \tau)}{1 - g_j} \right) \right] \quad (j = 1, 2) \\
D_j(b_j, d_j, \tau, \phi) &= \frac{b_j - \rho_H \sigma_H \phi i + d_j}{\sigma_H^2} \cdot \frac{1 - \exp(d_j \tau)}{1 - g_j \exp(d_j \tau)} \quad (j = 1, 2) \\
g_j &= \frac{b_j - \rho_H \sigma_H \phi i + d_j}{b_j - \rho_H \sigma_H \phi i - d_j} \quad (j = 1, 2) \\
d_j &= \sqrt{(\rho_H \sigma_H \phi i - b_j)^2 - \sigma_H^2 (2u_j \phi i - \phi^2)} \quad (j = 1, 2) \\
a &= \kappa_H \theta_H, \quad b_1 = \kappa_H - \rho_H \sigma_H, \quad b_2 = \kappa_H \\
u_1 &= 0.5, \quad u_2 = -0.5
\end{aligned}$$

- Schöbel and Zhu (Stein and Stein is obtained by setting $\rho = 0$):

$$\begin{aligned}
C(S, K, v, \tau) &= S_t F_1 - K e^{-r\tau} F_2 \\
F_j &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left(\frac{f_j(\phi) \exp(-i\phi \ln K)}{i\phi} \right) d\phi \quad (j = 1, 2) \\
f_1(\theta) &= \exp\{i\phi(r\tau + \ln S_t) - 0.5 \rho(1 + i\phi)[\sigma^{-1}v_t^2 + \sigma\tau]\} \\
&\quad \times \exp\{0.5 D(\tau, s_1, s_3)v_t^2 + B(\tau, s_1, s_2, s_3)v_t + C(\tau, s_1, s_2, s_3)\} \\
f_2(\theta) &= \exp\{i\phi(r\tau + \ln S_t) - 0.5 \rho i\phi[\sigma^{-1}v_t^2 + \sigma\tau]\} \\
&\quad \times \exp\{0.5 D(\tau, u_1, u_3)v_t^2 + B(\tau, u_1, u_2, u_3)v_t + C(\tau, u_1, u_2, u_3)\} \\
s_k &= h_k(1 + i\phi) \quad (k = 1, 2, 3) \\
u_k &= h_k(i\phi) \quad (k = 1, 2, 3) \\
h_1(x) &= -0.5 x^2(1 - \rho^2) + 0.5 x(1 - 2\kappa\rho\sigma^{-1}) \\
h_2(x) &= x\kappa\theta\rho\sigma^{-1} \\
h_3(x) &= 0.5 x\rho\sigma^{-1} \\
B(\tau, x, y, z) &= \frac{1}{\sigma^2\gamma_1} \left[\frac{(\kappa\theta\gamma_1 - \gamma_2\gamma_3) + \gamma_3[\sinh(\gamma_1\tau) + \gamma_2 \cosh(\gamma_1\tau)]}{\cosh(\gamma_1\tau) + \gamma_2 \sinh(\gamma_1\tau)} - \kappa\theta\gamma_1 \right] \\
C(\tau, x, y, z) &= -0.5 \ln [\cosh(\gamma_1\tau) + \gamma_2 \sinh(\gamma_1\tau)] + 0.5 \kappa\tau \\
&\quad + \frac{\kappa^2\theta^2\gamma_1^2 - \gamma_3^2}{2\sigma^2\gamma_1^3} \left[\frac{\sinh(\gamma_1\tau)}{\cosh(\gamma_1\tau) + \gamma_2 \sinh(\gamma_1\tau)} - \gamma_1\tau \right] \\
&\quad + \frac{(\kappa\theta\gamma_1 - \gamma_2\gamma_3)\gamma_3}{\sigma^2\gamma_1^3} \left[\frac{\cosh(\gamma_1\tau) - 1}{\cosh(\gamma_1\tau) + \gamma_2 \sinh(\gamma_1\tau)} \right] \\
D(\tau, x, y) &= \frac{1}{\sigma^2} \left[\kappa - \gamma_1 \frac{\sinh(\gamma_1\tau) + \gamma_2 \cosh(\gamma_1\tau)}{\cosh(\gamma_1\tau) + \gamma_2 \sinh(\gamma_1\tau)} \right] \\
\gamma_1 &= \sqrt{2\sigma^2x + \kappa^2}, \quad \gamma_2 = (\kappa - 2\sigma^2z)\gamma_1^{-1} \quad \gamma_3 = \kappa^2\theta - y\sigma^2
\end{aligned}$$

