DOES MODEL FIT MATTER FOR HEDGING? EVIDENCE FROM FTSE 100 OPTIONS

CAROL ALEXANDER
ANDREAS KAECK*

This study implements a variety of different calibration methods applied to the Heston model and examines their effect on the performance of standard and minimum-variance hedging of vanilla options on the FTSE 100 index. Simple adjustments to the Black–Scholes–Merton model are used as a benchmark. Our empirical findings apply to delta, delta-gamma, or delta-vega hedging and they are robust to varying the option maturities and moneyness, and to different market regimes. On the methodological side, an efficient technique for simultaneous calibration to option price and implied volatility index data is introduced. © 2011 Wiley Periodicals, Inc. Jrl Fut Mark

1. INTRODUCTION

Research on hedging vanilla options has important practical applications. In particular, market makers set prices so that their net profit from the trade, after deducting hedging costs, has positive expectation. More accurate hedging allows these traders to reduce their bid–ask spread, and thus increase their volume of trade. There is no constraint that the model used for hedging liquid options should be the same as the model used for pricing illiquid options.

We thank an anonymous referee for excellent comments on an earlier draft of this study.

*Correspondence author, ICMA Centre, Henley Business School at the University of Reading, UK. Tel: +44 (0)118 378 8239, Fax: +44 (0)118 931 4741, e-mail: a.kaeck@icmacentre.rdg.ac.uk

Received December 2010; Accepted May 2011

■ *Carol Alexander and Andreas Kaeck are at the ICMA Centre, Henley Business School at the University of Reading, United Kingdom.*

The Journal of Futures Markets, Vol. 00, No. 00, 1–30 (2011)

Published online in Wiley Online Library (www.onlinelibrary.com)
DOI: 10.1002/fut.20537

Indeed, there is a large literature on option-pricing models that focuses exclusively on their relative hedging performance, in particular, on their accuracy when hedging standard European (vanilla) equity index options.

This study implements a variety of calibration methods for the stochastic volatility model of Heston (1993) with the aim of assessing their effect on the performance of standard and minimum-variance (MV) delta, delta-gamma, and delta-vega hedging of vanilla options. As a benchmark we use simple adjustments to the Black–Scholes–Merton (BSM) hedge ratios (Black & Scholes, 1973; Merton, 1973).

We use the Heston model because the research on its performance for hedging equity index options is already considerable, and robust results have been obtained using various indices and sample periods. Kim and Kim (2004) show that Heston's model outperforms other stochastic volatility models in-sample, out-of-sample, and for hedging, and Kim and Kim (2005) confirm this when jumps are added to the price process. Bakshi, Cao, and Chen (1997) use the Heston model to conclude that "*once stochastic volatility is modeled, adding the SI (stochastic interest rate) or the random jump feature does not enhance hedging performance any further.*" Alexander, Kaeck, and Nogueira (2009) find that the Heston model performs particularly well in a MV delta-gamma hedging exercise. The Heston model has also become a frequently used model in simulation studies that investigate hedging performance, e.g. Psychoyios and Skiadopoulos (2006), Poulsen, Schenk-Hoppé, and Ewald (2009), Branger, Krautheim, Schlag, and Seeger (2011).

Although the pleasing hedging performance of the Heston model, coupled with its ease of calibration, has established it as a standard stochastic volatility model for the investigation of hedging performance, there are two important issues. First, several studies report that the Heston dynamics are grossly misspecified, e.g. Eraker, Johannes, and Polson (2003), Jones (2003), Broadie, Chernov, and Johannes (2007). Shortcomings of the model are related to the existence of possible jumps in price and volatility. In fact, some authors advocate to leave the class of affine pricing models altogether.[1] An important consequence is that model misspecification can translate into inconsistencies between parameter estimates that are solely based on time series of underlying prices and parameter estimates that are inferred from option quotes.[2] Parameter estimates based on the underlying price process are often regarded as inferior, because they are based on historical information rather than on

[1]For example, Christoffersen, Jacobs, and Mimouni (2010) report that non-affine models are superior for both in- and out-of-sample statistical tests.
[2]See, for example, Bakshi et al. (1997). Also, Broadi et al. (2007) give an illustrative example of how unrealistic the parameters of a misspecified model can be, when calibrated to a cross section of option prices (their Figure 3, p. 1464).

trader's beliefs about the evolution of the underlying price over the life of the option.[3] Altogether, the fact that different calibration techniques can lead to very different parameter values for the Heston model leads one to the important question of whether the data source for the model calibration matters.

Second, most of the hedging literature relies on daily calibration of the Heston model to a cross section of option prices. Although common, this approach is inconsistent, because within the model all parameters are assumed to be time-homogeneous. Thus, the model that has usually been tested in the empirical hedging literature is not the Heston model itself, but an extension where calibrated parameters are allowed to vary over time. A correct implementation of the model should separate out an estimation period and use one (constant) set of parameters throughout the out-of-sample hedging period.[4] Alternatively, if the calibrated parameters are systematically related to the underlying price process, adjustments of hedge ratios are necessary (see Alexander et al., 2009).[5]

Our work represents a new departure in empirical hedging studies in the literature. We focus not on the hedging model but on the model calibration technique itself. We provide evidence on the relative importance, for the purpose of option hedging, of the data source (options or underlying prices) and the quality of fit. Considering various calibration approaches, with daily and pre-sample calibration to option prices, we compare the hedging performance with that obtained when parameters are inferred purely from the underlying index. Our empirical results are very comprehensive: they cover a large sample (from 2 January 2002 to 31 December 2008) of delta, delta-gamma, and delta-vega hedging, using both standard and MV approaches. On the methodological side, we introduce a procedure to calibrate option pricing models simultaneously to option and implied volatility index data, which is very efficient because it considerably reduces the dimensionality of the optimization problem.

We proceed as follows: Section 2 presents an in-depth discussion of the Heston model and the simple adjustments to the BSM model that are used as benchmarks; Section 3 describes the three different calibration techniques applied to the Heston model; Section 4 defines the hedging strategies; Section 5 presents the empirical results; and Section 6 concludes.

[3]One could also infer parameter values from options prices and underlying prices simultaneously, as in Chernov and Ghysels (2000), Pan (2002), Eraker (2004), and others. However, such approaches are often extremely time consuming and are difficult to handle for a very large data set such as ours.
[4]Of course, such an approach compromises the fit of the pricing model. It may also deteriorate hedging performance; indeed one of our aims is to quantify how the daily recalibration of parameters improves the hedging performance of the Heston model.
[5]We do not include these hedge ratios in our study because each calibration requires a large number of integrals to be evaluated using Fourier transform methods, which is computationally infeasible on a large sample.

2. MODELS

2.1. Heston Model

The Heston model assumes that on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{Q}, (\mathcal{F}_t)_{t \in \mathbb{R}})$, where $\mathbb{Q}$ denotes the objective (real world) probability measure, the equity index $S(t)$ and its variance $V(t)$ solve

$$dS(t) = \mu S(t)\, dt + \sqrt{V(t)} S(t)\, dW^s(t), \tag{1}$$

$$dV(t) = [\theta - \kappa V(t)]\, dt + \xi \sqrt{V(t)}\, dW^v(t), \tag{2}$$

where $W^s(t)$ and $W^v(t)$ are two correlated Brownian motions with $E[W^s(t)W^v(t)] = \rho t$, μ is the constant drift of the index, κ the mean reversion speed of the variance, θ/κ is the long-term variance, and ξ is commonly referred to as the volatility-of-volatility (vol-of-vol).[6] We collect all structural parameters in the vector $\Phi = \{\mu, \kappa, \theta, \xi, \rho\}$.

This model is able to produce many stylized facts reported for equity indices. First, it accommodates increasing excess kurtosis in returns via increasing ξ and $|\rho|$. This feature is important, because it produces implied volatility graphs that have the well-known smile pattern. Also, a negative correlation between the underlying and its variance creates negative skewness in returns, an effect that is often termed the *leverage effect* (Black, 1976). It implies that falling prices tend to be accompanied by increasing variance and it leads to skewed or "smirk"-shaped implied volatilities with respect to strike. In addition, the model allows variance to tend to a long-term value. Thus, variance can fluctuate over time, but it will never wander away too much from its equilibrium value to which it is pulled back in the long run. This results in more realistic term structure behavior of volatility compared to simpler drift specifications. The model is, however, not completely consistent with the time-series of most equity indices. The main drawback is that sample paths generated by the Heston model are continuous and no sudden jumps in the price process can be created.

Although econometric research on the data generating process of equity indices favors models with jumps (see, for example, Eraker et al., 2003; Ignatieva, Rodriques, and Seeger, 2009), the empirical hedging literature finds little evidence that including jumps improves the hedging performance. Bakshi et al. (1997) argue that the likelihood of jumps (empirical estimates for equity indices indicate about 0.5 to two jumps per year on average) is often too low to

[6]The evolution of the variance can also be written as $dV(t) = \kappa[m - V(t)]\, dt + \xi\sqrt{V(t)}\, dW^v(t)$, thus explicitly providing a parameter (here $\sqrt{m}$) that accounts for the long-run value of volatility. We use the alternative specification because it simplifies notation when including risk premia.

have a discernible effect on out-of-sample option hedging exercises. This finding is confirmed by Branger, Krautheim, Schlag, and Seeger (2011) who use a simulation environment on hedging vanilla options to show that it is far more important to account for stochastic volatility than to include price jumps.