References

- [1] Gurdip Bakshi, Charles Cao, and Zhiwu Chen. Empirical performance of alternative option pricing models. *Journal of Finance*, 52:2003–2049, 1997.
- [2] Gurdip Bakshi, Charles Cao, and Zhiwu Chen. Pricing and hedging long-term options. Technical report, University of Maryland, Pennsylvania State University, and Ohio State University, 1998.
- [3] Simon Benninga and Joram Mayshar. Heterogeneity and option pricing. Technical report, Tel-Aviv University and Hebrew University, Jerusalem, 1997.
- [4] Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 3:637–654, 1973.
- [5] George M. Constantinides. Transaction costs and the volatility implied by option prices. Technical report, University of Chicago, 1998.
- [6] John C. Cox, Jonathan E. Ingersoll, Jr., and Stephen A. Ross. A theory of the term structure of interest rates. *Econometrica*, 53:385–407, 1985.
- [7] William L. Goffe, Gary D. Ferrier, and John Rogers. Global optimization of statistical functions with simulated annealing. *Journal of Econometrics*, 60:65–69, 1994.
- [8] Steven L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2):327–343, 1993.
- [9] John Hull and Alan White. The pricing of options on assets with stochastic volatilities. *Journal of Finance*, 42:281–300, 1987.
- [10] William R. Melick and Charles P. Thomas. Recovering an asset’s implied pdf from option prices: An application to crude oil during the gulf crisis. *Journal of Financial and Quantitative Analysis*, 32(1), 1997.
- [11] Mark Rubinstein. Implied binomial trees. *Journal of Finance*, 49:771–818, 1994.
- [12] Rainer Schöbel and Jianwei Zhu. Stochastic volatility with an Ornstein-Uhlenbeck process: An extension. Technical report, Eberhard-Karls-Universität Tübingen, 1998.
- [13] Louis O. Scott. Option pricing when the variance changes randomly: Theory, estimation and an application. *Journal of Financial and Quantitative Analysis*, 22:419–438, 1987.
- [14] Louis O. Scott. Pricing stock options in a jump-diffusion model with stochastic volatility and interest rates: Applications of Fourier inversion methods. *Mathematical Finance*, 7(4):413–426, 1997.
- [15] Elias M. Stein and Jeremy C. Stein. Stock price distributions with stochastic volatility: An analytic approach. *Review of Financial Studies*, 4(4):727–752, 1991.

- [16] James B. Wiggins. Option values under stochastic volatility. *Journal of Financial Economics*, 19:351–372, 1987.

Table 1:
Descriptive Statistics of the Sample

In each cell the entries are the median of the observed midpoint option prices, the median of the relative bid-ask spread (in brackets, computed as the difference between ask and bid, divided by the midpoint price), and the total number of observations (in square brackets). The option quotes and the contemporaneous index levels are recorded at 11.30 a.m. daily for the period from July 1 to December 31, 1996. The moneyness ratio M is measured as the current index level divided by the present value of the strike price.

		Time to maturity			All maturities
		< 2 months	2 - 6 months	> 6 months	
Moneyness	$M < 0.94$	2.28 (0.030) [248]	6.95 (0.052) [278]	29.58 (0.110) [92]	4.81 (0.042) [618]
	$0.94 \leq M < 0.98$	5.06 (0.051) [709]	23.45 (0.075) [865]	56.60 (0.071) [189]	15.60 (0.062) [1,763]
	$0.98 \leq M < 1.02$	34.68 (0.054) [757]	67.40 (0.054) [928]	102.42 (0.063) [179]	56.32 (0.055) [1,864]
	$1.02 \leq M < 1.06$	109.39 (0.040) [664]	133.22 (0.045) [835]	165.27 (0.053) [163]	128.40 (0.044) [1,662]
	$M \geq 1.06$	224.17 (0.030) [652]	243.85 (0.035) [1,064]	310.35 (0.036) [332]	246.61 (0.034) [2,048]
All moneyness groups		49.50 (0.044) [3,030]	93.52 (0.047) [3,970]	140.07 (0.054) [955]	83.35 (0.046) [7955]