One of the main reasons why the Heston model has become popular is its ability to express prices of standard vanilla options in nearly closed form. This requires the transition from the objective probability measure $\mathbb{Q}$ to the risk-neutral measure $\tilde{\mathbb{Q}}$. For a linear variance risk premium, the structure of the model remains identical after a change of measure:

$$dS(t) = rS(t)\,dt + \sqrt{V(t)}S(t)\,d\tilde{W}^s(t)$$

$$dV(t) = [\theta - \tilde{\kappa}V(t)]\,dt + \xi\sqrt{V(t)}\,d\tilde{W}^v(t),$$

where $\tilde{W}^s(t)$ and $\tilde{W}^v(t)$ are two Brownian motions under $\tilde{\mathbb{Q}}$, still having correlation ρ, and r denotes the risk-free interest rate. Structural parameters are now $\Phi^{rn} = \{\tilde{\kappa}, \theta, \xi, \rho\}$ where the variance risk premium is $\lambda = \kappa - \tilde{\kappa}$. Absence of arbitrage requires ξ, and ρ to be the same under both measures, thus no matter whether parameters are inferred from the evolution of the equity index (and thus under the objective measure $\mathbb{Q}$) or from its option prices (thus under the risk-neutral measure $\tilde{\mathbb{Q}}$), parameters should be consistent.

Vanilla option prices can be calculated by inverting the characteristic function as shown in Heston (1993), Carr and Madan (1999), Lewis (2000), or Duffie, Pan, and Singleton (2000). Adopting the methodology from Carr and Madan (1999), this leads to a time-t vanilla call price (with maturity T, residual time to maturity $\tau = T - t$, and strike K) given by the following integral, which can be evaluated by standard numerical integration methods or the fast Fourier transform:

$$C(t, T, K) = \frac{\exp(-\alpha \log K)}{\pi} \int_0^\infty e^{-iv\log K} \frac{e^{-r\tau} f(v - (\alpha + 1)i, t, T)}{\alpha^2 + \alpha - v^2 + i(2\alpha + 1)v}\,dv,$$

where α denotes a dampening factor (such that the $(\alpha + 1)$st moment of log $S(T)$ exists) and $f(u, t, T) = \tilde{\mathbb{E}}[e^{iu\log S(T)} | \mathcal{F}_t]$ is the characteristic function of the log stock price under the risk-neutral measure. This is given by:

$$\begin{aligned} f(u, t, T) = {} & \exp\{iu(\log S(t) + r(T - t))\} \\ & \times \exp\left\{\frac{\theta}{\xi^2}\left[(\tilde{\kappa} - \rho\xi iu - d)(T - t) - 2\log\left(\frac{1 - ge^{-d(T-t)}}{1 - g}\right)\right]\right\} \\ & \times \exp\left\{\frac{(\tilde{\kappa} - \rho\xi iu - d)(1 - e^{-d(T-t)})}{\xi^2(1 - ge^{-d(T-t)})} V(t)\right\}, \end{aligned}$$

where

$$d = \sqrt{(\rho\xi ui - \tilde{\kappa})^2 + \xi^2(iu + u^2)},$$

$$g = \frac{\tilde{\kappa} - \rho\xi iu - d}{\tilde{\kappa} - \rho\xi iu + d},$$

and $i = \sqrt{-1}$. To obtain the characteristic function under the statistical measure $\mathbb{Q}$, we just need to replace the risk-free rate r by μ and $\tilde{\kappa}$ by its counterpart κ.

2.2. Smile Adjustments and Sticky Models

We use three adjustments to the BSM model as benchmarks against which to compare the performance of the Heston model. This helps to understand the significance of the calibration effect on the Heston model's hedging performance. Moreover, we contribute to the discussion on whether stochastic volatility models can outperform simple hedging models, which do not specify explicit dynamics for the underlying price. This question has been addressed by several authors; for example, by Dumas, Fleming, and Whaley (1998), Vähämaa (2004), and Kim (2009). So far, empirical studies have been limited to relatively short samples and results have been inconclusive.

Smile-adjusted hedge ratios are based on the empirical observation that the implied volatility surface $IV(m|F)$, which is a function of moneyness $m = \log(K/F)/\sqrt{\tau}$, may have a dynamic dependence on the underlying price.[7] Here we use F, the index futures price with the same residual time to maturity τ as the option, for the underlying variable as it is F rather than the spot price S that is the hedging instrument for an equity index option. Applying the chain rule to differentiate the price of a vanilla option with moneyness m with respect to F, we obtain:

$$\delta_{adj}(m|F) = \delta_{BSM}(m|F) + \nu_{BSM}(m|F)\, IV_F(m|F). \tag{3}$$

In other words, to account for possible movements in the implied volatility surface when the underlying price moves, the option's implied BSM delta, δ_{BSM}, should be adjusted by an amount that depends on the option vega, ν_{BSM}, and the implied volatility-price sensitivity, IV_F.[8]

[7]One could alternatively use a moneyness definition that is independent of the option maturity and use the maturity of the option as an additional parameter. We have found that the hedging results for this formulation is virtually unchanged to the results reported here and therefore we use the simpler maturity-dependent moneyness definition.

[8]We follow Engelmann, Fenglet, and Schwendner (2006) by defining the vega in all sticky models as the implied BSM vega (i.e. the vega based on the option's implied volatility), so it represents a parallel shift in the volatility surface.

Differentiating again with respect to F yields the smile-adjusted gamma:

$$\gamma_{adj}(m|F) = \gamma_{BSM}(m|F) + \nu_{BSM}(m|F)\, IV_{FF}(m|F) + 2\eta_{BSM}(m|F)\, IV_F(m|F) + \kappa_{BSM}(m|F)\, IV_F(m|F)^2,$$

where η_{BSM} is the implied BSM vanna (cross derivative of the option price with respect to price and volatility) and κ_{BSM} is the implied BSM volga (second derivative of the option price with respect to volatility).

Two such approaches have been termed "sticky-strike" (SS) and "sticky-moneyness" (SM) (see Derman, 1999). Here the underlying price process is a geometric Brownian motion, but the constant volatility of that process depends on the characteristics of the option that is being priced. In the SS model, the volatility depends on the strike of the option and in the SM model the volatility depends on the option's moneyness. Since there is only one underlying price for all options, these models use deliberately inconsistent assumptions about the price process. For instance, if we were to use a binomial tree to price an option under the SM assumption, we would need to use many trees with different constant volatilities, jumping between them as the option's moneyness changes with the movement in F over the life of the option. Nevertheless, they offer a pragmatic approach to hedging and for this purpose have become very popular with option traders.

The effect of the SS assumption is that $IV_F = 0$, i.e. the implied volatility surface, does not move with F. Hence, SS price hedge ratios are just the implied BSM hedge ratios, i.e. hedge ratios that are computed using the BSM formula with an option-specific implied volatility. This is in stark contrast to the original BSM model where hedge ratios for all options are calculated with a unique volatility parameter.

Under the SM assumption the volatility of the GBM price process depends on its moneyness rather than its strike. Since an underlying price move now changes the implied volatility of an option with a given strike, the model is sometimes referred to as the "floating smile" model. The SM delta has

$$IV_F(m|F) = -\frac{1}{F\sqrt{\tau}}\, IV_m(m|F)$$

in (3). To compute IV_m we parameterize the implied volatility to be a cubic function of moneyness.[9] That is, at a given time t

$$IV(m, \alpha) = \alpha_0 + \alpha_1 m + \alpha_2 m^2 + \alpha_3 m^3, \tag{5}$$

[9]We also tested more general implied volatility-moneyness specifications that allow for additional dependence on the time-to-maturity of the option and/or mixed terms including both moneyness and time-to-maturity, but found that a simple cubic polynomial works as well as more complex specifications.

where the parameter vector α is time dependent (we have suppressed this in our notation, for brevity).[10] Then, IV_m and higher order derivatives in the smile-adjusted delta and gamma formulae (see Equations (3) and (4)) are simple derivatives of (5). Using these, the SM delta may be calculated directly.

In our empirical results we employ a more direct standard central difference approach. For instance, for the option delta we shift the futures price and the implied volatility consistently:

$$\delta_{SM} = \frac{BSM(F_{up}, K, \tau, IV(m_{up}, \alpha)) - BSM(F_{down}, K, \tau, IV(m_{down}, \alpha))}{F_{up} - F_{down}},$$

where $m_{up/down} = \log(K/F_{up/down})/\sqrt{\tau}$ and BSM denotes the BSM option price. The same central difference methodology can be applied to obtain a gamma hedge ratio.[11]

A third, but more complex smile-adjusted hedge, is labeled "sticky-tree" (ST) or "local volatility". This approach goes back to Derman and Kani (1994), Dupire (1994), and Rubinstein (1994). There is a wide literature on the implementation of local volatility surfaces, which often includes the numerical solution of the Dupire partial differential equation. However, Derman, Kani, and Zou (1996) and Coleman, Kim, Li, and Verma (2001) provide theoretical justifications to approximate the local volatility delta hedge ratio by setting

$$IV_F(m|F) = \frac{1}{K\sqrt{\tau}} IV_m(m|F)$$

in (3). Similarly, the gamma follows from (4) on differentiating the above with respect to F. For these hedge ratios, partial derivatives of the smile are taken from the interpolated volatility surface in Equation (5).

3. MODEL CALIBRATION

3.1. Local vs. Global Calibration

Here we describe the two main approaches to stochastic volatility model calibration that have been adopted in the empirical literature. Daily calibration to the cross section of option prices, where parameters and the latent state variable are simultaneously inferred from quotes on one trading day, consists of optimizing an objective function such as

[10]The parameters in the vector α are calibrated by minimizing the objective function (6) below. This guarantees consistency in the choice of the calibration objective function for all models considered in this study.
[11]A similar approach, but based on the local volatility surface, has been adopted by Engelmann, Fengler, and Schwendner (2006).

$$(\widehat{\Phi_t^{rn}}, \hat{V}(t)) = \arg\min \frac{1}{N_t} \sum_{i=1}^{N_t} [P^{mo}(t, K_i, T_i, \Phi_t^{rn}, V(t)) - P^{ma}(t, K_i, T_i)]^2, \tag{6}$$

where P^{mo} and P^{ma} denote the model and market prices of the options in the calibration set, $T_i(K_i)$ is the maturity (strike) of the ith option, N_t refers to the number of options used for the calibration at time t, and trading days in the calibration sample are denoted $t = 1, \ldots, H$. Repeated application of this calibration technique to panel data on option prices yields a time-series for each of the structural parameters $\widehat{\Phi_t^{rn}}$, and for the latent state variable $\hat{V}(t)$.

A detailed discussion on the choice of calibration objective can be found in Christoffersen and Jacobs (2004). Instead of using an absolute price metric, one could employ a relative price metric, or minimize the squared implied volatility errors. Weights that emphasize the errors on the most liquid options could also be applied. We opt for (6) to facilitate the use of standard deviation as a meaningful and consistent out-of-sample objective for our hedging experiment. Also, this approach has been adopted, for example, in Bakshi et al. (1997), Kim and Kim (2005), and Alexander and Nogueir (2007) and it is well-known that it produces a reasonable fit to option prices on any given day. In the following we refer to this calibration approach as local calibration.

However, the local calibration approach violates one of the basic assumptions of the model: that the structural parameters are time-homogeneous. If calibrations are stable over time this might not be a major concern, as minor variations can be attributed to price discreteness and other microstructure effects. However, this is rarely the case for equity indices. We, therefore, consider an alternative, which circumvents the problem of time-varying parameters, where we reserve a pre-sample period for estimating the latent state variable and one set of structural parameters.

We refer to this approach as global calibration. Here, the equivalent (unweighted price metric) optimization problem to (6) becomes:

$$(\widehat{\Phi_t^{rn}}, \{\hat{V}(t)\}_{t=1}^{H}) = \arg\min \frac{1}{\sum N_t} \sum_{t=1}^{H} \sum_{i=1}^{N_t} [P^{mo}(t, K_i, T_i, \Phi^{rn}, V(t)) - P^{ma}(t, K_i, T_i)]^2.$$

Several optimization techniques for global calibration have been adopted, among these the simulated method of moments (adopted, for example, in Bakshi et al., 2000) or the bootstrap procedure in Broadie et al. (2007). One of the main challenges to address with this optimization is the dimensionality of the problem, as we need to infer not only the structural parameters but also the spot variance over the whole estimation period. In addition, if the model is to be applied out-of-sample, a filtering algorithm needs to be implemented to extract the unobservable variance on the out-of-sample days.

We now introduce a new approach to global calibration that solves this dimensionality problem by utilizing some well-known results from the model-free volatility literature that link the (unobservable) state variable $V(t)$ with the (observable) underlying volatility index (VI) having some constant time to maturity τ_v. Applying the results of Britten-Jones and Neuberger (2000) to the Heston model, one obtains:

$$VI^2(t, \tau_v) = \frac{\theta}{\tilde{\kappa}}\left(1 - \frac{1 - e^{-\tilde{\kappa}\tau_v}}{\tilde{\kappa}\tau_v}\right) + \frac{1 - e^{-\tilde{\kappa}\tau_v}}{\tilde{\kappa}\tau_v} V(t). \tag{7}$$

This closed-form relation between volatility indices and the spot variance can be used in the calibration, and circumvents the need to estimate the spot variance directly. In addition, there is no need to implement a filtering algorithm for the out-of-sample period, because spot volatilities become directly observable once the structural parameters are known.[12]

Inverting (7) yields the following function for spot variance, in terms of the volatility index:

$$V(t) = \frac{e^{\tilde{\kappa}\tau_v} VI^2(t, \tau_v)\tilde{\kappa}^2\tau_v + \theta(e^{\tilde{\kappa}\tau_v}(1 - \tilde{\kappa}\tau_v) - 1)}{(e^{\tilde{\kappa}\tau_v} - 1)\tilde{\kappa}}. \tag{8}$$

The dimensionality of the optimization can now be reduced by using the volatility index directly in the optimization routine, thus using Equation (8) we obtain a much simpler optimization problem as follows:

$$\begin{aligned}\widehat{\Phi^{rn}} = \arg\min \frac{1}{\sum N_t} \sum_{t=1}^{H} \sum_{i=1}^{N_t} [P^{mo}(t, K_i, T_i, \Phi^{rn}, V(t, \Phi^{rn}, VI(t)))\\ - P^{ma}(t, K_i, T_i)]^2.\end{aligned} \tag{9}$$

3.2. Time-Series Consistency

A concern that is common to both the local and the global calibration approach is that they ignore the information from the time-series of the underlying. It is often argued that calibrating to option prices has the advantage of being a forward-looking strategy. However, when a pure calibration approach is adopted, the implied parameters are usually far away from their time-series counterparts

[12]It has to be pointed out that the use of this calibration technique relies on the availability of a volatility index. For most major equity indices such data are provided by option exchanges such as the CBOE (Chicago Board Options Exchange). If no volatility index is available, as for single stocks, one could use European options to construct a time-series before calibration (see Jiang & Tian, 2005). Alternative approaches that circumvent the direct estimation of the latent state variable include the two-step procedure of Christoffersen et al. (2010), which alternates between an optimization over the structural parameters and a filtering algorithm for the unobserved variance.

(Bakshi et al., 1997). This is problematic since for $\mathbb{Q}$ and $\tilde{\mathbb{Q}}$ to be absolutely continuous only variance drift parameters are allowed to change across measures.

To reduce the effect of model misspecification on risk premia estimates, this observation motivated Broadie et al. (2007) to implement a two-stage calibration methodology. First they fix parameters that are restricted across measures to the time-series estimates of Eraker et al. (2003) and then they calibrate the remaining parameters and the spot volatility to option prices.[13] Since a change of measure for pure jump processes is very flexible and requires no restriction on parameters, in a jump-diffusion framework there are still enough degrees of freedom to obtain a good fit to option prices even after fixing some parameters to their time-series values. However, in the Heston model we are left with only one free parameter. For this reason, we adopt a more pragmatic procedure: First we estimate the parameters in Φ from the time-series of FTSE 100 returns, and then we translate these estimates into risk-neutral parameters by making assumptions on the variance risk premium λ. This approach has also been adopted by Dotsis, Psychoyios, and Skiadopoulos (2007) for the VIX index. We have experimented with different values for the risk premia as will be described in more detail below, but choose to report results only for $\lambda = 2$.

Estimating the parameters of the Heston model from return observations of the FTSE index is not straightforward, because of the latent nature of volatility. Recently, Bayesian methods have become popular, mainly because of their advantages when it comes to estimating unobserved state variables such as jumps, jump sizes, or volatility. Jacquier, Polson, and Rossi (1994) pioneer Monte Carlo Markov Chain (MCMC) algorithms to estimate discrete time stochastic volatility models and Eraker et al. (2003) adopt a similar procedure to continuous-time finance models (including the Heston model). The authors show that MCMC applied to a time-discretized version of Equation (1) and (2) yields accurate inference about the parameters.

MCMC methods require the derivation of complete conditional distributions. Using a time-discretized version of the Heston process it is straightforward to derive these for μ, κ, and θ. For the correlation ρ and the volatility-of-volatility parameter ξ, two different algorithms are used in the literature. Jacquier, Polson, and Rossi (2004) reparameterize the variance process, which leads to known distributions for both parameters, whereas the methodology in Eraker et al. (2003) requires a Metropolis step to update the correlation. We have experimented with both algorithms but for brevity only reporting results based on the method of Jacquier et al. (2004), which circumvents the use of a Metropolis algorithm. Complete conditional distributions for the variance lead

[13]It should be noted that fixing some parameters to reasonable values before calibration is quite common (see Bates, 2000; Broadie et al., 2007 and others.)

to densities of unknown form, which can be sampled by a hybrid Metropolis-within-Gibbs step. For more details on the algorithm for the Heston model and the Metropolis step (and MCMC in general) we refer the interested reader to Eraker et al. (2003) or Johannes and Polson (2006). To start the procedure, it remains to fix the prior information. We use similar prior distributions and parameters as Eraker et al. (2003). There is, in general, little information in the prior distributions and hence we put as much weight as possible to the data. More specific details on our implementation of this algorithm, including the prior information, are given in the appendix.

4. HEDGING STRATEGIES

As usual, we define the standard delta and gamma of an option to be the first and second partial derivatives of its model price with respect to the underlying asset price. In a complete market with no frictions, continuously rebalanced delta hedging should remove all risk from the option position. But in practice transaction costs and trading limitations render continuous rebalancing impossible and one has to resort to using discrete hedging intervals. We shall rebalance daily, and so the resulting residual risk should be reduced by the addition of a gamma or vega hedge, which requires an investment in a second option on the same underlying. When the option is priced according to an underlying price process with constant volatility, the option vega is the first partial derivative of the option price with respect to this volatility. But in models which render the market incomplete, a standard delta hedge would leave residual risk even under continuous rebalancing. In the Heston model, where the residual risk arises because of stochastic volatility, a delta hedge has to be complemented by a vega hedge to remove risks completely. Under the Heston model, which has a stochastic volatility, we define the option vega to be the first derivative of the Heston model price with respect to the spot volatility.

Alternative tractable hedging strategies have been introduced that minimize the variance of the local hedging error. These hedge ratios have been labeled minimum variance (MV) or locally risk minimizing hedge; they have been studied, for example, in Bakshi et al. (1997) and Poulsen et al. (2009). In stochastic volatility models this hedge ratio takes into account all the risks arising from the movements in the Brownian motion $W^s(t)$ and also partially hedges risks due to a stochastic volatility that is correlated with $W^s(t)$. Minimum-variance gamma hedging has been proposed for stochastic volatility models in order to partly account for gamma risk that arises from the volatility process. Therefore, in our empirical hedging study we use MV delta and delta-gamma hedges as well as standard delta, delta-gamma and delta-vega hedging strategies. For a derivation of MV delta and gamma for the Heston model, see Alexander and Nogueira (2007).