Table 2:
Implied Parameter Estimates

The table shows the median of the parameter estimates for the four models, the 25 percent quantile (in brackets), and the 75 percent quantile (in square brackets). Each day t ($t = 1, \dots, 126$) the parameters are estimated using a simulated annealing algorithm to minimize the sum of squared differences between observed market prices and theoretical option values. The theoretical prices for the various models are given in the appendix. In the B&S, S&S, and S&Z models v represents the instantaneous volatility of the log returns of the underlying index, whereas in the Heston model y stands for the instantaneous variance. The parameters κ , θ , and σ , (κ_H , θ_H , σ_H) are, respectively, the speed of adjustment, the long-run mean, the volatility of the instantaneous volatility v_t (the instantaneous variance y_t). The correlation between log index returns and volatility (variance) changes is denoted by ρ (ρ_H). SSE is the sum of squared errors.

	v	κ	θ	σ	ρ	SSE
B&S	0.1213 (0.1152) [0.1321]					1,399.48 (1,053.85) [1,884.57]
S&S	0.0944 (0.0839) [0.1117]	3.8058 (2.0429) [6.8497]	0.0868 (0.0608) [0.1217]	0.2968 (0.2293) [0.3529]		899.54 (483.66) [1,327.23]
S&Z	0.0920 (0.0775) [0.1074]	6.1176 (2.4097) [9.4030]	0.1157 (0.0921) [0.1362]	0.2575 (0.1879) [0.3534]	-0.4936 (-0.5779) [-0.3954]	227.49 (111.03) [486.51]
	y	κ_H	θ_H	σ_H	ρ_H	SSE
Heston	0.0095 (0.0075) [0.0138]	7.1525 (4.1130) [11.0323]	0.0200 (0.0168) [0.0255]	0.4355 (0.3372) [0.5331]	-0.5747 (-0.6986) [-0.4688]	223.95 (114.36) [501.43]

Table 3:
Pricing Errors (Out of Sample)

The entries in the table are the median absolute DM pricing error and the median absolute percentage pricing error (in brackets). The absolute DM pricing error for option i on day t is defined as $|e_{it}|$ with $e_{it} = C_{it} - \hat{C}_{it}(S_t, r_t, K_{it}, \tau_{it}; \hat{\Psi}_{t-1})$, where $\hat{C}_{it}(S_t, r_t, K_{it}, \tau_{it}; \hat{\Psi}_{t-1})$ represents the theoretical price of the option using the implied parameter estimates from day $t - 1$. The absolute percentage pricing error is defined as $100 \cdot |e_{it}/C_{it}|$.

			Time to maturity							
			Model	< 2 months	2 - 6 months	≥ 6 months	All maturities			
Moneyness	$M < 0.94$	B&S	1.54	(76.3)	3.15	(40.7)	6.78	(23.3)	2.41	(48.3)
		S&S	1.56	(71.7)	4.87	(55.6)	10.12	(33.2)	2.82	(58.9)
		S&Z	1.76	(89.4)	1.91	(27.6)	4.44	(15.7)	1.94	(41.3)
		Heston	1.76	(89.4)	1.82	(25.4)	4.63	(17.4)	1.93	(40.5)
	$0.94 \leq M < 0.98$	B&S	2.61	(49.2)	5.13	(22.1)	5.20	(9.7)	4.14	(26.9)
		S&S	1.84	(33.8)	4.49	(19.4)	8.16	(14.0)	3.51	(22.6)
		S&Z	1.45	(31.8)	1.67	(7.2)	3.57	(6.9)	1.68	(11.4)
		Heston	1.56	(29.6)	1.60	(7.0)	4.08	(7.0)	1.69	(11.4)
	$0.98 \leq M < 1.02$	B&S	4.33	(12.6)	3.17	(4.4)	4.19	(4.2)	3.72	(6.4)
		S&S	1.59	(4.6)	2.32	(3.5)	4.00	(4.0)	2.03	(4.0)
		S&Z	1.47	(4.5)	1.68	(2.7)	4.00	(3.9)	1.69	(3.5)
		Heston	1.60	(4.8)	1.78	(2.8)	3.87	(4.0)	1.81	(3.5)
	$1.02 \leq M < 1.06$	B&S	1.34	(1.2)	3.37	(2.6)	7.45	(4.6)	2.44	(2.1)
		S&S	1.43	(1.4)	3.54	(2.7)	4.55	(2.9)	2.58	(2.1)
		S&Z	1.18	(1.1)	1.55	(1.2)	4.03	(2.2)	1.43	(1.2)
		Heston	1.21	(1.1)	1.53	(1.2)	3.38	(2.2)	1.47	(1.2)
	$M \geq 1.06$	B&S	1.10	(0.5)	3.86	(1.6)	6.49	(1.8)	2.66	(1.0)
		S&S	0.88	(0.4)	2.26	(0.9)	4.29	(1.2)	1.83	(0.7)
		S&Z	0.88	(0.3)	1.61	(0.6)	4.50	(1.4)	1.46	(0.6)
		Heston	0.81	(0.3)	1.58	(0.6)	4.21	(1.3)	1.42	(0.6)
	All moneyness groups	B&S	1.99	(5.2)	3.79	(3.6)	5.95	(4.2)	3.16	(3.9)
		S&S	1.40	(3.1)	3.12	(3.2)	5.12	(3.4)	2.41	(3.2)
		S&Z	1.26	(2.7)	1.66	(1.9)	4.09	(3.0)	1.61	(2.3)
		Heston	1.30	(2.8)	1.62	(1.9)	4.07	(2.7)	1.61	(2.2)