To fix notation assume that a hedge is rebalanced at discrete intervals of length Δt. In a delta hedging strategy, at every point in time t, the short position in a call or put option (with price here denoted simply by O) is complemented by an investment in $X_S(t)$ shares of the underlying, where $X_S(t)$ is equal to the delta hedge ratio of the pricing model. As a hedging instrument we employ the futures with the maturity T' closest to the maturity T of the option (and in most cases, $T' = T$). Thus, we translate the sensitivity $X_S(t)$ into $X_{F_{T'}}(t)$, where $F_{T'}(t)$ denotes the price of the closest futures contract. This requires the calculation of hedge ratios with respect to the hedge instrument, which can be easily derived from chain rule of calculus. Furthermore, we discount the gain/loss in the futures position in order to account for the fact that such gain/loss is only realized at maturity of the futures. In addition to the futures, we add an investment of $X_r(t)$ in the risk-free rate, so that the hedge portfolio has a value of zero at initiation.

After one time step Δt we obtain a hedging error of

$$\begin{aligned} &e^{-r(T'-t-\Delta t)}X_{F_{T'}}(t)[F_{T'}(t+\Delta t) - F_{T'}(t)] \\ &\quad + [e^{r\Delta t} - 1]X_r(t) - [O(t+\Delta t, T) - O(t, T)] \end{aligned} \tag{10}$$

for every option and every rebalancing day in our sample. Similarly, delta-gamma and delta-vega hedging performance can be assessed for every option price in our sample, by defining a hedging error that includes the position in a second option $X_{\overline{O}}(t)$:

$$\begin{aligned} &e^{-r(T'-t-\Delta t)}X_{F_{T'}}(t)[F_{T'}(t+\Delta t) - F_{T'}(t)] + [e^{r\Delta t} - 1]X_r(t) \\ &\quad + X_{\overline{O}}(t)[\overline{O}(t+\Delta t, T) - \overline{O}(t, T)] - [O(t+\Delta t, T) - O(t, T)], \end{aligned} \tag{11}$$

where $X_{\overline{O}}(t)$ is determined by the ratio of the gamma (or vega) of the hedging option and the option to be hedged. Such a hedge can be implemented with any option that exhibits a nonzero gamma (or vega).

5. EMPIRICAL RESULTS

In this section we first describe the data, then we present the calibration results for the Heston model, and finally we report the hedging results for the Heston model and the three smile adjustments of Section 2. Hedging results are for standard delta, delta-gamma, and delta-vega hedging under all models, and for MV delta and delta-gamma hedging under the Heston model.

5.1. Data

We use daily closing prices of FTSE 100 options obtained from Euronext-LIFFE (London International Financial Futures Exchange) from 2 January 2002 until 31 December 2008. These contracts are European style options on the FTSE 100 equity index (ESX). The options can be classified into two distinct categories depending on their expiration month. The first category are options in the quarterly cycle, where the expiration day is the third Friday in March, June, September, or December. In addition to the quarterly cycle, serial options are made available for trading such that the quarterly cycle is complemented by short-term options that expire in non-cycle months. Prior to 2004, the next four cycle-months and two short-term serial options are available for trading. For instance in the beginning of March, expiry months are March, April, May, June, September, and December.[14] From January 2004 on, also long-term options with time to maturity of up to two years were added to the data set and the number of serial options was raised to three.[15] As regards the strike range, option quotes between far out-of-the money and far in-the-money (ITM) are available where the quoted strikes typically range between around 70 and 120% of the futures price.

We apply standard filters to our data. First, we remove quotes that violate standard no-arbitrage conditions such as in Bakshi et al. (1997). Second, we remove data outside the moneyness range -0.5 and 0.3 as these options are infrequently traded. For the same reason we discard options with more than one year to maturity. Third, we refrain from using option quotes with less than five days to maturity as their prices are often heavily influenced by market-microstructure effects. We also plot all options in implied volatility-moneyness space to filter out obvious recording errors. After applying these filters we are left with more than 700,000 prices. We further divide the sample into an in-sample calibration period (January 2002 until June 2005) and an out-of sample period (July 2005 until December 2008). This leaves more than 356,000 observations for our out-of-sample hedging exercise.

To explore the sample graphically, and because we shall need this series when implementing our global calibration procedure, Figure 1 depicts the volatility index derived from these option prices from 2002 until the end of 2008.[16] It is apparent that the in-sample and out-of-sample periods are generally quite similar in nature, including both tranquil and more volatile market regimes. Volatile periods are observed between mid 2002 and mid 2003, and

[14]The expiry date is always the third Friday in the month.

[15]The long-term options extend the quarterly expiration months beyond one year and consequently the number of different maturities reaches up to 12.

[16]This index was constructed according to the VIX methodology. We are extremely grateful to Stamatis Leontsinis for constructing this series.

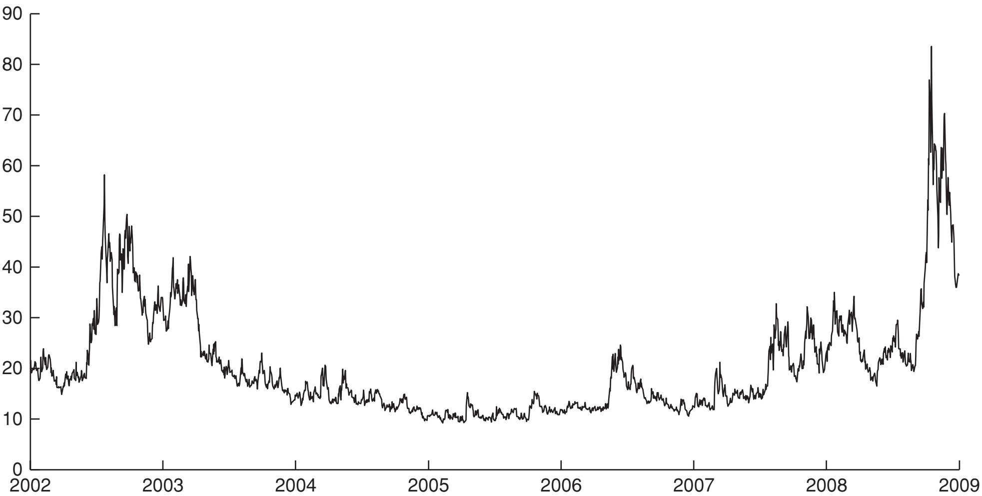

FIGURE 1
VFTSE Index. This figure plots the evolution of the VFTSE index for 30 days to maturity for the sample period. The index is constructed using the VIX methodology and is normalized to a yearly standard deviation.

especially during the last quarter of 2008 when, with the advent of the banking crisis, FTSE 100 volatility reached more than 80%. Before this, such high levels of implied volatilities had only been witnessed after the global stock market crash of 1987. This extreme event poses a challenge for any hedging model and provide a means to compare the performance of alternative hedges during both normal and crash market regimes.

Finally, we require a substitute for the risk-free interest rate. No standard proxy has emerged in the literature to date. While some authors prefer using maturity matched interest rates by interpolating the yield curve (for example, Broadie et al., 2007; Alexander et al., 2009), others employ a short-term interest rate as a substitute for the unobservable instantaneous short rate. For example, Bliss and Panigirtzoglou (2004) argue that interpolated interest rates are unlikely to represent realistic borrowing and lending rates for traders in this market, and thus use the three-month LIBOR rate (11 a.m. fixing as reported by Bloomberg) as an approximate risk-free rate. In addition, short-term rates are heavily affected by central bank interventions and are less liquid compared with the three-month rate. As most of the options in the sample are short-term options, interest rates have little effect on prices. Nevertheless, we experimented with both maturity-matched rates and a constant interest rate. Since short-term interest rates can become highly erratic under central bank interventions, we prefer to follow Bliss and Panigirtzoglou (2004) in using the three-month LIBOR rate as a substitute for the risk-free rate.

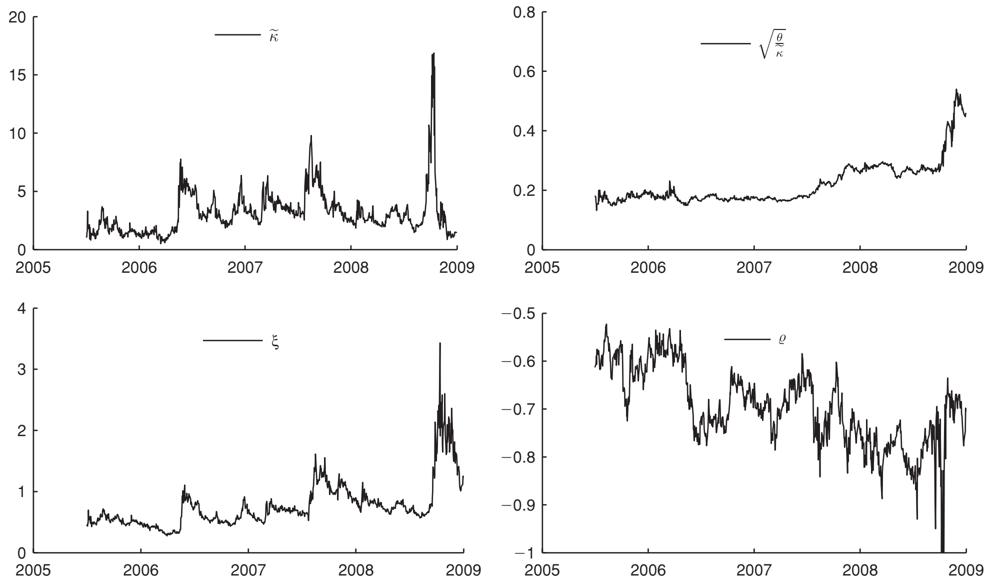

FIGURE 2
Local optimization. This figure plots the evolution of the structural parameters for daily calibrations of the Heston model to vanilla option prices on the FTSE 100 index, from June 2005 to December 2008. On each day in the sample, the parameters are obtained by minimizing the objective function (6).

5.2. Calibration Results

The Heston model is calibrated using local, global, and time-series methods. For the global calibration we employ our new procedure as described in Section 3. For the time-series calibration we use the in-sample calibration period for calibration from 2 January 2002 to 30 June 2005. After estimating the spot index dynamics we apply a variance risk premium of $\lambda = 2$.[17]

The parameter estimates obtained from local calibration of the Heston model to our FTSE 100 option data set are depicted in Figure 2. The calibrated parameters change considerably from day to day, especially during the banking crisis. The day after the Lehman Brothers default the correlation coefficient hit its boundary value of minus one, thus creating a spike with exceptionally large and negative skewness in the risk-neutral density. There also seems to be a positive correlation between the mean reversion and vol-of-vol. The estimation of the long-run volatility is the most stable. It remained around 20% for most of the first half of the sample, but during the banking crisis its value adjusted to around 50%.

[17]Results for other values of the variance risk premium and for other in-sample calibration periods have been calculated but are not reported for brevity, since they do not change the qualitative nature of our results. They are available from the authors on request.

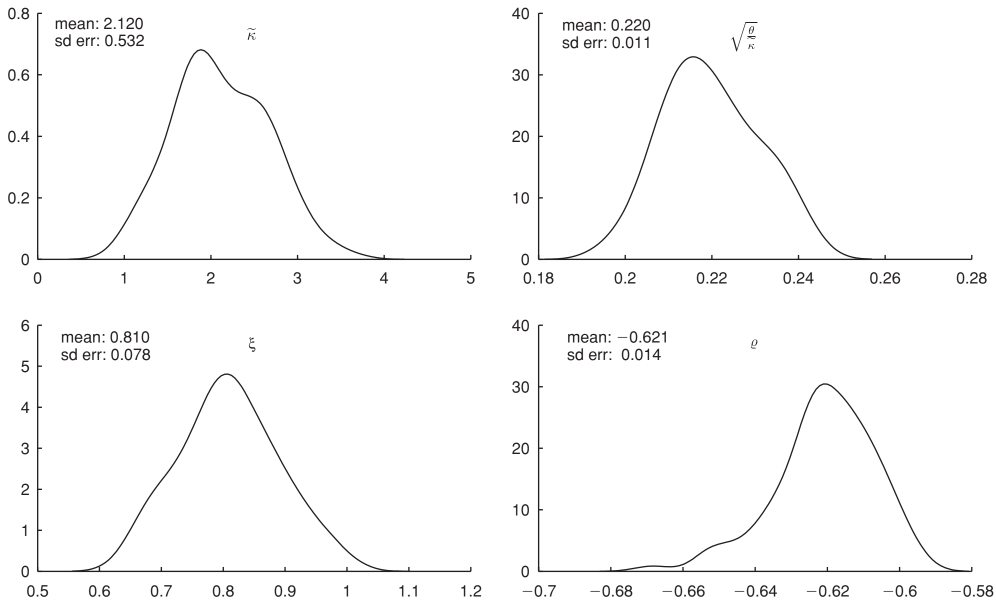

FIGURE 3

Global optimization. This figure plots Kernel densities of the parameter estimates obtained from 100 random samples of 50 Wednesdays in our calibration period. Parameters are obtained by optimizing Equation (9). Point estimates are obtained by using the mean over the 100 samples and standard errors are calculated from the standard deviation of the parameter estimates.

For our global calibrations of the Heston model, based on (9), we use VFTSE index shown in Figure 1 to derive the spot volatilities using (8). Then, we follow the procedure employed by Broadie et al., 2007 and choose representative option data from 50 random Wednesdays during the calibration sample to estimate the structural parameters in Φ^{rn}. This procedure is repeated for 100 samples, which has the advantage of providing bootstrapped standard errors for the parameters, thus revealing how accurate the parameter estimates can be. The results are depicted by the kernel densities shown in Figure 3. In our hedging study we use the mean value, across all 100 samples, as our point estimate for each parameter. A comparison of the results with Figure 2 reveals that these mean values for all structural parameters are similar to the averages of the locally calibrated values during the out-of-sample period. Of course, on a given day the locally calibrated parameters can differ substantially from these averages. However, this comparison is instructive because it confirms that the out-of-sample and the in-sample periods are similar in nature and potential differences between the hedging performance of alternative calibration techniques is unlikely to be due to fundamentally different parameter estimates in- and out-of sample.

TABLE I
Heston Model Calibration from Spot Data

Parameter	μ	κ	θ	ξ	ρ	*LT Vol*
From January 1, 2002 to June 30, 2005						
Posterior mean	0.0007	3.6079	0.1112	0.3919	−0.7098	17.8203
Posterior mode	0.0010	3.5269	0.1096	0.3914	−0.7166	17.5875
Posterior std. dev	0.0606	1.2251	0.0343	0.0421	0.0700	2.5575
5% percentile	−0.0988	1.7545	0.0578	0.3241	−0.8114	14.0799
95% percentile	0.1012	5.7775	0.1705	0.4620	−0.5846	22.3825

This table provides estimates of the structural parameters for the Heston model calibrated to FTSE 100 Spot data in the period from January 2002 until June 2005. We discarded the first 20,000 runs of the chain as a burnin and summarize the posterior with 20,000 draws.

Finally, for the time series calibration we use the daily returns on the FTSE spot index. For consistency, we have used the same sample as for the global calibration.[18] Applying the MCMC procedure, we use the posterior mean as our point estimate. This is in line with a quadratic loss function and is standard in the related literature. A possible alternative is to use the mode of the distribution but, as apparent from the results reported in Table I, they are extremely close for all parameters. The results imply an estimate for the mean reversion parameter κ of 3.61, a long-term volatility value of 17.82%, a vol-of-vol parameter ξ of 0.39, and an index-variance correlation parameter of -71%. Of course, these are slightly different from the results obtained by Eraker et al. (2003) for the S&P index, due to the different behavior of the FTSE and the different sample period (Eraker et al., 2003 include the market crash of 1987).

Our time-series calibration results exhibit similar stylized facts to those previously found when comparing estimates of S&P 500 dynamics that can be obtained from spot and option calibrations. That is, the vol-of-vol parameter differs significantly from its value obtained from a pure calibration to option prices. This observation was first made by Bakshi et al. (1997) and is often attributed to the misspecification apparent in the Heston model. Our option price implied parameter is 0.81 for the global calibration and has a similar average value over the local calibration sample. This is more than twice times the estimate obtained from spot data. Thus, using different calibration techniques might lead to very different results for the out-of-sample hedging exercise.

[18]We also calibrated parameters using an extended sample covering 15 years (i.e. starting in January 1990) to satisfy ourselves that the calibration sample was large enough to produce robust parameter estimates. For example, if volatility were relatively slow to mean revert, we might find that the parameters related to mean reversion of variance were quite different when calibrated to a much longer time series. However, the longer calibration sample had little effect on parameter values, and even less effect on our hedging results. Hence, these results are omitted for brevity, but are available from the authors on request.

5.3 Hedging Results

On each trading day t between July 1, 2005 and December 30, 2008 we suppose that one option of every available strike and maturity has been written. Then we gamma hedge all these options using a single put option of relatively short maturity (typically in the region of 45 days) that is closest to at-the-money (ATM). Similarly, for the vega hedge we use a single put option of relatively long maturity (around 120 days), also closest to ATM. This is because near ATM options have relatively large gammas and vegas, and puts are usually more liquid than calls.

Consistent with much of the empirical hedging literature we employ two different rebalancing intervals, one day and one week.[19] After this time the hedge portfolio is closed and the hedging error (10) or (11) is recorded. Then, another portfolio of options is written and the procedure is repeated until all the data are exhausted. Under daily (weekly) rebalancing a total of 361,913 (73,078) options are written. We assume transactions costs have a similar effect on all hedges as our data includes only closing prices, no bid–ask spreads.

Our reported performance measure is the standard deviation of the hedge error, as this risk metric is consistent with our calibration objectives and the aim of hedging is to minimize risk.[20] To examine whether hedge performance is affected by market regime, we compute this standard deviation over two sub-samples: from July 1, 2005 to June 30, 2008, and from July 1, 2008 to December 31, 2008. Finally, to assess whether the performance depends on the moneyness of the option, we report results for all options and also disaggregated results for sub-groups of options with moneyness buckets [−0.5, −0.3), [−0.3, −0.1), [−0.1,0.1) and [0.1,0.3], where moneyness is defined as before.

5.3.1 Daily rebalancing

Empirical hedging results for the daily rebalancing frequency are provided in Table II. We investigate delta hedging first. There are three main results considering all options in the sample. First, delta hedging results in a large reduction of risk with standard deviations of the hedging error being at least three to four times smaller than the standard deviations of the unhedged positions. Second, for all calibration techniques, using a MV hedge ratio rather than a standard partial derivative for the delta reduces the risk considerably. This finding

[19]We follow standard practice in the empirical option literature and use nonoverlapping Wednesday data for the weekly exercise.

[20]We have also computed average hedging errors and the mean absolute hedging error but have not included these to reduce the size of the results tables. Again, they are available from the authors on request.