Table 4:
Frequency of Overpricing and Violations of Bid-Ask Bounds
(Out of Sample)

In each cell the entries are the relative frequency with which the four models overprice the options in the sample and the relative frequency with which the theoretical prices generated by the models fell outside the band between the observed bid and the observed ask prices (in brackets). The DM pricing error for option i on day t defined as $e_{it} = C_{it} - \hat{C}_{it}(S_t, r_t, K_{it}, \tau_{it}; \hat{\Psi}_{t-1})$, where $\hat{C}_{it}(S_t, r_t, K_{it}, \tau_{it}; \hat{\Psi}_{t-1})$ represents the theoretical price of the option using the implied parameter estimates from day $t-1$. A model overprices an option if the theoretical price is greater than the observed midpoint price, i.e. if $e_{it} < 0$. An observation is counted as outside the spread, if $\hat{C}_{it}(S_t, r_t, K_{it}, \tau_{it}; \hat{\Psi}_{t-1}) < C_{it}^b$ or $\hat{C}_{it}(S_t, r_t, K_{it}, \tau_{it}; \hat{\Psi}_{t-1}) > C_{it}^a$, where C_{it}^b (C_{it}^a) denotes the observed bid (ask) price for option i on day t .

			Time to maturity			All maturities
			< 2 months	2 - 6 months	≥ 6 months	
Moneyness	$M < 0.94$	B&S	27.8 (95.2)	80.9 (93.5)	77.2 (91.3)	59.1 (93.9)
		S&S	35.5 (98.0)	90.6 (96.8)	88.0 (94.6)	68.1 (96.9)
		S&Z	7.7 (94.4)	45.7 (85.3)	55.4 (78.3)	31.9 (87.9)
		Heston	5.2 (93.5)	45.0 (80.2)	51.1 (76.1)	29.9 (85.0)
	$0.94 \leq M < 0.98$	B&S	76.3 (97.2)	92.8 (92.6)	66.7 (81.0)	83.3 (93.2)
		S&S	69.0 (95.1)	93.2 (91.3)	86.2 (89.4)	82.7 (92.6)
		S&Z	27.6 (89.3)	53.5 (70.0)	69.3 (70.4)	44.8 (77.8)
		Heston	26.8 (88.6)	51.3 (67.4)	63.0 (72.0)	42.7 (76.4)
	$0.98 \leq M < 1.02$	B&S	92.2 (86.5)	66.9 (66.0)	31.3 (54.7)	73.8 (73.2)
		S&S	55.5 (69.8)	44.1 (59.1)	59.2 (56.4)	50.2 (63.1)
		S&Z	53.2 (67.0)	48.3 (48.9)	63.1 (55.3)	51.7 (56.8)
		Heston	61.0 (69.4)	53.5 (50.0)	63.1 (58.1)	57.5 (58.6)
	$1.02 \leq M < 1.06$	B&S	44.7 (27.4)	14.6 (55.2)	16.6 (68.1)	26.8 (45.4)
		S&S	19.0 (31.3)	6.1 (58.7)	32.5 (50.3)	13.8 (46.9)
		S&Z	53.6 (25.5)	47.4 (23.2)	61.4 (40.5)	51.3 (25.8)
		Heston	56.8 (24.5)	46.5 (22.9)	68.7 (46.0)	52.8 (25.8)
	$M \geq 1.06$	B&S	8.1 (10.6)	3.9 (41.2)	13.6 (44.3)	6.8 (31.9)
		S&S	16.4 (6.4)	19.5 (22.6)	40.1 (31.6)	21.8 (18.9)
		S&Z	42.2 (6.0)	54.1 (11.0)	65.4 (34.9)	52.1 (13.3)
		Heston	41.9 (4.1)	47.8 (10.0)	66.0 (38.6)	48.9 (12.7)
	All moneyness groups	B&S	54.7 (60.4)	45.7 (64.8)	34.0 (62.1)	47.7 (62.8)
		S&S	40.6 (55.9)	43.5 (58.9)	56.1 (57.0)	43.9 (57.5)
		S&Z	41.2 (52.2)	50.6 (40.5)	64.1 (50.9)	48.7 (46.2)
		Heston	43.4 (52.0)	49.4 (39.5)	63.9 (53.7)	48.9 (45.9)