TABLE II
Hedging Results—One Day Rebalancing

	TS		*Local*		*Global*					
	Partial	*MV*	*Partial*	*MV*	*Partial*	*MV*	*ST*	*SM*	*SS*	*UH*
Delta Hedge										
All options	10.27	7.90	10.67	7.09	10.12	7.71	6.72	10.68	7.36	41.76
Std. err	(0.05)	(0.05)	(0.05)	(0.04)	(0.05)	(0.04)	(0.04)	(0.05)	(0.04)	
1%-perc.	10.15	7.80	10.55	7.01	10.01	7.62	6.64	10.55	7.26	
99%-perc.	10.38	8.00	10.81	7.20	10.25	7.81	6.81	10.81	7.47	
−0.5 to −0.3	5.39	5.43	6.42	5.06	5.88	5.37	4.60	6.15	4.99	41.72
−0.3 to −0.1	7.84	7.34	10.16	6.93	9.44	7.13	6.61	10.11	7.22	39.20
−0.1 to 0.1	11.13	8.32	12.98	8.11	12.19	8.70	7.83	12.93	8.55	37.69
0.1 to 0.3	13.71	9.40	10.73	7.29	10.63	8.53	6.78	11.00	7.49	48.53
Delta-Gamma Hedge										
All options	8.44	3.88	7.92	2.53	7.37	3.00	5.62	9.43	4.70	41.76
Std. err	(0.13)	(0.02)	(0.06)	(0.02)	(0.08)	(0.01)	(0.09)	(0.15)	(0.05)	
1%-perc.	8.13	3.84	7.78	2.49	7.2	2.97	5.43	9.08	4.59	
99%-perc.	8.73	3.91	8.07	2.57	7.56	3.04	5.83	9.79	4.81	
−0.5 to −0.3	7.01	3.15	4.62	2.76	4.97	2.93	5.89	5.44	3.92	41.72
−0.3 to −0.1	8.73	3.10	5.29	2.56	6.11	2.68	6.06	6.35	4.58	39.20
−0.1 to 0.1	8.41	2.84	7.42	2.05	7.83	2.34	4.99	8.09	4.81	37.69
0.1 to 0.3	9.16	5.76	11.84	2.81	9.29	3.94	5.66	14.63	5.22	48.53
Delta-Vega Hedge										
All options	3.69		2.44		2.69		2.86	2.37	2.31	41.76
Std. err	(0.01)		(0.02)		(0.02)		(0.02)	(0.02)	(0.02)	
1%-perc.	3.66		2.40		2.65		2.80	2.33	2.26	
99%-perc.	3.73		2.48		2.73		2.91	2.40	2.35	
−0.5 to −0.3	2.97		2.52		2.65		3.23	2.60	2.46	41.72
−0.3 to −0.1	3.51		2.64		2.65		2.89	2.51	2.55	39.20
−0.1 to 0.1	2.69		2.11		2.33		2.15	2.04	2.02	37.69
0.1 to 0.3	5.16		2.53		3.13		3.23	2.39	2.25	48.53

This table reports the hedge error standard deviation for daily rebalancing, followed by its bootstrapped standard error (in parentheses) and 1 and 99% percentiles. Moneyness is defined in terms of $m = \log(K/F)/\sqrt{\tau}$ and given in the first column. The models are abbreviated as TS (time-series calibration with $\lambda = 2$), Local (parameters from local calibration), Global (parameters from global calibration), ST (sticky tree), SM (sticky moneyness), SS (sticky strike), and UH (unhegded position).

confirms the results of Alexander and Nogueira (2007) and Poulsen, Schenk-Hoppé, and Ewald (2009) and is due to the residual risk arising from uncertainty in volatility. This result is also theoretically expected because our performance measure is designed to favor hedges that have small variances. We therefore concentrate on these hedge ratios for the Heston model in the remainder. Third, our results imply that the hedging performance also systematically depends on the data used to calibrate the model. Using only spot data results in the highest standard deviation of 7.90.[21] Using information from

[21]This finding is also independent of the sample period used for the parameter estimation and of the size of the risk premium. We omitted detailed results for expositional clarity but they are available from the authors upon request.

option quotes, however, leads to a reduced standard deviation of the hedge error of 7.71 for the MV hedge (using the out-of-sample parameters). Yet the best performing strategy is the one that utilizes parameters from the daily cross section of option prices, with a standard deviation of 7.09. Therefore, in our data set daily recalibration improves the hedging performance. The differences reported are also statistically significant, as bootstrapped confidence intervals for the hedge performance statistics derived from each hedging model do not overlap. The difference between hedge error standard deviations using (a) the spot and (b) the global calibration parameters are only significant at a 5% level; but all other differences are very highly significant.

This last result is at odds with evidence from local volatility hedging, reported by Dumas, Fleming, and Whaley (1998). These authors find that overfitting local volatility can lead to an increase in hedging errors. In our sample, the closer the parameter values are to those of the local calibration (and thus the better the fit), the better the hedging performance.

Comparing the performance of the Heston model with the smile-adjusted hedge ratios reveals that simple adjustments of the BSM model can outperform the more sophisticated Heston model. Our local volatility implementation performs best among all considered strategies with a standard deviation of 6.72. The practitioner BSM model (using the implied volatility surface to calculate hedge ratios) provides slightly higher risk with 7.36. Not surprisingly, since Alexander and Nogueira (2007) show that the two hedge ratios are theoretically the same, the sticky moneyness assumption and the standard delta hedge performance for the local calibration yield very similar results, and in agreement with several other studies, these models provide the highest hedge errors on average (see for instance Alexander, Rubinov, Kalepky, & Leontsinis, 2011). We remark that the Heston model can only outperform the practitioner BSM model (in terms of standard deviation) when we consider a local calibration procedure. This result highlights the fact that a calibration technique may have important implications on hedging performance. It also casts doubt on whether we should attribute the good hedging performance of the Heston model that is sometimes reported in the literature to the dynamics of the model. Indeed, our result indicate that its good hedging performance is merely a result of a good fit to option prices.

Many of our conclusions are robust across moneyness categories, i.e. the rankings of the hedging performance are similar in all moneyness buckets: the ST (local volatility) hedge always ranks first and the local calibration of the stochastic volatility model performs better than the SS hedge for all but the first moneyness category (-0.5 to -0.3 moneyness). An exception is the global calibration compared with the spot estimation of the stochastic volatility hedge. It is evident that the reason why the global calibration performs better on average is that it is better for deep ITM calls, i.e. options in the moneyness bucket (0.1–0.3).[22] In the other moneyness categories it appears that using option prices to fit parameters leads to no better hedging performance than we obtain when calibrating parameters to a time series on the spot prices of the underlying. Another result that is not robust across moneyness is that the improved performance gained from using a MV hedge, based on the spot estimates for parameters, is only visible for options with positive moneyness.

Results for the delta-gamma strategy (middle section of Table II) show a slightly different picture. Again, within the Heston model implementations, the local calibration provides the best risk reduction, the smile adjustments are now consistently outperformed by the Heston model. So although simple smile-adjustments to the BSM delta hedge can work very well (and the local volatility delta hedge is exceptionally good), they are less effective for gamma hedging. Even the ST smile-adjusted gamma hedge, which corresponds to the local volatility hedge, cannot outperform the stochastic volatility MV delta-gamma hedge. While the Heston model gamma hedge reduces risk considerably (with a standard deviation about half the size of the delta-only hedge error standard deviation), the improvement offered by a gamma hedge for the smile adjusted hedging models is at best moderate. Moreover, with the addition of a gamma hedge, the SS model consistently outperforms the ST model. These results are also fairly robust across moneyness categories.

Finally, we turn to the results for delta-vega hedging strategies. Again, Heston model implementations yield results that are consistent with our earlier findings: the in-sample fit of the pricing model improves the out-of-sample hedging error statistics with a hedge standard deviation of 2.44 for the local calibration model. As for the delta-gamma hedge, result for smile adjustments are novel, and very interesting: the ST (local volatility) implementation, which performed best in a delta hedge, turns out to be the worst once a vega hedge is added. SM and SS adjustments perform very well, even improving on the local calibration implementation slightly (with a standard deviation of 2.37 and 2.31, respectively). This finding is surprising as one would expect the Heston model to manage volatility risk more accurately than a theoretically inconsistent model where the underlying price could follow one of the multiple GBMs with different, constant volatilities, depending on the option's strike or moneyness. Interestingly, simulation results in Branger, Krauthem, Schlag, and Seeger (2008) indicate that if the true data-generating process is the Heston model augmented with jumps, SS hedges perform very well for short maturities but deteriorate with the maturity of the options. Most option data sets are

[22]Option prices in this category are higher and thus this influences the results more than other categories. It is thus important to consider the results of all buckets separately. We have also investigated the hedging performance of put and calls separately and found no qualitative difference in the results.

biased toward short-term options (such as ours) and hence our empirical results appear to be consistent with this finding.

Overall, the delta-vega hedging strategy also performs better than a delta-gamma hedge, hence it appears to be more important to hedge volatility risk before considering second-order price risk for a daily rebalancing strategy. As far as the different moneyness buckets are concerned, our results are again robust. The only noteworthy exception is, as before, the last moneyness category (0.1–0.3), where the improvements from using the options quotes for the calibration is largest, whereas in the other categories the gain in hedging accuracy from using options for calibration can be fairly small. It is likely that the different behavior of this bucket is due to the fact that it is the least liquid category.

5.3.2. *Weekly rebalancing*

Results for the weekly rebalancing strategies are given in Table III. Overall, the qualitative results from the daily hedging exercise carry over to the weekly rebalancing frequency. In particular, in the delta hedging exercise the ST hedging errors are the least variable. The second best performance is from the locally calibrated Heston model, followed by the global and the time series calibrations. The SS hedge, and even more so the SM hedge, perform worse than any Heston model implementation. Bootstrapping results again confirm that the difference between the models is highly significant.

For the delta-gamma strategy, all smile adjustments underperform the Heston model, just as they did with daily rebalancing. Again, the locally calibrated Heston model has the best performance, with a hedging error standard deviation that is about one third of its value based on a pure delta hedge. But under weekly rebalancing, delta-gamma hedging now reduces risk more efficiently and has a similar performance to delta-vega hedging, implying that the importance of a convexity adjustment increases with longer hedging horizons. Under delta-vega hedging both the practitioner BSM and SM models again outperform the Heston model, for all moneyeness buckets. So there is little change in the qualitative nature of our results, compared with those obtained with daily rebalancing.[23]

5.3.3. *Robustness checks*

Table IV provides hedge error standard deviations over the two sub-samples: January 2005–June 2008, representing a relatively tranquil period, and July 2008–December 2008, an extremely volatility period at the height of the

[23]The hedging error is roughly proportional to the square root of the rebalancing frequency, so with weekly rebalancing hedging error increases by more than a factor of two. This merely reflects the increased uncertainty and is consistent with the reported increase from daily to weekly hedging errors reported in Bakshi et al. (1997).

TABLE III
Hedging Results—One Week Rebalancing

	TS		*Local*		*Global*					
	Partial	*MV*	*Partial*	*MV*	*Partial*	*MV*	*ST*	*SM*	*SS*	*UH*
Delta Hedge										
All options	22.95	16.97	23.62	14.73	22.74	15.88	13.95	23.51	17.02	84.29
Std. err	(0.15)	(0.14)	(0.18)	(0.12)	(0.17)	(0.12)	(0.12)	(0.17)	(0.14)	
1%-perc.	22.61	16.66	23.23	14.45	22.38	15.57	13.69	23.13	16.70	
99%-perc.	23.28	17.3	24.04	15.01	23.15	16.17	14.21	23.93	17.33	
−0.5 to −0.3	12.53	12.17	14.27	11.40	13.48	11.98	10.70	13.69	11.75	84.35
−0.3 to −0.1	18.75	16.58	23.09	15.73	22.22	16.09	14.88	23.01	17.36	78.42
−0.1 to 0.1	26.17	18.59	29.68	17.38	28.56	18.36	16.31	29.43	20.45	76.34
0.1 to 0.3	28.34	18.44	21.52	12.22	20.89	15.03	11.92	21.84	15.40	98.22
Delta-Gamma Hedge										
All options	15.23	7.69	16.08	5.00	14.30	5.77	9.98	16.90	9.42	84.29
Std. err	(0.22)	(0.05)	(0.16)	(0.07)	(0.15)	(0.05)	(0.24)	(0.29)	(0.11)	
1%-perc.	14.74	7.58	15.72	4.85	13.95	5.66	9.41	16.23	9.18	
99%-perc.	15.75	7.82	16.46	5.15	14.62	5.89	10.51	17.61	9.68	
−0.5 to −0.3	12.87	6.57	9.51	6.01	9.11	6.18	9.80	10.08	8.35	84.35
−0.3 to −0.1	15.32	6.68	10.48	5.61	10.89	5.74	10.85	11.26	9.06	78.42
−0.1 to 0.1	15.39	5.51	14.94	3.60	15.63	4.23	9.31	14.97	9.56	76.34
0.1 to 0.3	16.58	11.07	24.24	4.91	18.40	6.99	9.96	25.90	10.35	98.22
Delta-Vega Hedge										
All options	7.41		5.17		5.43		5.57	4.70	4.69	84.29
Std. err	(0.06)		(0.07)		(0.06)		(0.07)	(0.05)	(0.05)	
1%-perc.	7.29		5.01		5.30		5.40	4.59	4.57	
99%-perc.	7.57		5.34		5.58		5.74	4.82	4.83	
−0.5 to −0.3	6.23		5.69		5.75		6.39	5.56	5.52	84.35
−0.3 to −0.1	7.51		6.04		5.85		6.28	5.27	5.40	78.42
−0.1 to 0.1	5.26		3.97		4.42		3.99	3.63	3.59	76.34
0.1 to 0.3	9.98		5.08		5.82		5.74	4.50	4.36	98.22

This table reports the hedge error standard deviation for weekly rebalancing, followed by its bootstrapped standard error (in parentheses) and 1% and 99% percentiles. Moneyness is defined in terms of $m = \log(K/F)/\sqrt{\tau}$ and given in the first column. The models are abbreviated as TS (time-series calibration with $\lambda = 2$), Local (parameters from local calibration), Global (parameters from global calibration), ST (sticky tree), SM (sticky moneyness), SS (sticky strike), and UH (unhegded position).

credit and banking crisis. As expected, delta hedging results in a large reduction of risk, especially during the less volatile period (Panel A). Naturally, it was much more difficult to eliminate risk during the banking crisis and standard deviations of the hedging error were much greater in all categories. However, once an additional option is used in the hedge, the differences between the two periods reduce markedly. Clearly, gamma or vega hedging is particularly important during crisis periods, when such hedges improve upon the delta hedge by a factor of three or four, whereas in tranquil market periods the reduction in hedging error standard deviations is merely of the order of two or less.

TABLE IV
Robustness

	D	*MVD*	*DG*	*MVDG*	*DV*	*UH*	*D*	*MVD*	*DG*	*MVDG*	*DV*	*UH*
	Panel A: 1 January 2005 to 30 June 2008						*Panel B: 1 July 2008 to 31 December 2008*					
TS	7.42	5.04	5.03	3.54	3.37	35.75	21.54	18.05	20.06	5.73	5.46	71.48
Local	7.19	4.64	5.40	2.18	2.03	35.75	23.61	16.00	17.46	4.28	4.37	71.48
Global	7.01	5.15	4.81	2.66	2.30	35.75	21.93	17.20	16.71	4.79	4.62	71.48
ST	4.53		2.68		2.19	35.75	14.90		14.42		5.68	71.48
SM	7.15		4.76		2.07	35.75	23.73		23.82		3.91	71.48
SS	4.68		3.00		1.97	35.75	16.84		10.81		3.98	71.48

This table reports the hedge error standard deviation over two subsamples. The models are abbreviated as TS (time-series calibration with $\lambda = 2$), Local (parameters from local calibration), Global (parameters from global calibration), ST (sticky tree), SM (sticky moneyness), SS (sticky strike), and UH (unhegded position).

Finally, to investigate whether the Heston model is better for hedging long-term or short-term options we also divided the option data set along the maturity dimension. The rationale is that jumps in price or volatility may affect the short-term smile more than the long-term smile, in which case a pure stochastic volatility model could perform better for long-term options than it does for short-term contracts.[24] Yet, when hedging S&P 500 options, Bakshi et al. (1997) report that including jumps yields no discernible improvement. Our findings support this result, since the results in Table IV are robust to the option's maturity. The detailed results disaggregated by option maturity are not included, for brevity, but are available from the authors upon request.

6. CONCLUSION

This study has considered three main issues that have not been fully addressed in the previous option-hedging literature:

1. *What effect does the calibration procedure have on the hedging performance of a stochastic volatility model*? Calibration can affect the hedging results considerably. Dumas, Fleming, and Whaley (1998) find that over-fitting of a model can increase the variability of hedging errors, and frequent recalibration to option prices is not even consistent with most stochastic volatility models; yet we find that daily recalibration of the Heston model to option prices clearly improves its hedging performance. This is true for all

[24]The stochastic volatility component—because of its diffusive nature—cannot generate enough negative skewness and excess kurtosis in the short run. Price jumps on the other hand affect skewness and kurtosis even in the short run (see Bakshi et al., 1997).

hedging strategies considered (delta, delta-gamma, and delta-vega) and for both rebalancing frequencies (one day and one week). The results are also robust with respect to different sub-samples.

2. *Can MV hedging improve upon simple adjustments to the BSM model?* The standard Heston model hedge ratios perform worse than any other model (except the delta-only hedge sticky-delta model and this is theoretically equivalent to the standard Heston delta hedge). MV hedge ratios improve the performance very considerably, but they are still not able to outperform the simple smile-adjustments of the BSM hedges that are so popular with traders, at least when delta and delta-vega hedging. Only when delta-gamma hedging does it seem to be worthwhile using the Heston model.
3. *Is delta-vega hedging always better than delta-gamma hedging (or vice versa) or does it depend on the options being hedged, or the hedging model?* For a daily rebalancing strategy it appears to be more important to hedge volatility risk before considering second-order price risk, since delta-vega definitely leads to a greater risk reduction than delta-gamma hedging. However, with weekly rebalancing the vega and gamma hedges have a similar performance. We remark that the excellent performance of the simple sticky strike and sticky moneyness delta-vega hedges is a little surprising, because they assume the underlying follows a GBM in which the constant volatility depends on the option being hedged; from theoretical viewpoint one would expect the Heston model to capture volatility risk more accurately.

It should be emphasized that our conclusions have been reached entirely on the basis of hedging FTSE 100 index options. Having said this, they are obtained using a very long sample period. Our findings are robust to the market regime, whether it be unusually volatile or tranquil; they are also reasonably robust across moneyness categories, the only noteworthy exception being the very high strike options that are the least liquid and the most difficult to hedge. It is for these options that we find the greatest improvements from using options prices for the Heston model calibration; in the other categories, the differences can sometimes be rather small.

APPENDIX A: MCMC ALGORITHM

In Bayesian statistics, one aims at recovering the posterior distribution of the unknown variables given the observed data $p(\Phi, V|Y)$. Learning about this distribution is achieved by updating prior beliefs with the information obtained from observing the data. Thus, formally we have

$$p(\Phi, V|Y) \propto p(Y|\Phi, V)p(\Phi, V), \tag{A1}$$

where $p(Y|\Phi, V)$ denotes the *likelihood* and $p(\Phi, V)$ the *prior density*. However, even though Equation (A1) gives a simple formula for the distribution of unobservable parameters, obtaining the density in closed form is impossible for most practical applications. Instead, a sampling algorithm has to be applied to produce random draws from this distribution. Then by sampling many times, we recover sufficient information about the shape of the density function and related quantities such as their marginals.

In order to overcome the problem of sampling from complex multivariate distributions, one can iteratively draw from the so-called complete conditional densities. For the problem at hand, these densities are $p(\phi_i|\Phi_{-i}, V, Y)$ and $p(V_t|\Phi, V_{-t}, Y)$, where the notation Φ_{-i} indicates the parameter vector without the ith element (same notation applies to the variance V) and $\phi_i \in \Phi$. Once these distributions are derived, they can be used in a recursive procedure, hence replacing the simulation of the multivariate posterior distribution by simulation of several lower dimensional (here univariate) distributions. This procedure is known as the Gibbs sampler. Loosely speaking, it produces a Markov chain with invariant distribution equal to the posterior. Thus, the chain can be used to obtain samples from the posterior distribution once convergence of the chain is achieved.