Table 5:
Absolute DM Hedging Errors

The entries in the table represent the median absolute DM hedging error. On each day t a hedge portfolio is set up for an option sold short by going long the underlying index and, in the case of the extended B&S model and all the three stochastic volatility models, another option (with different strike price and/or different time to maturity). In the case of the standard B&S model only the underlying index is bought. The number of units for the components of the hedge portfolio are chosen such that over a time period of infinitesimal length, the portfolio should have no exposure to any risk factor. The option price sensitivities needed to compute these quantities are calculated using the parameters estimated on day $t - 1$. If the proceeds from the short sale do not cover the costs of the hedge portfolio the funds are borrowed at the risk-free rate observed on day t . On day $t + 1$ the complete portfolio is liquidated. The absolute DM hedging error is then defined as the absolute value of the DM amount that is left after liquidation.

			Time to maturity				
			Model	< 2 months	2 - 6 months	≥ 6 months	All maturities
Moneyness	$M < 0.94$	B&S	1.14	1.52	3.36	1.38	
		ext. B&S	1.54	1.17	2.64	1.45	
		S&S	1.41	1.11	2.78	1.39	
		S&Z	1.96	1.60	2.59	1.90	
		Heston	2.16	1.43	2.63	1.90	
	$0.94 \leq M < 0.98$	B&S	1.44	1.54	1.91	1.51	
		ext. B&S	1.36	1.24	2.47	1.37	
		S&S	1.40	1.28	3.15	1.40	
		S&Z	1.23	1.65	1.29	1.55	
		Heston	1.65	1.24	3.26	1.53	
	$0.98 \leq M < 1.02$	B&S	1.41	1.39	1.77	1.42	
		ext. B&S	0.81	0.95	2.50	0.93	
		S&S	0.83	1.04	2.65	0.98	
		S&Z	0.80	0.95	3.35	0.94	
		Heston	0.77	0.97	3.90	0.92	
	$1.02 \leq M < 1.06$	B&S	1.10	1.17	1.58	1.18	
		ext. B&S	0.56	0.54	1.78	0.58	
		S&S	0.56	0.57	1.91	0.59	
		S&Z	0.51	0.55	1.89	0.56	
		Heston	0.49	0.54	2.08	0.56	
	$M \geq 1.06$	B&S	1.02	1.41	2.37	1.33	
		ext. B&S	0.59	0.97	3.52	0.92	
		S&S	0.50	0.96	4.12	0.85	
		S&Z	0.49	0.86	3.71	0.74	
		Heston	0.49	0.84	4.00	0.73	
All moneyness groups	B&S	1.26	1.38	1.95	1.36		
	ext. B&S	0.82	0.92	2.78	0.94		
	S&S	0.81	0.96	3.10	0.94		
	S&Z	0.84	0.90	3.08	0.93		
	Heston	0.83	0.89	3.33	0.92		

Figure 1:

Implied Volatilities for Short-Term Options

The graphs in the figure show the implied volatilities for individual options as a function of their moneyness for the four models. Implied volatilities on day t are computed by equating the theoretical price based on the structural parameters estimated on day $t - 1$ (none for B&S). Short-term options are options with a time to maturity up to two months.




Figure 2:

Implied Volatilities for Long-Term Options

The graphs in the figure show the implied volatilities for individual options as a function of their moneyness for the four models. Implied volatilities on day t are computed by equating the theoretical price based on the structural parameters estimated on day $t - 1$ (none for B&S). Long-term options are options with a time to maturity between six and nine months.