For the MCMC implementation of the Heston model, we need to specify prior distributions of the parameters μ, κ, θ, ξ, and ρ. Since we follow common practice and implement our algorithm on daily percentage returns (thus we use returns scaled by a factor 100), we need to reflect this scaling in the choice of our prior distributions. Our goal is to induce little knowledge through the priors and put the emphasis on the information contained in the observed return time series. Furthermore, to facilitate sampling from the complete conditional distributions we opt for prior distributions of the conjugate class. In particular, we choose $\mu \sim \mathcal{N}(0, 0.1)$, $\kappa \sim \mathcal{TN}^+(0, 1)$, and $\theta \sim \mathcal{TN}^+(0, 1)$, where $\mathcal{N}$ denotes the normal distribution with mean m and standard deviation s and $\mathcal{TN}^+$ denotes the truncated normal distribution (with truncation to the positive real axis and the two parameters now reflect the mean and standard deviation of the underlying normal density).[25] This choice entails very little information considering typical parameter values obtained in the literature. We transform the parameters (ξ, ρ) to (ψ_1, ψ_2) with $\psi_1 = \xi\rho$ and $\psi_2 = \xi^2(1 - \rho^2)$. This is motivated by the observation that the transformed parameters can be easily sampled as they are linear regression parameters for the variance time series (with slope ψ_1 and heteroscedastic error term ψ_2) for which sampling is standard. We use the uninformative priors $\psi_1 \sim \mathcal{N}(0, 1)$ and $\psi_2 \sim \mathcal{IG}(0.05, 1)$ where $\mathcal{IG}(m, s)$

[25]This truncation was also applied in Li, Wells, and Yu (2008) who provide exact formulae for the complete conditional distributions for the Heston model.

denotes an inverse gamma distribution with mean m and standard deviation s. Having sampled from ψ_1 and ψ_2 we can transform back to $\xi = \sqrt{\psi_1^2 + \psi_2}$ and $\rho = \psi_1/\xi$. For more details on this transformation we refer to Jacquier, Polson, and Rossi (2004).

Sampling the variance vector is less straightforward for two reasons. First, block updating for the parameter vector is not available and one has to cycle through the variance vector one by one. And second, posterior distributions for the individual variances are nonstandard. By the Markov property and the Bayes formula we obtain (with appropriate adjustments for the first and last variance)

$$p(V_t|V_{-t}, Y, \Phi) \propto p(Y_{t+1}|V_t, \Phi)p(V_{t+1}|V_t, Y_{t+1}, \Phi)p(V_t|V_{t-1}, Y_t, \Phi),$$

which is rather involved as the variance parameter enters in several place on the right-hand side. We have tested two updating algorithms, the random walk Metropolis and the adaptive rejection Metropolis sampling (ARMS) algorithm of Gilks, Best, and Tan (1995). We found that both produce accurate results for this model, indeed estimated parameters for the FTSE index were virtually the same for both algorithms. We ended up reporting results from the ARMS algorithm.

BIBLIOGRAPHY

Alexander, C., & Nogueira, L. M. (2007). Model-free hedge ratios and scale-invariant models. Journal of Banking and Finance, 31(6), 1839–1861.

Alexander, C., Kaeck, A., & Nogueira, L. M. (2009). Model risk adjusted hedge ratios. Journal of Futures Markets, 29(11), 1021–1049.

Alexander, C., Rubinov, A., Kalepky, M., & Leontsinis, S. (2011). Regime-dependent smile-adjusted delta hedging. Journal of Futures Markets, DOI: 10.1002/fut.20517.

Bakshi, G., Cao, C., & Chen, Z. (1997). Empirical performance of alternative option pricing models. Journal of Finance, 52(5), 2003–2049.

Bakshi, G., Cao, C., & Chen, Z. (2000). Pricing and hedging long-term options. Journal of Econometrics, 94(1–2), 277–318.

Bates, D. S. (2000). Post-'87 crash fears in the S&P 500 futures option market. Journal of Econometrics, 94(1–2), 181–238.

Black, F. (1976). The pricing of commodity contracts. Journal of Financial Economics, 3(1–2), 167–179.

Black, F., & Scholes, M. S. (1973). The pricing of options and corporate liabilities. Journal of Political Economy, 81(3), 637–654.

Bliss, R. R., & Panigirtzoglou, N. (2004). Option-implied risk aversion estimates. Journal of Finance, 59(1), 407–446.

Branger, N., Krautheim, E., Schlag, C., & Seeger, N. (2011). Hedging under model mis-specification: All risk factors are equal, but some are more equal than others Journal of Futures Markets, DOI: 10.1002/fut.20530.

Britten-Jones, M., & Neuberger, A. (2000). Option prices, implied price processes, and stochastic volatility. Journal of Finance, 55(2), 839–866.

Broadie, M., Chernov, M., & Johannes, M. (2007). Model specification and risk premia: Evidence from futures options. Journal of Finance, 62(3), 1453–1490.

Carr, P., & Madan, D. (1999). Option valuation using the fast Fourier transform. Journal of Computational Finance, 2(4), 61–73.

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(3), 407–458.

Christoffersen, P., & Jacobs, K. (2004). The importance of the loss function in option valuation. Journal of Financial Economics, 72(2), 291–318.

Christoffersen, P., Jacobs, K., & Mimouni, K. (2010). Volatility dynamics for the S&P 500: Evidence from realized volatility, daily returns, and option prices. Review of Financial Studies, 23(8), 3141–3189.

Coleman, T., Kim, Y., Li, Y., & Verma, A. (2001). Dynamic hedging with a deterministic local volatility function model. Journal of Risk, 4(1), 63–89.

Derman, E. (1999). Regimes in volatility. Risk, 12(4), 55–59.

Derman, E., & Kani, I. (1994). Riding on a smile. Risk, 7(2), 32–29.

Derman, E., Kani, I., & Zou, J. Z. (1996). The local volatility surface: Unlocking the information in index option prices. Financial Analysts Journal, 52(4), 25–36.

Dotsis, G., Psychoyios, D., & Skiadopoulos, G. (2007). An empirical comparison of continuous-time models of implied volatility indices. Journal of Banking and Finance, 31(12), 3584–3603.

Duffie, D., Pan, J., & Singleton, K. (2000). Transform analysis and asset pricing for affine jump-diffusions. Econometrica, 68(6), 1343–1376.

Dumas, B., Fleming, J., & Whaley, R. E. (1998). Implied volatility functions: Empirical tests. Journal of Finance, 53(6), 2059–2106.

Dupire, B. (1994). Pricing with a smile. Risk, 7(1), 18–20.

Engelmann, B., Fengler, M. R., & Schwendner, P. (2006). Hedging under alternative stickiness assumptions: An empirical analysis for barrier options. Journal of Risk, 12(1).

Eraker, B. (2004). Do stock prices and volatility jump? Reconciling evidence from spot and option prices. Journal of Finance, 59(3), 1367–1404.

Eraker, B., Johannes, M., & Polson, N. (2003). The impact of jumps in volatility and returns. Journal of Finance, 53(3), 1269–1300.

Gilks, W. R., Best, N. G., & Tan, K. K. (1995). Adaptive rejection metropolis sampling within Gibbs sampling. Applied Statistics, 44(4), 455–472.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. Review of Financial Studies, 6(2), 327–343.

Ignatieva, K., Rodrigues, P., & Seeger, N. (2009). Stochastic volatility and jumps: Exponentially affine yes or no? an empirical analysis of S&P500 dynamics. Goethe University Frankfurt, Working Paper.

Jacquier, E., Polson, N. G., & Rossi, P. E. (1994). Bayesian analysis of stochastic volatility models. Journal of Business & Economic Statistics, 12(4), 371–389.

Jacquier, E., Polson, N. G., & Rossi, P. E. (2004). Bayesian analysis of stochastic volatility models with fat-tails and correlated errors. Journal of Econometrics, 122(1), 185–212.

Jiang, G. J., & Tian, Y. S. (2005). The model-free implied volatility and its information content. Review of Financial Studies, 18(4), 1305–1342.

Johannes, M., & Polson, N. (2006). MCMC methods for financial econometrics. In Yacine Aït-Sahalia and Lars Hansen (Eds.), Handbook of Financial Econometrics. New York: Elsevier.

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

Kim, I. J., & Kim, S. (2004). Empirical comparison of alternative stochastic volatility option pricing models: Evidence from Korean KOSPI 200 index options market. Pacific-Basin Finance Journal, 12(2), 117–142.

Kim, I. J., & Kim, S. (2005). Is it important to consider the jump component for pricing and hedging short-term options? Journal of Futures Markets, 25(10), 989–1009.

Kim, S. (2009). The performance of traders' rules in options market. Journal of Futures Markets, 29(11), 999–1020.

Lewis, A. L. (2000). Option valuation under stochastic volatility. Finance Press.

Li, H., Wells, M. T., & Yu, C. L. (2008). A Bayesian analysis of return dynamics with Levy jumps. Review of Financial Studies, 21(5), 2345–2378.

Merton, R. C. (1973). Theory of rational option pricing. Bell Journal of Economics, 4(1), 141–183.

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

Poulsen, R., Schenk-Hoppé, K. R., & Ewald, C.-O. (2009). Risk minimization in stochastic volatility models: Model risk and empirical performance. Quantitative Finance, 9, 693–704.

Psychoyios, D., & Skiadopoulos, G. (2006). Volatility options: Hedging effectiveness, pricing, and model error. Journal of Futures Markets, 26(1), 1–31.

Rubinstein, M. (1994). Implied binomial trees. Journal of Finance, 49(3), 771–818.

Vähämaa, S. (2004). Delta hedging with the smile. Financial Markets and Portfolio Management, 18(3), 241–255.